

The End of HPC

Patrick Geoffray
patrick@myri.com

07 September 2010
CCGSC – Flat Rock, NC

The end of an era

- Founded in 1994, spin-off from Caltech.
- Original DARPA project, technology from Caltech Mosaic and Cosmic Cube, licensed in Intel Delta, Paragon.
- Myrinet network fueled the Cluster Computing expansion:
 - Performance, Scalability, Cost.
- Firmware-based architecture:
 - Flexibility, constant innovation.
- Convergence with 10 Gb/s Ethernet in 2005:
 - Open door to enterprise data center.
 - Enterprise customers with well-defined problems, willing to pay to solve them, pragmatic.
- Slow move into high-performance data-center networking.



**SO LONG
&
THANKS FOR ALL THE FISH**

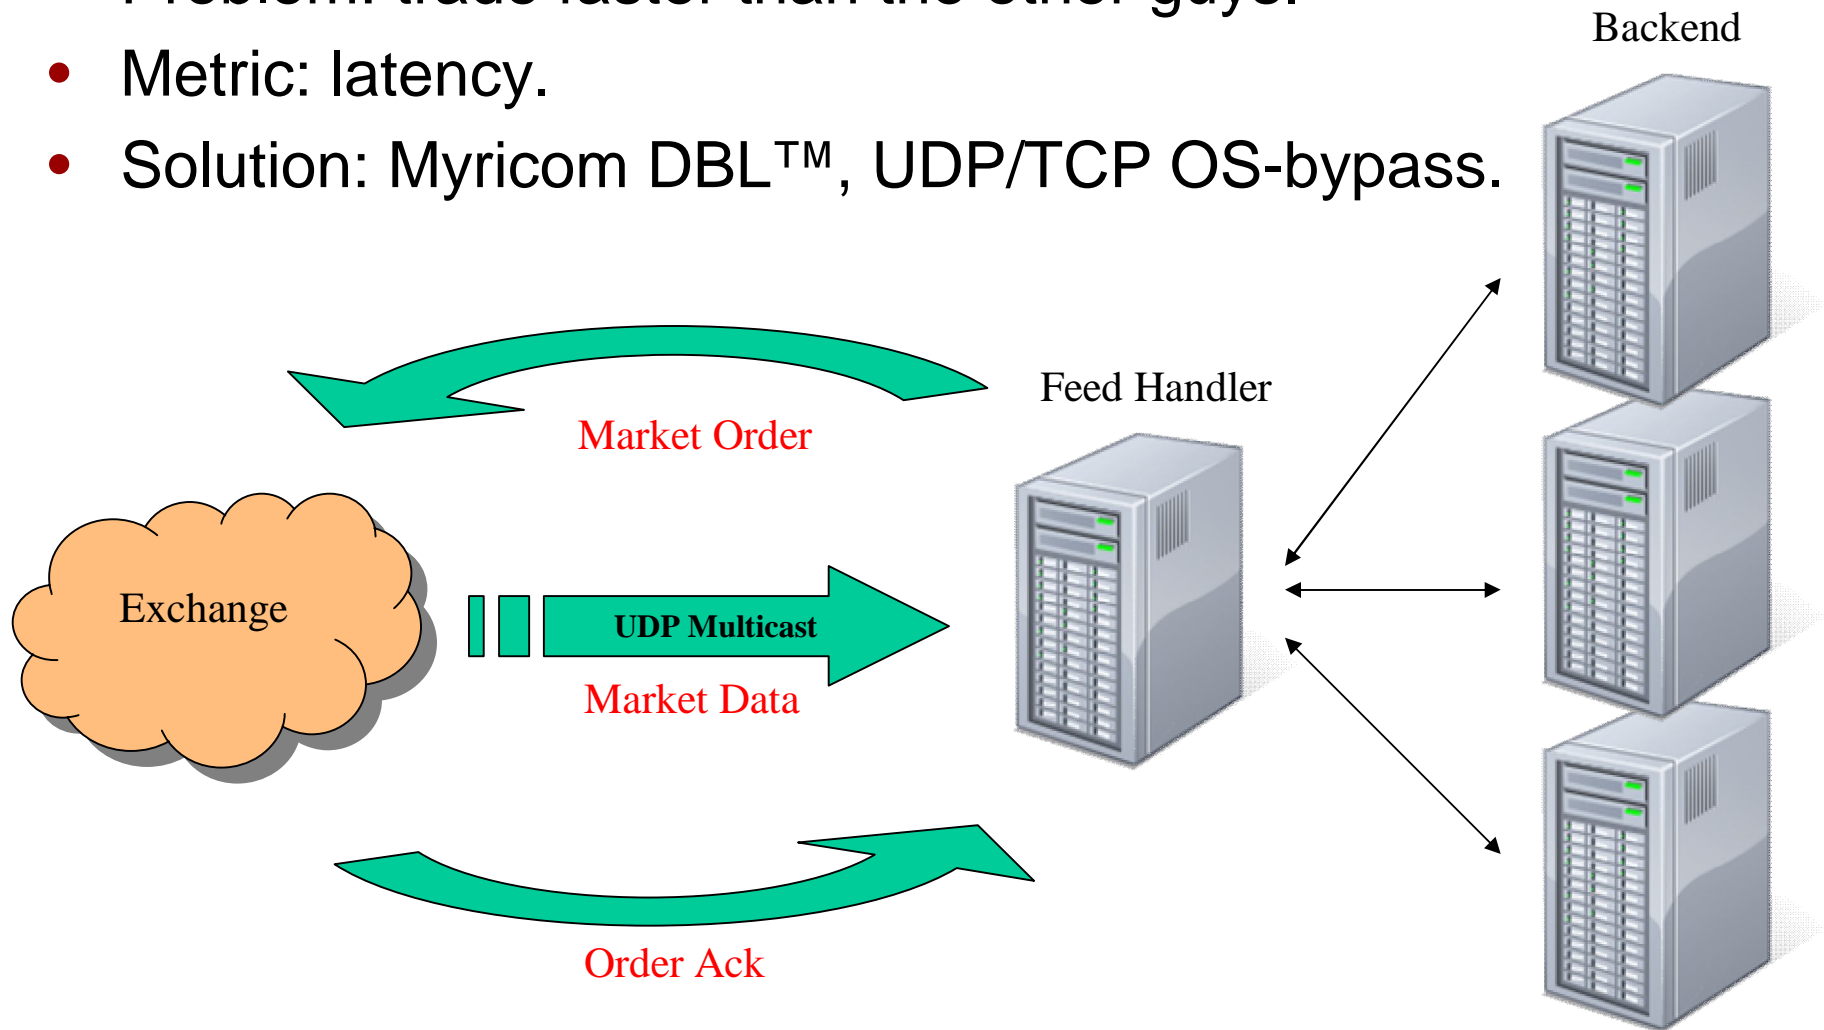


High-Performance Ethernet


- Specialized NIC firmwares for vertical markets:
 - Financial trading
 - Video streaming, IPTV
 - Network security
 - Voice-over-IP
 - Virtualization
 - Storage
- Commodity hardware, added-value software.
- Not offload, processing is better done in the host.
- Leverage 15-year old techniques designed for HPC: OS-bypass, zero-copy, demultiplexing.

Financial Trading

- Problem: trade faster than the other guys.
- Metric: latency.
- Solution: Myricom DBL™, UDP/TCP OS-bypass.



High Frequency Trading



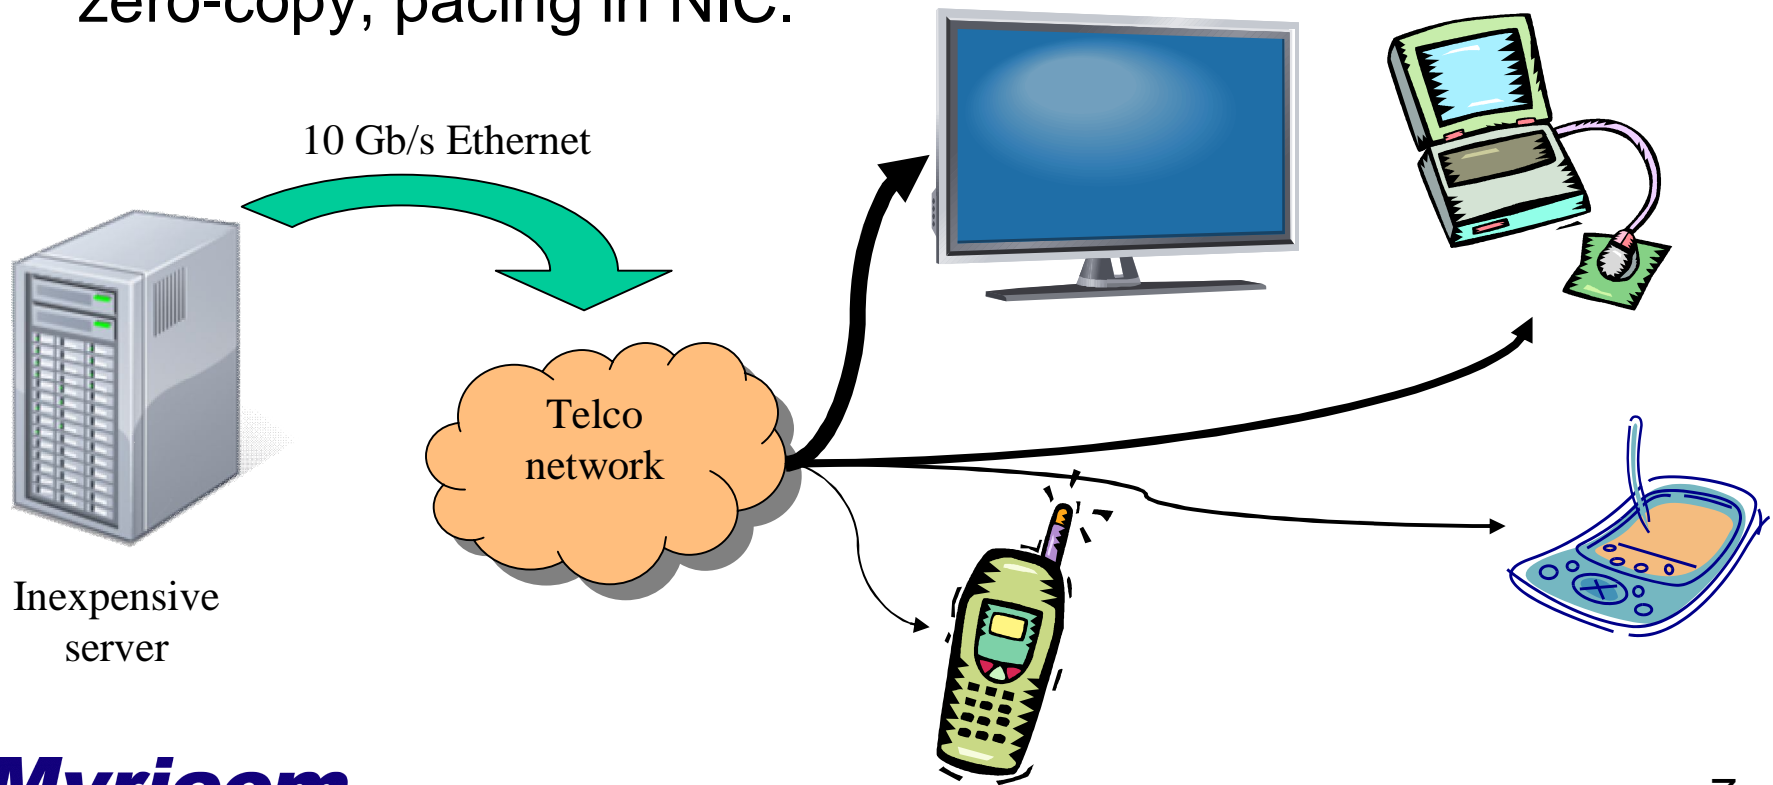
	Sell 1M XYZ	
Buy 100 XYZ @ 10.01	Buy 500K XYZ	XYZ @ 10.01
Buy 100 XYZ @ 10.02		XYZ @ 10.02
...		...
Buy 100 XYZ @ 7.00		Buy 100 XYZ @ 12.01
Buy 480K XYZ @ 8.00		Sell 480K XYZ @ 12.00



Friendly HFT Algo



Video Servers (IPTV, VOD)

- Problem: stream video (HD, SD) to large set of clients (TV, laptop, PDA, Cell), very cost sensitive.
- Metric: throughput, CPU overhead, low-jitter.
- Solution: Myricom VideoPump™, UDP OS-bypass, zero-copy, pacing in NIC.



Cyber Security

- Problem: capture network traffic, inspect it, eventually re-inject packets.
- Metric: packet rate, CPU overhead
- Solution: Myricom Sniffer10G™, line-rate packet capture and injection, OS-bypass, zero-copy.



Firewall
Intrusion Detection System
Network Monitoring
Intelligence Collection
DDoS Protection
Traffic Recording
Packet Generation

14.8 Mpps (10 Gb Ethernet, 64B pkts)

The end of HPC ?

- HPC as a R&D engine, trickle down technology transfer?
 - Nobody adopted MPI, wonder why.
 - Grid computing is still looking for a home.
- Enterprise/consumer markets leading innovation, not traditional HPC crowd.
 - Trickle in the other direction: GPU, SSD, ...
- Right motivation: money, not science.
- HPC efforts are not problem-oriented:
 - 5% efficiency on most applications.
 - Success story: D.E. Shaw “Anton” FFT machine.
- HPC funding is political, Linpack is the yardstick:
 - Japan (old) / China (new) is taking over.
 - Posturing.

Questions ? Comments ?

***Feel bad about your HPC job ?
That's ok, have a drink tonight***