Mark 6 Next-Generation VLBI Data System

Alan R. Whitney
MIT Haystack Observatory

5 March 2012
IVS General Meeting
Madrid, Spain

Data rates for VLBI2010

- VLBI2010 data rates are dictated by
  - Small antennas (12m class)
    - Antennas must be able to move very quickly around sky
  - Weak sources
    - Sources need to be ~uniformly distributed in the sky and have simple or no structure, which severely constrains available sources
  - Short observations
    - VLBI2010 on-source observations will be ~10 secs each
    - antenna must move around sky quickly to map temporal atmospheric flucuations
    - most of observation period is spent moving antenna from source-to-source

All these factors conspire to dictate very high data rates (both instantaneous and average)

VLBI Data Rates and Volume are not for the faint of heart!

- VLBI2010 at 4 Gbps/station average, 4 to 20 stations
  - ~5-40 TB/station/day
  - Global 10-station experiment @ 4 Gbps/station up to ~400 TB/day
  - Single 10-day experiment can produce up to ~4 PB
- Higher data rates (8-32 Gbps) are already on the horizon; higher data rates → more sensitivity
- Available disk supply can support only few days of observations at these rates

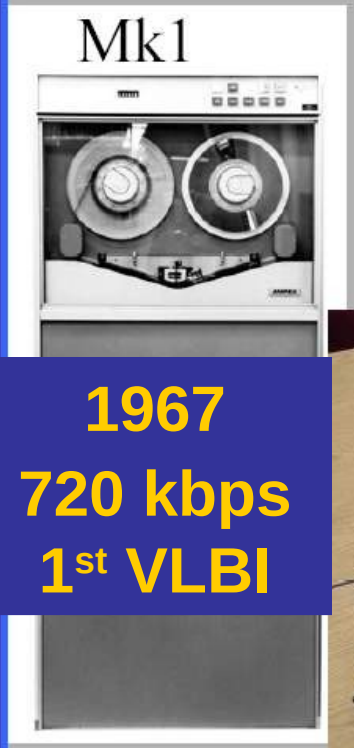

1967
720 kbps
1st VLBI

1971
4 Mbps

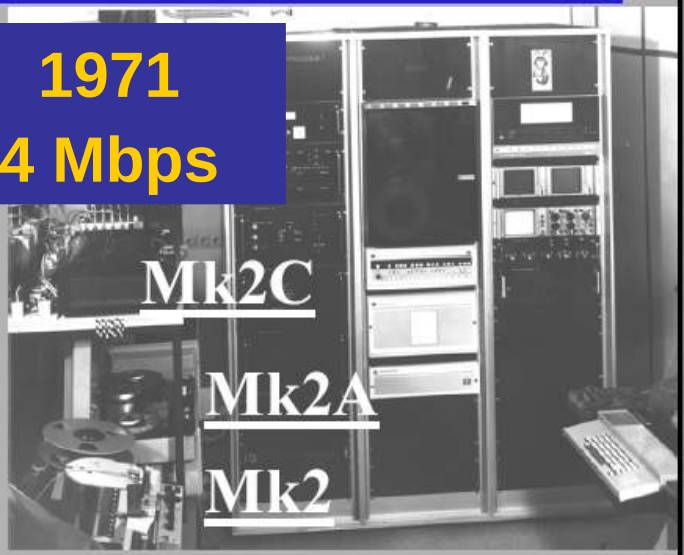

2002
1 Gbps
2006
2 Gbps
2011
4 Gbps

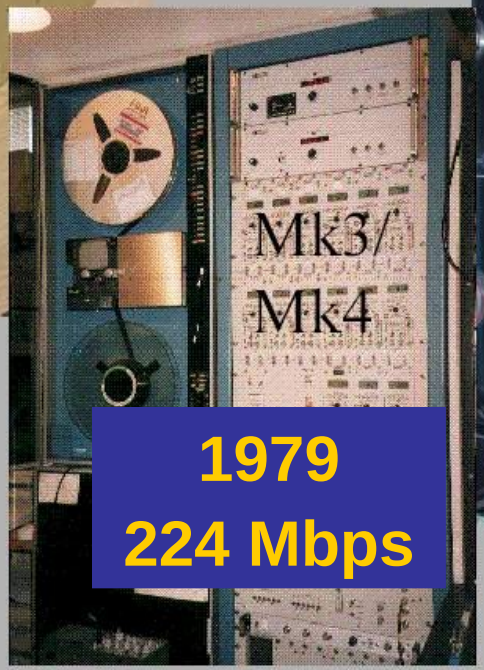

1979
224 Mbps

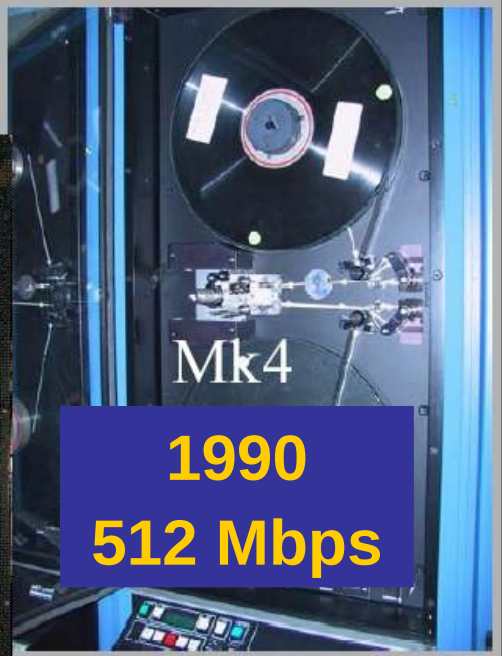

1990
512 Mbps

Mark 5 Data Acquisition System

(Mark 5A/B/B+/C all look the same)

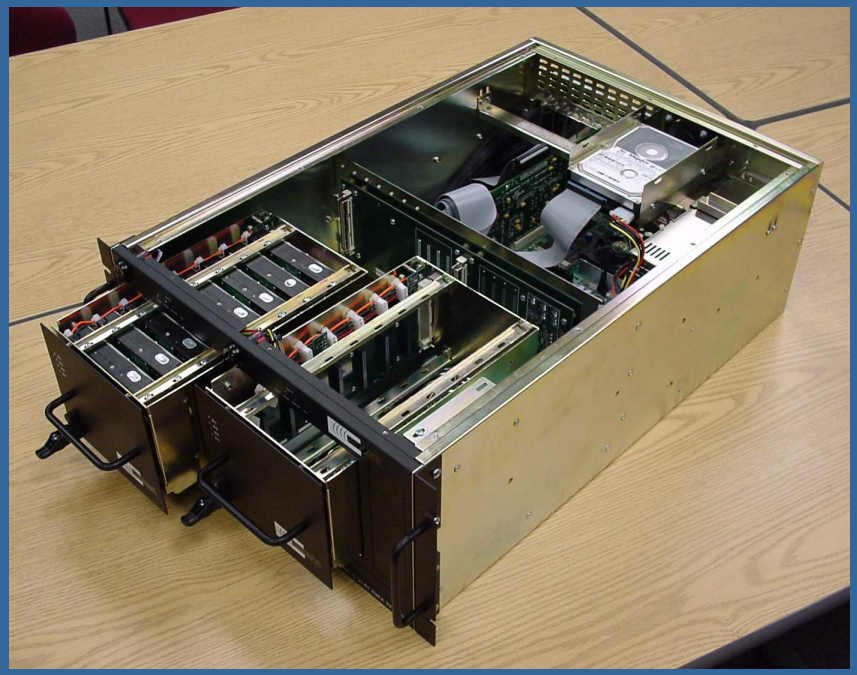

	Year introduced	Record rate (Mbps)	Interface	Cost (USk$)	#deployed
Mark 5A	**2002**	**1024**	**Mk4/VLBA**	**21**	**~130**
Mark 5B	**2005**	**1024**	**VSI-H**	**22**	**~40**
Mark 5B+	**2006**	**2048**	**VSI-H**	**23**	**~30**
Mark 5C	**2011/12**	**4096**	**10GigE**	**21**	**~20**

Mark 5 includes a significant amount of proprietary technology

Next up – Mark 6

- 16Gbps sustained record/playback and burst-mode capability
- 4 x 10GigE input ports
- Based on inexpensive high-performance COTS hardware
- Linux OS w/open-source software
- Resilient File System to manage slow and failed disks
- e-VLBI support
- Fully VDIF/VTP compliant
- Goal is to preserve as much investment in existing Mark 5 systems and disk libraries as possible
- Mark 6 collaboration:
  - Haystack Observatory – all software and software support
  - NASA/GSFC High-End Network Computing group – consultation on high-performance COTS hardware
  - Conduant Corp – Mark 6 disk module

Basic Mark 6 System

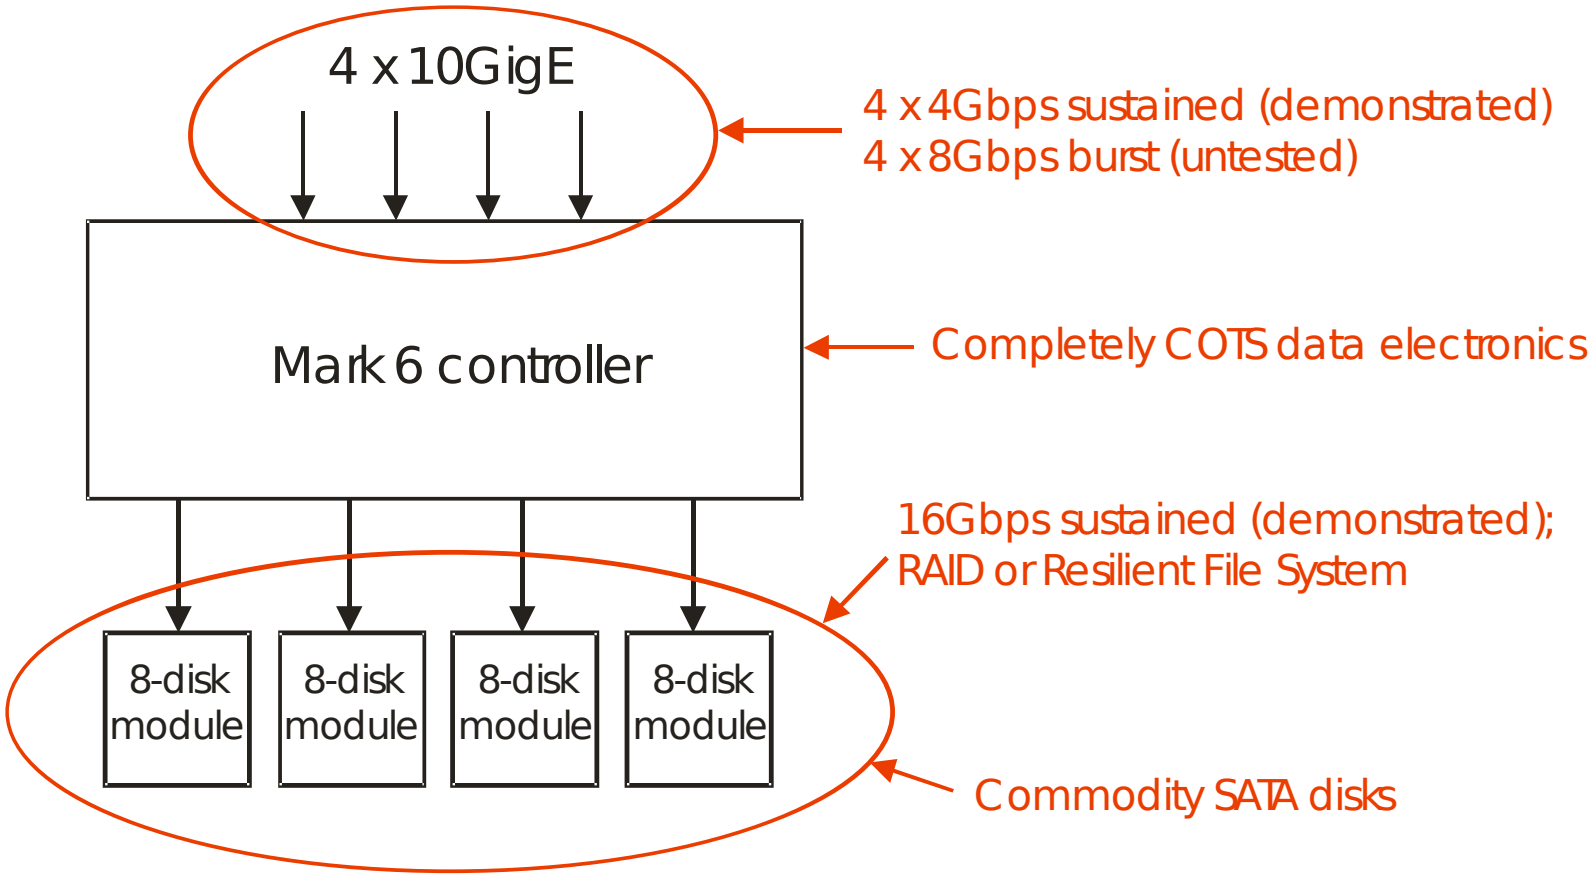

16 Gbps VLBI demonstration with Mark 6

24 October 2011

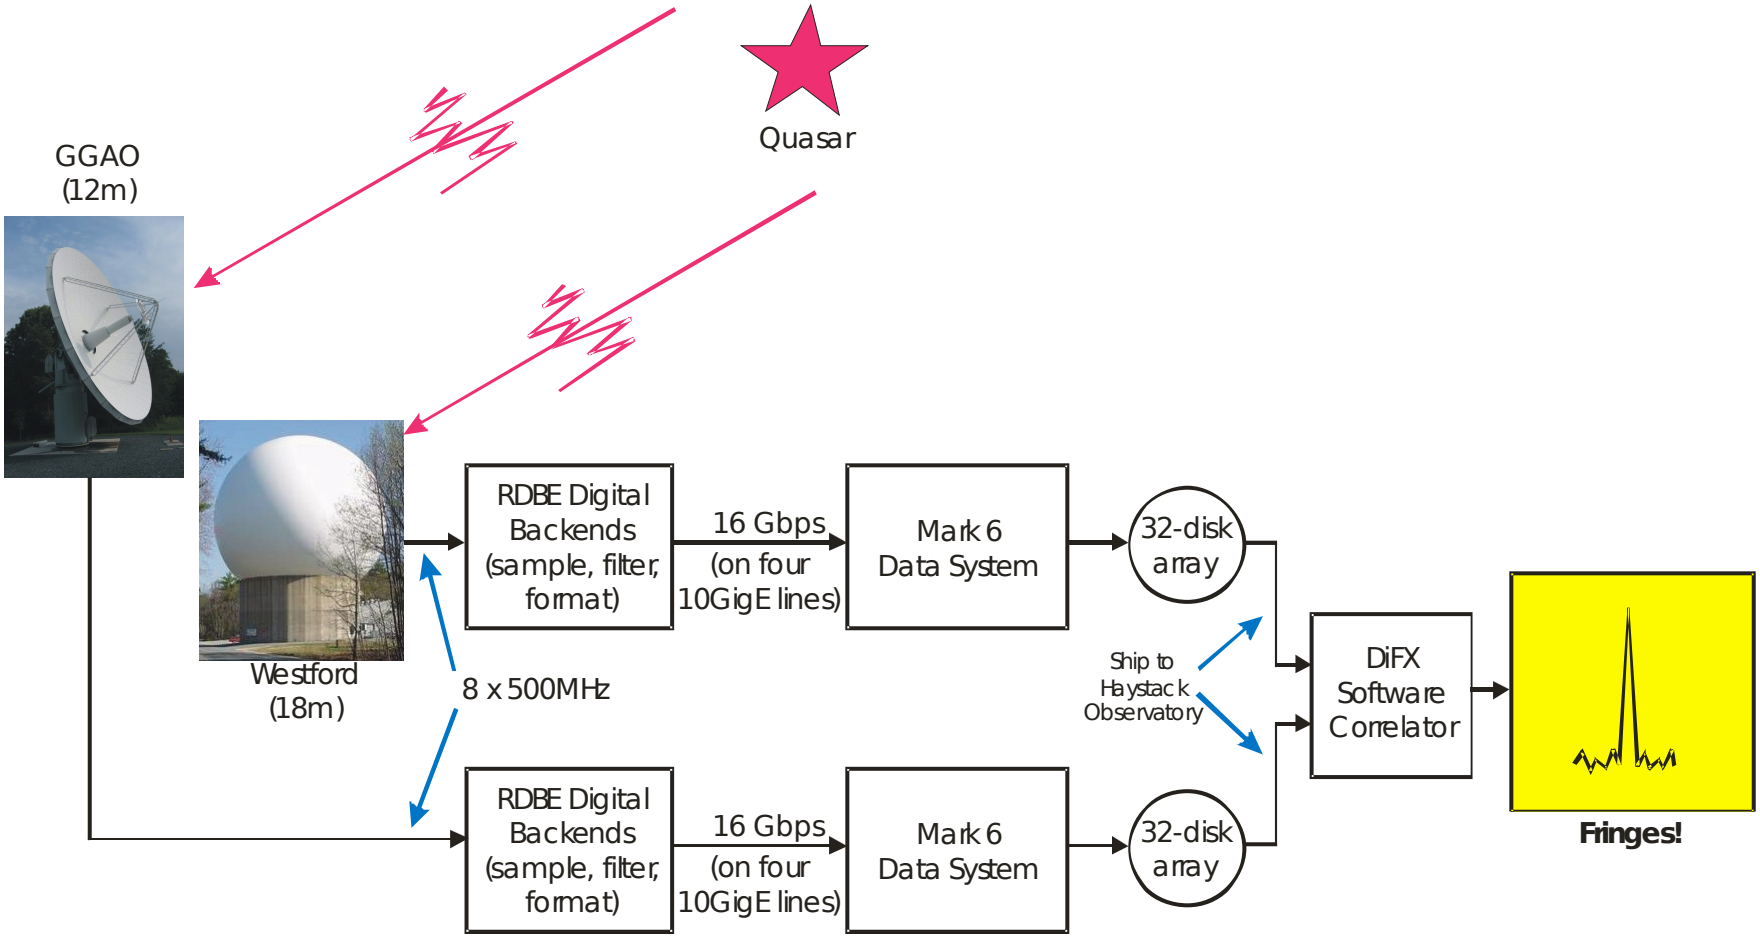

Correlation results (single 500MHz channel)

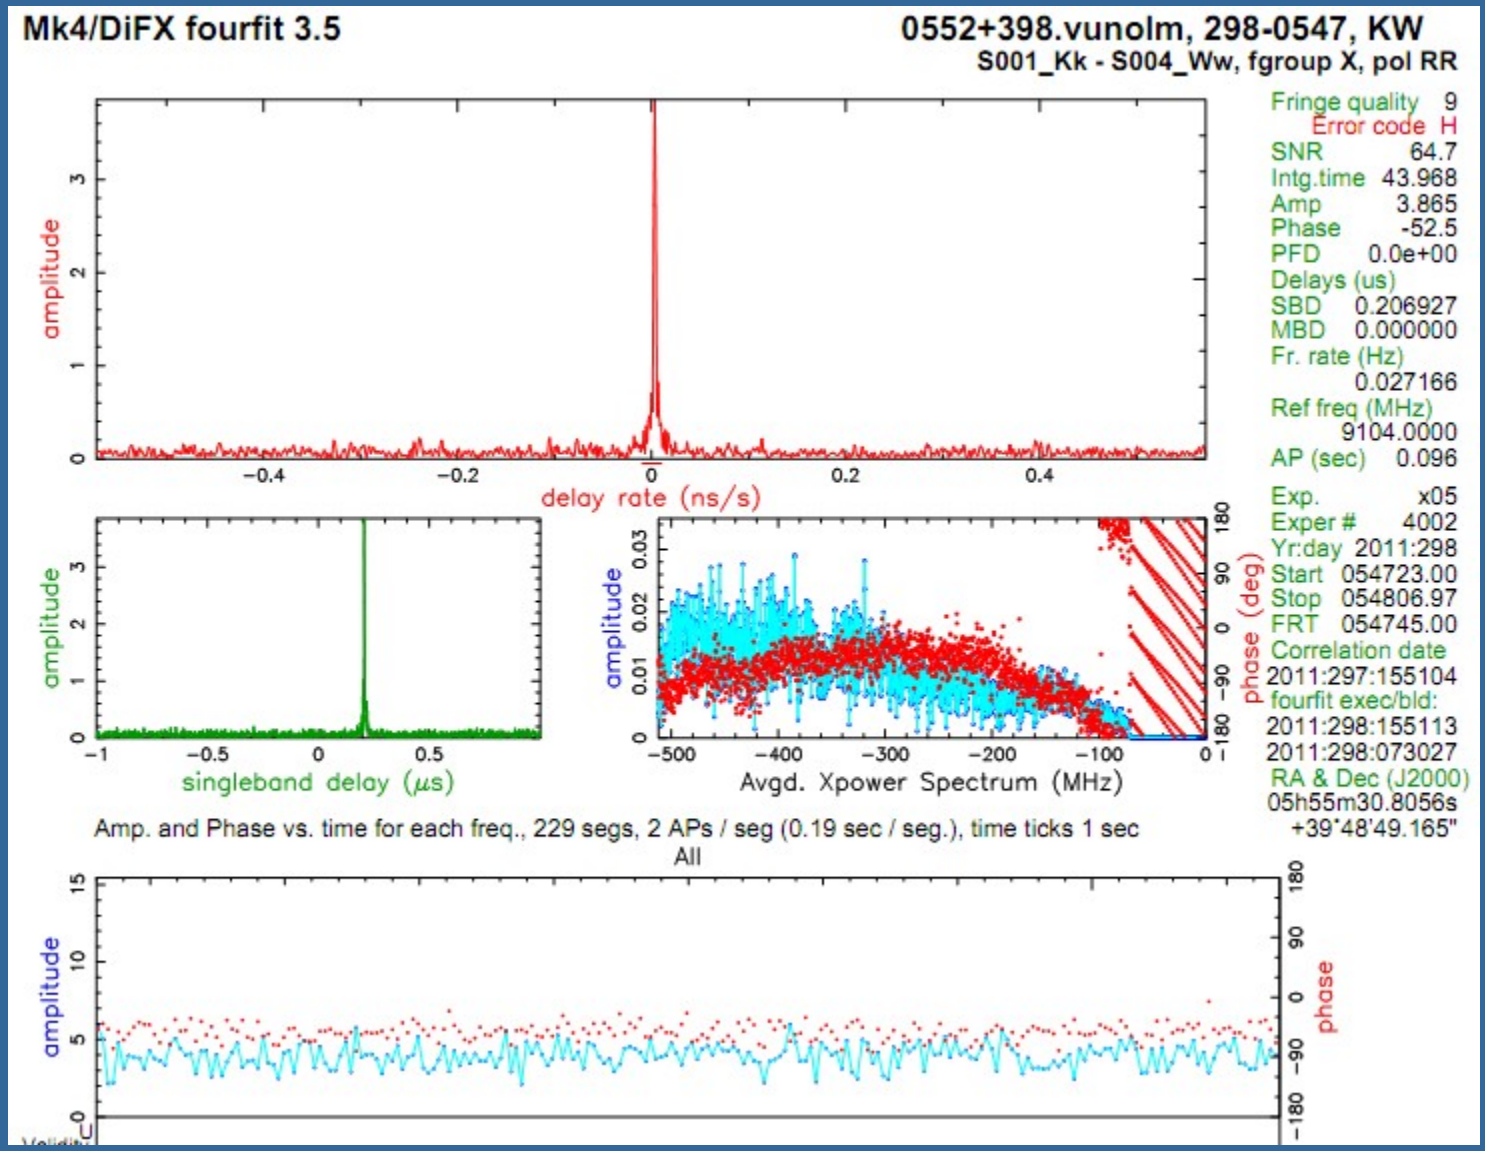

Mark 6 Project Status

- Sustained 16Gbps from four 10GigE interfaces to disk has been demonstrated
- To be completed:
  - Full VSI-S command set
  - Resilient File System (current implementation is RAID-based)
  - Playback as standard Linux files
- Prototype Mark 6-specific hardware pieces arrived at Haystack last week from Conduant
  - New Mark 6 chassis-backplane power-management boards
  - Mark 5-to-Mark 6 SATA disk module upgrade kit
  - New cable-management panel

Prototype Mark 6 hardware

High-speed data connections to module front-panel via two standard SAS cables

Existing Mark 5 chassis is upgradeable to Mark 6

New chassis backplanes for disk power management

Cable-management panel (unused cables retract into panel)

Existing Mark 5 SATA disk modules are upgradeable to Mark 6 (new backplane and front panel)

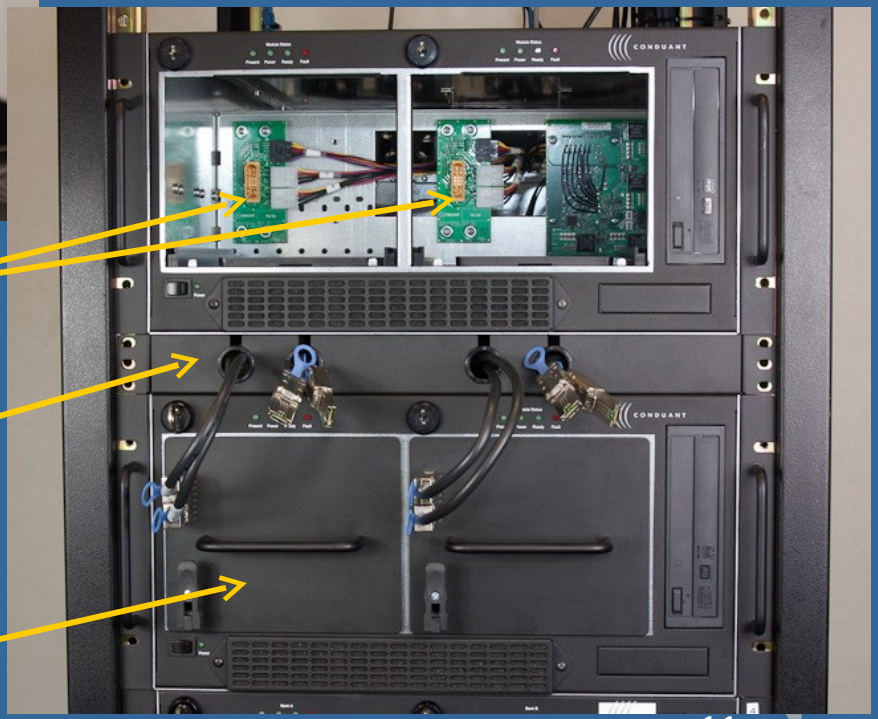

Projected Mark 6 schedule

- Mar 2012 – GGAO/Westford Mark 6 test with broadband VLBI2010 system (dual-pol with 2GHz BW/pol)
- Mar 2012 – Test Conduant prototype hardware; integrate complete hardware system; begin integration with Field System
- mid/late 2012 – System complete and fully tested; (new complete Mark 6 system <$15k)

VDIF
(VLBI Data Interchange Format)

- Standardized format for raw time-sampled VLBI data
- Compatible with both VLBI data-recording systems and e-VLBI data transmission
- Highly flexible to accommodate a large variety of channel and frequency configurations, including mixed sample-rate data
- VDIF being implemented for all new VLBI2010 systems
- Accompanying VLBI Transport Protocol (VTP) specifies e-VLBI data-transmission protocol for VDIF-formatted data stream

For details: www.vlbi.org

VTP
(VLBI Transmission Protocol)

- Companion specification to VDIF
- Specifies e-VLBI data-transmission protocol for VDIF-formatted data streams
  - Normally must use UDP or UDP-like protocol to maintain necessary data rate
  - Addition of Packet Serial Numbers (PSNs) helps to keep packets organized and identify missing packets (a few missing packets are not normally a problem)

Thank you

Backup slides

Thank You's

Haystack/Westford –
Chris Beaudoin, Pete Bolis, Roger Cappallo, Shep Doeleman, Geoff Crew, Rich Crowley, Dave Fields, Alan Hinton, David Lapsley, Arthur Niell, Mike Poirier, Chet Ruszczyk, Jason SooHoo, Ken Wilson

NASA/GSFC VLBI Group –
Tom Clark, Ed Himwich, Chopo Ma

NASA/GSFC GGAO –
Roger Allshouse, Wendy Avelar, Jay Redmond

NASA/GSFC High-End Computer Networking Group –
Bill Fink, Pat Gary (recently deceased), Paul Lang

Conduant –

Phil Brunelle, Greg Lynott, Ken Owens

Recording-rate cost vs. time

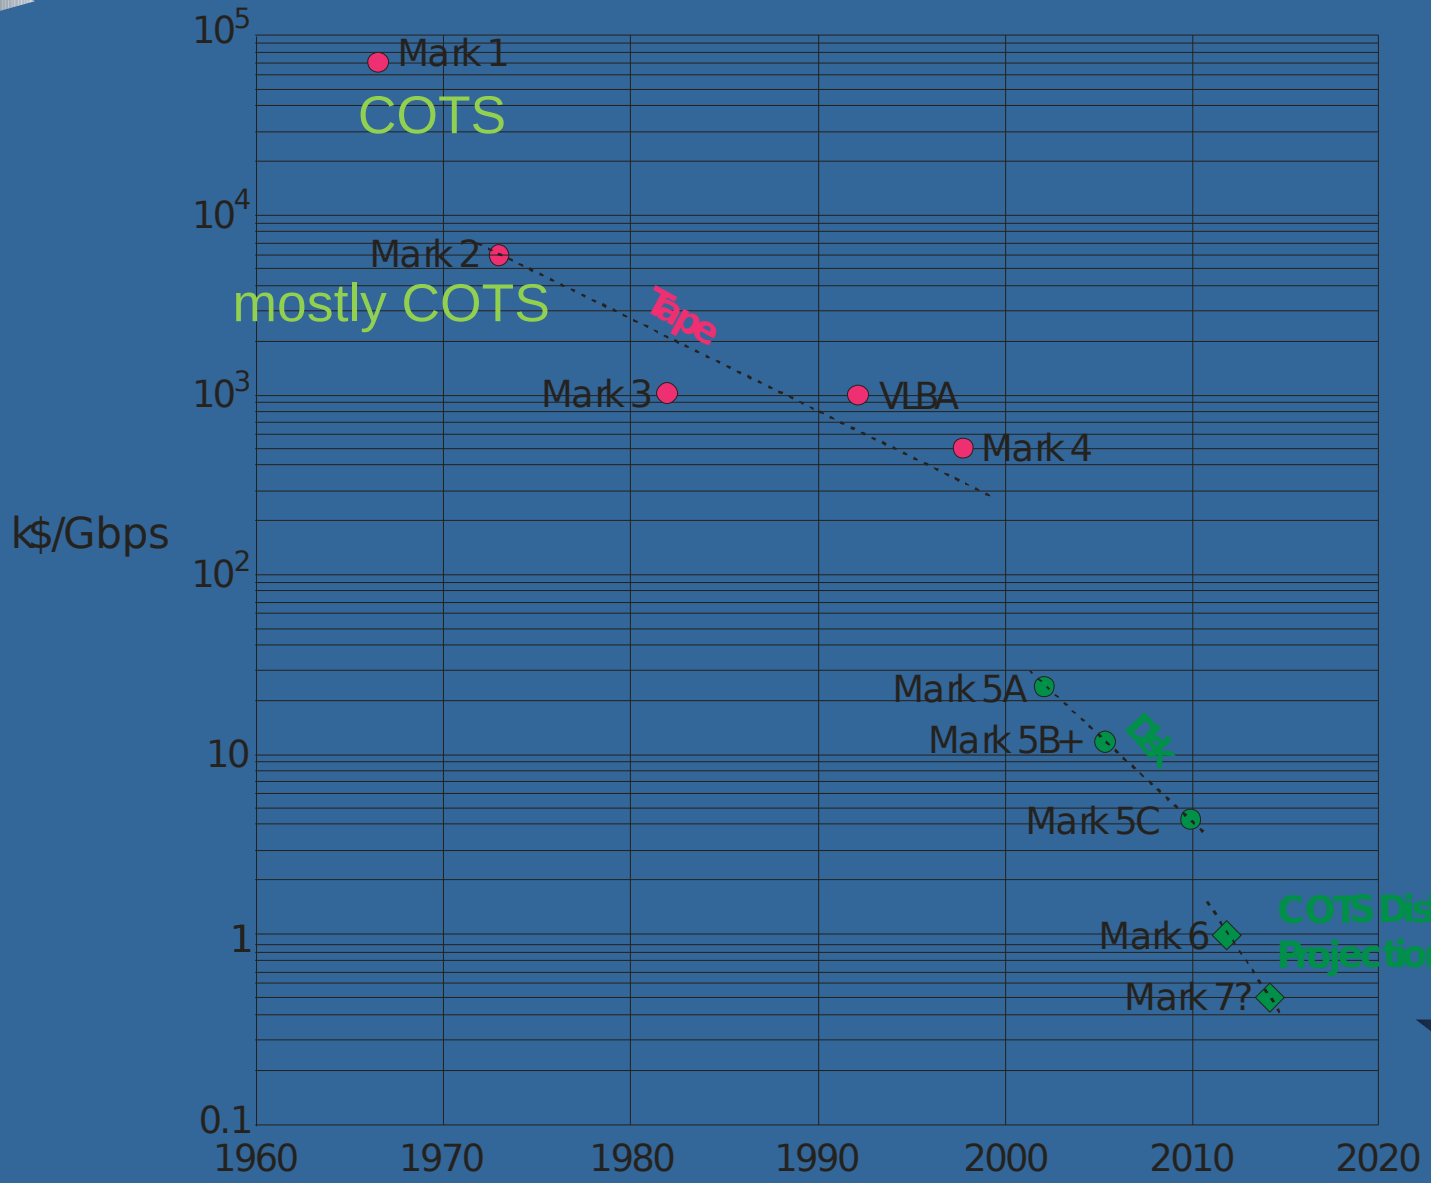

Backend-bandwidth cost vs. time

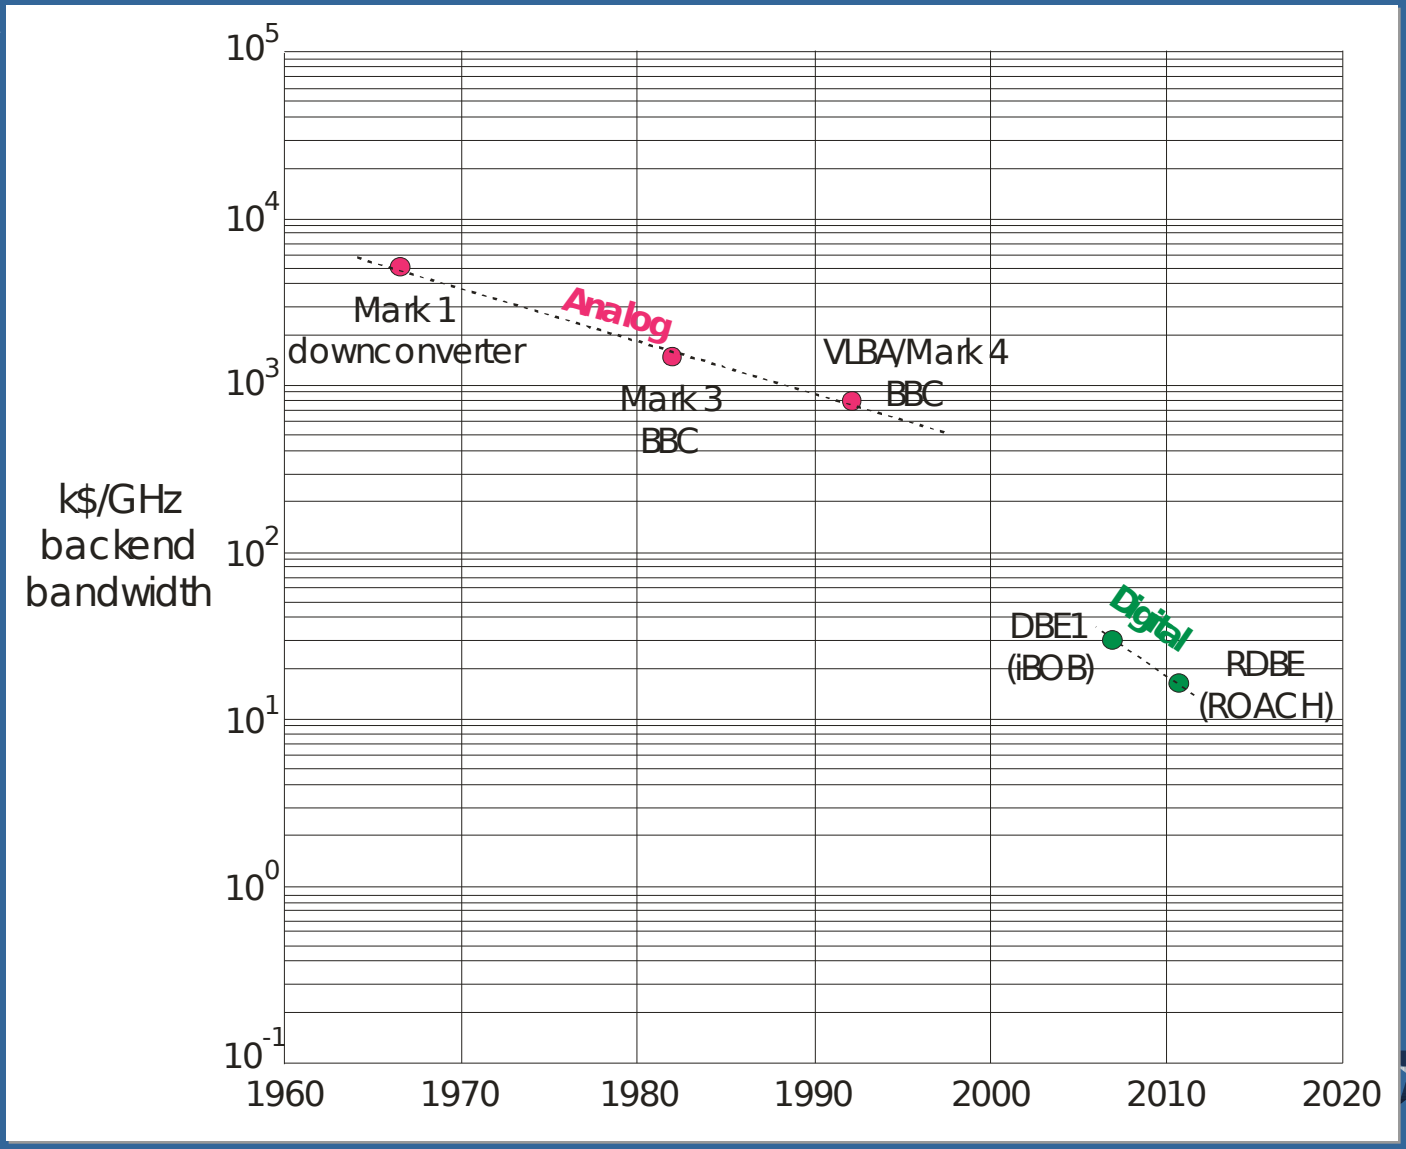

Recording-rate capability vs time

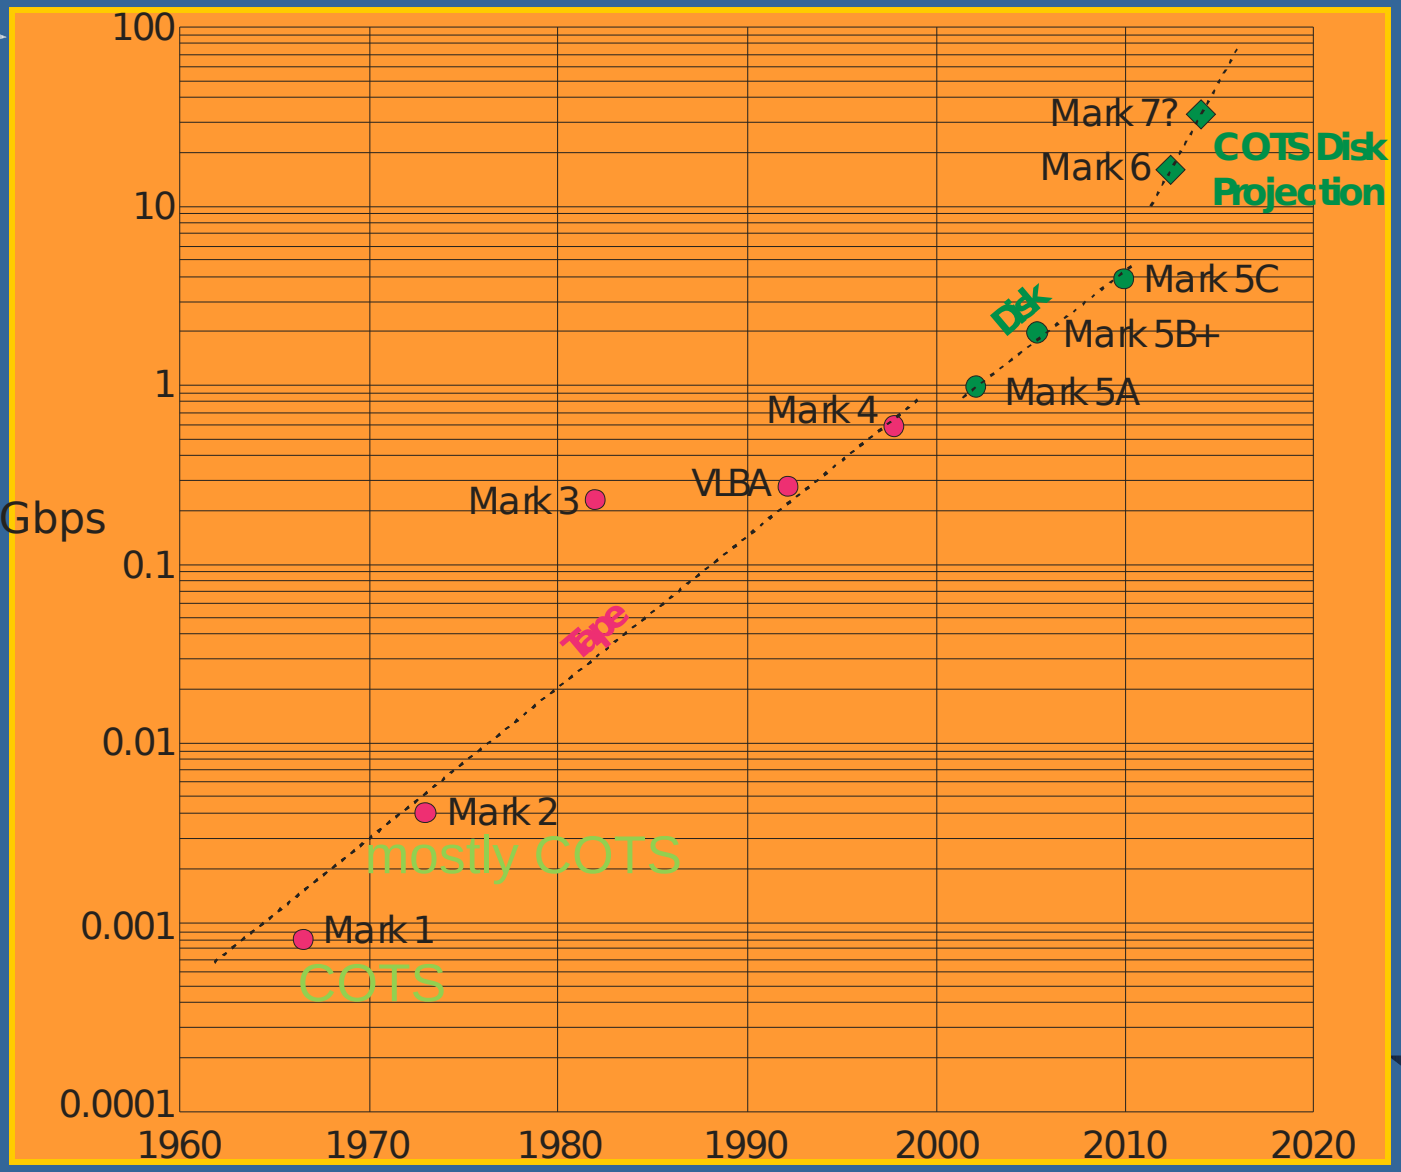

Mark 6 M&C and concepts

- VSI-S command set
- Recording units are defined as ‘volumes’, each of which consists of one or more physical disk modules
  - Multi-module volumes are required for recording rates >~4Ggps
  - Multi-module volumes retain identity thru correlation processing, then are returned to single-module volumes
- Volumes are managed on an ordered ‘Volume Stack’ that allows multiple volumes to be mounted simultaneously
  - Allows volumes to be queued in specific order for usage
  - Supports automated switchover to next volume in Volume Stack when current module becomes full; switchover takes place between scans
- Disk statistics gathered during recording allow easy identification of slow/failing disks by disk serial number

Mark 6 16Gbps demonstration system

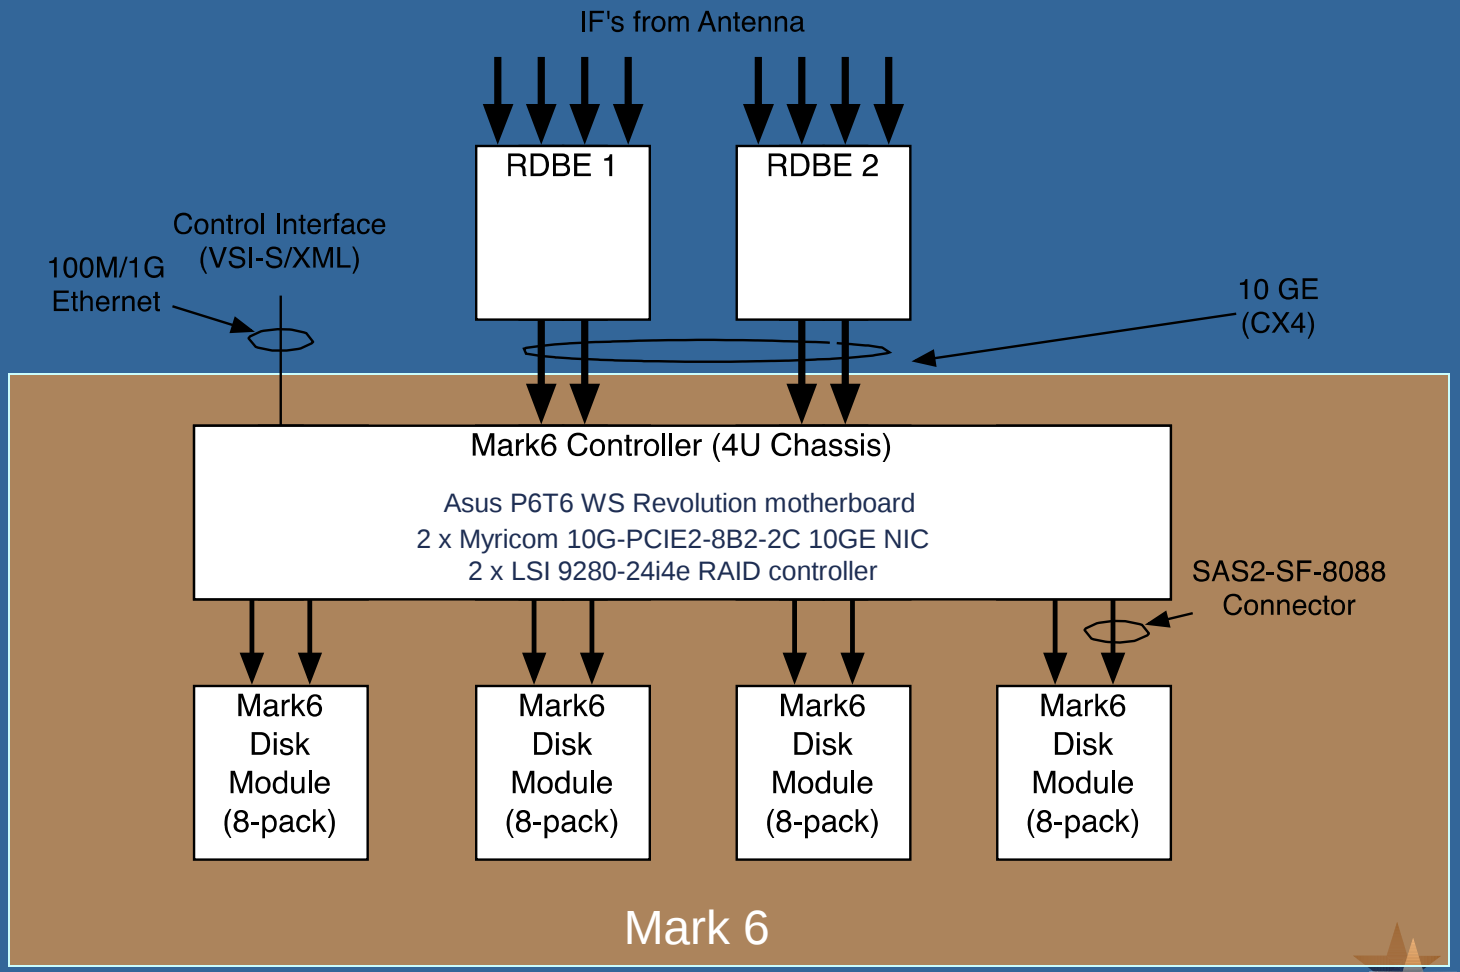

How will Mark 6 be available?

- Several options:
  - Purchase new Mark 6 system from Conduant
  - Upgrade existing Mark 5 system (either yourself or with kit from Conduant)
  - Upgrade Mark 5 SATA-modules (with upgrade kits from Conduant)
  - Purchase Mark 6 modules (with or without disks)