Haystack Technology Roadmap

Russ McWhirter, John Barrett, Pedro Elosegui, Chet Ruszczyk, Ganesh Rajagopalan, Dan Hoak, Alex Burns, Haystack IT Dept.

MIT Haystack Observatory

Outline

- Haystack Research and Collaboration Network (RECN)
- Development of Digital Backend Version 5 (DBEv5)
- Haystack data storage development

Prior Haystack e-transfer architecture

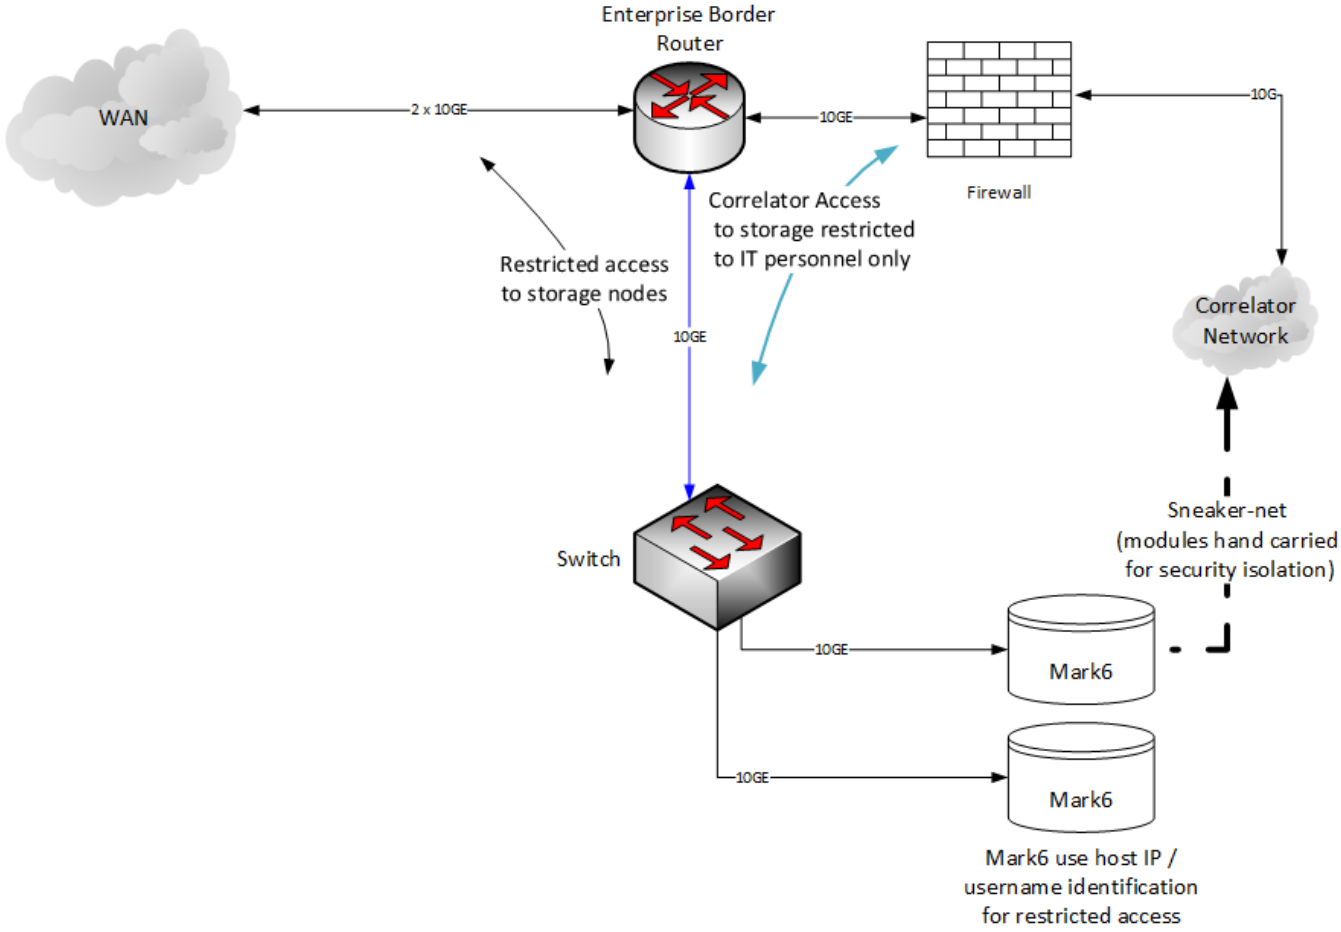

Haystack Research and Collaboration Network (RECN)

- Highlights
  - 400Gbps connectivity now available
    - One 100Gbps fiber lit, with connectivity to Internet2 thru campus
  - Science DMZ
    - Funded by NASA SGP
    - Renamed - Science Isolated Storage (SIS) zone connected and available

SIS Objectives

- Network architecture expandable for other research areas storage nodes, e.g. Astronomy
- Ability to read in data for correlation directly or via sneaker net
  - Consideration is the number of stations participating in the experiment
  - Minimize human interaction for correlation time
    - Transfer of modules via sneakernet
- Evaluate processes to reduce e-transfer to correlation time
  - Critical for Geodesy where 16 stations presently observing
  - 3 weeks to get data for correlation
- Provide 100Gbps network to Westford for e-transfer

SIS Storage

- Network attached storage (NAS) nodes
  - Three NAS's
    - 2.6 PB of total storage
    - Two flavors of processor
      - Intel / AMD
      - Evaluating performance interacting with correlator to directly pull data from NAS nodes
- Mark6+
  - To continue sneaker net based upon number of stations
  - 32 Gbps average input recording rates

RECN Network Diagram

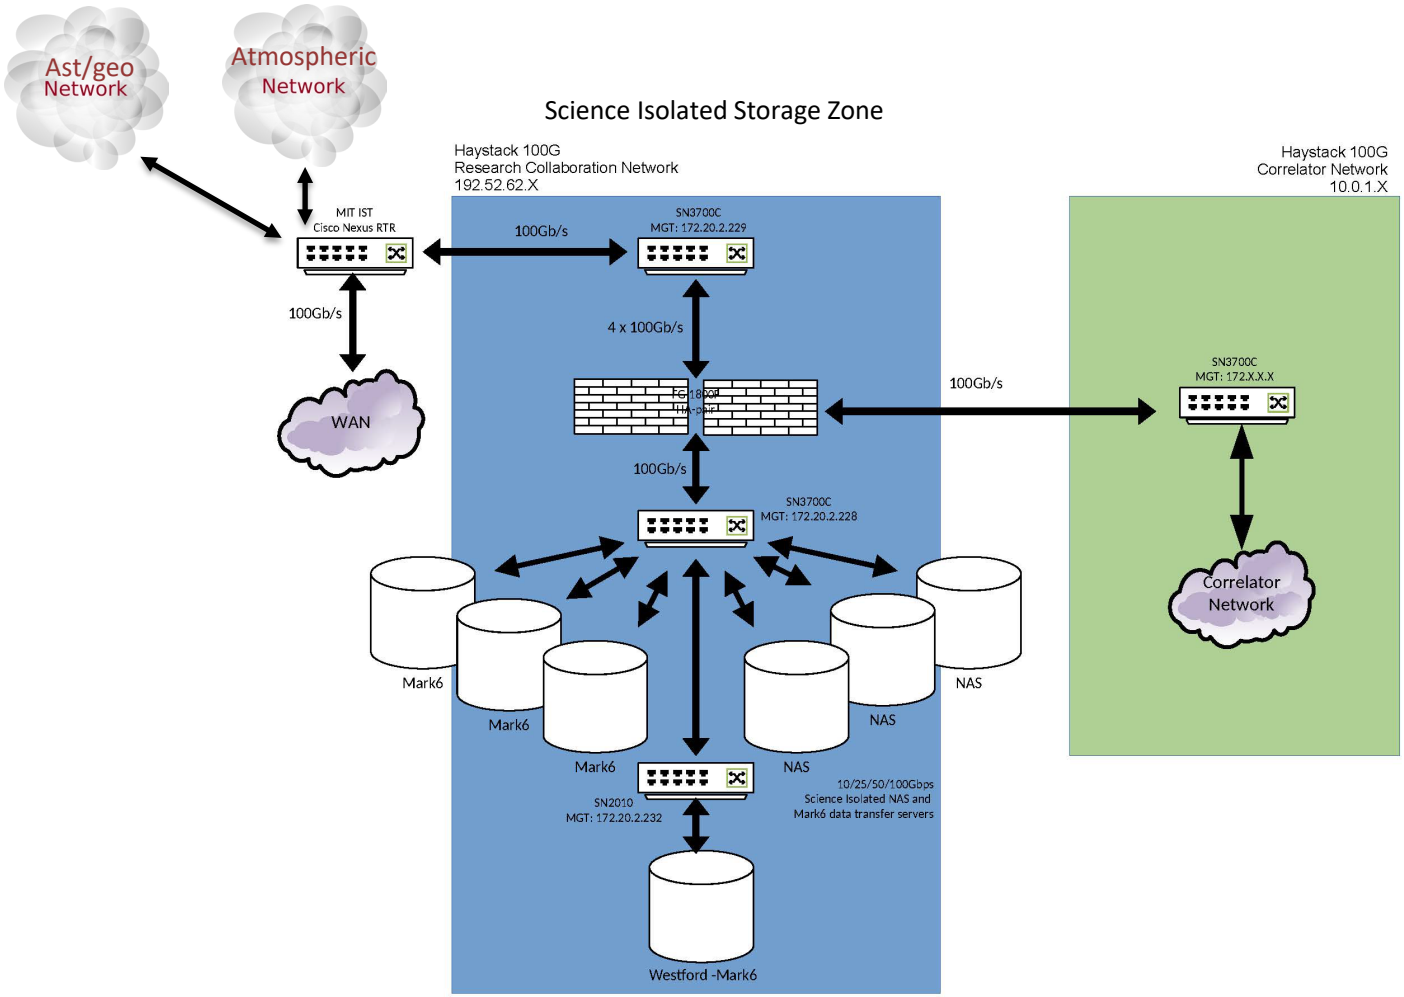

MIT VGOS Signal Chain

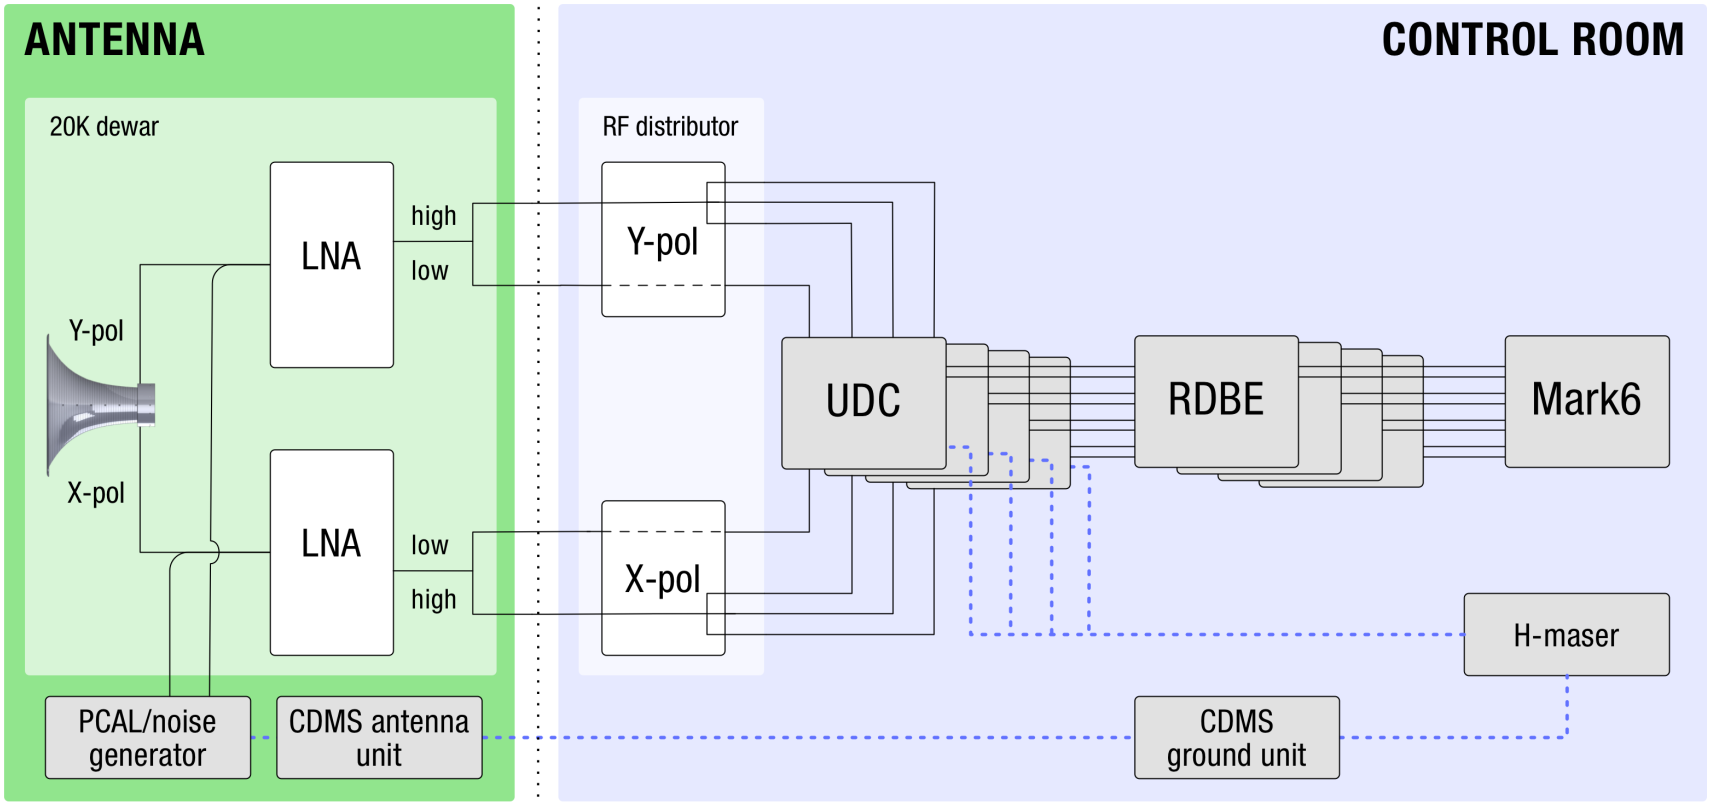

DBEv5 on CASPER architecture

- Goal was a drop in replacement for the R2DBE's
  - Same personality / features
  - 2 GHz IF processing
  - Port existing CASPER based personality
    - Utilizes VHDL and CASPER blocks to maximize performance
    - The CASPER filterbanks channelize the input signals, which are then re-quantized to 2-bits, and the data are output in VDIF protocol through the 100-GigE outputs
  - Extend for GRITSS optional data path (Marco Midon presentation)
    - Software configurable between VGOS and GRITSS
- Leverages hardware commercially available
  - 4x2 RFSOC model and features here
    - Max 6 GHz Input Freq
    - 14 bit ADC
    - Onboard synthesizer
    - FPGA has 4272 DSP slices (double of Roach2)
    - 100G NIC egress for VLBI data

CASPER Realities

- Simulations of ported R2 Personality exposed critical bug
  - 100 Gbps NIC could not be clocked at 256 MHz
    - Buffer overflows
    - Works up to 255 MHz then fails
  - Mitigation strategy was to reduce the internal clock from 256 to 128 MHz
    - Reduces the Poly phase filter bank (PFB) block from processing 2GHz bandwidth to 1 GHz
      - Meets VGOS requirements

CASPER Realities

- Reduction of PFB clock to 128 MHz results in
  - 32 points of the Fast Fourier Transform (FFT) design versus present 64 points of the original design
- 2nd bug discovered
  - CASPER FFT block does not work below 64 points
- Interfacing with Casper @ UC Berkeley to resolve
  - 100 Gbps network interface is not their own
    - Difficult to debug and correct – not an option
  - FFT block they will look into
    - Mitigation strategy is to use Xilink FFT block (verified)

Update Signal Chain Architecture

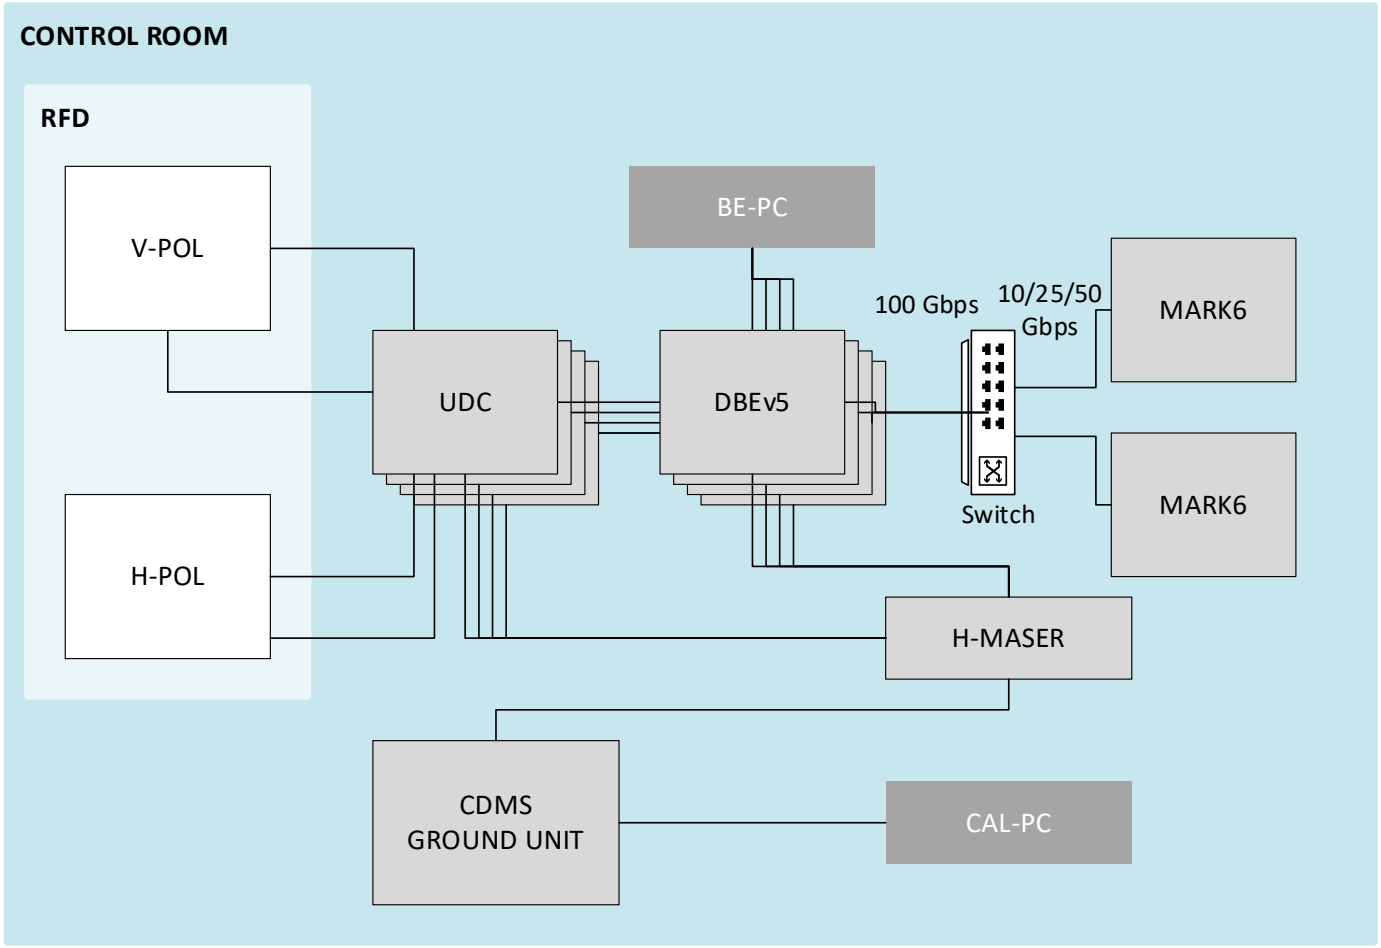

Haystack Data Storage Development

- Mark6 have been deployed since 2015
  - Workhorse for Geodesy and EHT
  - Extended to Mark6+
    - 32 Gbps capability
    - And higher rates (see John Barretts presentation Weds)
  - Reality science and backend technology are pushing higher data rates
  - Working on new storage stategy

Storage Strategy

- Two tier approach
- Tier 1 is an “edge router buffer”
- Tier 2 is any legacy storage node in service
  - Mark6’s
  - NAS (e.g. flex buff)
  - Any thing else used
- Resource broker that configures
  - Data flows based on ingress data rate to Tier 1
  - Egress rates and configurations available to Tier 2 nodes

Storage Architecture

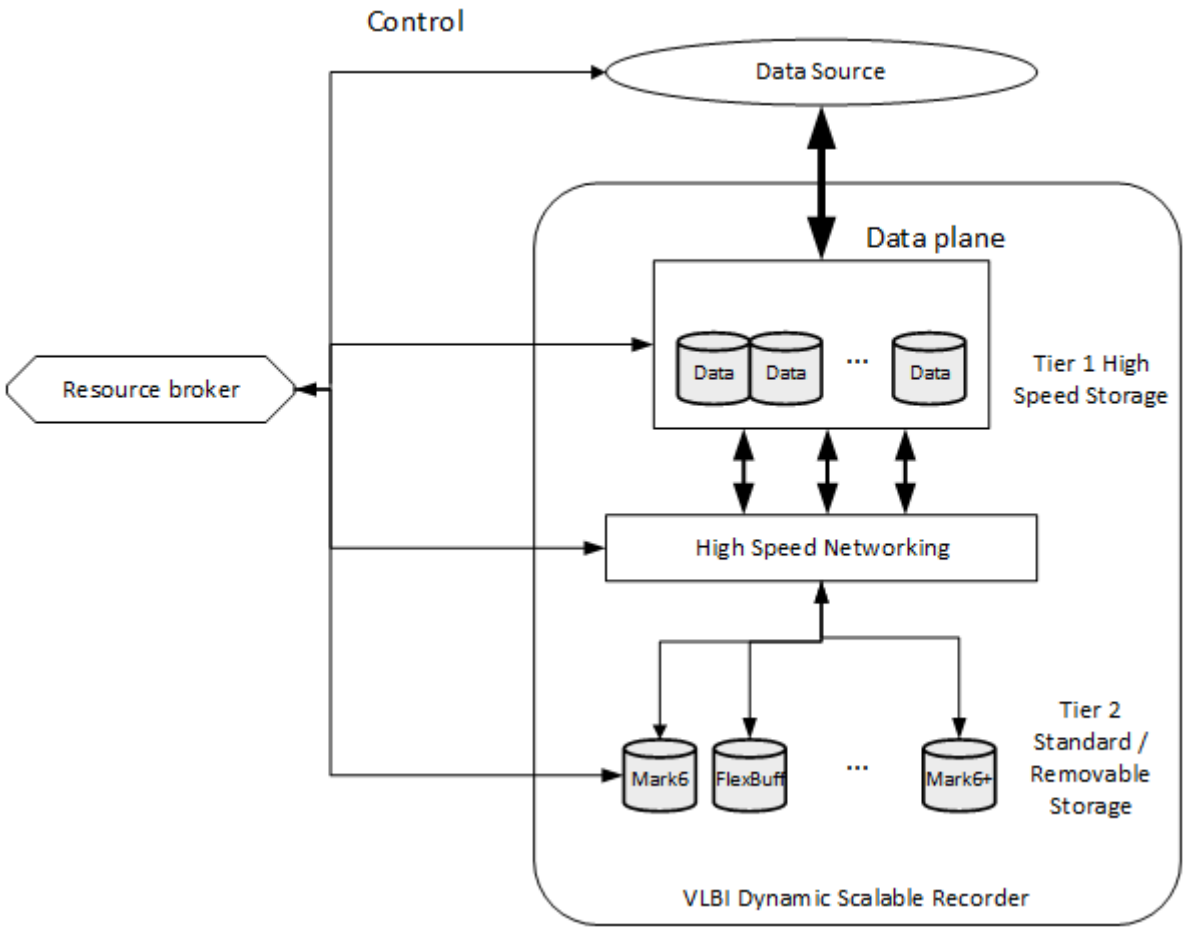

Tier 1

- HERB 1
  - Sorry broke the “MarkX” name since it is not a data recorder but a high speed storage buffer
  - Name can be
    - Tribute to the original founder of Haystack “Herb Weis”
      - See origin story on the outside wall
    - Or Haystack Edge Router Buffer
    - Your call. ☺
  - Hardware base system identified (ordered)
    - Development plan established

Summary

- Haystack has 100 Gbps to Science Isolated Storage for improved e-transfer to correlation times.
- MIT/NASA VGOS network switching to R2DBEs
- MIT/NASA GRITSS DBEv5 is pathfinder to sustainable support
- Introduction to new storage architecture and the “H1”