

Experiments with Differentiated Services at iCAIR

Presenters: Dilip Kandlur, IBM

Joel Mambretti, iCAIR

kandlur@us.ibm.com, j-mambretti@nwu.edu

First Joint Internet2/DOE QoS Workshop
"QBone: Early Experiences and the Road Ahead"

Feb 9, 10, 2000



In Collaboration With

- iCAIR & IBM
 - Brian Carpenter
 - Doug Freimuth
 - Ashish Mehra
 - Dinesh Verma
 - Tim Ward
 - Jeremy Weinberger
 - Dan Weaver
- U Penn
 - Roch Guerin
 - Wael Ashwami (supported in part by NSF contract #ANI 99-06855)
- MREN
 - Linda Winkler et al
- ANL
 - Linda Winkler
 - Alain Roy
 - Sander Volker
- CERN
 - Joop Joosten
 - Paolo Moroni
- EMERGE (supported by DOE Science Grid funding)
 - PI Tom DeFanti

Outline

- Objectives
 - large scale experiments with high quality video streaming and other delay sensitive applications, using DiffServ across heterogeneous networks
 - evaluate impact of different enforcement mechanisms at the application level, i.e., video quality et al
- Experimental Apparatus and Mechanisms
 - Different Server QoS mechanisms and Policy control
 - EF service enforcement design
- Experiments and experiences
- Related activities at iCAIR

DiffServ Context

- IETF DiffServ (Co-Chair, Brian Carpenter)
- IETF RFC 2768 on Middleware (National Science Foundation-Funded Forum Dec 1998 at Northwestern, hosted by Cisco, IBM, ANL MCS and iCAIR)
- IETF RFCs: 2474, 2475, 2598 (EF)

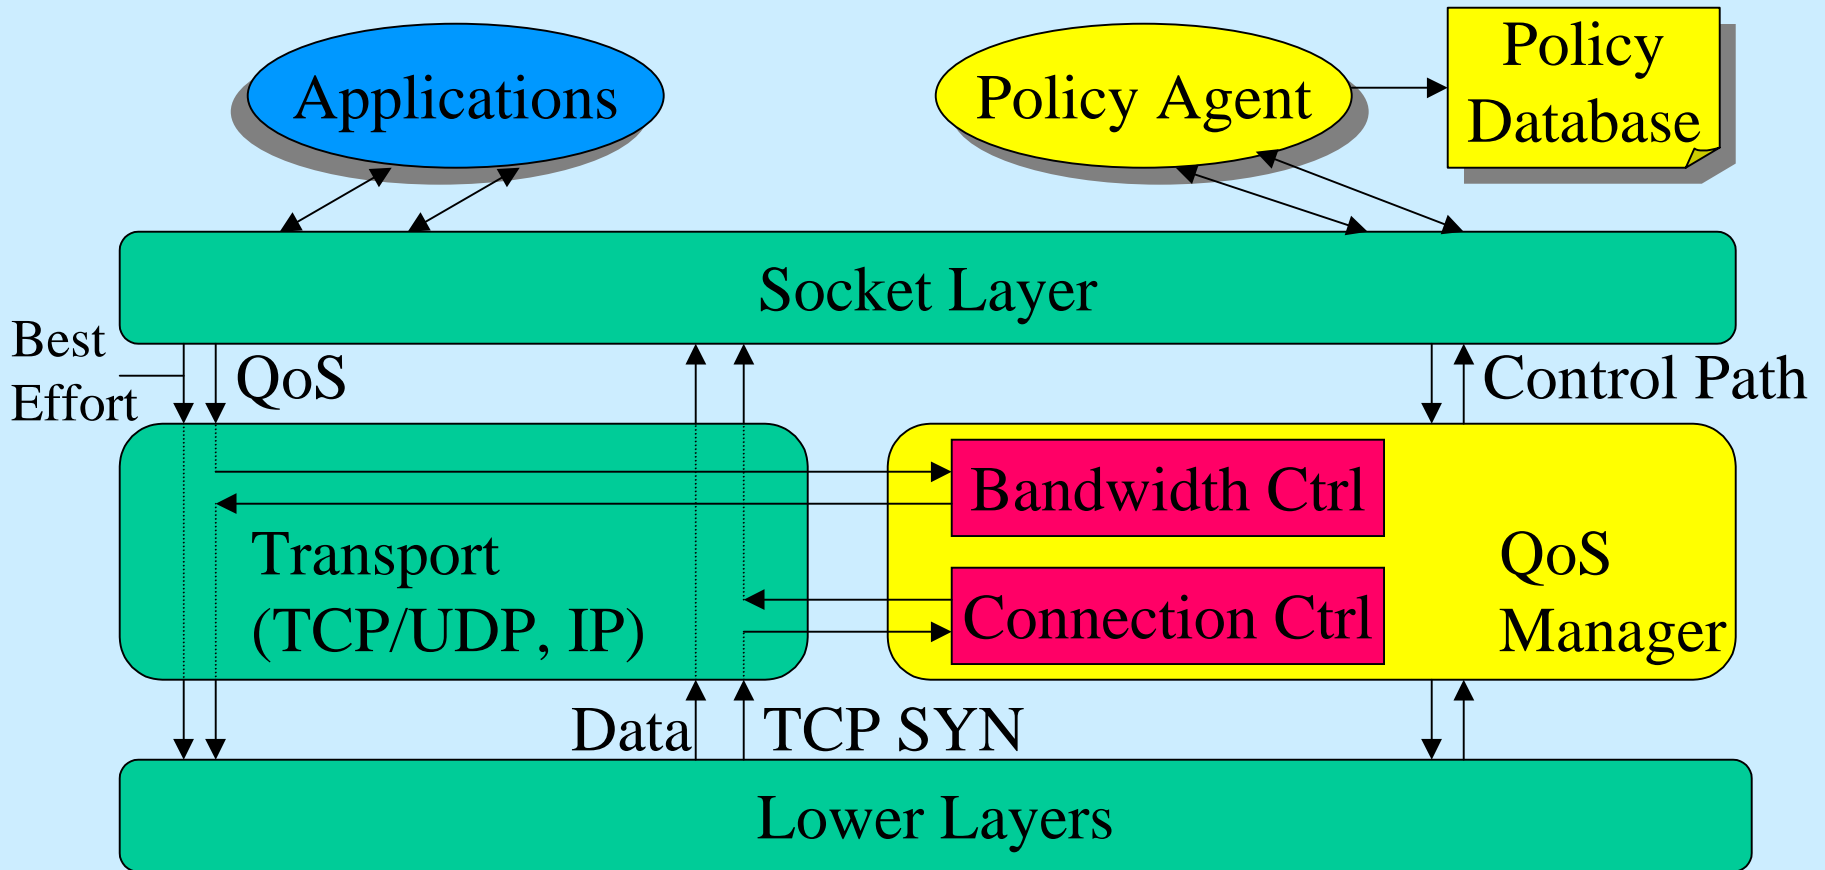
QoS and DiffServ Experiments

- Applications - [Digital Video](#), VoIP, VR, ...
- DiffServ Instrumentation
 - Host Policy (eg, via LDAP)
 - Host Packet Marking
 - Router enforcement via ATM PVCs and ingress CAR
 - Monitoring using RTFM
 - Performance Instrumentation
- Experiments
 - Sensitivity to EF policing at different domain boundaries
 - how to configure parameters in the absence of shaping
 - Tracking of video quality as function of network parameters and policing options
 - Quantitative and qualitative evaluations of video quality
 - Dropping vs best-effort forwarding of excess traffic

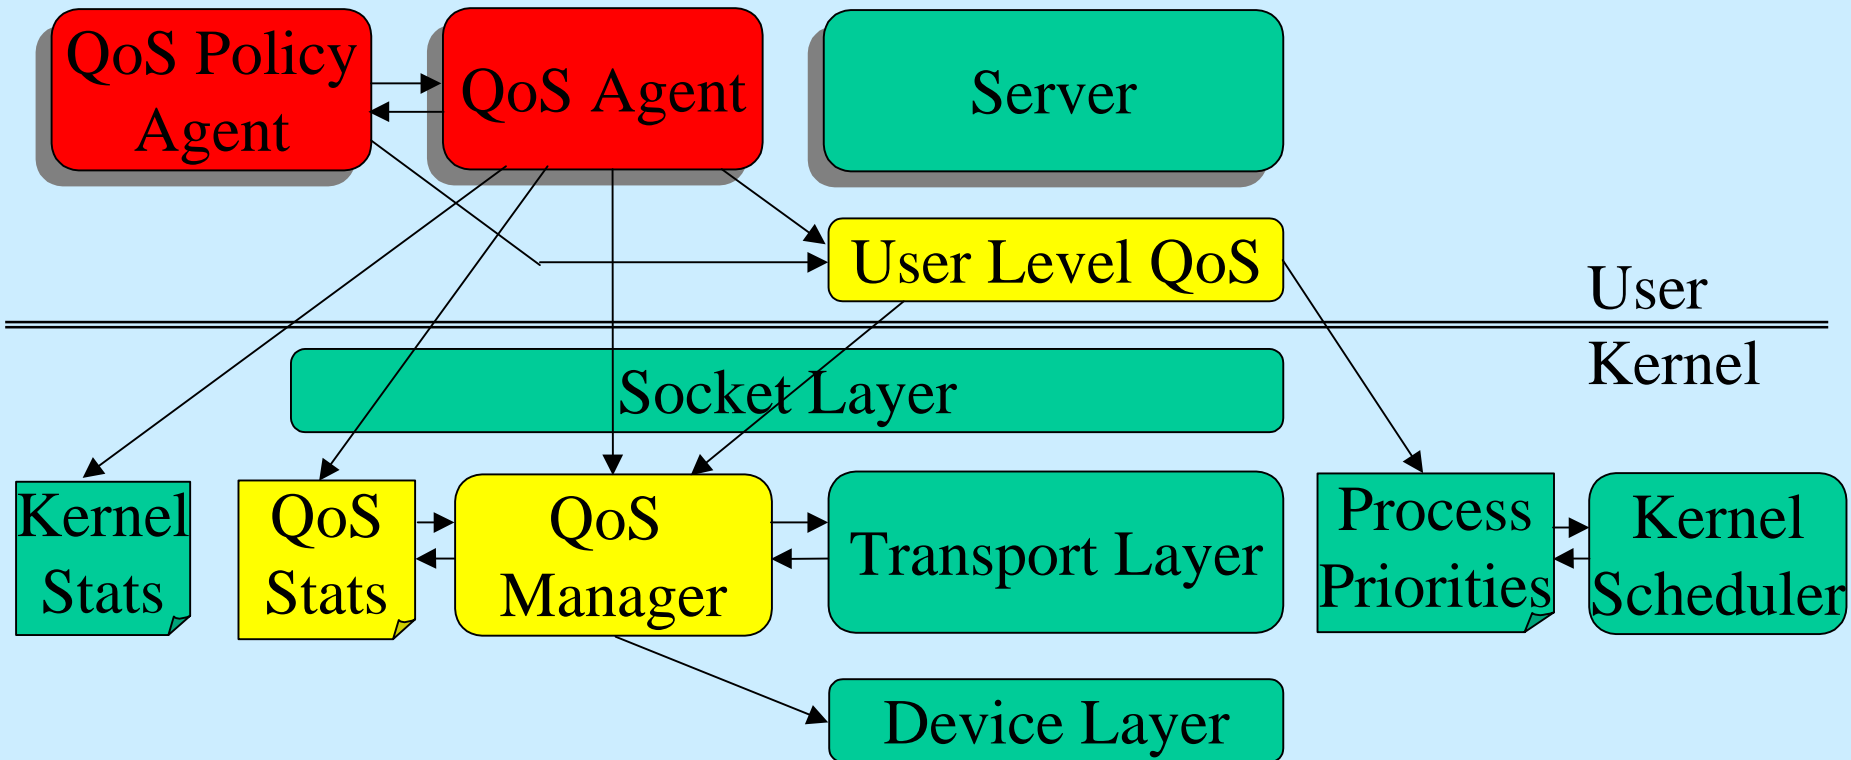
Server QoS Mechanisms

- Policy-based QoS support in AIX 4.3.3 release
 - does not require application modification
 - IETF DiffServ and IntServ support
 - traffic control: policing, marking, shaping
- Enhanced outbound traffic conditioning
 - *Scalability*: large aggregates, adaptive shaping timers
 - *Performance*: search optimizations, multiple triggers
 - *Fairness*: logical partitions for fair token sharing (TCP)

Policy Based QoS Architecture



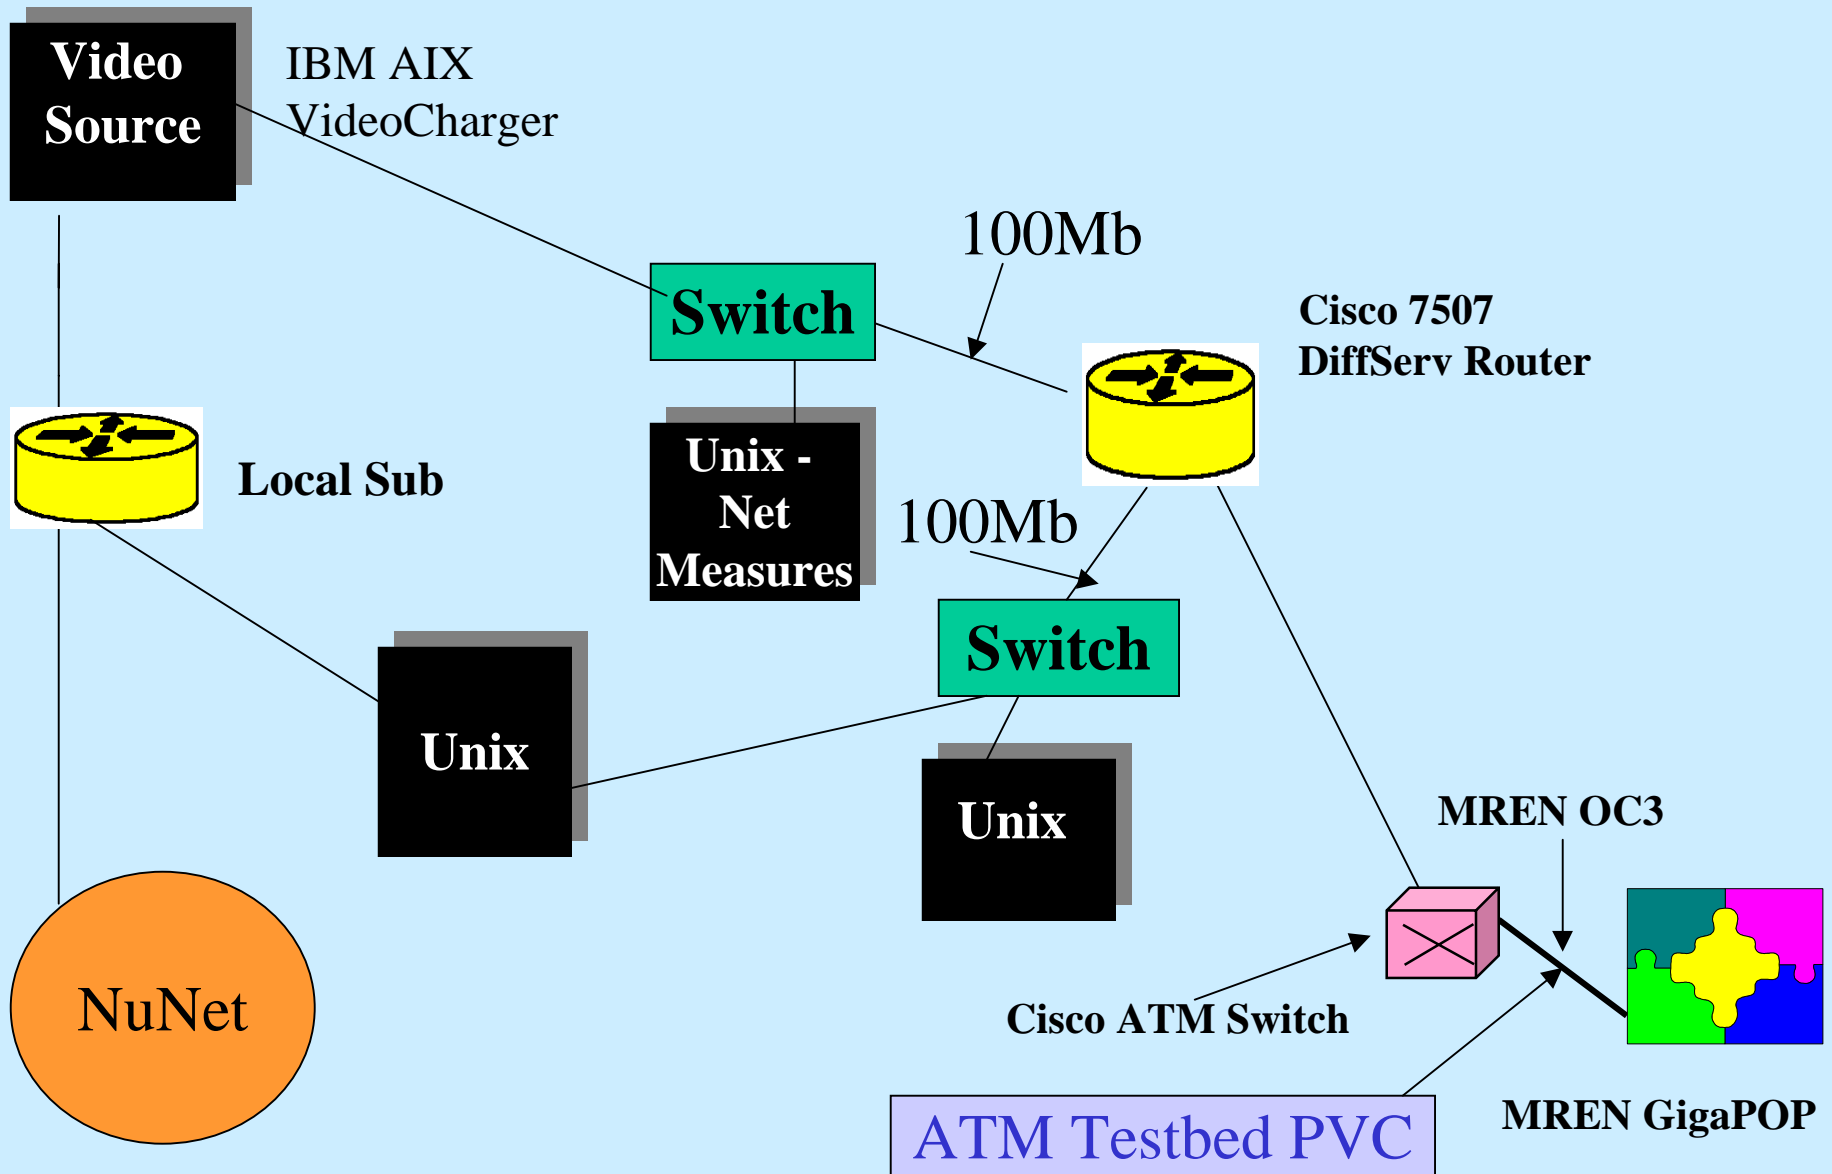
Server QoS Architecture



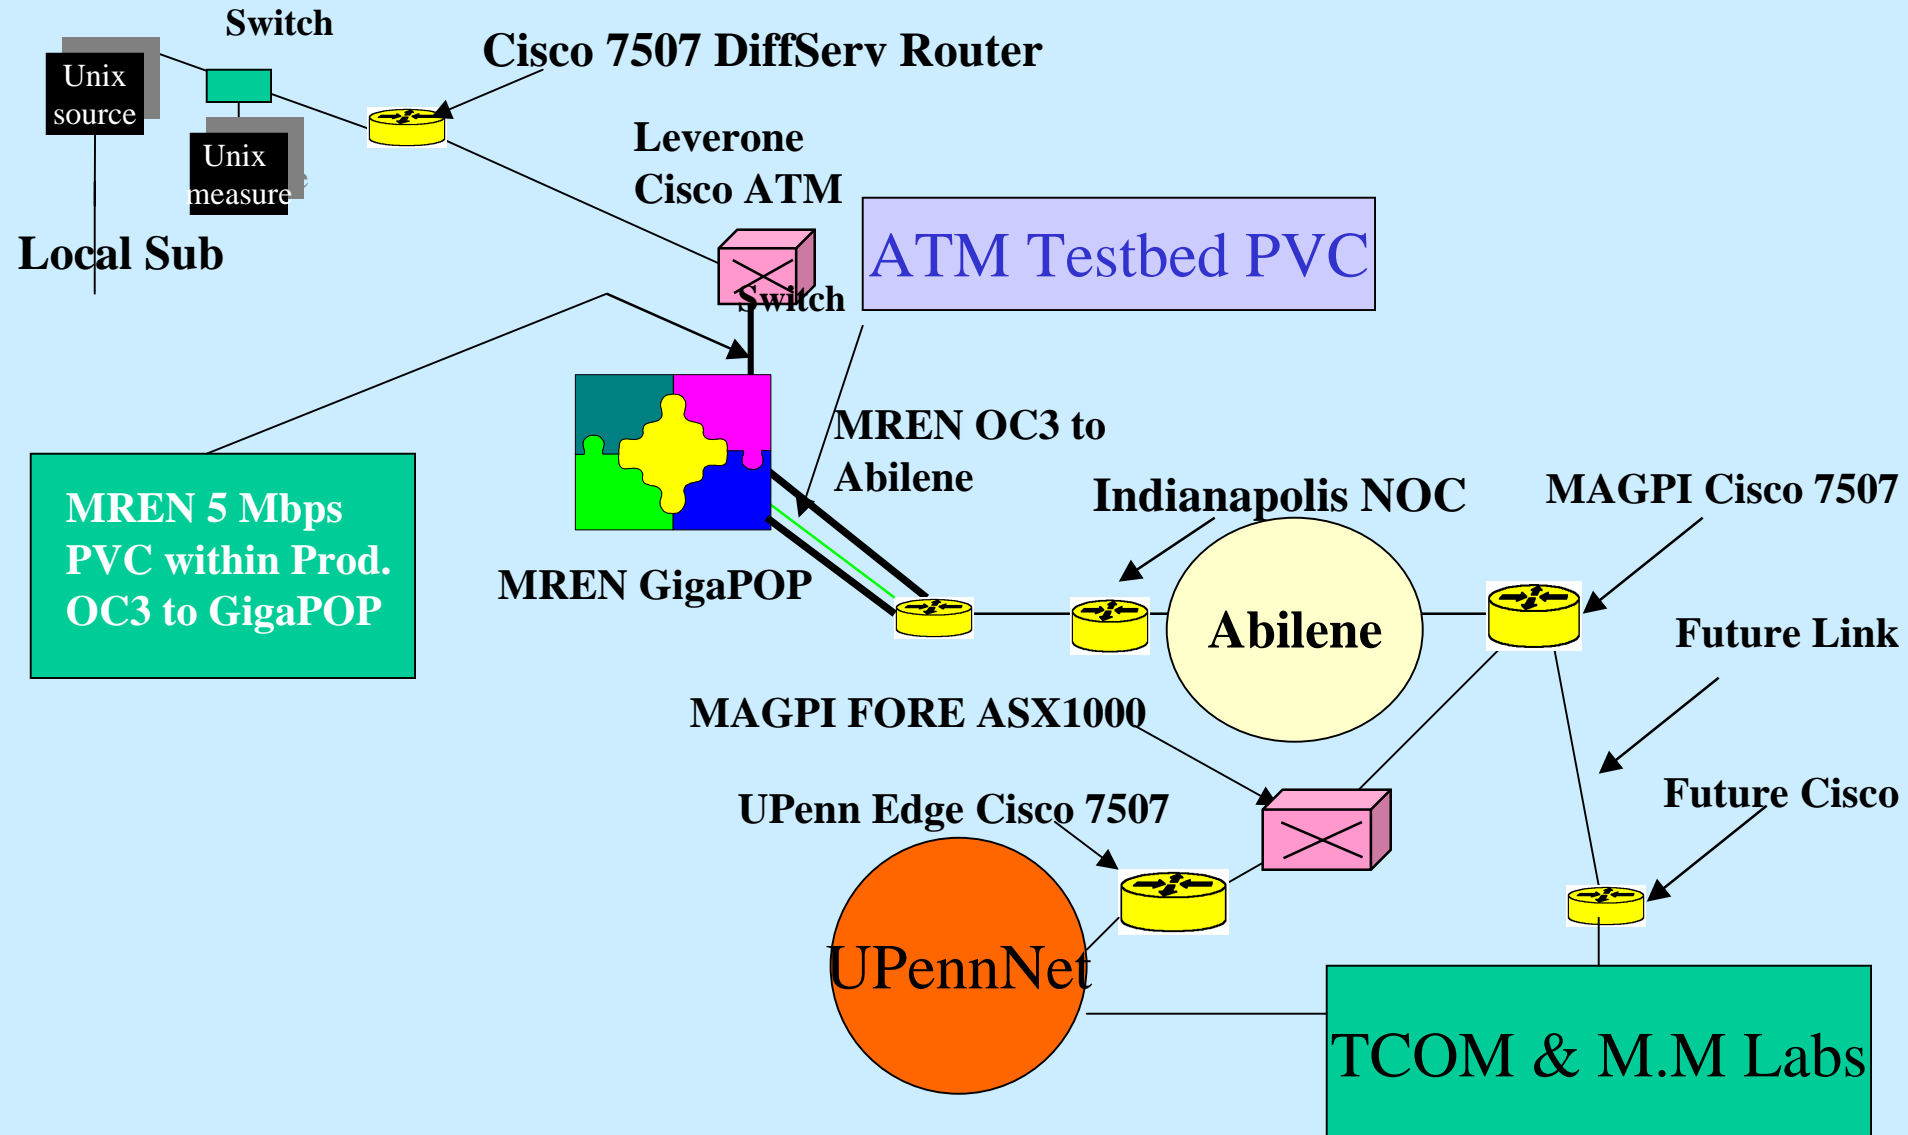
Pitfalls and Issues

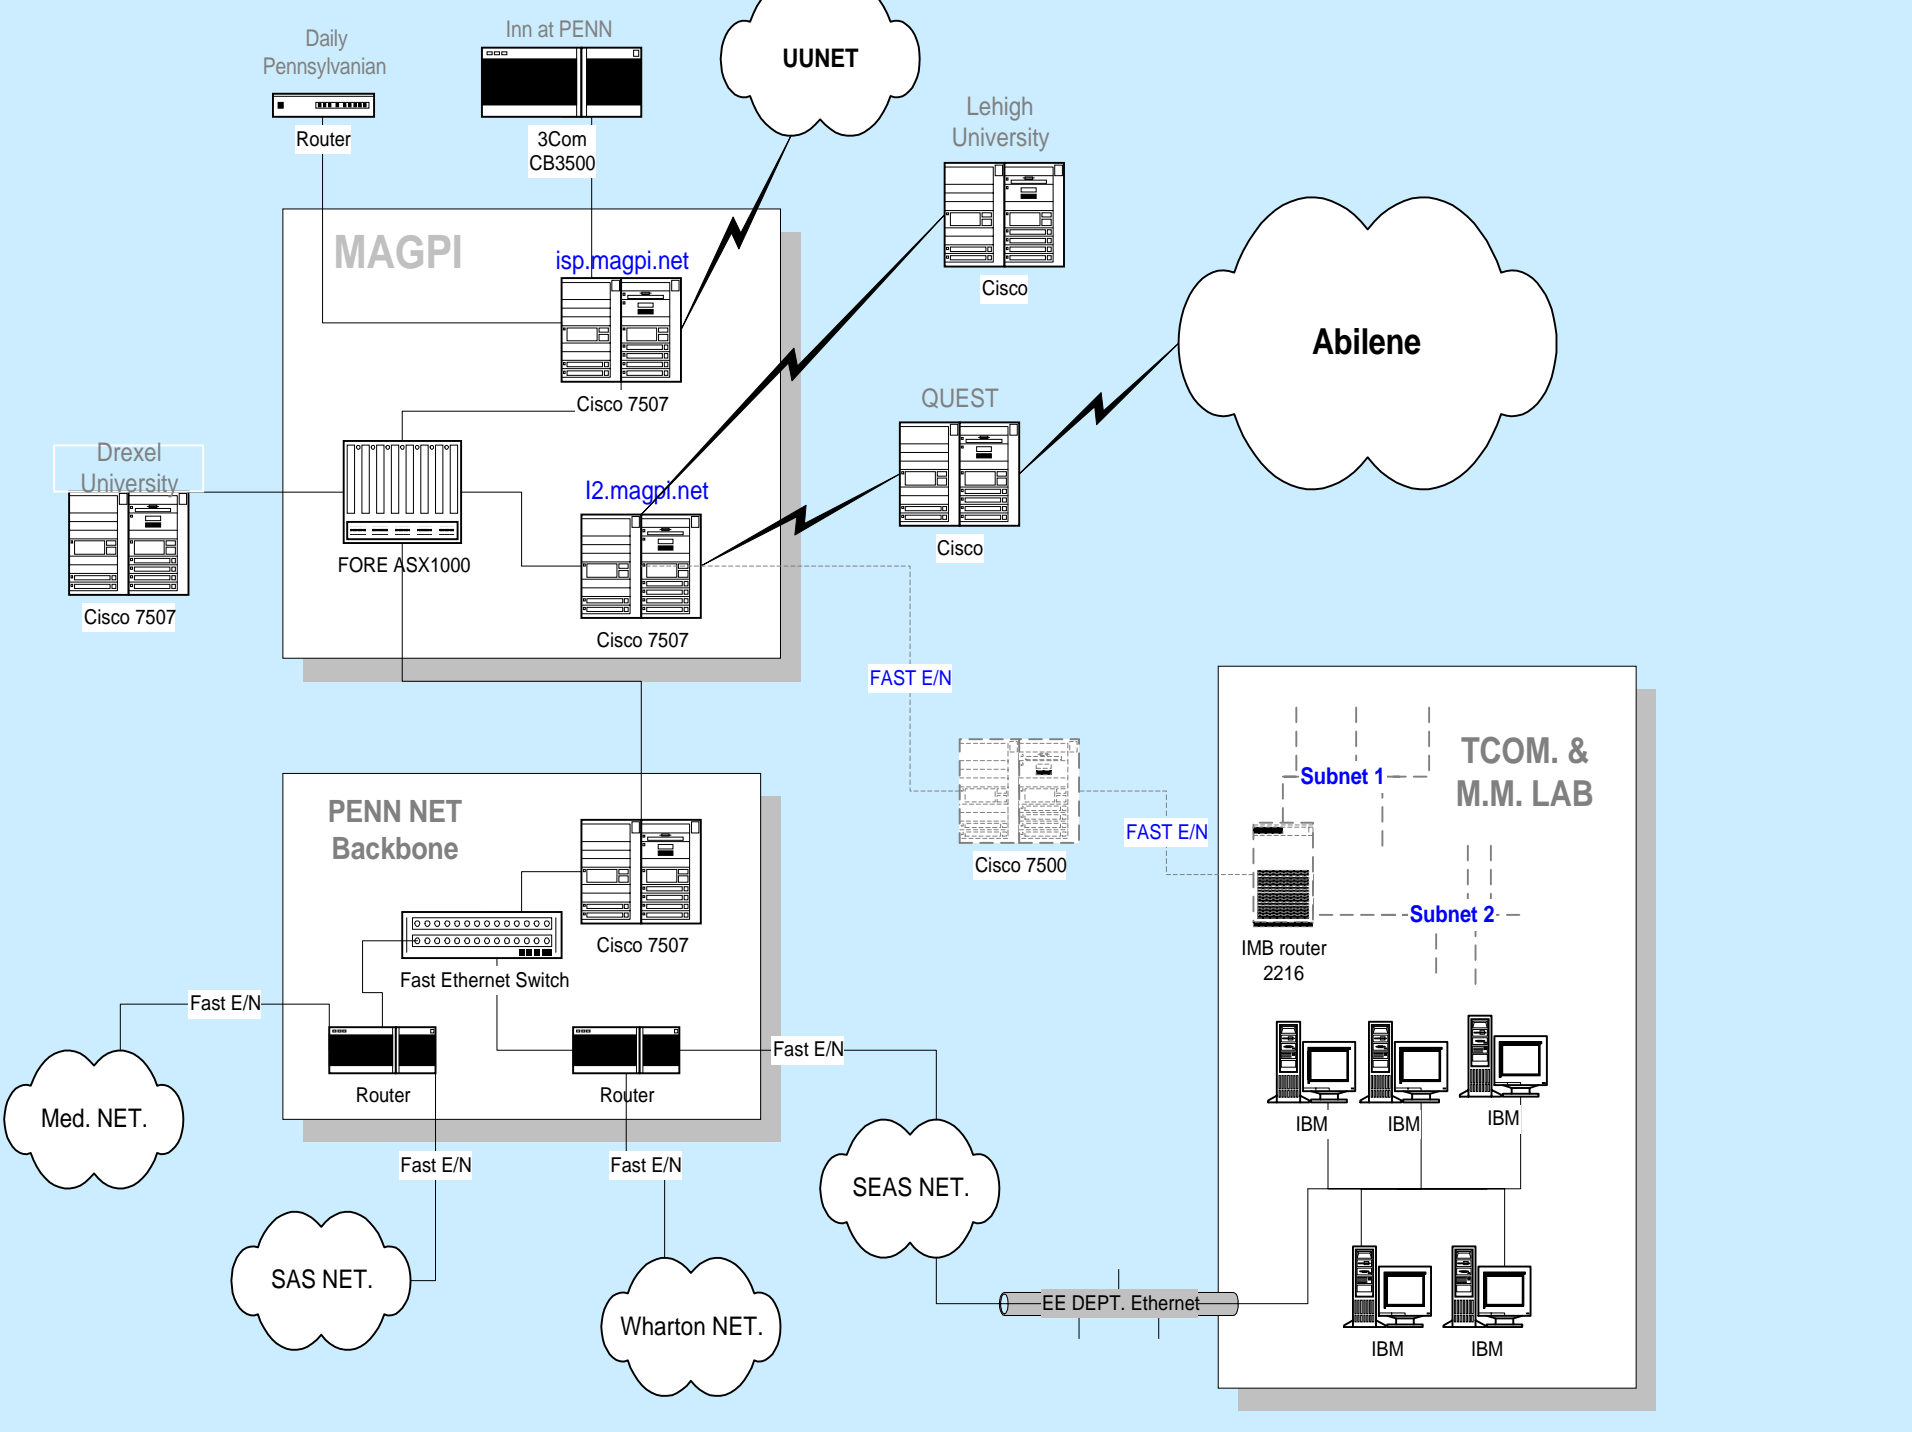
- Impact of policy control of access to Diff-Serv
 - Asymmetric routes
 - Selective advertising of routes
 - Dual-homed systems
- => Lot's of added complexity and places where things can go wrong (and they have...)
- Heterogeneity of network technologies
 - ATM and IP interactions
 - Impact of ATM segmentation and mapping of traffic contracts
 - Different levels of classification capabilities in routers
 - CARR based on IP precedence vs DSCP vs SA/DA in different IOS releases
- Deployment logistics
 - interactions between QBone and production networks
 - interactions between QBone and NGI networks
 - "Experimental" vs stable IOS versions

DiffServ Testbed @ iCAIR



iCAIR <=> UPenn





Policy Implementation Mechanisms

- Random Early Discard (RED)
- Queuing Implementations
- Priority Queuing
- Class-Based Queuing
- Weighted Fair Queuing (WFQ)
 - All Flows Equally Weighted
 - Packets with Same Characteristics Identified as in the Same Flow

CB-WFQ = Enhanced QoS Granularity

- Extends WFQ to Allow for Support for User Defined Traffic Classes
- Can Define Class Based on Protocol, Ingress Interface, Access Control List etc
- Packets Fitting Specified Criteria Constitute Class
- Classes Linked to Queues
- Classes Are Given BW, Weight, Maximum Packet Limit, Queue Limit etc.

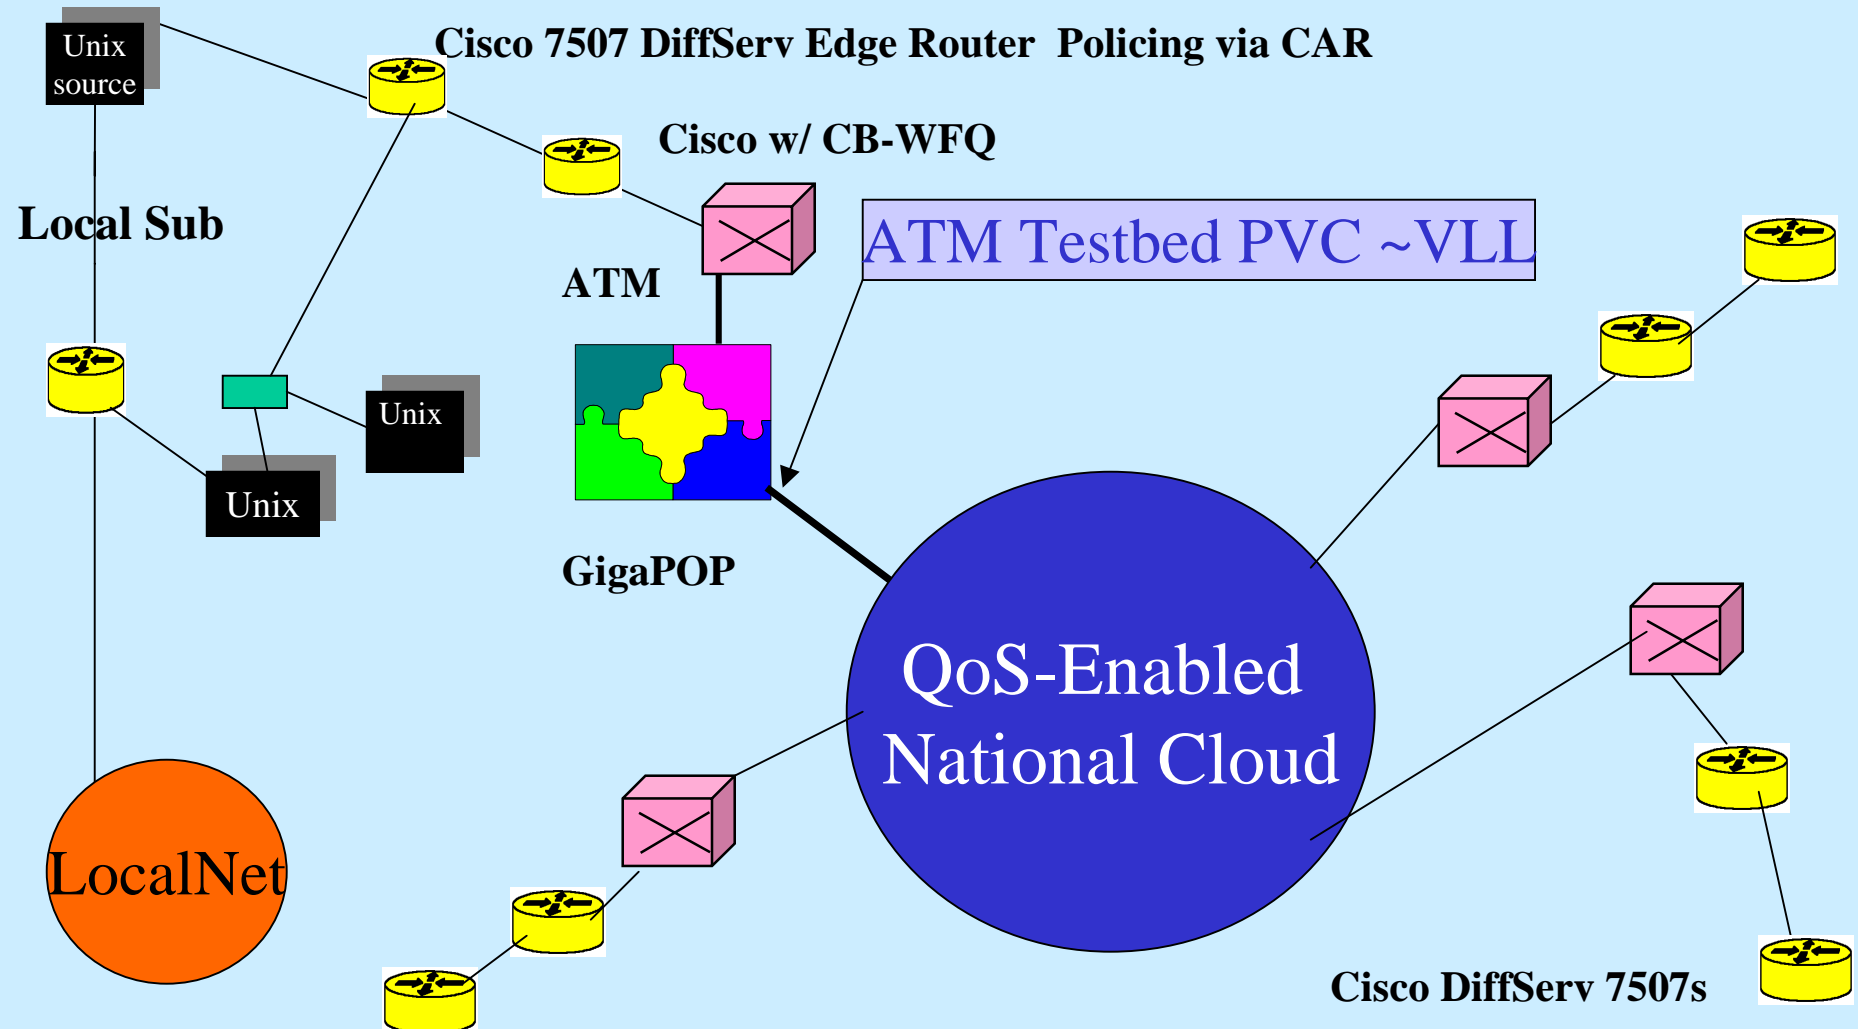
iCAIR & EMERGE: ESnet/MREN Regional Grid Experimental NGI Testbed

- Funded by DOE NGI Program
- Globus Services, GARA etc.
- A to B Path Defined Non-Dynamically via a Dedicated PVC (~VLL, VBB?)
- L2 -- ATM over SONET:
- iCAIR ~ 17 Mile Link to MREN GigaPop, 35 Mile Link further to ANL MCS
- Various Flow Control Mechanisms
- Cisco 7507s IOS 12.0 (x) XE beta images



DiffServ QoS Testbed (eg, EMERGE)

Local DiffServ TestBed

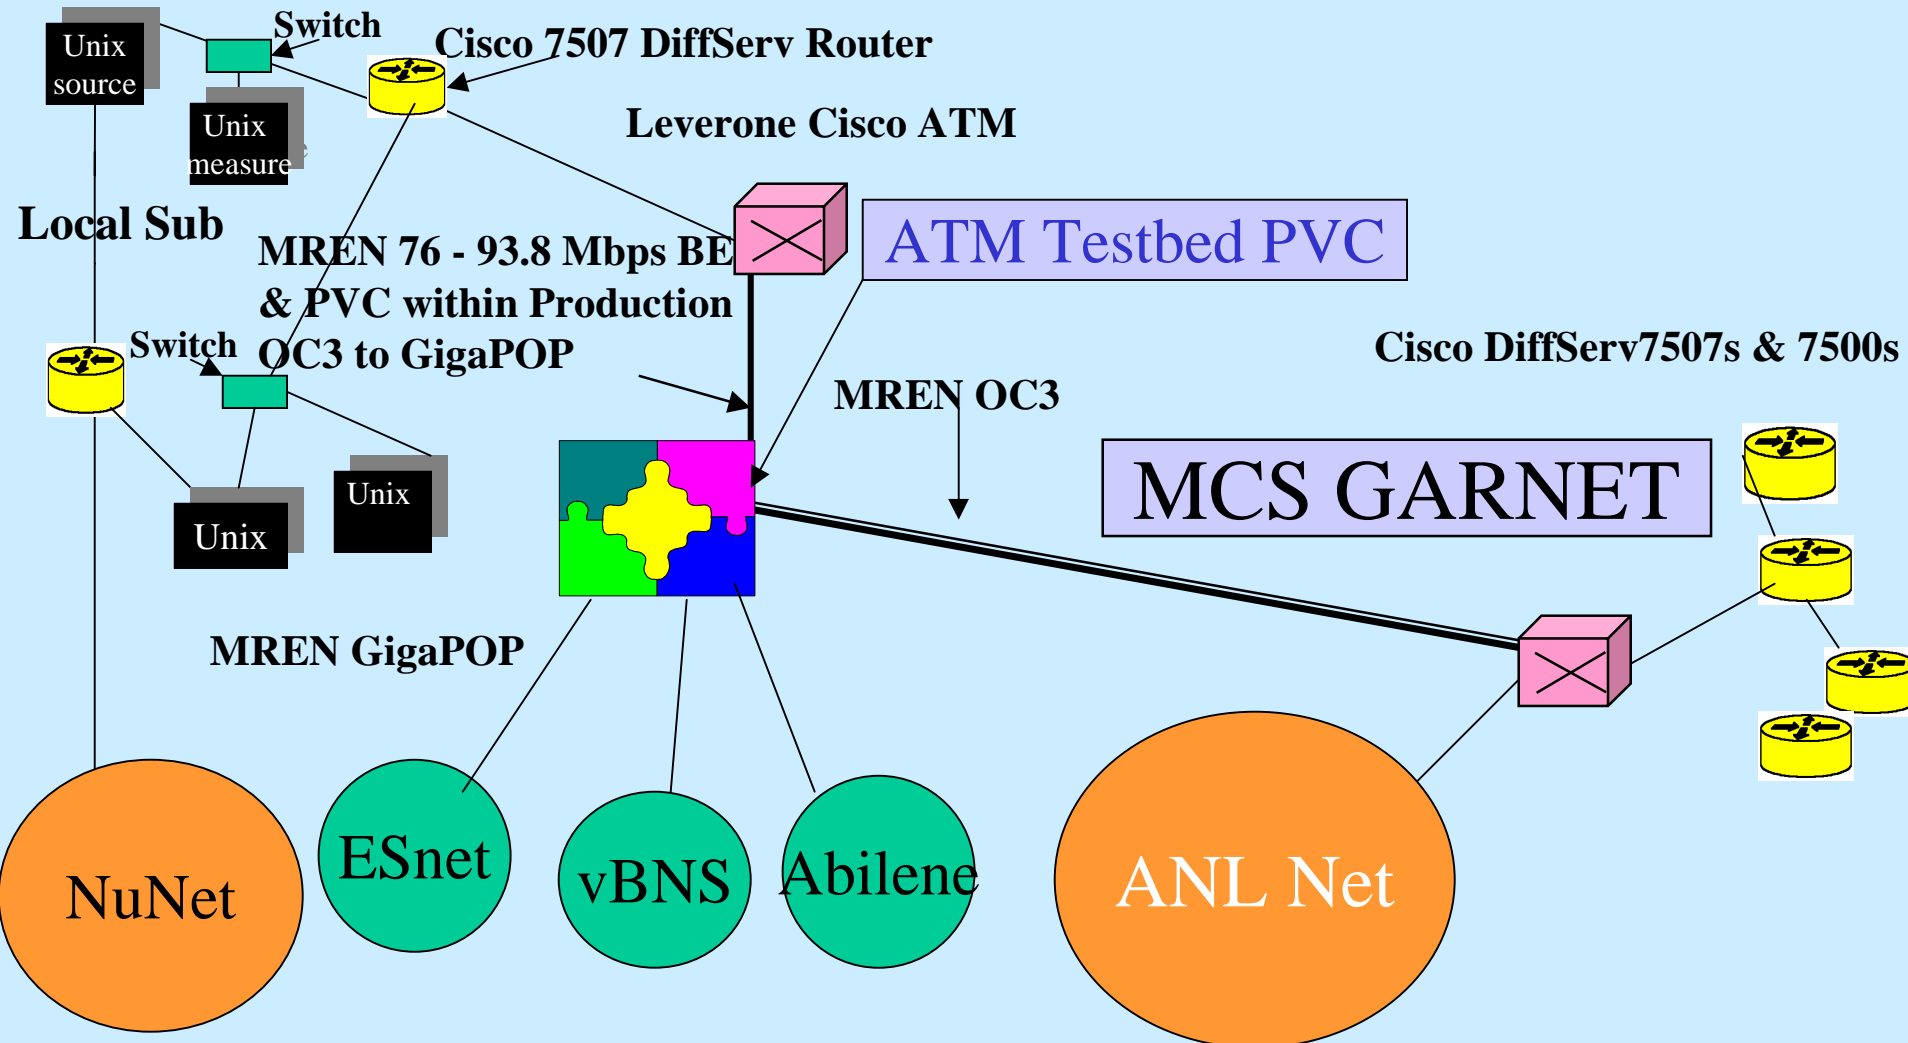


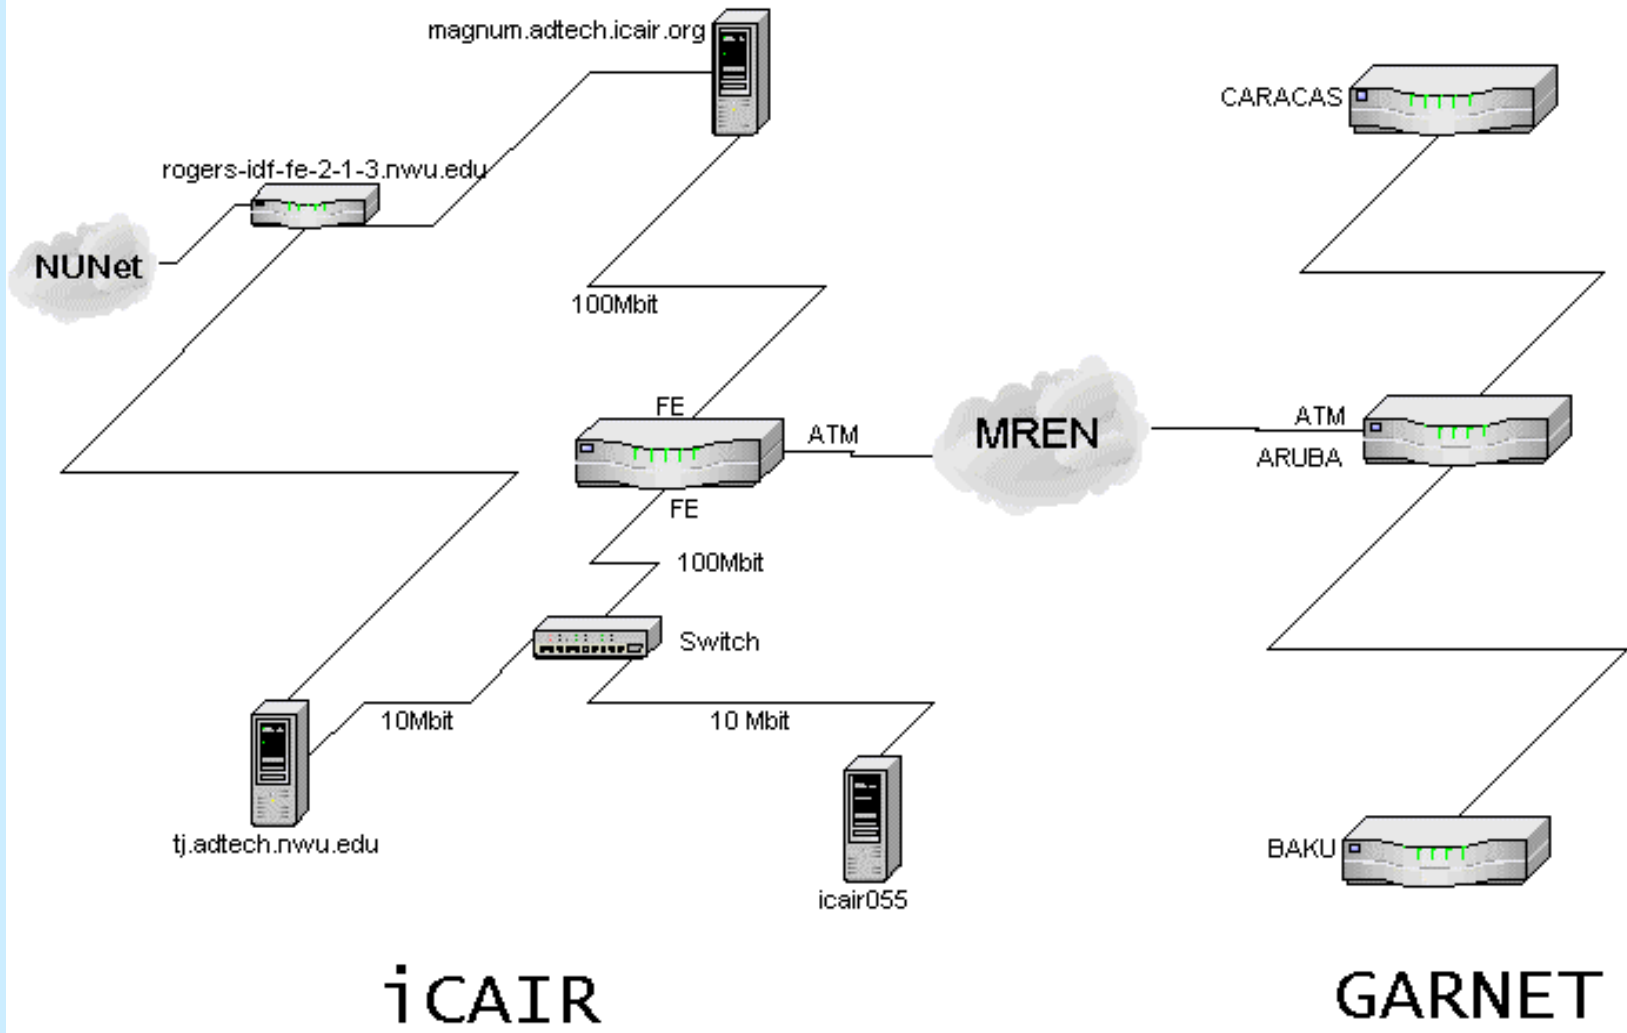


iCAIR EMERGE Testbed:

iCAIR<=>ANL

iCAIR DiffServ EMERGE TestBed

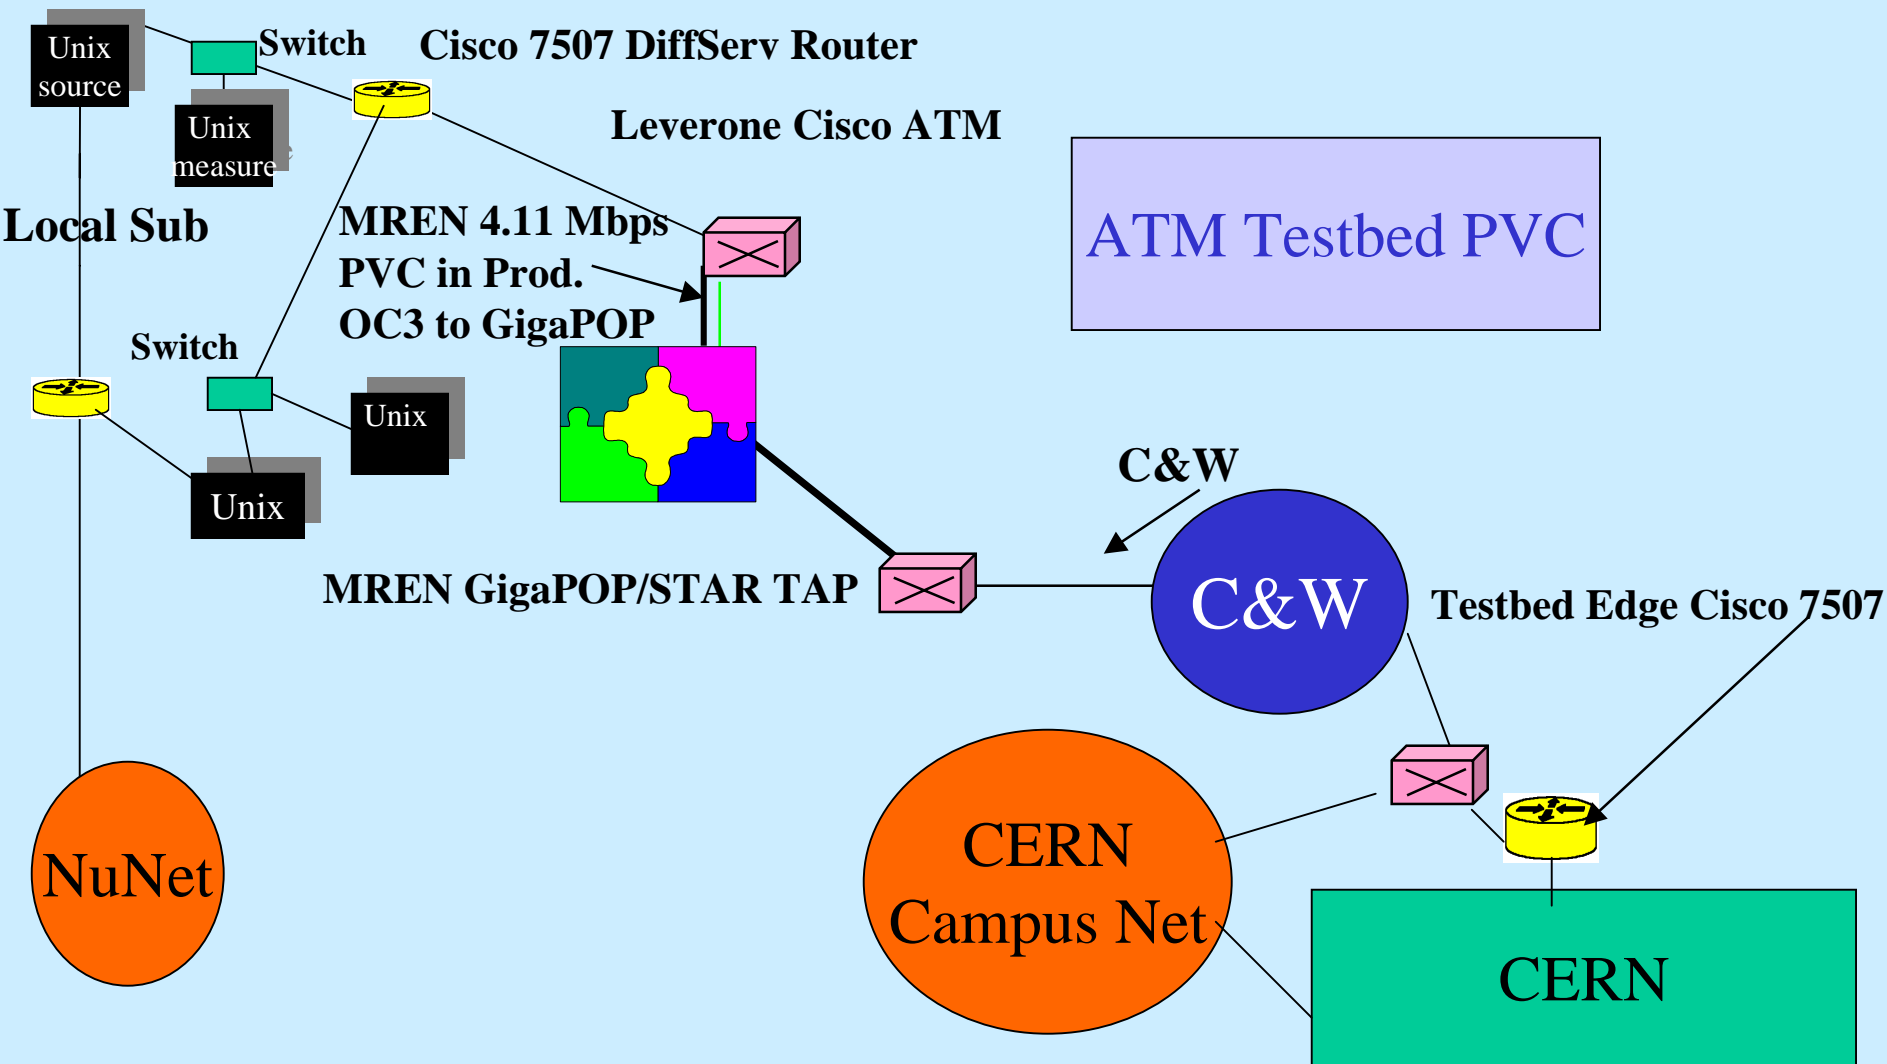






iCAIR<=>CERN

iCAIR DiffServ CERN TestBed



Extensions and Enhancements

- Integration with AI X Workload Manager
 - enhanced traffic classification
 - dynamic bandwidth shares
- Expansion to Linux
- Integration with Congestion Manager
 - common congestion notification for flows within a class
 - service-based monitoring and control
- Additional Testbed Experiments, Upenn, EMERGE, CERN, etc.

Pointers for Information

- <http://magnum.adtech.icair.org/quick>
- <http://www.icair.org>
- <http://www.research.ibm.com/compsci/communications/projects/qos>
- <http://www.evl.uic.edu/cavern/EMERGE>