Internet Protocols for Network-Attached Peripherals

Steve Hotz, Rod Van Meter,
Greg Finn
{hotz,rdv,finn}@isi.edu

NASA GSFC MSST & IEEE Symp. MSS
March 25, 1998

Goals of This Talk

Show that:

- IP is a good match for a network-attached peripheral (NAP) wire protocol
- IP can be fast and efficient, and easily implemented ***in a NAP***

Outline

- Netstation project
- Protocols for NAPs
  - Networking problems
  - IP as solution
- Implementation
  - Issues
  - Experience
- Conclusions

A Netstation

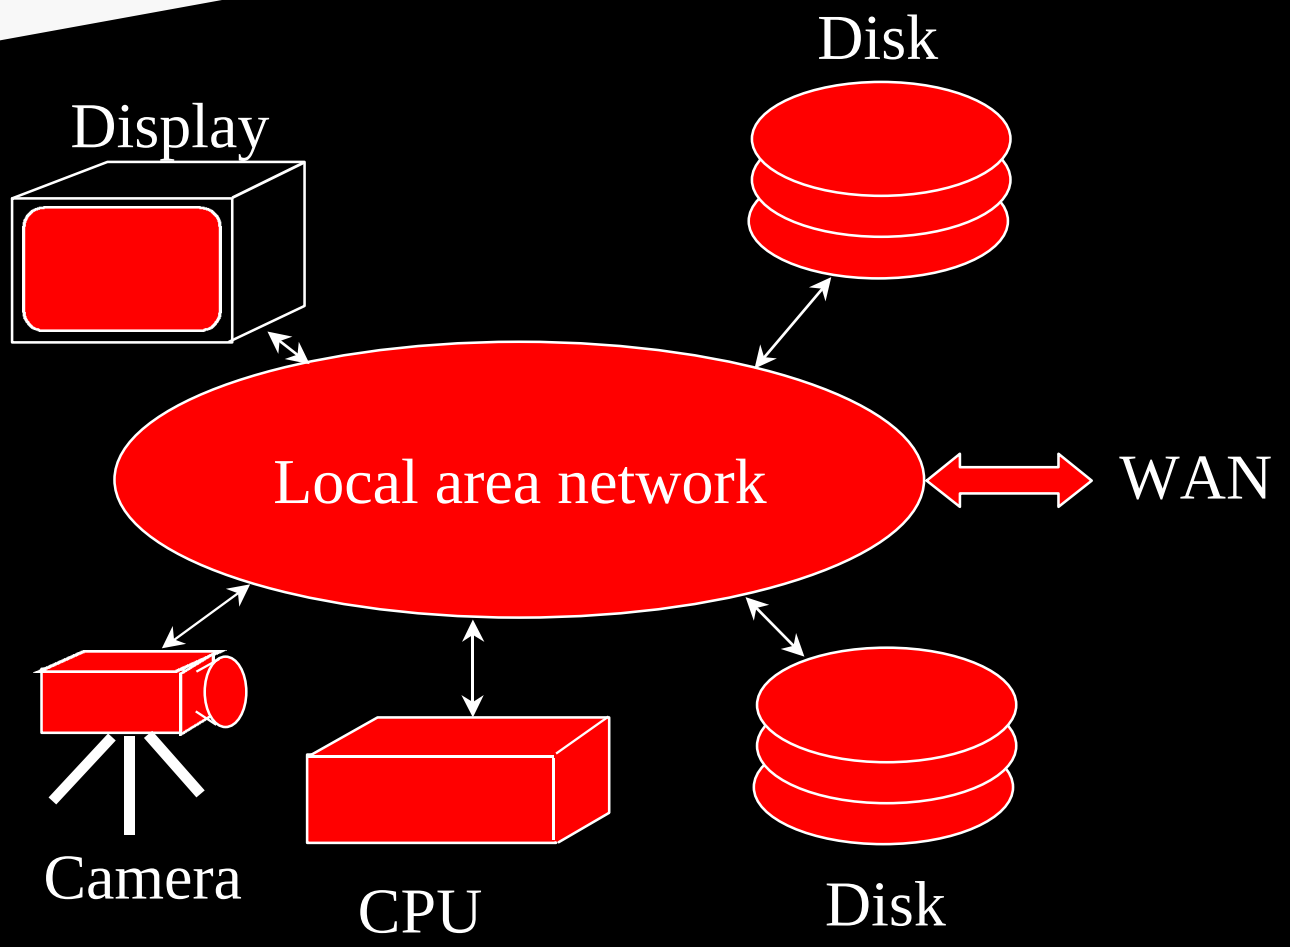

The Netstation Project

- Replace I/O bus with a gigabit network
- Buses not scaling in:
  - # devices connected
  - aggregate bandwidth
  - distance

Netstation Research Areas

- Network protocols
- Network-Attached Peripheral (NAP) hardware & firmware
- NAP-capable OS
- Security and access models

Netstation Protocols

- Originally, custom DTP over Myrinet
- Moved to TCP/IP and UDP/IP
- Disk: SCSI cmds over UDP
- Display: TCP/IP

Netstation Theses

1) IP good as wire protocol
2) NAP IP implementation can be efficient
3) Host IP implementation can be efficient
4) Security-enhanced derived virtual device is a good NAP access model

...concentrate on 1 & 2 here

Outline

- Netstation project
- Protocols for NAPs
  - Networking problems
  - IP as solution
- Implementation
  - Issues
  - Experience
- Conclusions

Networking

ISO Model	Netstation
Application	SCSI
Presentation	
Session	
Transport	TCP, UDP
Network	IP
Link	ethernet, Myrinet
Physical	ethernet, Myrinet

Networking Issues for NAPs

as I/O nets get larger & more complex...

- Media bridging
- Congestion control
- Flow control
- Demultiplexing
- Legacy systems

- Latency
- Security
- Reliability
- Heterogeneity

...all become bigger problems!

Net Technologies for NAPs

- HiPPI
  - 800
  - 6400
- SSA
- Fibre Channel
  - fabrics
  - arbitrated loop
- ATM
- 1394 (FireWire)
- Myrinet
- Switched SCSI (SPI)
- ESCON
- Gigabit ethernet
- ServerNet
- WAN channel extenders

Solution: IP

Strengths:

- Heterogeneity
- Scalability

Weaknesses:

- Complexity
- Performance

Complexity

- You get what you pay for
- IP very simple
- Support functions minimal

Functionality

Baseline

- IP
- ICMP
- ARP
- TCP and/or UDP

Optional

- DNS
- RARP
- SNMP

Performance

- Wire performance
- Node OS/architecture/app issues
- Node protocol implementation

Transport

- Fast path TCP is efficient
- UDP available as low-overhead alternative
- Other transport protocols possible

Outline

- Netstation project
- Protocols for NAPs
  - Networking problems
  - IP as solution
- Implementation
  - Issues
  - Experience
- Conclusions

Implementation Issues

- Per-packet costs
- Data-touching operations

...both can be solved

Per-Packet Costs

- Header processing
- Interrupt handling
- Task switching

Answers:

- Integrated Layer Processing (ILP)
- Bigger packets
- Lightweight embedded OS

Data-Touching Operations

- Checksumming
- Data copies

Answers:

- Hardware assist for checksum
- Zero-copy TCP/IP stacks
  (easy in embedded, real-memory OS)

Implementation Experience

- TCP/IP built for Netstation display NAP
- Based on INRIA user-level TCP
- Modest implementation effort
  - a few man-months
  - IP, ARP, ICMP + embedded OS work
- Device drivers hardest part!

Performance

- 200 instructions per packet through IP & TCP fast path receive
  - plus <200 for OS & device driver
- Estimate >100K pkts/sec.
  - zero-copy stack, HW checksum, 50 MIPS
- Hardware/driver problems prevented good bandwidth measurements

TCP/IP Code for Display NAP

Component	C Code (lines)	Memory (bytes)
ICMP	90	184
ARP	570	2804
IP	1210	6444
TCP	2700	12636
shared	390	1248
user lib	630	3908
TCP/IP Total	5590	27224

Conclusions

- IP solves networking problems for NAPs
- TCP/IP can be implemented efficiently in a NAP
- TCP can be efficient, but transport issues still open
- http://www.isi.edu/netstation/