
Status of ASDC and CERES SCF

CERES-II Science Team Meeting No. 7

Mike Little

Michelle Ferebee

April 24, 2007

Overview

- SCF supplies analysis and development support
 - Development of Data Production & Analysis Codes
 - Cal/Val and analysis of data
 - General Support
- ASDC supplies production data support
 - Ingest, production, archive and delivery
 - User services
- Climate Data Records demand max system stability
 - Changes to computing environment effect data products
- ASDC and SCF supply a tightly coupled computing environment for CERES
 - Must evolve in a coordinated manner

Evolving CERES Compute Environment

- 40 years of computation to produce CDRs
 - Must accommodate evolving computing environment
- Four Factors Drive Evolution
 - Need to decrease Cost
 - Increase Capacity to handle re-processing
 - Complete re-processing within 2 years
 - Accommodate improvements/corrections to algorithms
 - Changes to input data (MODIS, GMAO, GEO satellites, etc.)
 - Increasing volume of instrument-months of data
 - Obsolescence of computing hardware and software
 - Recognize that hardware/OS/apps
 - Help from Management
 - IT Security
 - MS Active Directory User Authentication
 - Software Engineering
- NASA Consolidation, Standardization and Interop

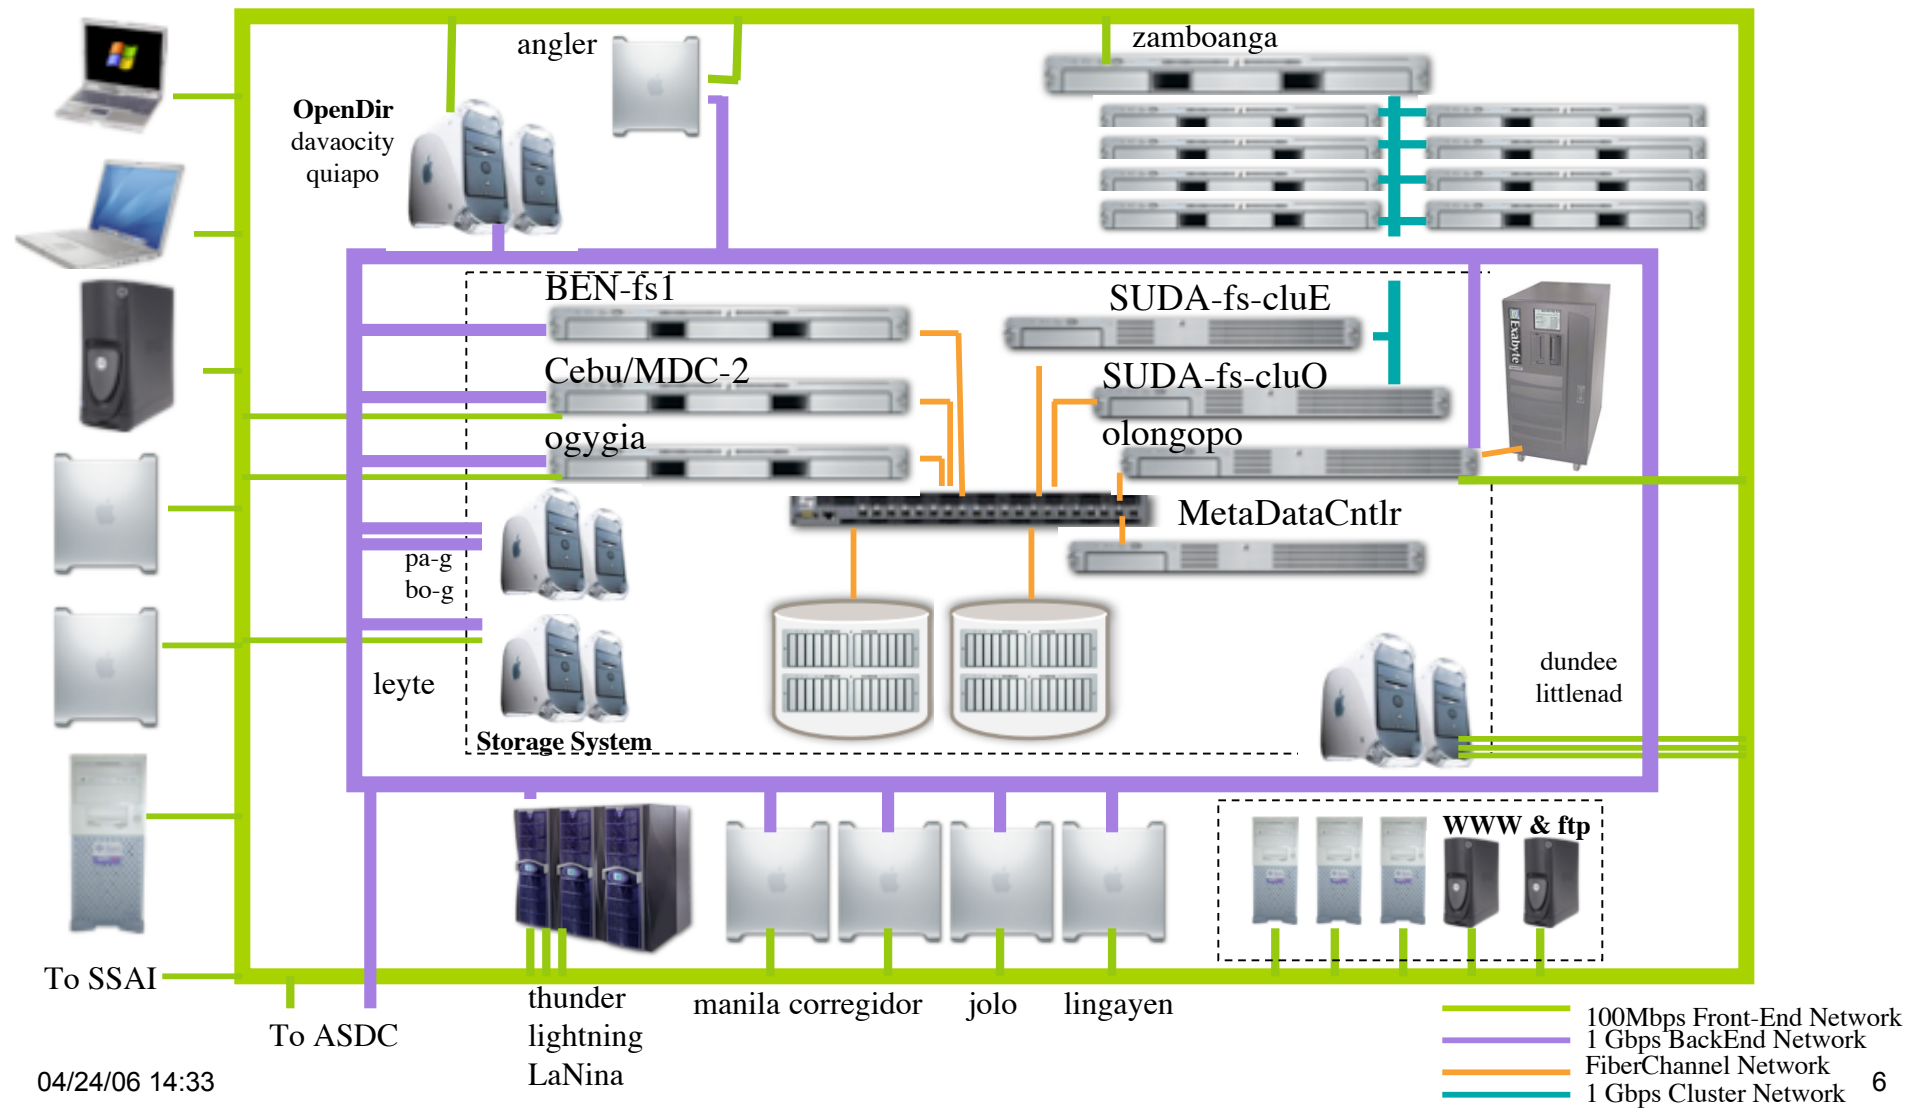
SCF/ASDC Initiatives

- Convert Codes to run on multiple platforms
 - SGI, Mac G5, IBM 970, Mac Intel, AMD/Intel LINUX
- Convert from multiprocessor to cluster computing
 - Big endian multiprocessor to big endian cluster
 - Big endian cluster to little endian cluster
 - Scientifically identical data products
- Increase capacity of shared storage (SAN)
 - Max data available for spontaneous computation
 - Increase reliability of data availability
- Improve SysAdmin Support to Science Team
 - Faq-o-matic
 - Replace help desk support software (e-mail & web)
- Increase capacity/reliability of Web and ftp servers
- Evolve ASDC Storage to ANGe
 - Modernize the ingest/archive/delivery tools
- ASDC implementing CERES Automation
 - Reduce production costs

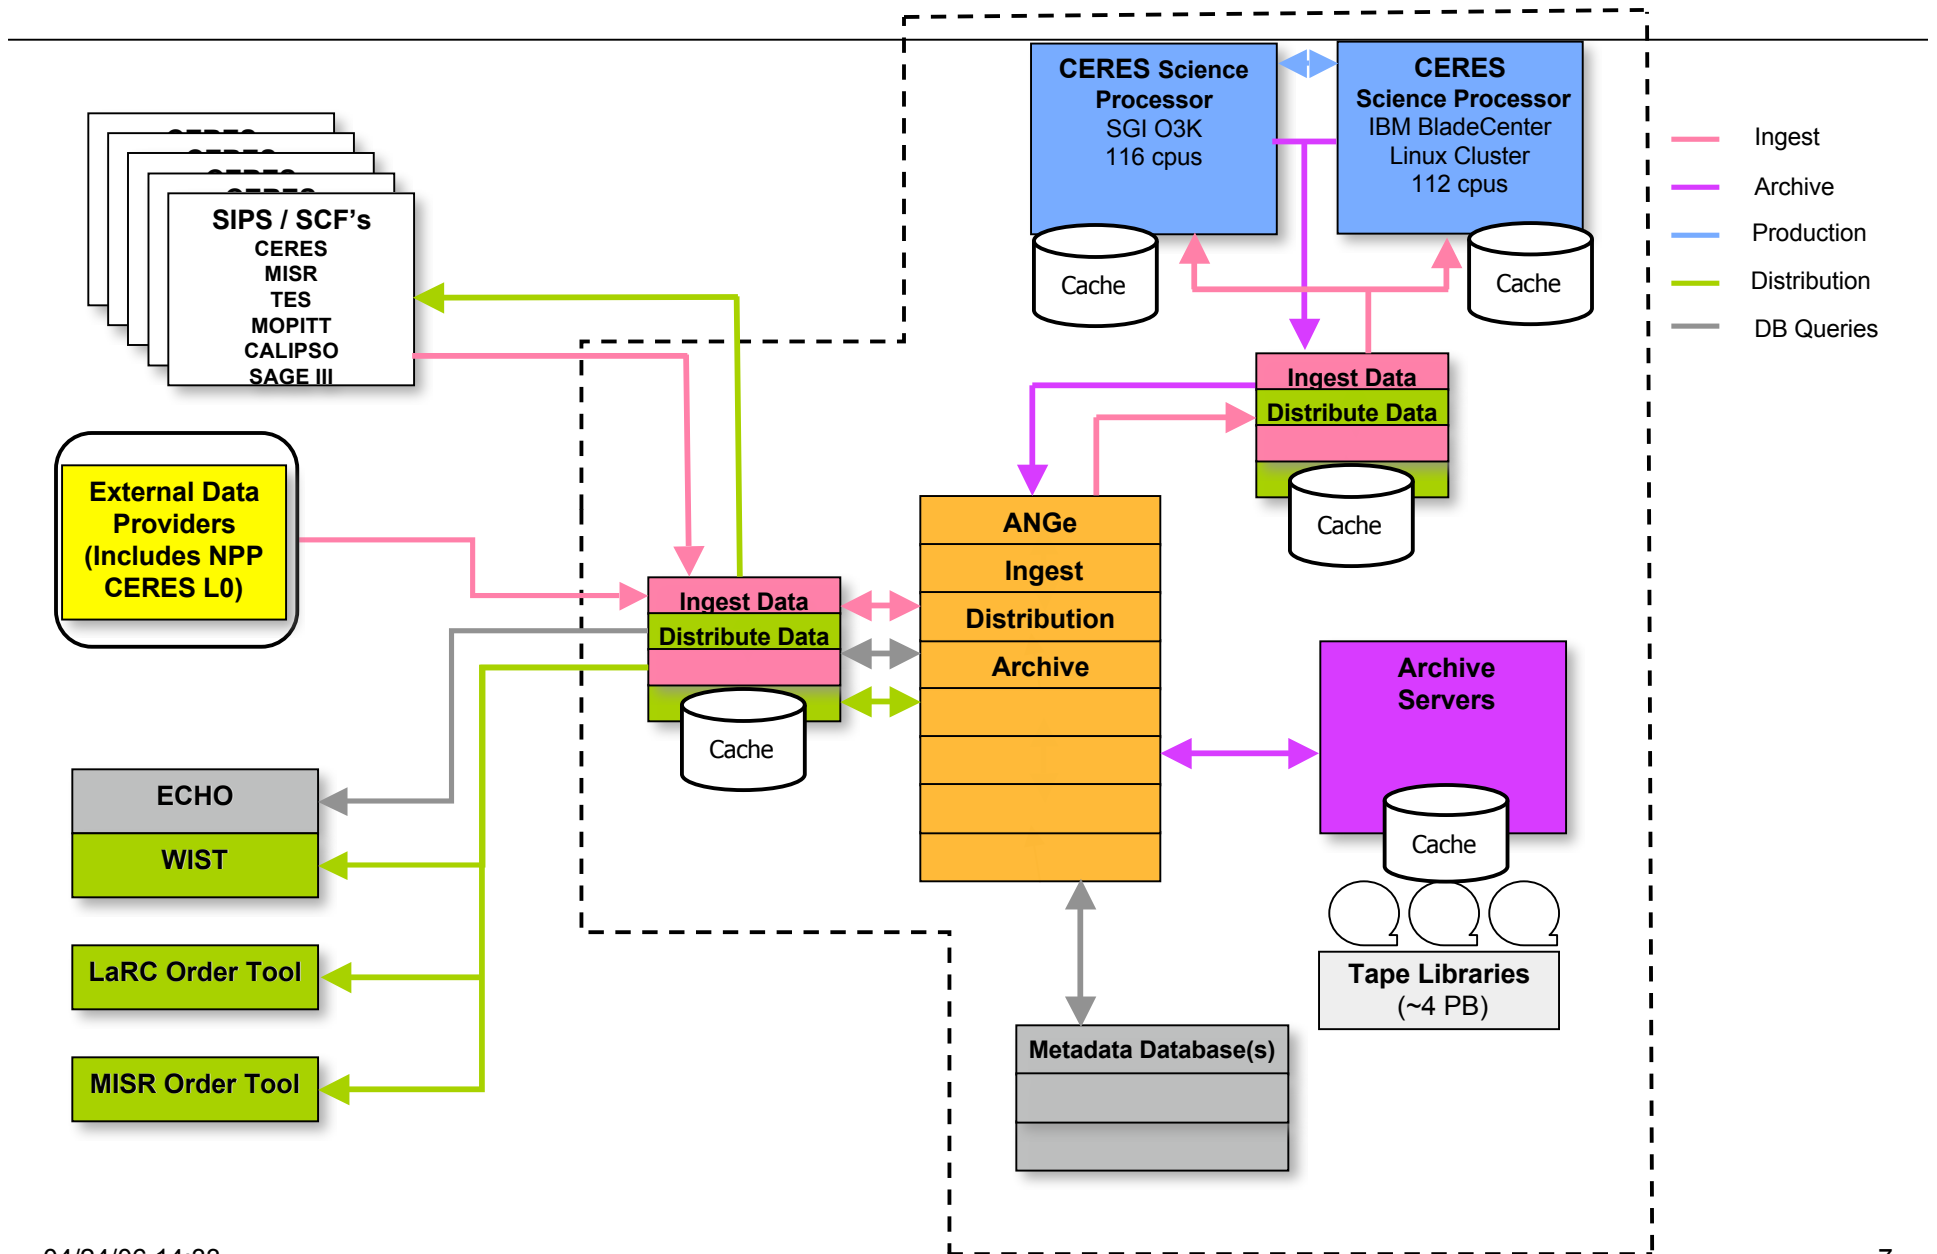
Externally Required Change

- IT Security Certification and Authorization
 - 800 pages of plans
 - New Protective Measures for Systems
 - System Configuration Management processes
 - Increase data collection/retention
 - Contingency Planning and Disaster Recovery
 - Successful test by early summer
- Conversion of e-mail to NASA NOMAD
 - Many processes report results by e-mail
 - All users/correspondents impacted
- Guidance to Minimize GFE to offsite Contractors
 - Modernize SSAI off-site workstations & support
- All-ODIN (whatever that means)
- Standardization of all Windows configurations
 - Hold on WinVista until approved by OMB

CERES SCF Target Architecture



ASDC Operational Environment (2009)



Backup Data

SGI 3800 and Attached Disk

- thunder, lightning (32 processors)
 - SCSI Disk Drives (10TB) experience INCREASING failure rates
 - Not configured as RAID
 - SCSI disks are harder to replenish (not on maintenance)
 - Not all directories are backed up
 - Aging components are more susceptible to power fluctuations
 - Current maintenance expires in early 2008
 - Continuation is estimated to be \$10k per month
 - Will keep in service as long as possible without maintenance
- LaNina (40 processors)
 - FibreChannel disk drives covered by maintenance agreement
 - Generally has been reliable to date
 - Maintenance agreement expires in April, 2008
 - Continuation is estimated to be <\$7k per month for next year

Open Directory

- Purpose: Ensure universal account authorization, file ownership, group access
- Current Status
 - Deployed on single Mac G5 (davaocity) with connectivity into each network in Room 124E
 - Backup system as a mirror
 - Attempting to integrate with NASA and Center Active Directory System
 - All computers mounting SAN via NFS will need to use OD to authenticate to ensure file ownership and group access is correct
 - AFP and Samba do this correctly already

SCF G5 Cluster

- Purpose: provide computational capacity for the CERES (and related) Science and Instrument staff
 - Code Development for data product generation
 - Analysis of data and data product validation studies
- Current Status
 - 47 Nodes (94 processors) managed by SunGridEngine
 - 2 user nodes (manila, corregidor) for compilation and testing
 - 2 More online by 01/01/07 (xxx, xxx) with 4 processors each
 - Nodes will be designated for specific groups (primary and backup)
 - Authenticates users via Open Directory (davaocity)
 - Users submit jobs from manila, corregidor
 - From 2 more on 01/01/07
 - Queues set up based on purchasing project, but can use idle nodes
 - CALIPSO - 2
 - AWIN - 4
 - CERES - 41
 - Designated for Instrument - 4
 - CERES General Purpose - 37
- Near Term work
 - Train users and support gaining experience
 - Assist in science/analysis code conversion
 - Add electrical capacity to computer room (~\$4k)

Storage Area Network (SAN)

- Purpose: to provide universal access to large capacity of data storage with confidence in its reliability
- Current Configuration: 2 SAN, RAID-5, 2 TapeLib (LTO2)
 - inSANity (84.17TB) currently 1 volume, 7 servers, 3 controllers
 - Cebu (FEN), ogygia (FEN), fs-g, fs2-g (BEN), olongopo (FEN, BU),
 - clu-fs(cluster)E+O
 - Insanity: 84.17 TB
 - quickSAN (25.04TB+11): currently 1 volume, 2 servers, 2 controllers,
 - You've got Backup (currently 100TB, soon 200TB)
- Allocations based on SS usage on SGI's & LaCie Disks
 - Predictions by SS reps on DMT
 - Storage Tsar to monitor usage and re-allocate quotas (Scott Zentz)

CERES SCF Comm Architecture

- Purpose: Provide access to the data with low latency and capacity commensurate with computer system performance
- Current Status
 - 6 main ethernet switched networks
 - Front-end Net (larcnet) (connects users to systems) (100Mbps)
 - Back-end Net (connects compute servers for file transfer) (1Gbps)
 - Cluster Net (connects nodes, including file servers) (1Gbps)
 - SUDA metadata Net (connects SAN servers for user access) (100Mbps)
 - ARCH metadata Net (connects SAN servers for archive) (100Mbps)
 - Storage Component Monitoring Net (permits config of RAIDS) (100Mbps)
 - 2 Fiberchannel switched networks (move data among disks/servers)(4Gbps)
 - SUDA Network
 - ARCH Network
- Near Term Plans
 - Clean-up and documentation
 - Tune BEN to max effectiveness in NFS transfers

Code Conversion

- Purpose: Ensure codes run on all platforms available
 - NOW SGI, G5-OS X, IBM-Linux
 - Next Intel-Linux
- Scientifically equivalent data products
 - Needed to permit changing computer systems over long-duration missions without introducing bias or discontinuities into long time series data products
 - Would like to avoid re-processing just to introduce a new HW environment
- Limited Workforce, Not to Interfere with Science/Ed3 work
 - Instrument, SARB, MOA, Synoptic SARB, TISA Grid converted
 - Jim Donaldson currently working on ERBE-like
 - Needed for Instrument Validation
 - Victor Sothcott working on Inversion
 - Clouds and TISA proceeding very slowly due to conflicting science needs
 - Would need to redirect other resources to work Clouds and TISA conversion
 - Temporarily Diverted to test Upgrade IRIX & Toolkit on SGI's for ASDC

System Admin Support

- Purpose: ensure system performance is maintained by detecting problems and correcting them
 - Monitoring systems
 - Preventive and corrective maintenance
 - Adjust system configuration as we gain experience with usage
 - Train & assist users in maximizing effective use of the SCF tools
 - Other user support functions
- Current State
 - Limited Mac experience
 - New Hire plus training for remainder of staff
 - Limited SGI experience
 - Supplemented by ASDC personnel when necessary
 - Not much automation
 - Ben Loyall, part time to help out with this
 - Short 4 personnel (1 trainee has been hired to start in Jan07)

System Logistics Support

- Purpose: Provide information on demand
- Planned Technologies
 - faq system
 - Property and system configuration management Support
 - Web and ftp support

Open SCF Issues

- Humans
 - Recruit 4 qualified SysAdmin personnel
 - Repair lack of documentation
 - Inexperience of user community in using a cluster environment requires more hand holding when the SA resources are short
 - Increase use of automation to improve service, reduce workload
- Facility
 - Power in computer room to be upgraded upon opening of P-card system
 - Currently available 200Amps at 208V, 3 Phase
 - Upgrade/tune comm to ASDC to 1Gbps with jumbo frames
 - Upgrade access to ASDC disk array via fiberchannel
- Cluster/SAN
 - Confirm Maintenance Procedures and Documentation
 - Work through conversion to use Intel based xServes
 - Tape Backup system online
 - Re-mix RAIDS to expand capacity of insanity volume on SUDA0
 - Increase allocations of space to users
 - Clean up and reconfigure ceresarchive
- Computer Security
 - Certification and Authorization (C&A) process imposed by OMB

Status of Conversion by SubSystem

Subsystem	Status	Planned Effort	Comments
1-Instrument	Completed		Current Mac version lost in disk failure; Donaldson rebuilding
2,3 ERBE-like	Converted		
4 Clouds	WIP (95%)		Main processor running.90/10
4 Convolution	Converted		
4 Inversion	Converted		Beta Ed3 for both platforms. Scripts in csh, not perl, but work.
5 SARB	Completed		No IDL on cluster for 5.4P2
6 TISA Grid	Converted		Beta Ed3 for both platforms
7.2.1 Syn SARB	Completed		
7.1 TISA Ave	Not Started	Feb/Mar-June	Depending on Doelling's priorities
8 TISA Ave	Not Started	Feb/Mar-June	Depending on Doelling's priorities
9 TISA Grid	Converted		Beta Ed3 for both platforms
10 TISA Ave	Completed with next delivery		
11 GGEO	WIP	Apr/May	Dependent upon clouds conversion
12 MOA	Completed		
cereslib	Completed		

CERES SAN Status

SubSystem or Project Quota	Pri Grp ID	POC	Actual Usage 3/19	Insanity Alloc 3/19	QSAN Alloc 3/30	Total Alloc 3/30	Need 3/30	Legacy Remain 3/19
CERES (Gen'l) incl archive	400	Scott	4.8GB	200GB	9.76TB	9.96TB		
CERES CM	470	Tammy	77.83GB	150GB	0	150GB		
Instrument	430	Denise	7.05TB	13TB	0	13TB	15TB	
ERBE-like	440	Dale	462GB	1TB	0	1TB		
Inversion	450	Victor	8.63TB	11TB	0	11TB	19TB	
SOFA	1024	Kratz	400GB	1TB	0	1TB		
SARB/MOA	460	Tom	475GB	1TB	0	1TB		
Clouds-DM	421	Sunny	0	2.26TB	0	2.26TB		
Clouds-Science	420	Sunny	2.64	3TB	0	3TB	7TB	
Clouds-Convolve	425	Walt	290GB	1TB	0	1TB		
TISA-DMT	410	Cathy, Raja	1.7TB	5TB	0	5TB		
TISA-Sci (Dave)	495	Dave	194GB	1TB	0	1TB		
RadApp	710	Stackhouse	3.24TB	6.5TB	0	6.5TB		
RFA	730	Stackhouse	756MB	512GB	0	512GB		
Bing	1026	Bing	3.96TB	9.76TB	0	9.76TB		
AWIN	1025	Louis Nguyen	0	0	19.52TB	19.52TB		
NEWS	1027	Sunny	0	10TB	4.64TB	14.64TB		
Laptop Security	1027	Little	0	0	10TB	10TB		