
Experiences with Active Kernel Modules at Washington University

John DeHart
Washington University
jdd@arl.wustl.edu
<http://www.arl.wustl.edu/arl/projects/ann.html>

October 14,15 1999

OpenSig'99

 Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

1

Questions: Who, What, Where

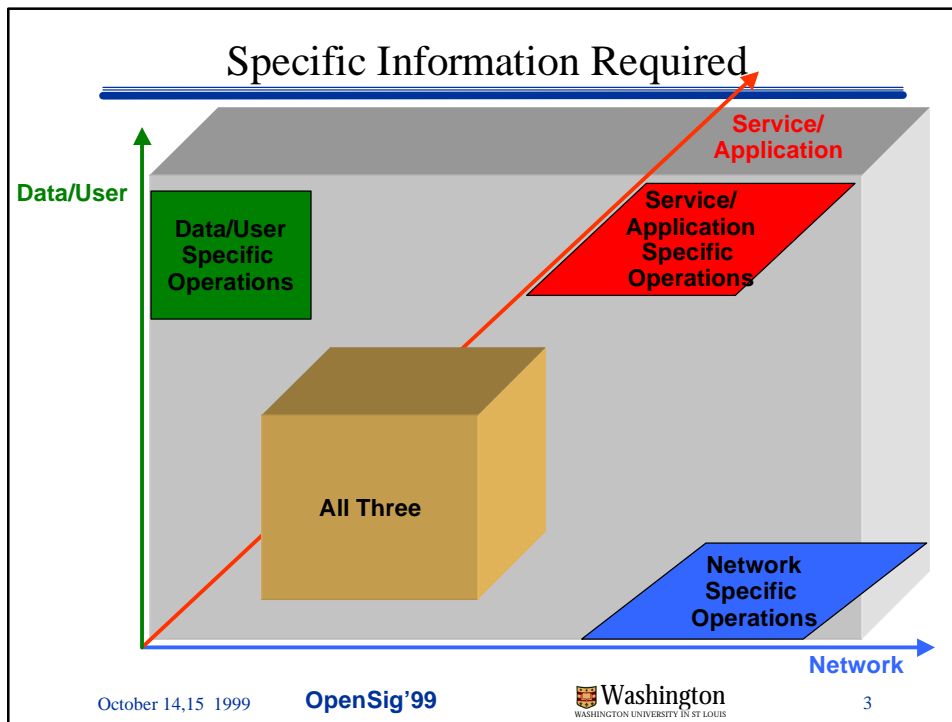
- Who has the **information** and **access** to be able to write the Active Code?
–
- What level of **performance** is required to make the active code effective?
–
- Where in the network are the **resources** to execute the code?

October 14,15 1999

OpenSig'99

 Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

2



- ## Who Writes the Active Code
- End User - “classic” active networking
 - **“Its my data, here’s how I want to manipulate it ...”**
 - Application/Service Developer
 - **“Its my application, I understand the format of it best, here’s some capabilities for manipulating it ...”**
 - e.g. Vendor X provides Active congestion control for its MPEG video conferencing system.
 - Network Provider/Developer
 - **“Its my router, I’ll give you these capabilities for manipulating your data ...”**
 - e.g. Alternate Routing and Topology based decisions.
- October 14,15 1999 **OpenSig'99** Washington
WASHINGTON UNIVERSITY IN ST. LOUIS 4

Where is the Active Code Executed

Greater Flexibility

- User Space Execution Environments
 - ANTs from MIT
 - CANES from Georgia Tech
 - PLAN from UPenn
 - ...
- Kernel Space Execution Environments
 - Scout from Arizona
 - ANN (Active Network Node) from WU
 - ...
- Hardware (e.g. FPGAs)
 - P4 from UPenn
 - PPX from WU
 - ...

Greater Performance

October 14,15 1999

OpenSig'99

Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

5

“Taxonomy” of Active Functionality

- Generic Network
 - Routing
 - **Congestion Control**
 - Network Monitoring and Control (Traffic Data collection, Topology discovery, ...)
 - Protocol Prototyping and Deployment
- Generic Data Manipulation
 - Encoding/Decoding/Transcoding
 - Encryption/Decryption
 - **Congestion Control**
- Specific Data Manipulation
 - Encoding/Decoding/Transcoding
 - **Congestion Control**
 - Application Specific Computation

October 14,15 1999

OpenSig'99

Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

6

Sample of Active Applications

- Everyone: Ping
- ActiveBridge - UPenn
- Universal Search Interface - Arizona
- MPEG video congestion control - Georgia Tech
- Wavelet video congestion control - WU
- Active Error Recovery - TASC/Umass
- Wireless
- Distributed Interactive Simulation
- ...

October 14,15 1999

OpenSig'99

 Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

7

Requirements for WU AN Platform

- Code should be deployed automatically and on-demand to router
 - however this happens rarely!
- Simple security model we understand
- Fast enough for data-path applications
 - 155 Mbits/s in software
 - 1.2 Gbits/s with hardware support (FPGA)

October 14,15 1999

OpenSig'99

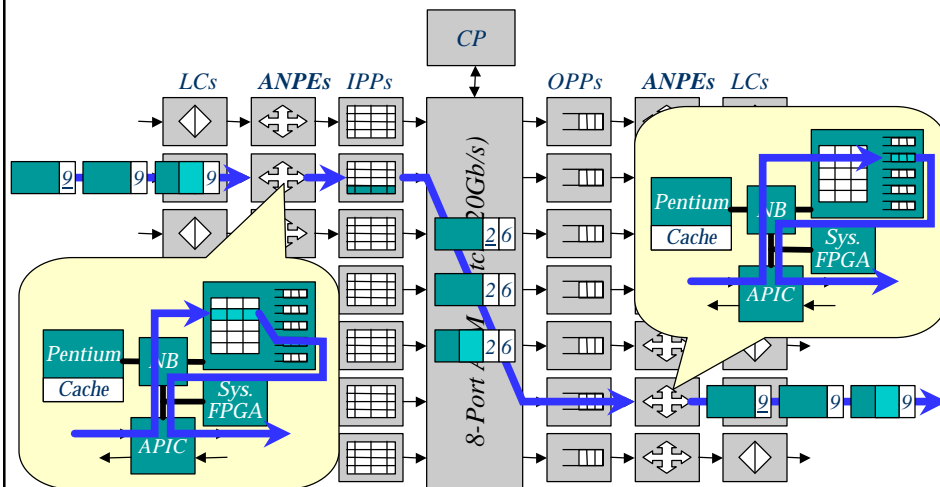
 Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

8

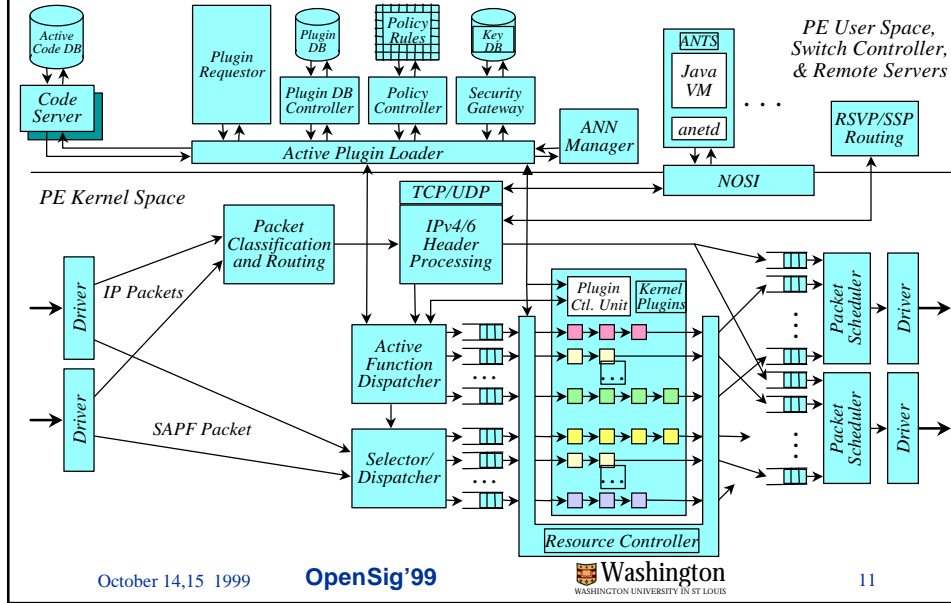
Router Plugins

- Modular kernel software architecture developed at Washington U. and ETH Zurich
- Allows for
 - dynamically adding code modules called Plugins at run-time into the IP forwarding loop
 - binding plugins to individual IP flows
- Provides fast packet classification

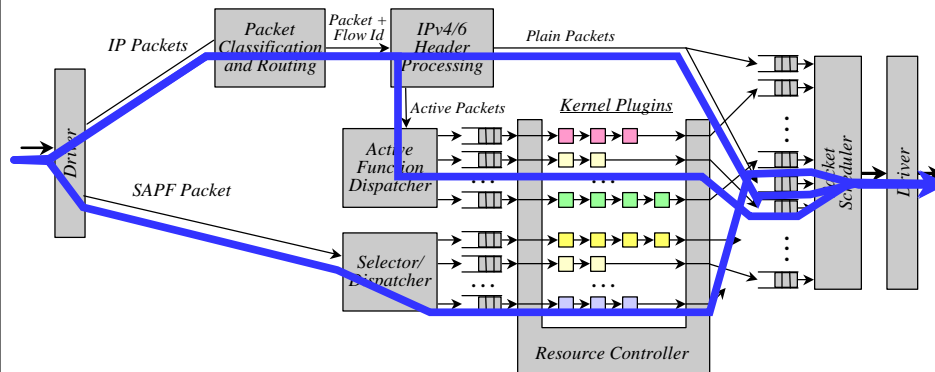
Packet Processing in the WU Active Network Node



Software Architecture



Principal Data Flows Through ANPE Kernel



- Standard processing for “plain” IP packets.
 - classification and routing, header processing, output queueing
- Active packets move through configured kernel plugins.
 - active function dispatcher passes packets to instances of plugin objects
 - instantiates objects or triggers download of plugin class, as needed
 - streamlined processing of SAPF packets using pre-established state

October 14,15 1999

OpenSig'99

Washington WASHINGTON UNIVERSITY IN ST. LOUIS

12

Active Congestion Control for Video

- Based on wavelet-based encoding method (WaveVideo)
 - up to hundreds of distinct frequency sub-bands (33)
 - low frequencies: general image definition
 - high frequencies: image details
- Active WaveVideo plugin
 - detects congestion
 - discards packets in high frequency sub-bands first
- Better performance than layered multicast
 - requires large number of multicast groups
 - reacts too slowly to congestion (several seconds)

October 14,15 1999

OpenSig'99

Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

13

WaveVideo Scaling



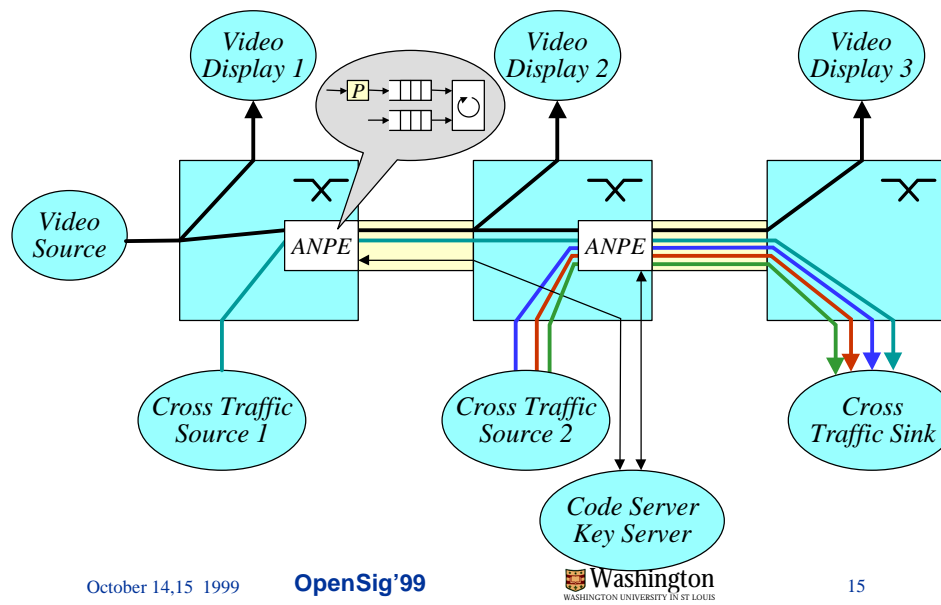
October 14,15 1999

OpenSig'99

Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

14

Demonstration Network Configuration



Results

- 500 ms to download and install plugins
- 22.3 us to scale individual packet
- 45,000 packets can be scaled per second
- 360 Mbit/s throughput (1KB packets)
- + 15 dB PSNR active vs non-active



- original no congestion ctrl Active CC

October 14,15 1999

OpenSig'99

Washington
WASHINGTON UNIVERSITY IN ST. LOUIS

16

Experiences and Reflections

- Us: Netbsd, Rest of world: → Linux
- Mbufs, Gates, Flow Ids, Filters, etc.
 - lots of details
- Performance
- Security
 - Per packet vs. Per plugin
- Applications
 -