

# Haystack correlation and post-correlation



**ROGER CAPPALLO**  
**HAYSTACK DIFX MEETING**  
**2011.12.5**

# Software Correlation



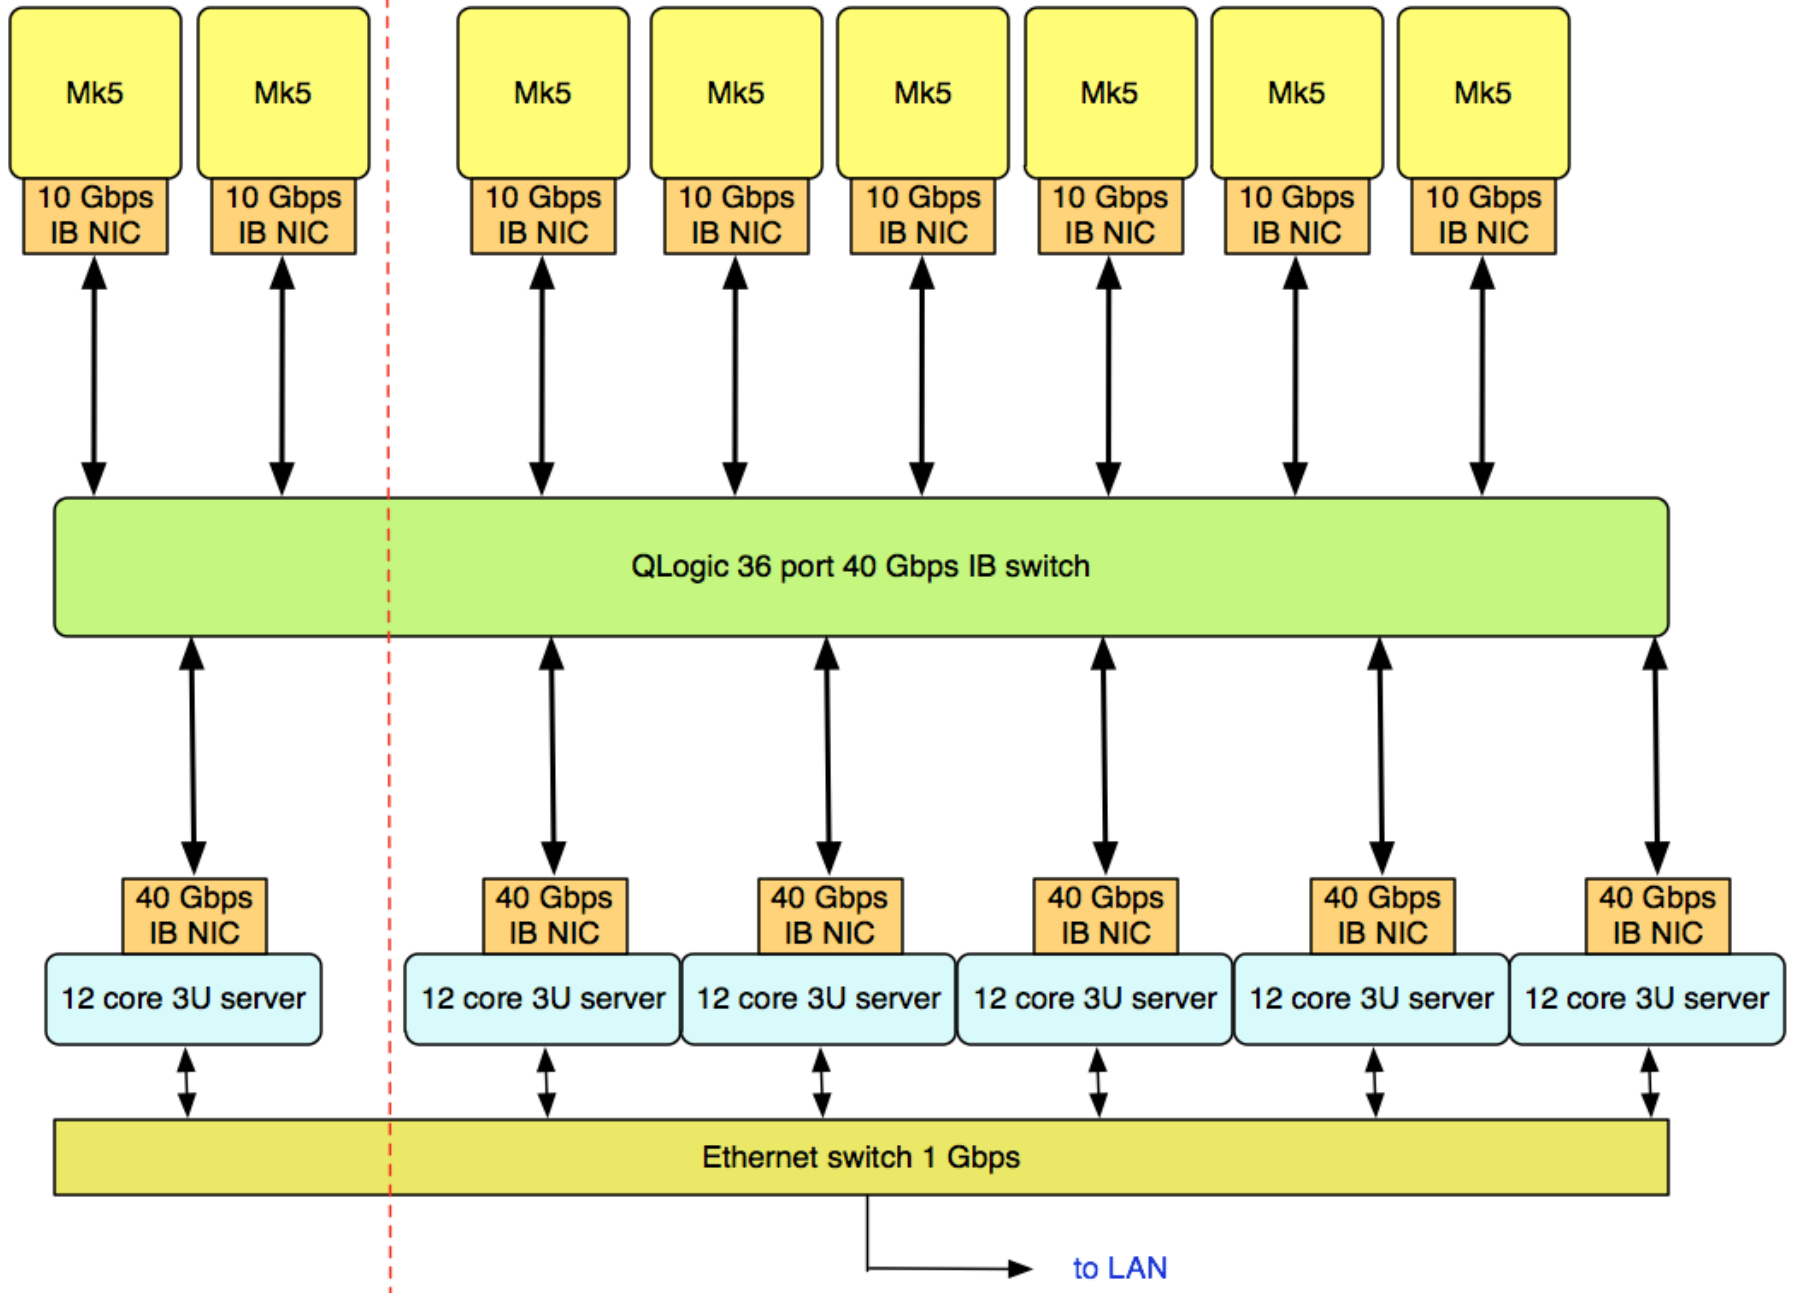
- Gradual transfer over from Mk4 hardware correlator
  - A little more gradual than I'd like!
- Use *difx2mark4* file conversion software to interface DiFX with the HOPS/Mk4 backend suite (*fourfit*, *aedit*, etc.)
- Continuing revision of code as various experiments and problems are encountered, principally from correlations at MPIfR and VLBA
- DiFX Plans:
  - Integrate Mk6
  - Modify to support ALMA phasing
- Recently acquired and integrated COTS hardware for a 72+ node DiFX-based software correlator

- Six Supermicro 3U server chassis, each having:
    - X8DAH+-F motherboard
    - 2 hexacore Intel Xeon X5650 CPU's
    - 24 GB RAM
    - 1 TB hard disk
    - 40 Gb/s IB/HCA
  - QLogic Infiniband switch
    - 36 ports
    - 40 Gb/s
  - Eight 10 Gb/s Infiniband adapter cards for the Mk5's
- Total cost ~\$34K



Phase 1

Phase 2



# Haystack Server Cluster



computer	chips	cores/ chip	h-t cores	Xeon model	clock GHz	cache MB
sc01	2	4	8	E5405	2.00	6 – L2
sc02	2	4	8	E5405	2.00	6 – L2
sc08	1	4	4	E5462	2.80	12 – L3
sc09	1	4	8	E5630	2.53	12 – L3
corro1	2	6	24	X5650	2.67	12 – L3
corro2	2	6	24	X5650	2.67	12 – L3
corro3	2	6	24	X5650	2.67	12 – L3
corro4	2	6	24	X5650	2.67	12 – L3
corro5	2	6	24	X5650	2.67	12 – L3
corro6	2	6	24	X5650	2.67	12 – L3

# Mk5's on Haystack Correlators

#		OS	vers	IPP	SDK	openmpi	gigE	10gE	IB	difx
00	5A	Debian	4.1.1-21	v7	8.2	1.4.3	X			X
01	5A	Debian	4.1.1-21	v5	8.3					
02	5A	Debian	4.1.1-21	v7	8.3	1.4.3	X			X
03	5A	Debian	4.1.1-21	v5	8.2					
04	5A	Debian	4.1.1-21	v5	8.2					
05	5A	Debian	4.1.1-21	v5	8.2					
06	5A	Debian	4.1.1-21	v5	8.2	1.4.3	X			X
07	5B	RedHat	3.2.2-5	v5	7.x	1.4.3				
08	5B	RedHat	3.2.3-47	v5	7.x	1.4.3	X			X
09	5B	Debian	5.0.8	v7	9.0	1.4.3	X	X	X	X
10	5B	Debian	4.0	v7	8.3	1.4.3	X			X
11	5B	Debian	4.0	v7	8.3	1.4.3	X			X
12	5B	Debian	5.0.8	v7	9.0	1.4.3	X		X	X
13	5B	Debian	5.0.5	v7	9.0	1.4.3	X	X		X

# InfiniBand Status



- QLogic switch requires management (*opensm*)
  - Using the pml cm (which seems to require identical hardware) *mpispeed* achieves up to 23 Gb/s.
  - With pml ob1 the speed drops to 13 Gb/s over the same path
  - Haven't yet been able to transfer data with *mpispeed* from mk5's to servers
    - 10 Mb/s Mellanox HCA → 40 Mb/s QLogic switch
    - Need to rebuild openMPI to get missing BTL components
  - Hope to get cluster ready for operations this week
- 
- pml – point to point management layer
  - cm – (Conner MacLeod) connection manager uses exposed matching fabrics
  - ob1 – (Obi-Wan) high-perf pml implementation in openmpi; uses RDMA if possible
  - btl – byte transfer layer