

# Second OASIS/I3 Retreat

(in conjunction with the ROC Retreat)


10-12 January 2005

Randy H. Katz, Ion Stoica, Anthony Joseph  
Computer Science Division


Electrical Engineering and Computer Science Department  
University of California, Berkeley

Berkeley, CA 94720-1776


# Retreat Goals & Technology Transfer



People  
Project Status  
Work in Progress  
Prototype Technology




UC Berkeley Project Team



Industrial Collaborators  
Friends

Early Access to Technology  
Promising Directions  
Reality Check  
Feedback



# Who is Here (Industry & University Friends)

- **Align Technology**
  - Chris Overton
- **Boeing**
  - Tom Henderson
- **Cisco**
  - David Jaffe
  - Dan Lenoski
- **Cypress Semiconductor**
  - David Chu
- **DOCOMO Labs**
  - Ulas Kozat
- **Hewlett-Packard Labs**
  - Nina Bhatti
  - Wai-Tian Dan Tan
- **HIIT**
  - Andrei Gurtov
- **IBM**
  - Kirsten Hildrum
  - Sandeep Uttamchandani
- **Intel**
  - Kevin Fall
  - Timothy Roscoe
- **Microsoft**
  - Sharad Agarwal
  - Venkat Padmanabhan
  - Doug Terry
  - Helen Wang
- **Nortel Networks**
  - Tal Lavian
- **NTT MCL**
  - Jianping Pan
- **Riverbed Systems**
  - Steve McCanne
- **Sun Microsystems**
  - Christoph Schuba
- **Telcordia Technologies**
  - Bryan Lyles
- **UC Davis**
  - S. J. Ben Yoo
- **University of Waterloo**
  - S. Keshav

# Who is Here (Berkeley)

- Professors
  - Anthony Joseph
  - Randy Katz
  - Ion Stoica
  - (Dave Patterson, ROC)
  - (Armando Fox, Stanford, ROC)
- Technical & Admin Staff
  - Bob Miller
  - Keith Sklower
- Grad Students
  - Gautam Altekar
  - Marco Barreno
  - Weidong Cui
  - Rodrigo Fonseca
  - Dennis Geels
  - Brighten Godfrey
  - Ling Huang
- Grad Students
  - Dilip Antony Joseph
  - Jayanthkumar Kannan
  - Karthik Lakshminarayanan
  - Boon Thau Loo
  - Sridhar Machiraju
  - Steven Martin
  - Ana Sanz Merino
  - Blaine Nelson
  - George Porter
  - Sean Rhea
  - Sriram Sankararaman
  - Mukund Seshadri
  - Anil Sewani
  - Mel Tsai
  - Li Yin
  - Fang Yu
  - Shelley Zhuang

# Retreat Purpose

## Second OASIS/I3 Retreat

- Extensions to the Internet Architecture
  - » I3: Internet Indirection Infrastructure
  - » New Methods for Naming and Addressing: P2P Systems and DHTs
  - » OASIS: Emerging technology of PNEs, applied to Network layer observation, analysis, and action
  - » Application for Reliable Adaptive Distributed Systems


## NSF Cybertrust Center Proposal: Center for Adaptive Trustworthy Systems (CATS)—Unsuccessful!

- Separate ROC-RADS and Net-RADS proposals
- "Protecting Networks with COPS: Checking, Observing, and Protecting Systems," PIs: Katz, Shenker, Stoica



**Overlays and  
Active  
Services for  
Inter-networked  
Storage**





says: "The Network is the Computer"

*We say: "The Computer is the Network"*

# Proliferation of Network Appliances



**Packeteer PacketShaper**  
*Traffic monitor and shaper*



**Network Appliance NetCache**  
*Localized content delivery platform*




**F5 Networks BIG-IP LoadBalancer**  
*Web server load balancer*



**Ingrian i225**  
*SSL offload appliance*



**Cisco SN 5420**  
*IP-SAN storage gateway*



**Nortel Alteon Switched Firewall**  
*Checkpoint firewall and L7 switch*



**NetScreen 500**  
*Firewall and VPN*



**Extreme Networks SummitPx1**  
*L2-L7 application switch*




**Cisco IDS 4250-XL**  
*Intrusion detection system*

In-the-Network Processing: the Computer IS THE Network

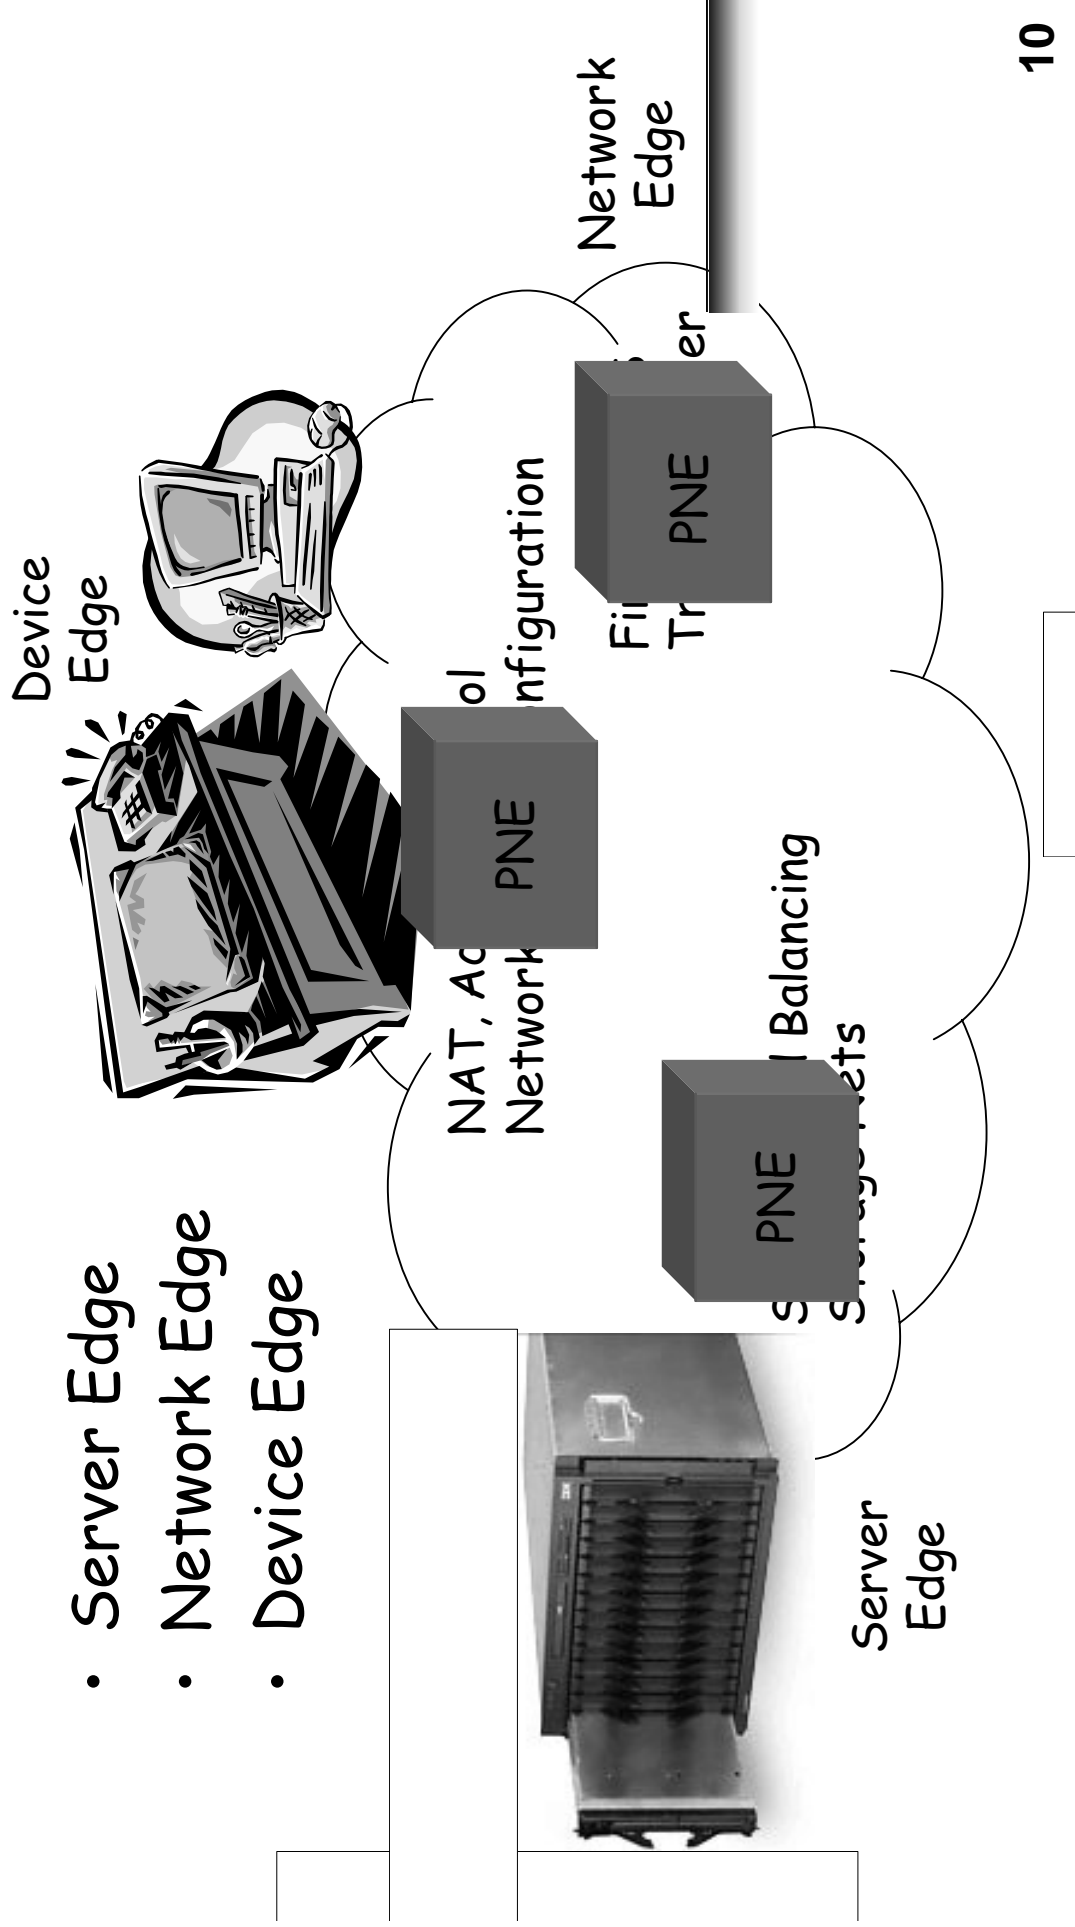


# Generic PNE Architecture



# Adaptive Edge Networks

- Server Edge
- Network Edge
- Device Edge



# OASIS Vision

- Specification/control environment for diverse network elements to realize full power of “inside the network” services and applications
- Via virtualized architecture for PNEs (aka RouterVM), retarget for diverse appliance-specific architectures
- Focus on stream extraction, intrusion detection, network monitoring, iSCSI acceleration
- Sys admins “program” the network through service specification and composition
- Open framework for multi-platform appliances, enabling third party service development

# Reliable Adaptive Distributed Systems


Dramatically improve the trustworthiness of networked systems

- Observe: design observation points throughout system
- Analyze: infer via statistical learning
  - Respond: detect anomalous behavior vs. baseline
  - Learn: use observations to modify responses to future observations
- Act:
  - Reactive: use control points in system for rapid recovery if detect something wrong
  - Proactive/protective: prophylactically act on system to prevent predicted impending failure

# Brittle Distributed Systems

- Fragile, easily broken, poor dependability and security
  - E.g., Amazon: yearly revenue \$3.1B, downtime costs \$600,000/hr
- Design for rapid detection, diagnosis, recovery
  - Rapid application and server recovery, agile network rerouting, proactive protective actions ...
  - No distinction between "normal operation" and "recovery"
- Elements of our solution
  - Programming paradigms for robust recovery
  - Crash-only software design for rapid server recovery
  - *Network protocols designed for observation to allow rapid detection of behavioral violations*
  - *Instrumentation and SLT for on-line analysis, anomaly detection, diagnosis of failure*
- Adaptation benchmarks to measure progress
  - What you can't measure, you can't improve
  - Collect real failure data to drive benchmarks

# Reliable Adaptive Distributed Systems




# Summer04 Retreat Feedback

- **Retreat Organization:**
  - Cramped space, no place to lounge and talk
  - Separate sessions should be partitioned by topic
  - Presentations in PDF as well as PPT
- **Technical Comments:**
  - Reference Architecture for Protocols, Services, Applications, Configurations needed!
  - "Intelligent" adaptive network management, network storage management represent huge challenges and opportunities
  - Bladed systems with Net+Processing+Storage represents a real opportunity as an architectural platform
  - Importance of building measurement capabilities into the system from the ground up
  - Progress needed on the integration of Observe, Analyze, Act
  - I3 needs to get out!
  - Integration of other research activities with PNEs

# Plan for the Retreat

- Monday, 10 January 2005
  - 0730 Bus to Tahoe
  - 1200 Lunch
  - 1330 Introduction to Retreat  
"Retreat Overview and Plan, OASIS Update"—Randy Katz  
"I3 Update"—Ion Stoica  
*Highlight Talk: "OpenDHT"*—Sean Rhea
  - 1500 Break
  - 1530 Introduction to New COPS NSF Proposal—Randy & Ion
  - 1430 Break
  - 1700 "State of Funding in Computer Science Research"—Dave Patterson
  - 1800 Dinner
  - 1930 Evening Session: "The Case for RADS"  
—Armando Fox





# Plan for the Retreat

- Tuesday, 11 January 2005
  - 0730 Breakfast
  - 0830 Parallel Sessions
    - » Programmable Network Elements—Randy
      - "Instrumenting 3 Tier Systems for Performance and Reliability"—George Porter
      - "Predicting PNE Performance from Router VM Specifications"—Mel Tsai
      - "Observe-Analyze-Act Paradigm for Storage System Optimization"—Yin Li
      - "Multimatch Classification Using SRAM and TCAM"—Fang Yu
    - » Potpourri Topics—Ion
      - "Beacon Vector Protocol"—Rodrigo Fonseca
      - "Semi-Supervised Learning on Email Characteristics for Novel Worm Detection"—Steven Martin and Anil Sewani
      - "Analyzing Countermeasures to SLT-based Techniques"—Blaine Nelson
      - "Using the Time-Series Nature of Data to Improve Prediction"—Ling Huang
  - 1000 Break
  - 1030 Parallel Sessions
    - » Performance and Dependability—Randy
      - "Towards More Dynamic Internet Routing"—Mukund Seshadri
      - "Binder: Extrusion-based Break-in Detection"—Wedong Cui
      - "Active Probing for Available Bandwidth Detection"—Sridhar Machiraju
    - » Peer-to-Peer and Overlay Networks—Ion
      - "Securing Forwarding Infrastructures"—Karthik Lakshminarayanan
      - "The Cost of Inconsistency in DHTs"—Shelley Zhuang
      - "On the Effect of Heterogeneity in Distributed Systems"—Brighten Godfrey
      - "Declarative Networks"—Boon Loo
  - 1200 Lunch
  - 1300 Ski Break
  - 1700 "Debugging Deployed Routing Overlays"—Dennis Geels
  - 1730 Poster Previews
  - 1800 Dinner



# Plan for the Retreat

- **Wednesday, 12 January**
  - 0730 Breakfast
  - 0830 Feedback on NSF Proposal
  - 1000 Break and Checkout
  - 1030 Industry Feedback Session
  - 1200 Lunch
  - 1300 Bus to Berkeley



# Recent OASIS/I3-Related Publications

- Y. Chen, D. Bindel, H. Song, R. H. Katz, "An Algebraic Approach to Practical and Scalable Overlay Network Monitoring," ACM SIGCOMM Conference, Portland, OR, (August 2004).
- F. Yu, T. V. Lakshman, R. H. Katz, "Multi-class Classification using TCAM," Hot Interconnects 12 Symposium on High Performance Interconnects, Stanford, CA, (August 2004). Best Paper Award.
- F. Yu, T. V. Lakshman, R. H. Katz, "Gigabit Rate Pattern-Matching using TCAM," International Conference on Network Protocols (ICNP), Berlin, Germany, (October 2004). Best Paper Award.
- A. Fox, E. Kiciman, D. Patterson, M. Jordan, R. H. Katz, "Combining Statistical Monitoring and Predictable Recovery for Self-Management," Proceedings of 2004 Workshop on Self-Managed Systems (WOSS'04) in conjunction with ACM SIGSOFT FSE-12, Newport Beach, CA, (October 2004). Earlier version presented at 2nd Bertinoro Workshop on Future Directions in Distributed Computing (FuDiCo II): Survivability: Obstacles and Solutions, (June 2004).
- S. Machiraju, R. H. Katz, "Verifying Global Invariants in Multi-Provider Distributed Systems," Proceedings HotNets-III Workshop on Hot Topics in Networking, San Diego, CA, (November 2004).
- S. Zhuang, R. H. Katz, I. Stoica, D. Geels, "On Failure Detection in Overlay Networks," Proc. IEEE Infocomm Conference, Miami, FL, (March 2005).
- S. Uttamchandani, L. Yin, G. Alvarez, J. Palmer, G. Agha, "CHAMELEON: a self-evolving, fully-adaptive resource arbitrator for storage systems", Proc. USENIX 2005 Technical Conference, Anaheim, CA, (March 2005).
- W. Cui, R. H. Katz, D. Tan, "BINDER: An Extrusion-based Break-In Detector for Personal Computers," Proc. USENIX 2005 Technical Conference, Anaheim, CA, (March 2005).

# Oasis/I3 Retreat Overview

Randy H. Katz

Univ. of California  
Berkeley, CA  
94720-1776

