WOMBAT
FINANCIAL SOFTWARE

Wombat Data Fabric: Bringing the Power of Native IB To Financial Services…

(and beyond?)

Ken Barnes
VP, Middleware Division
ken@wombatfs.com

Mike Schonberg
Senior Developer/Architect
mls@wombatfs.com

Shameless Chest Thumping

- Experts in high performance connectivity
- Focused on the “Algorithmic Trading” Market
- “Twice as fast” – Jerome Downey, MD Bear Stearns, WS&T magazine
- All of the top 13 banks are customers (US, Europe, Asia)
- 2 to 3X revenue growth YoY last 3 years

Gordon Moore Eat your Heart Out…

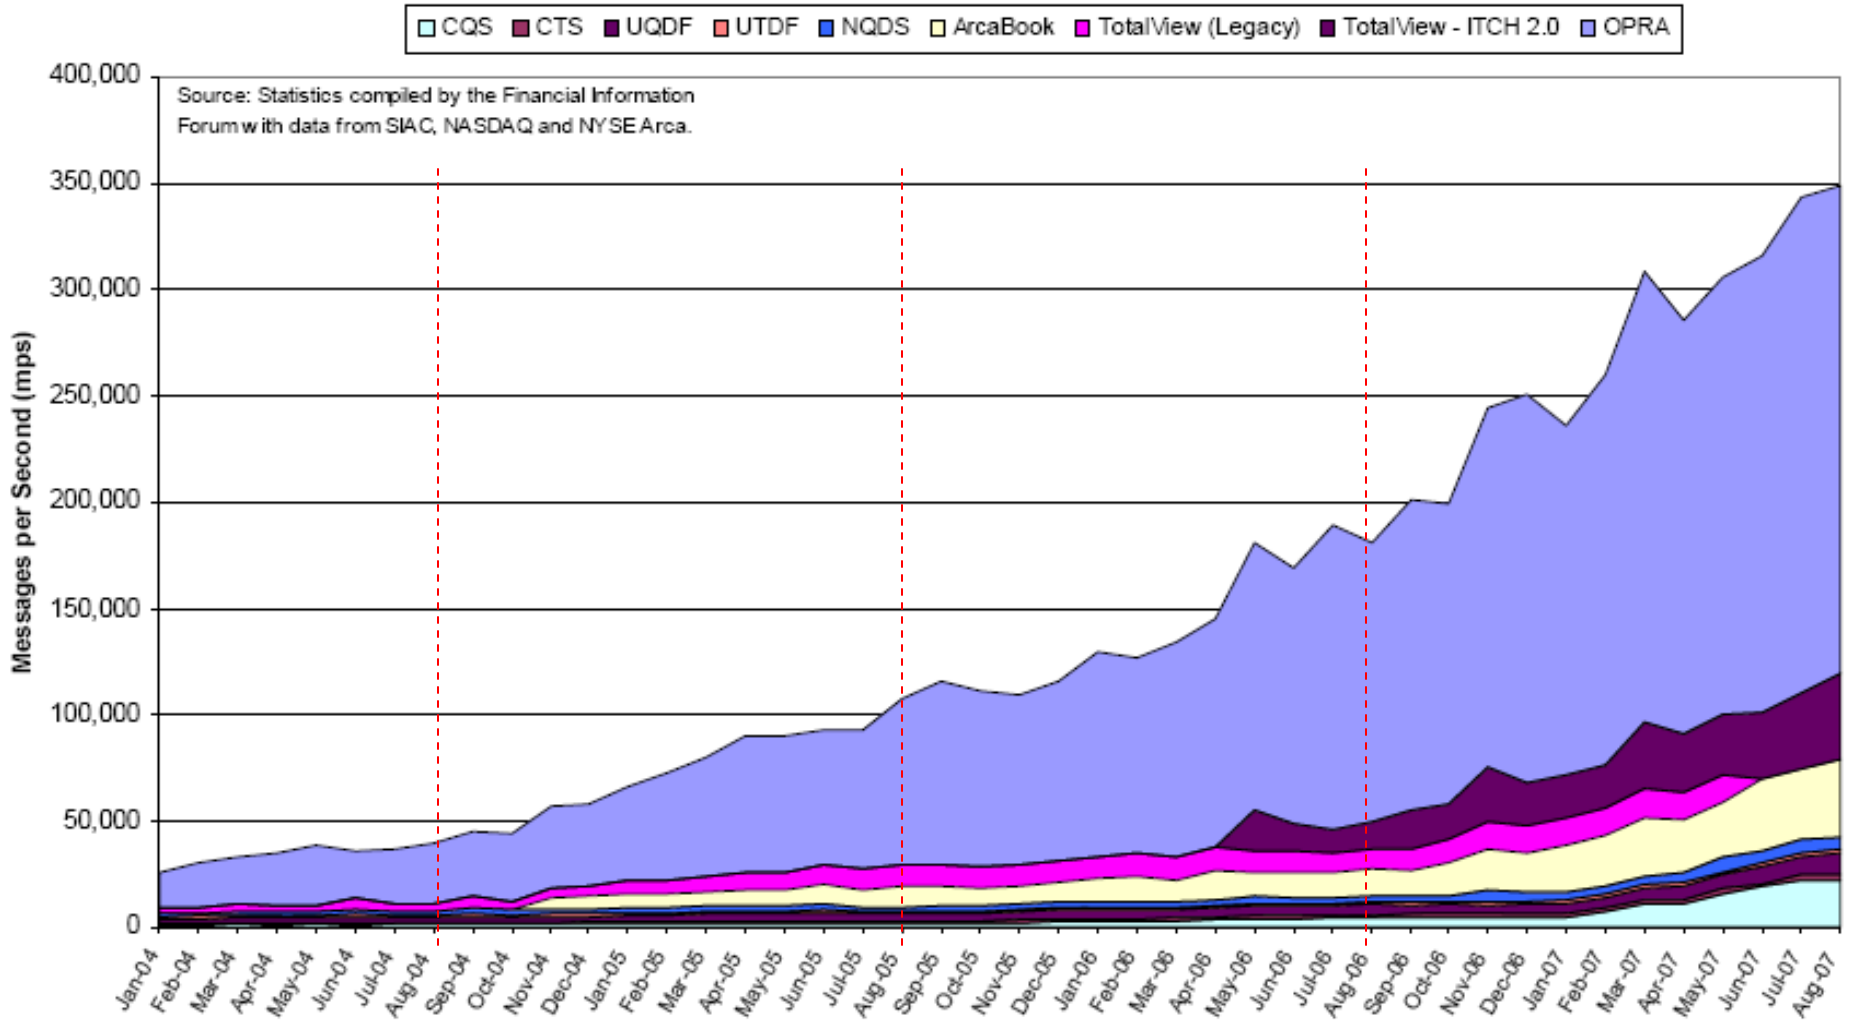

Innovation Driving Spikes

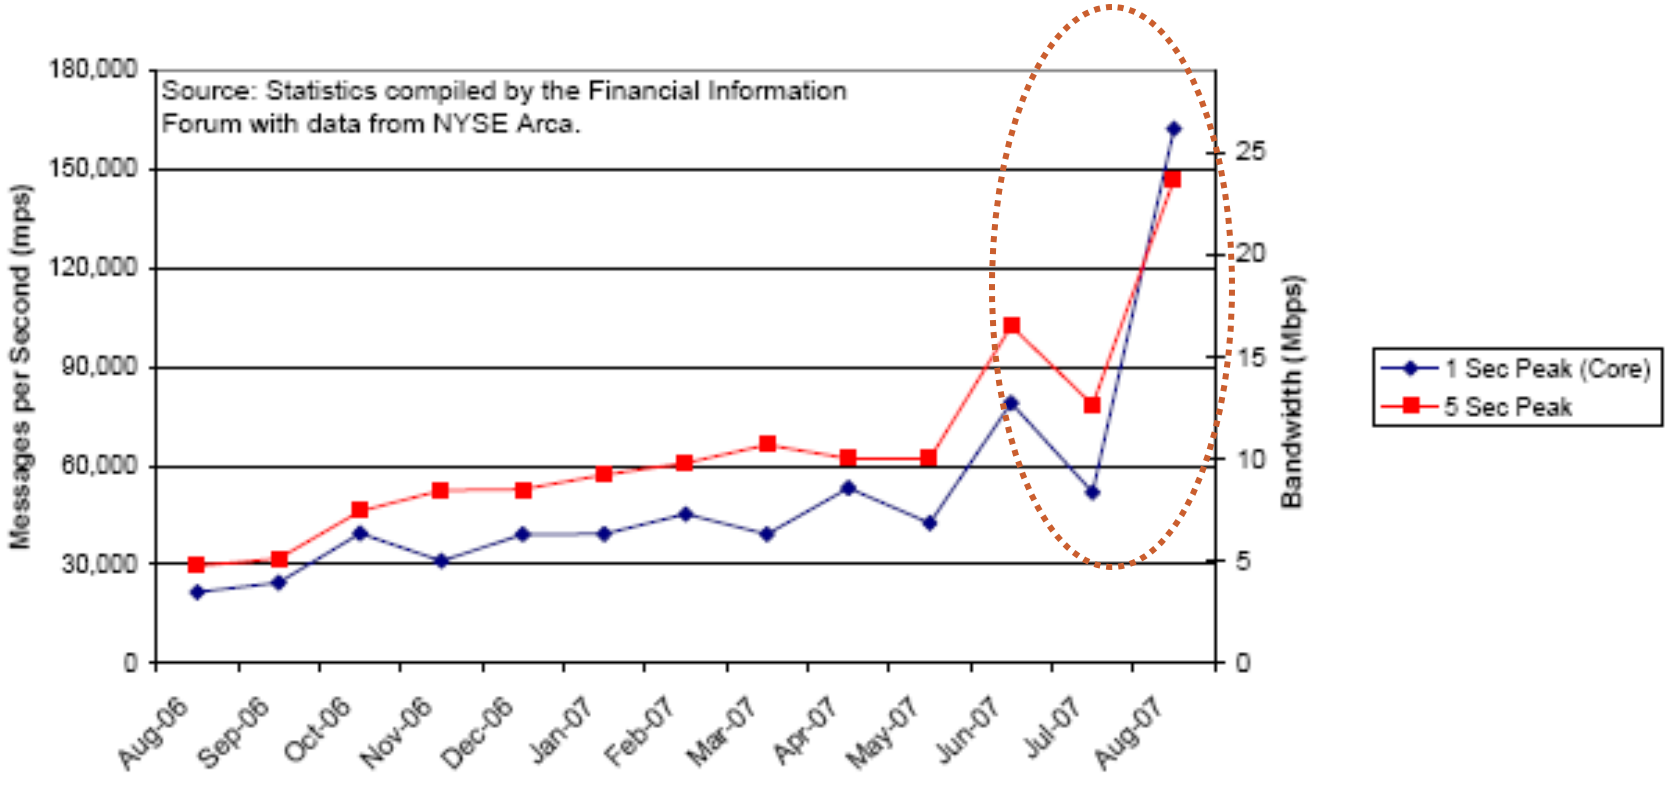

A Compounded Problem

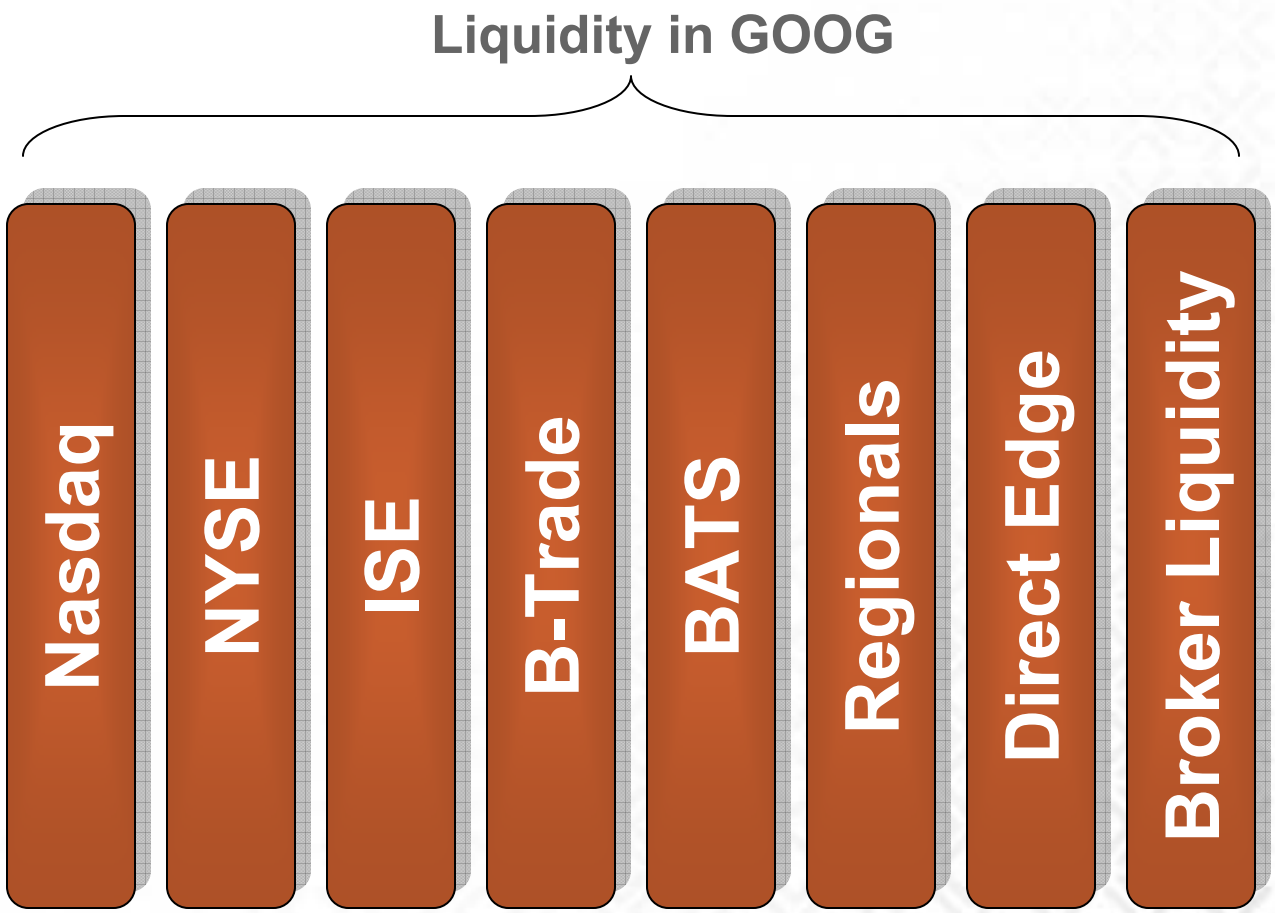

SuperBook- Generates Insights…and MORE data still

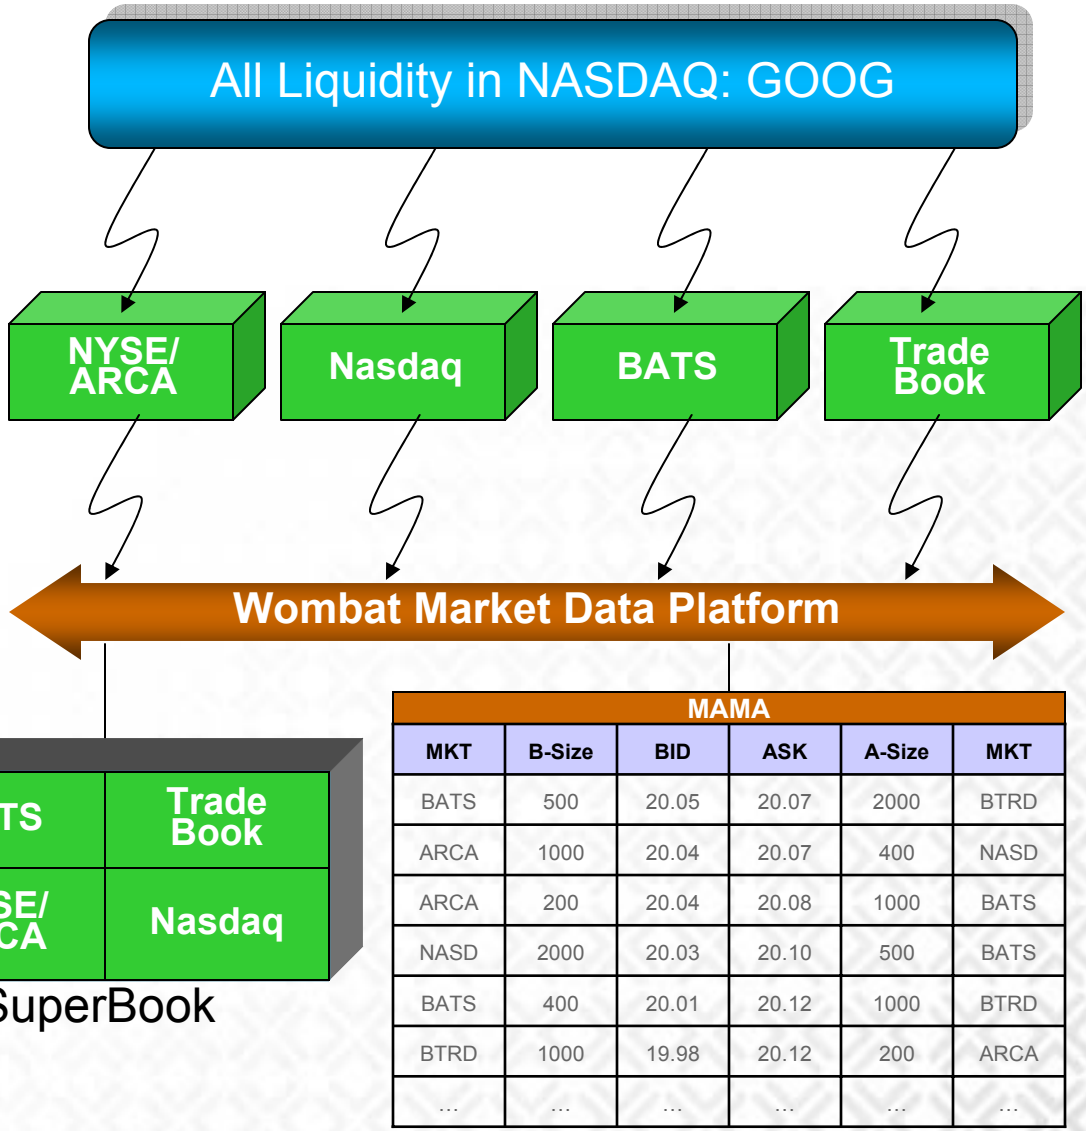

MAMA

MKT	B-Size	BID	ASK	A-Size	MKT
BATS	500	20.05	20.07	2000	BTRD
ARCA	1000	20.04	20.07	400	NASD
ARCA	200	20.04	20.08	1000	BATS
NASD	2000	20.03	20.10	500	BATS
BATS	400	20.01	20.12	1000	BTRD
BTRD	1000	19.98	20.12	200	ARCA
…	…	…	…	…	…

Net Effect: Server Sprawl

- Hardware
- Rack Space
- Power
- Cooling
- Management

Retransmissions and the Crybaby Consumer

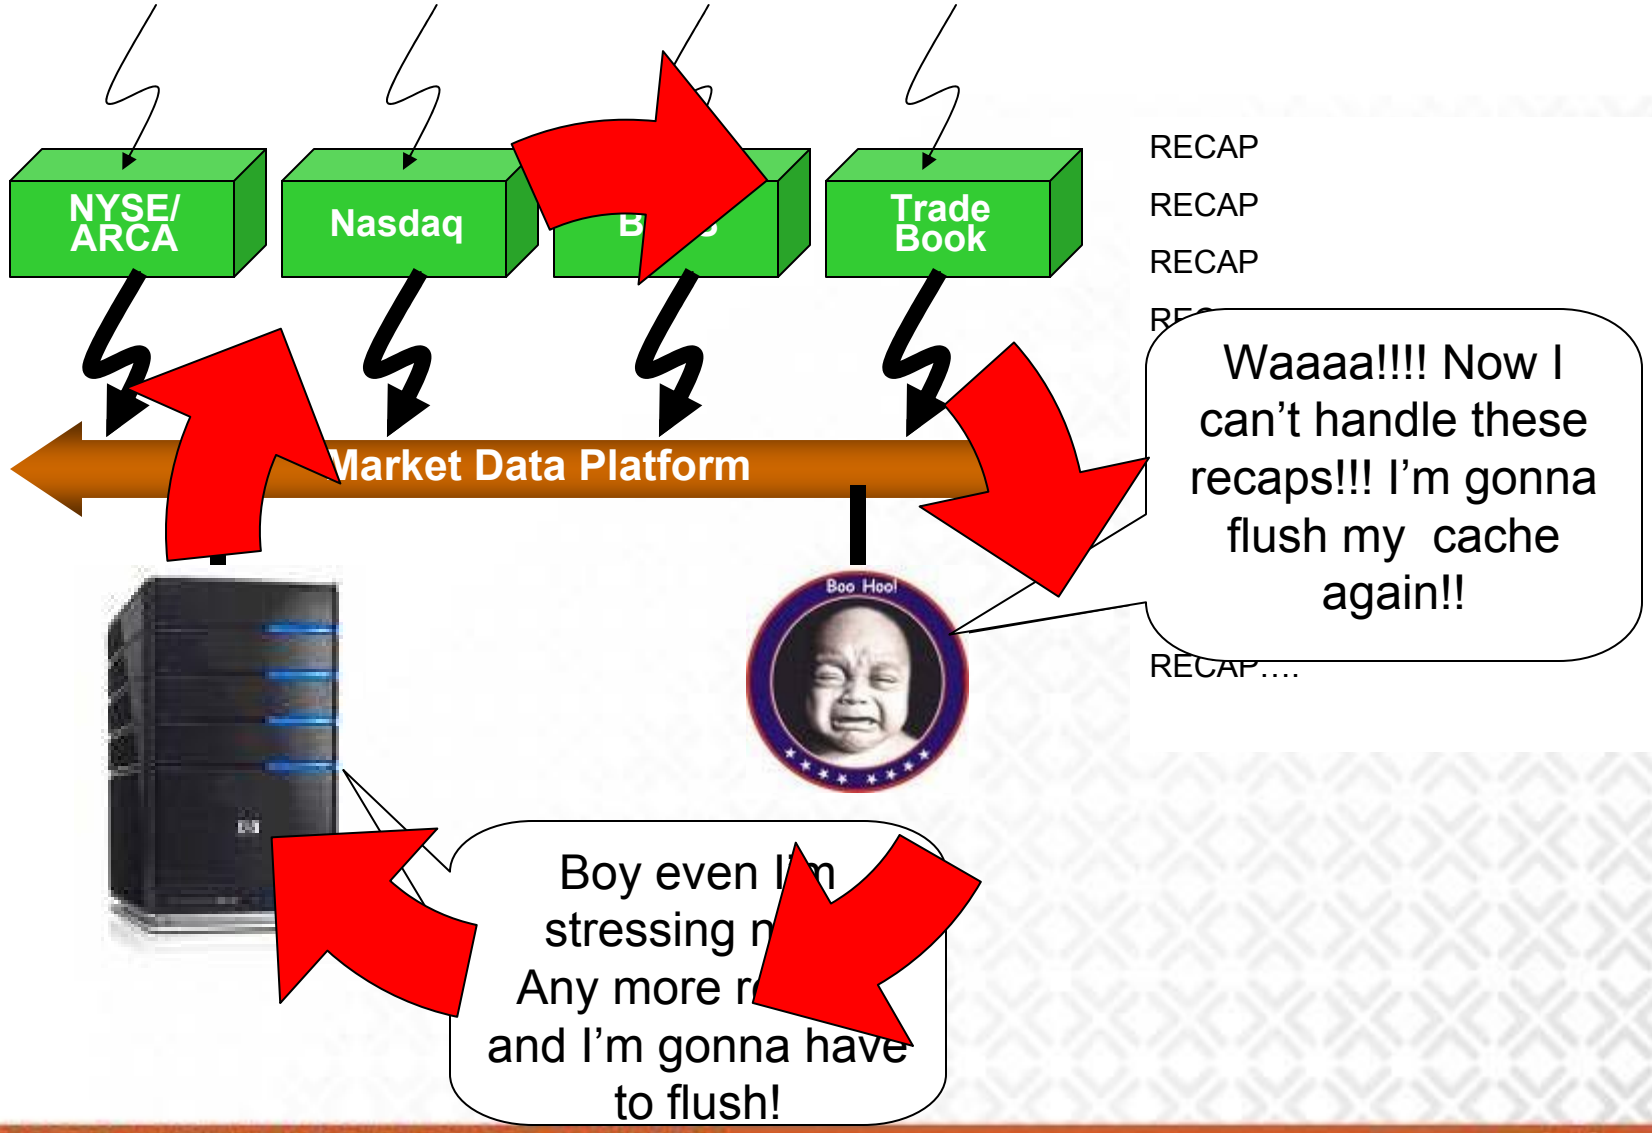

Market Data Middleware Evolution

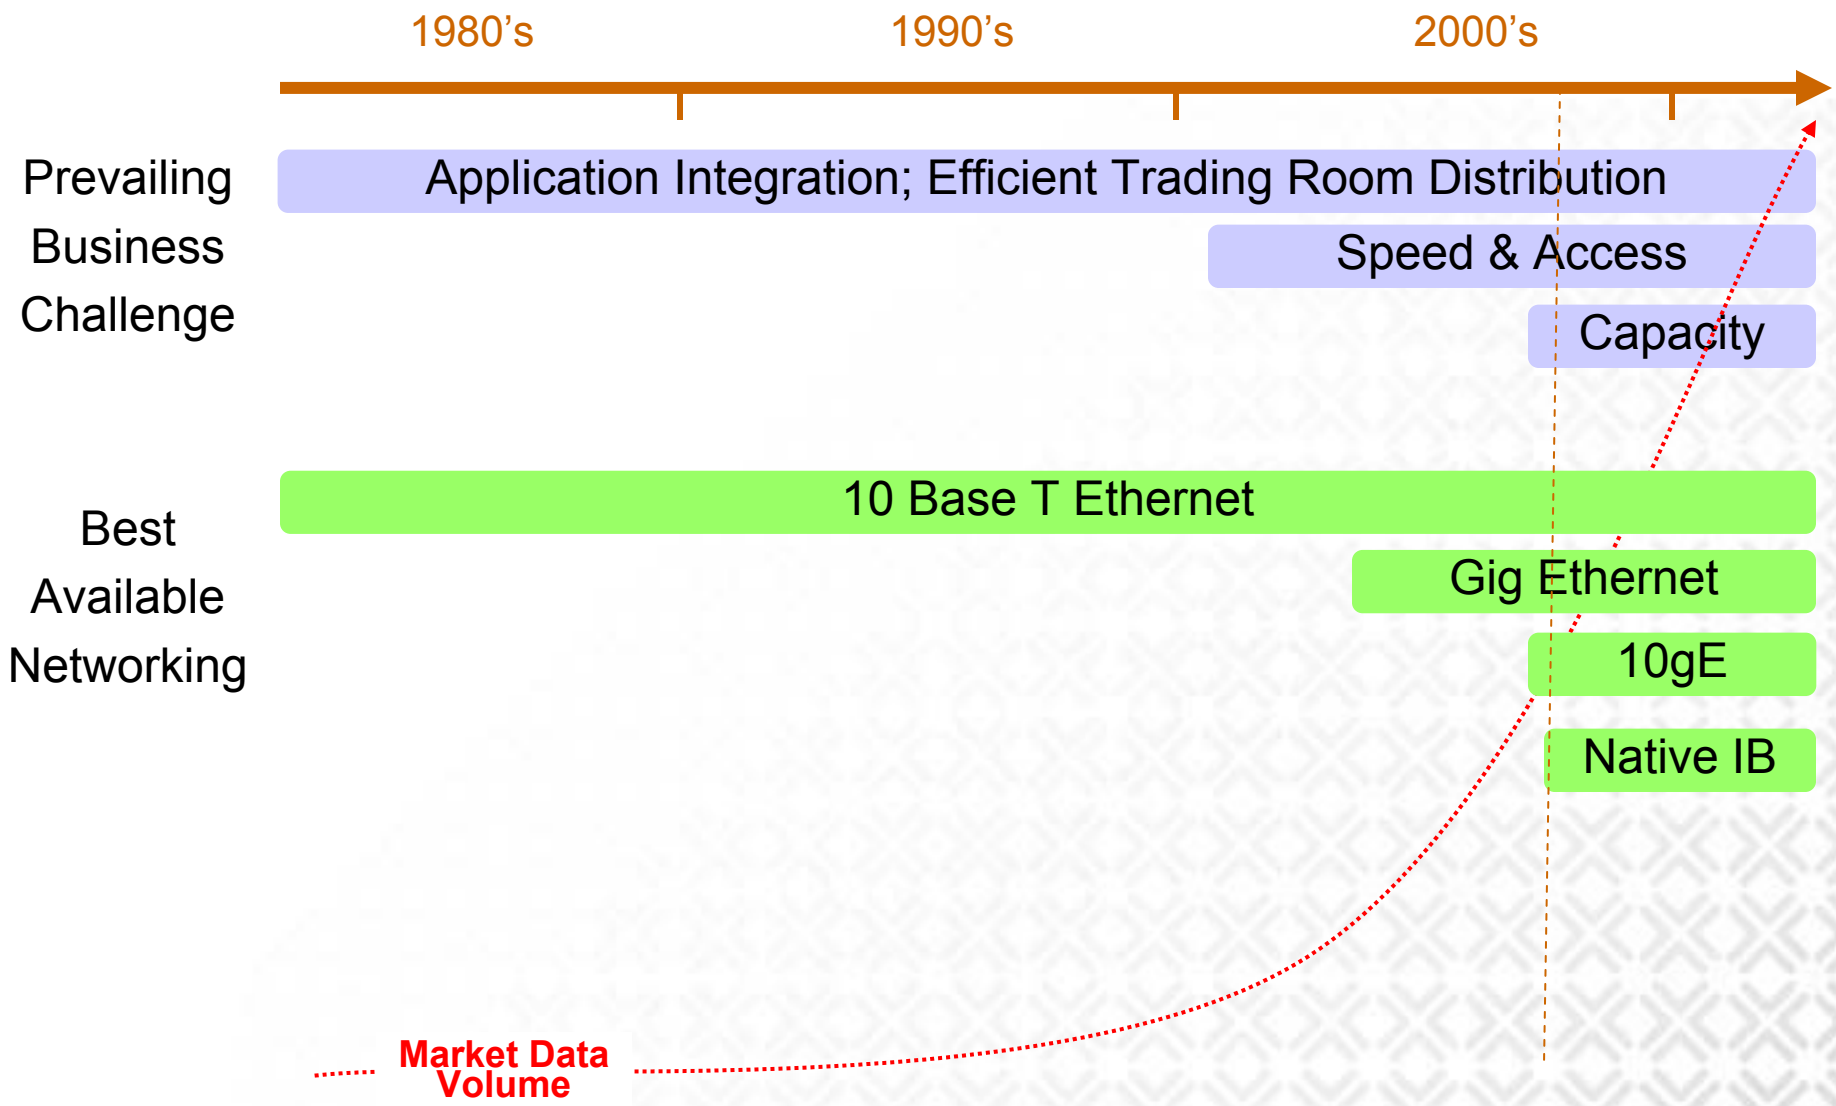

Wombat Middleware Innovation: Addressing Market Data Business Issues

Design Goals

- Deliver world class performance (*reliable* speed)
- Minimize sprawl
- Enable business agility

Compute Consumption: The Problem's in the I/O

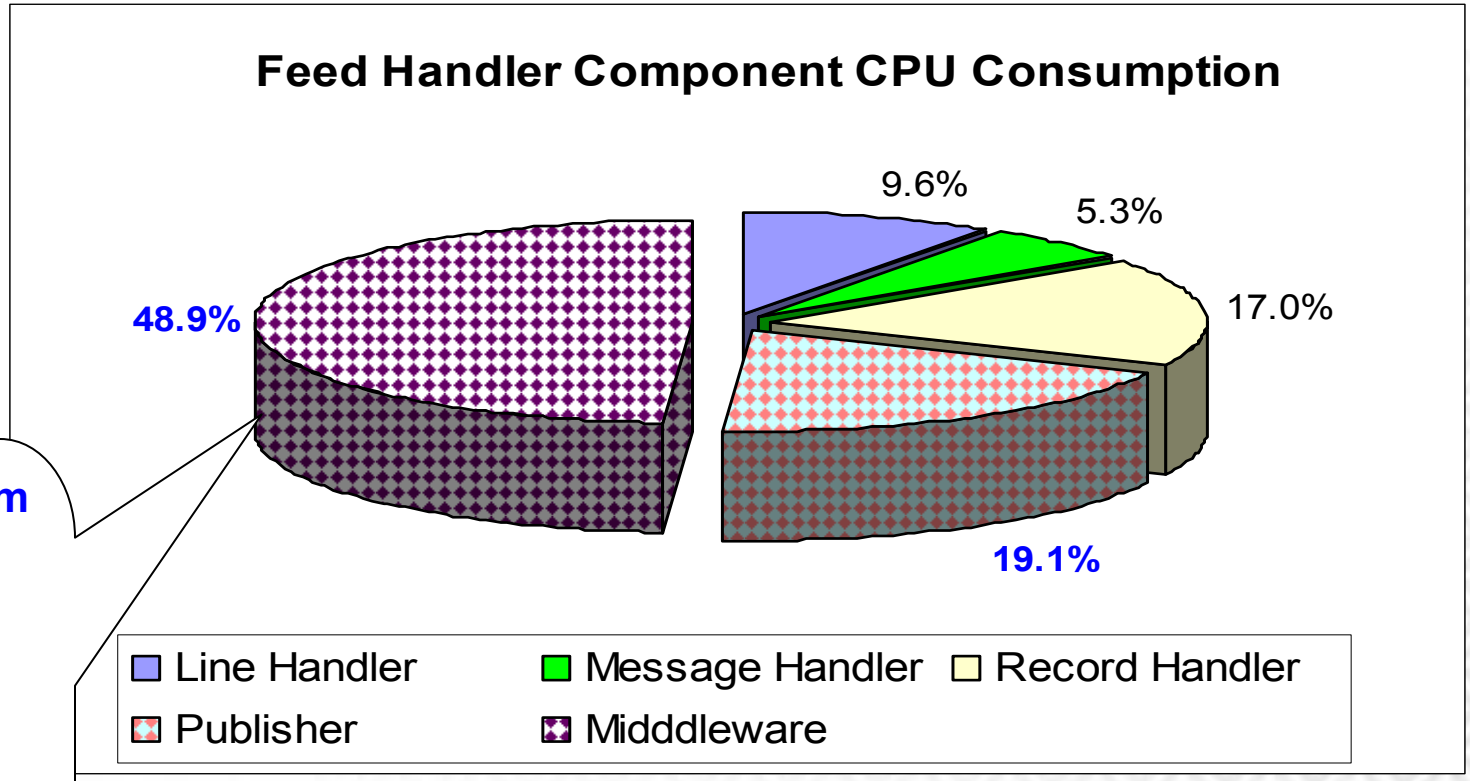

Offload-able from CPU!

- General I/O
- Unnecessary QoS code
- Multicast channel overkill

Latency Induction:
It's in the Copies and CPU Interrupts

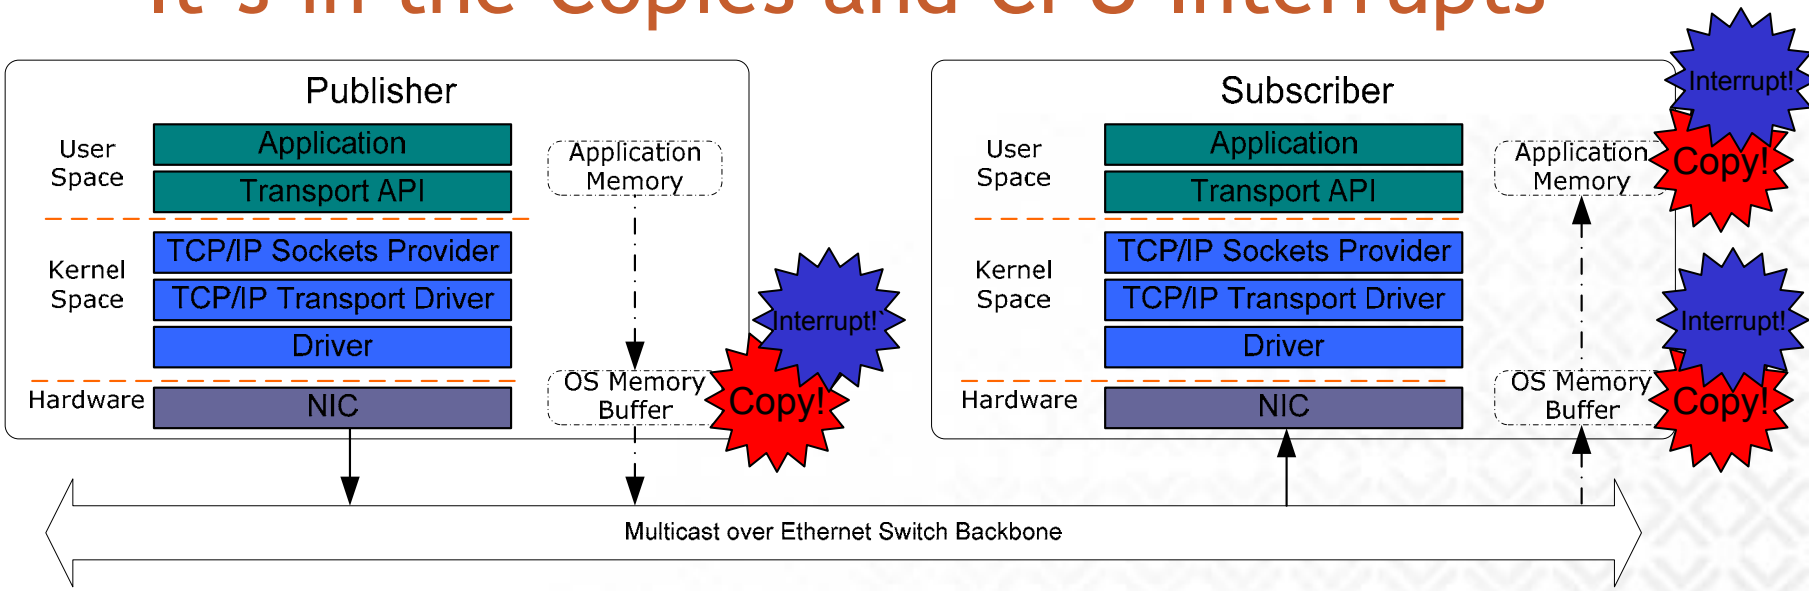

Ethernet Delivery vs RDMA over InfiniBand

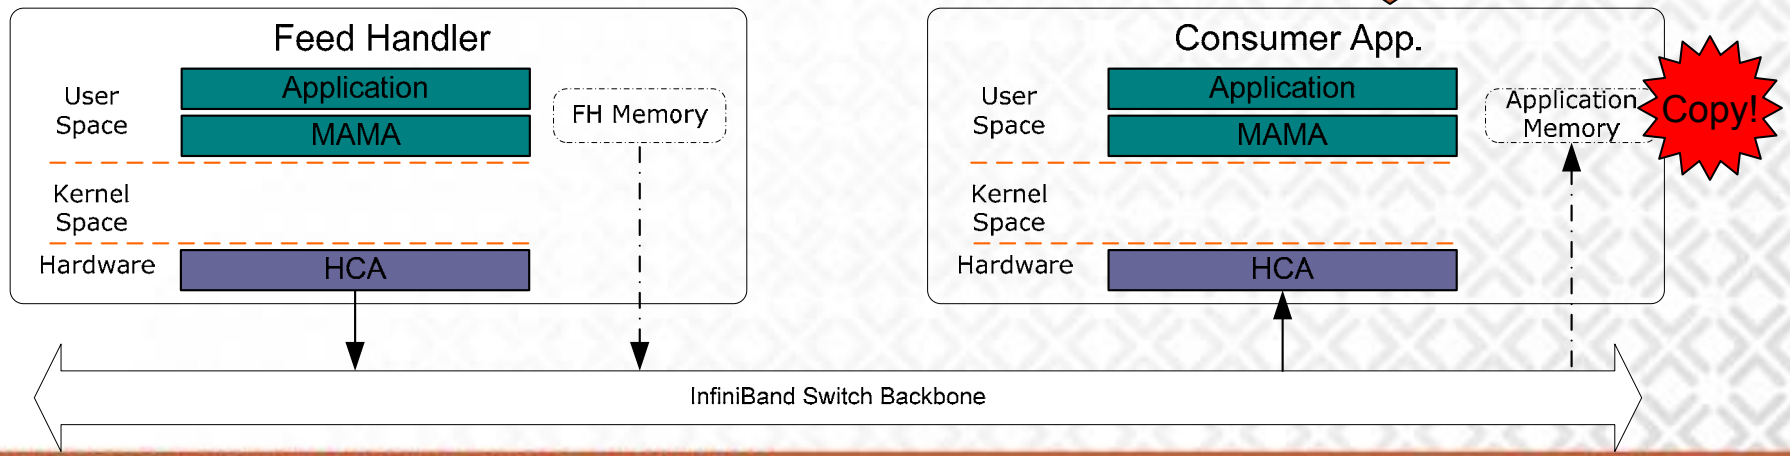

Wombat's IB Journey

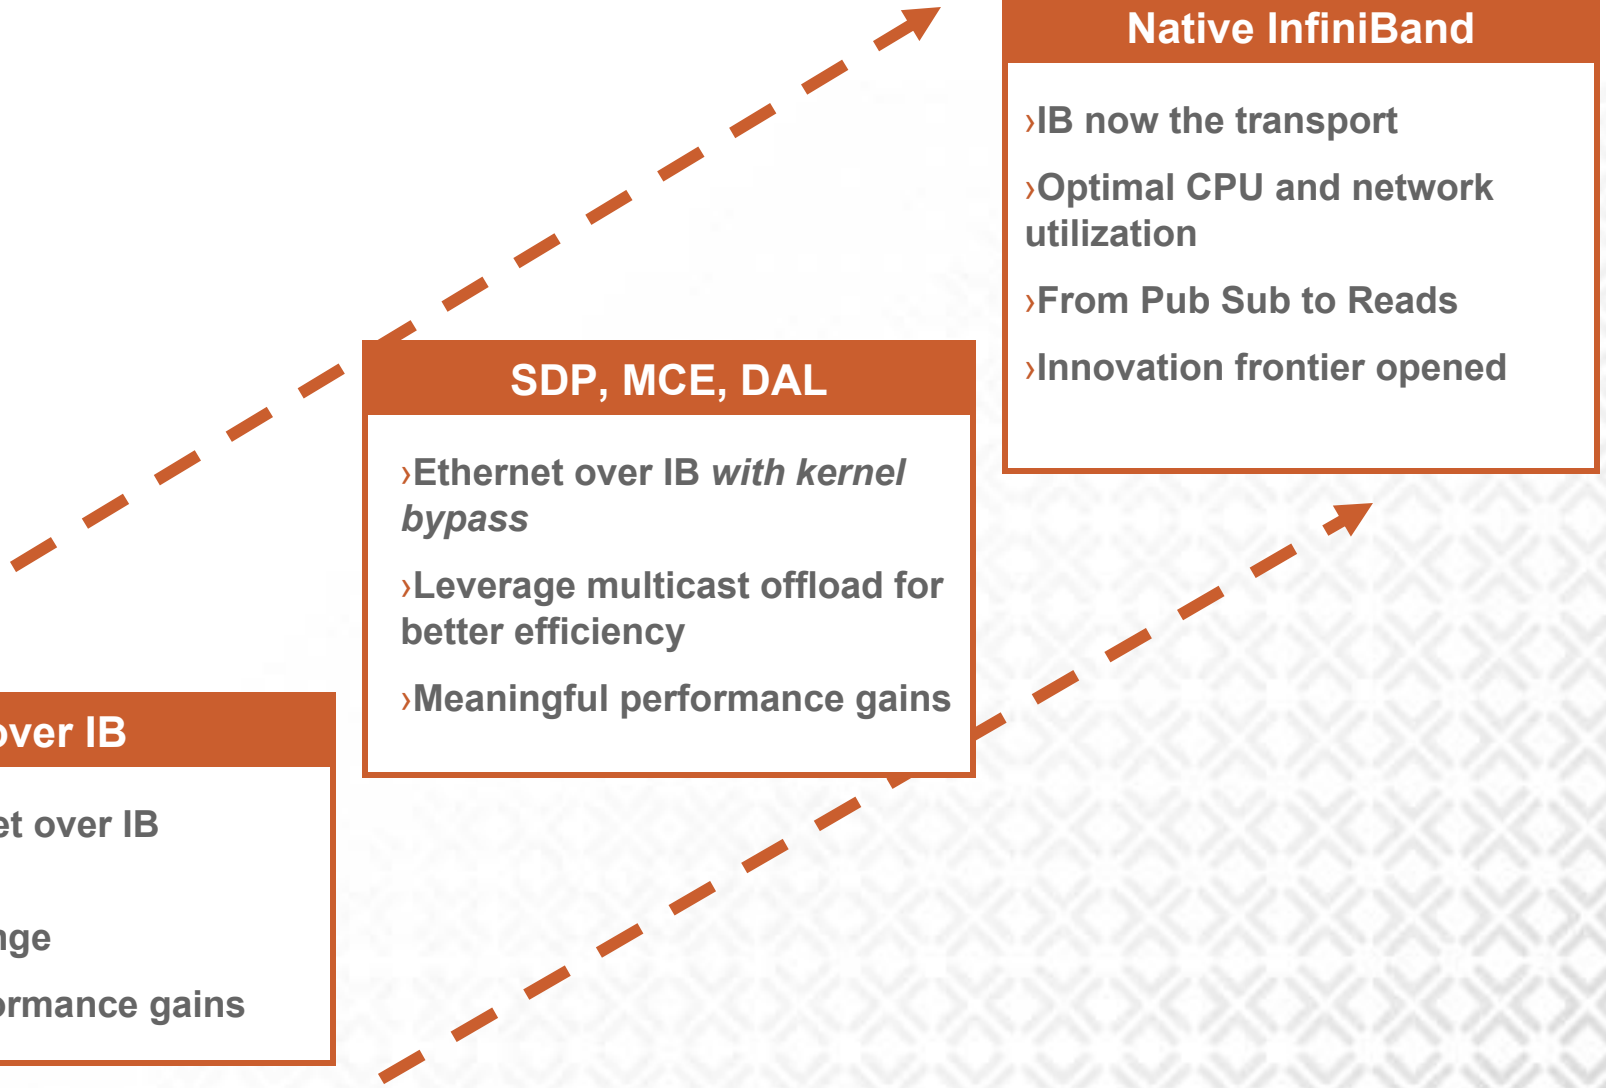

InfiniBand:
Fast, Guaranteed Delivery

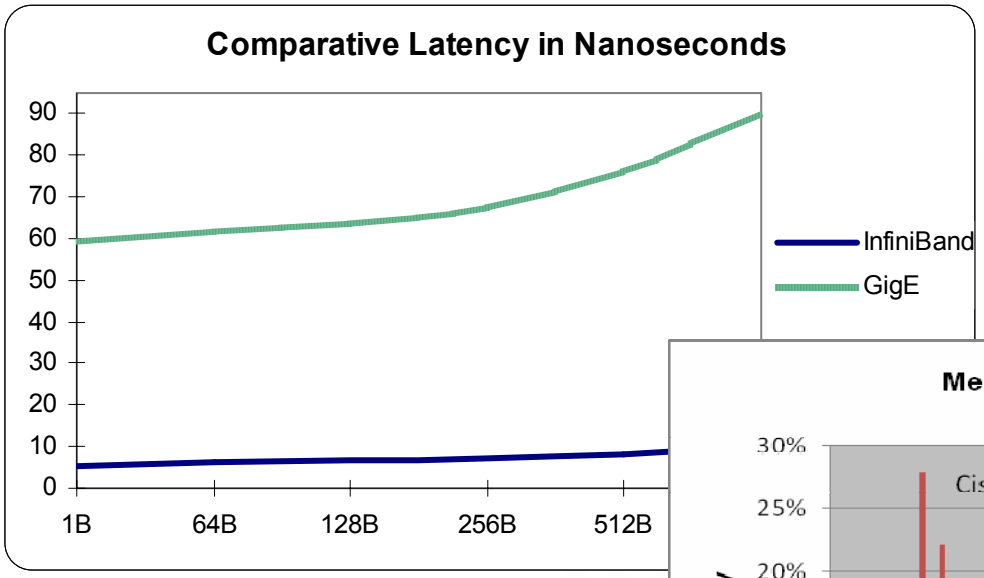

Source: NCSA

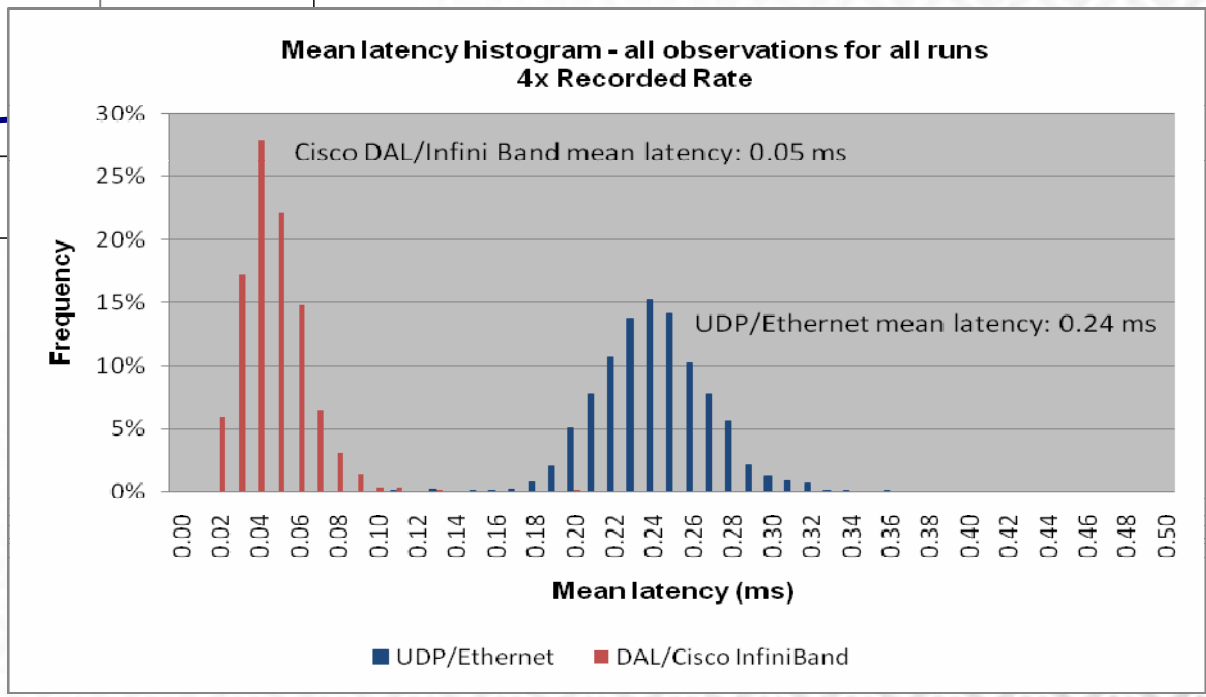

Source: STACresearch.com

Wombat Data Fabric:
The Market Data over RDMA Solution

- World's only Native IB/RDMA Middleware
- From publish to reads
- Throughput: 1 million messages/sec and beyond…
- Reduced CPU/Bandwidth consumption
  - Minimal CPU interrupts
  - Zero copies
  - I/O offload to the HCA
- No retransmissions / crybabies
- Commodity x86 servers; the flexibility of software
- New horizons for innovation

Just One 8-core 1U Server…

Can Deliver more than
500,000 messages/sec

Under 100u-secs

Sockets vs Shared Memory

Sockets - Ethernet

- Transient messages
- Retransmission handling
- Heavy context switching in the OS
- Multiple copies in the I/O stack
- 1 GB per sec
- Latency ~0.5ms
- Latency jitter (long tails)
- Saturation 625,000 messages per sec
- Streaming data required to maintain state

Shared Memory - Infiniband

- Data persisted in shared memory
- Consume at will
- OS bypass
- “Zero copy”
- 20-40 Gb per sec
- Latency ~0.065ms
- Tight latency distribution
- Saturation 10,000,000 per sec (per port)
- Zero latency snapshot
- Tick series in the fabric

Enterprise-Wide Enablement

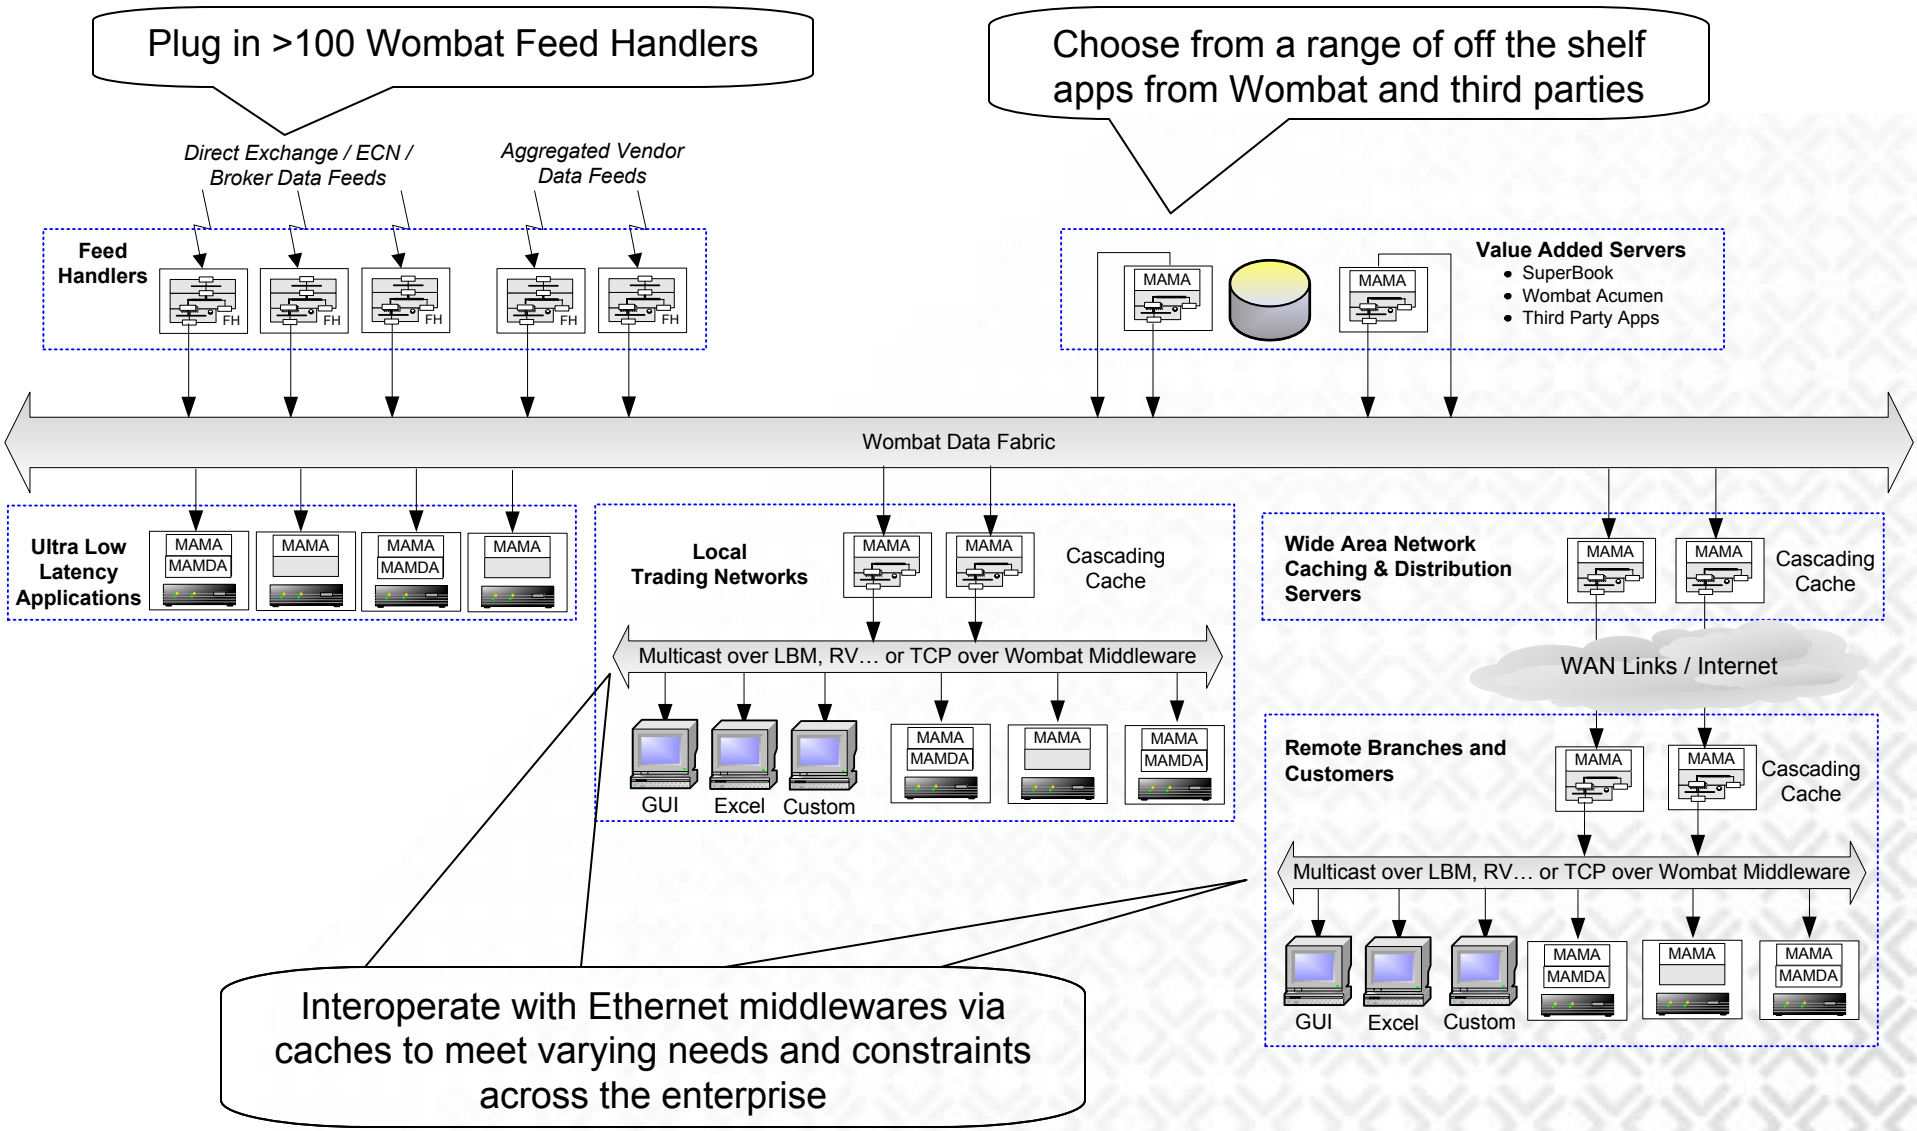

Unification via MAMA 4.0

- Binds the Wombat Platform & abstracts underlying middleware
- Low overhead
- Multi Impl.
- Dynamically loadable
- Standard programming model

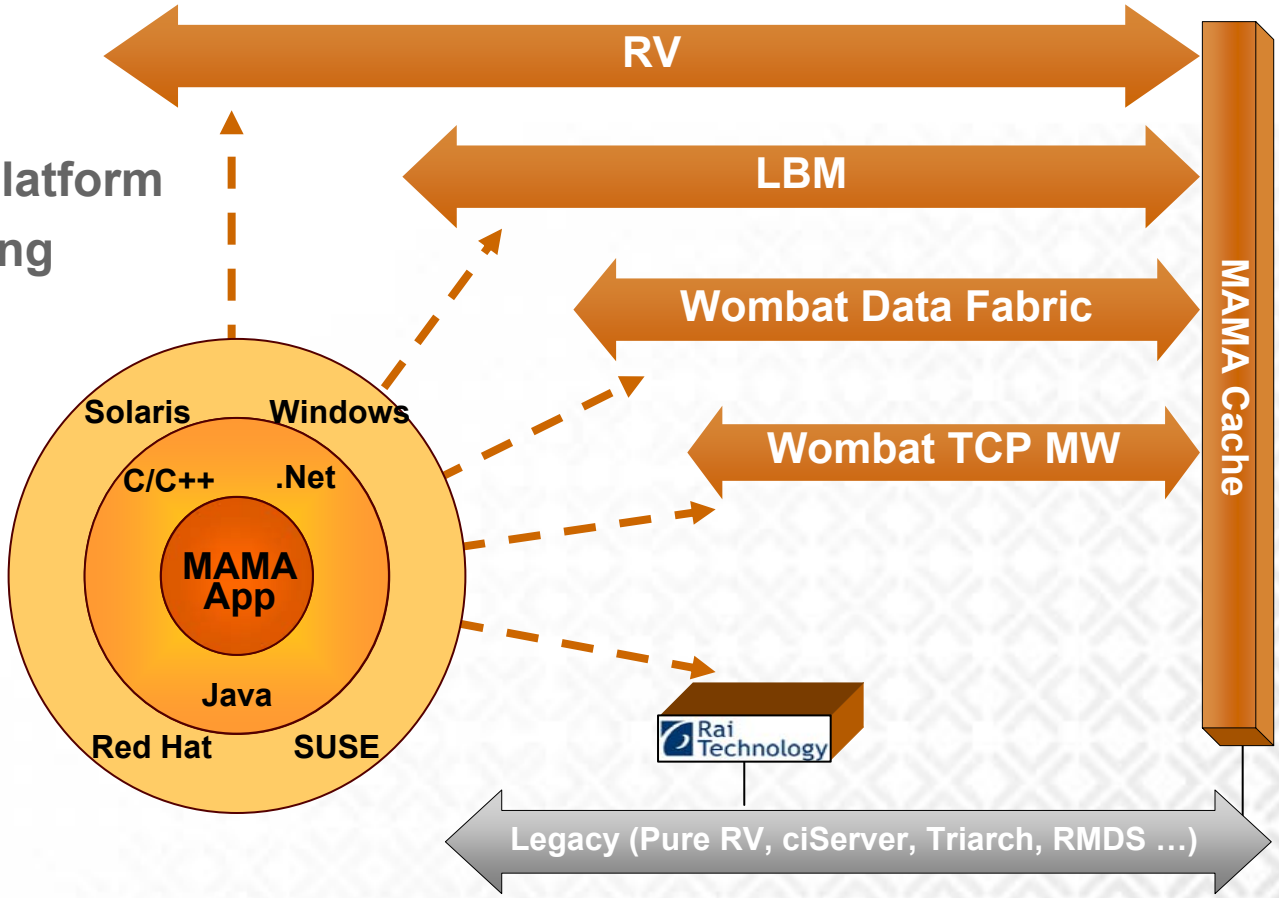

Future Innovations:
Market Data Object RDMA

Order Books with Traditional Middleware

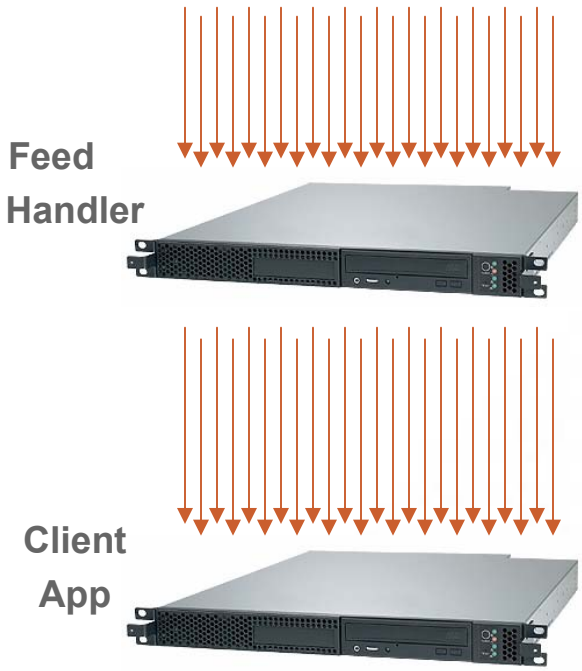

Book state must be maintained in subscriber

Order Books with Data Fabric

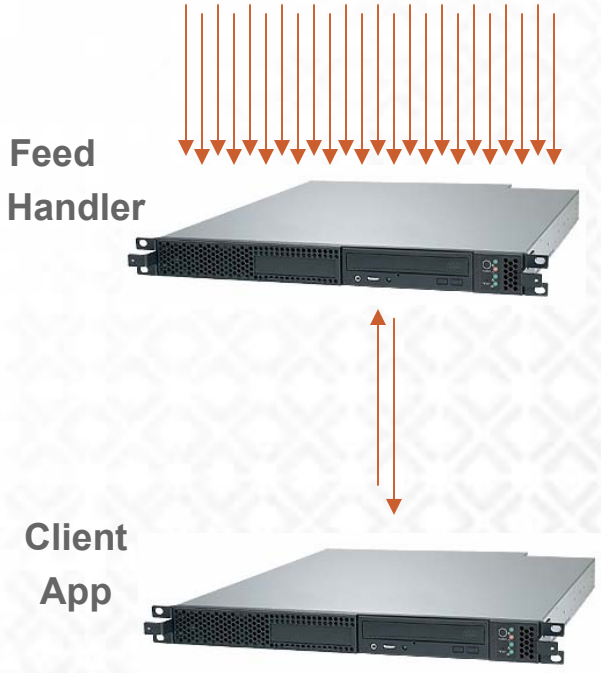

Book state maintained in feed handler & requested on demand in “zero” latency

Future Innovations:
Time Series in the Fabric

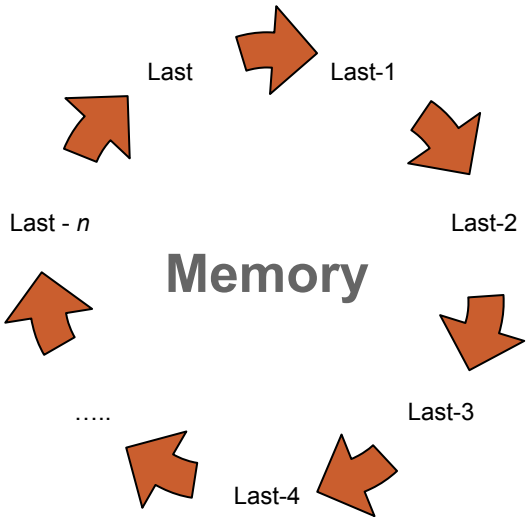

- **Traditional middleware model: Last value erased on update**
- **Fabric model: create ring buffers**

Future Innovations:
Local DMA with MAMA

- Port apps from remote machines to inside the publisher
- No recoding needed: still MAMA

RDMA

Feed Handler

MAMA

Client App Process

"LDMA"

Feed Handler

MAMA

Client App Process

In Summary

- IB is crossing the capital markets chasm now
- Wombat Data Fabric wraps the benefits of IB & RDMA under a pub sub-like middleware
- Now Alpha. Launch Dec 2007 in capital markets
- Looking for opportunities to add value beyond…

Supporting Materials

FPGA?

- A single node solution- not a platform solution
  - The handler may be faster- but what about your apps?
- Possible value but less than WDF provides
- VHDL programming- DANGER
  - How to handle exchange data changes?
  - C to VHDL tools are not proven to be mature
- Not supported by enterprise server community yet
- Yet still under Wombat exploration

Simple Physical Architecture

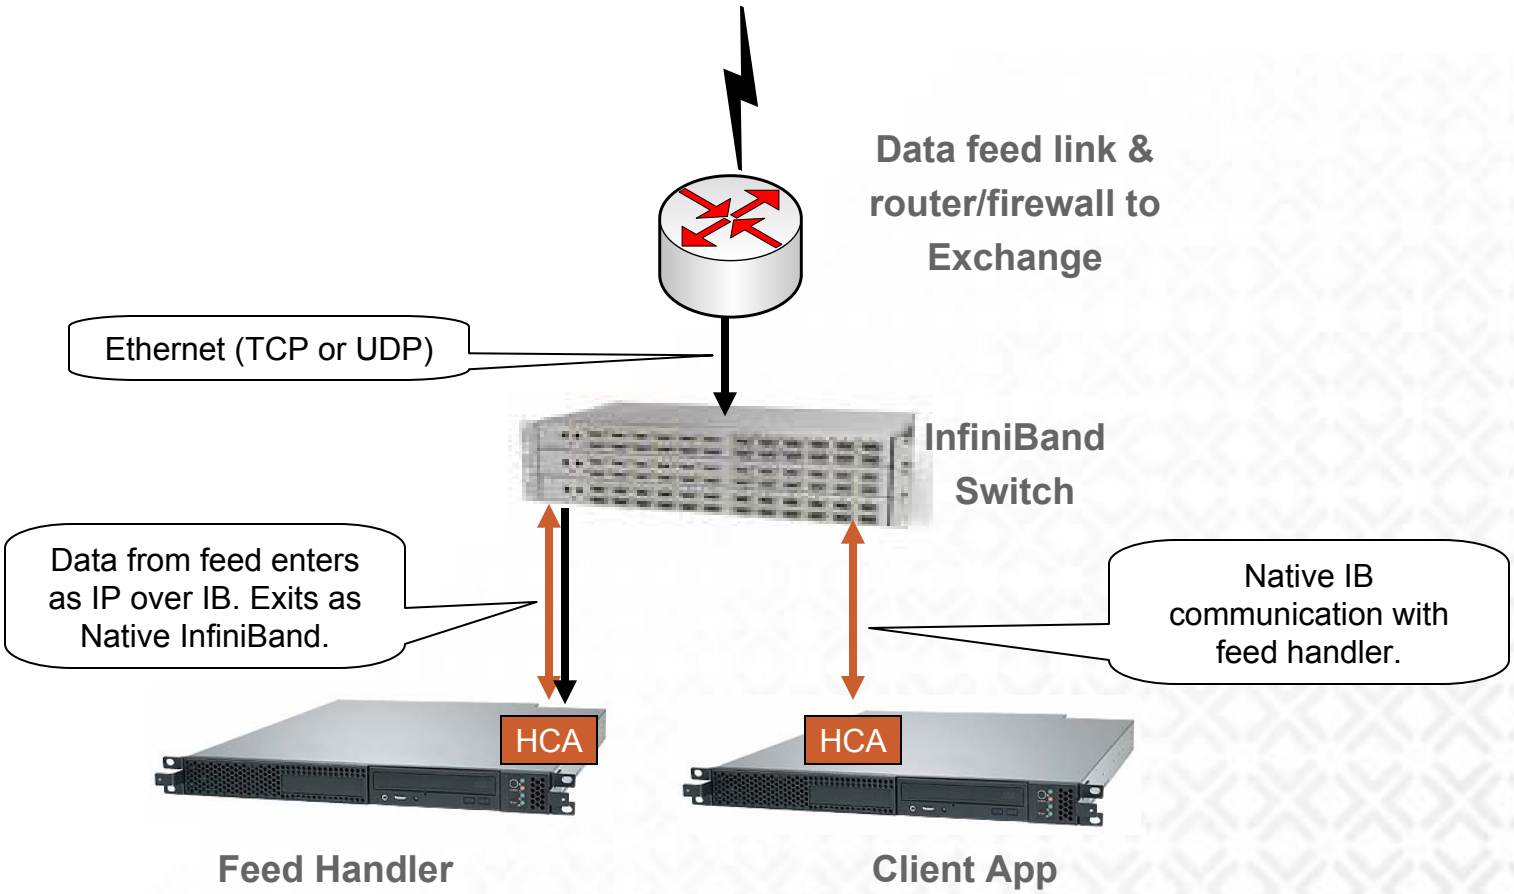

Rendezvous Hops