EVLA: Data Management

- EVLA has sub-contracted data management to NRAO Data Management group
- End-to-end processing needs being addressed by DM End-to-end (e2e) project
- Data reduction needs being addressed by AIPS++ project

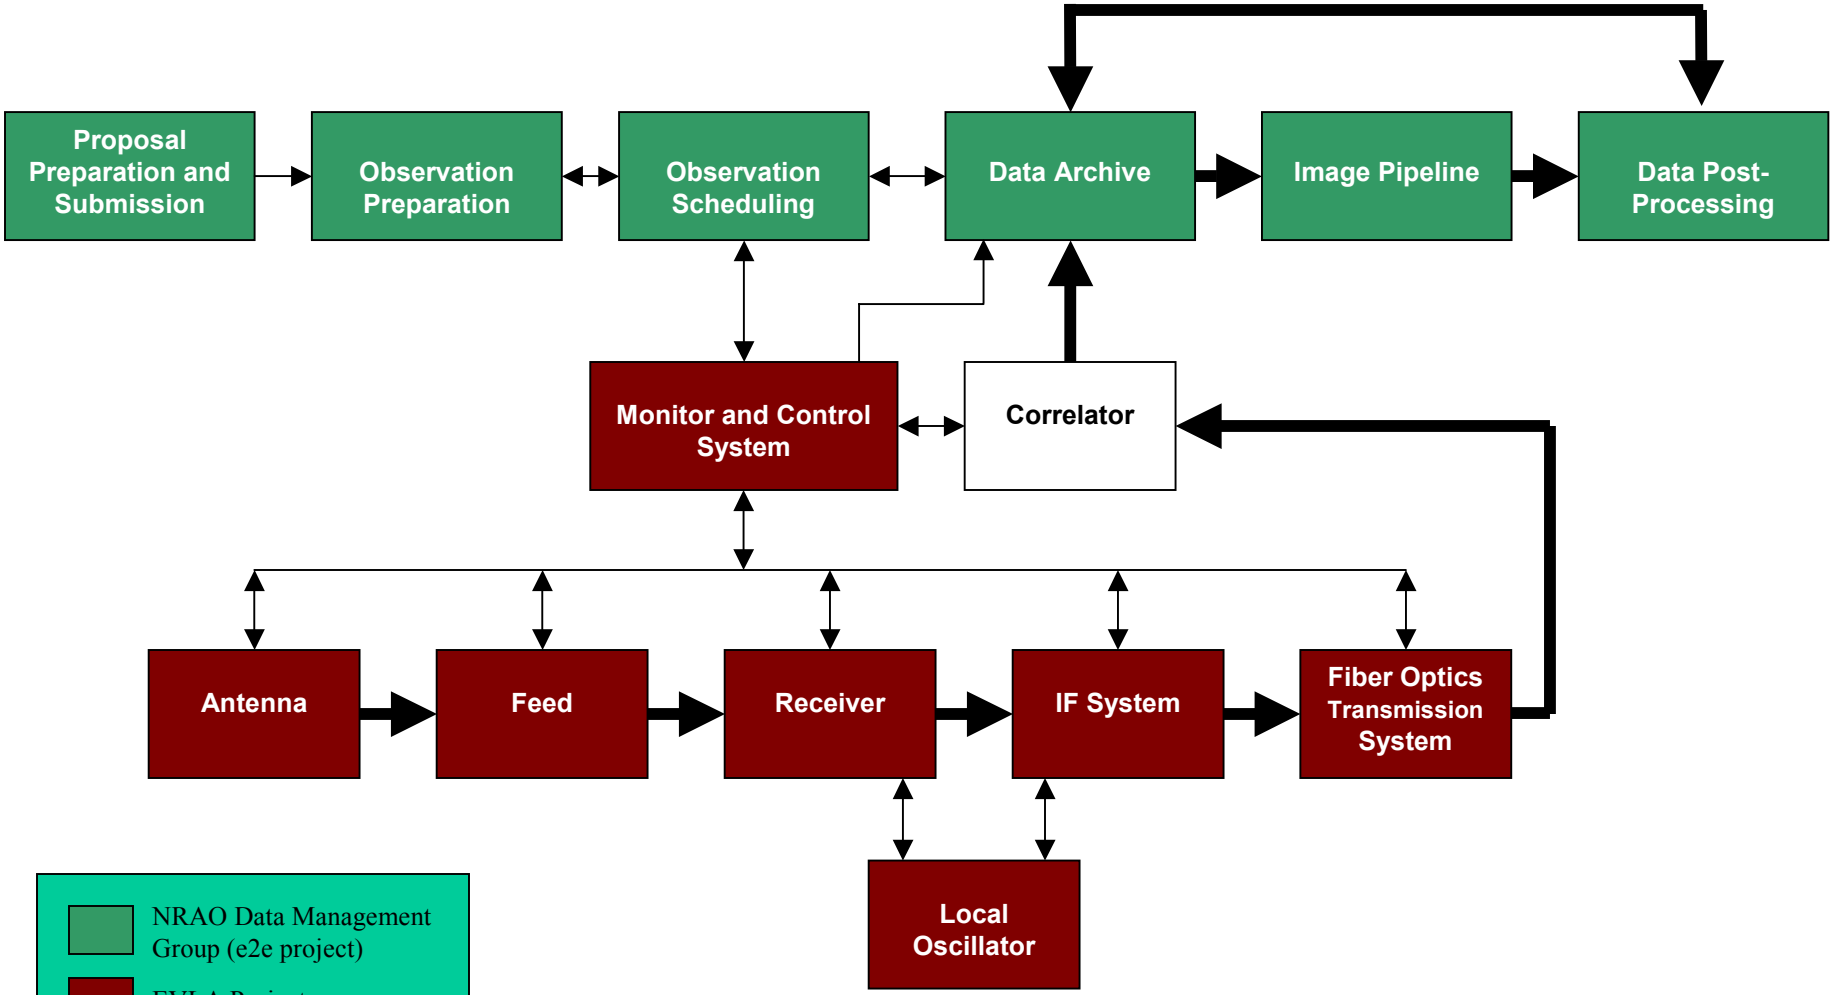

Principal EVLA Subsystems

End-to-end goals

- Streamline observer access to NRAO telescopes
  - *End to end* management from proposal to science
  - Cross-Observatory consistency
- Greatly improve data products to users of NRAO radio telescopes
  - Provide original, calibrated, and auxiliary data, default images and processing scripts
  - Improve monitoring of instrument behavior
- Greatly improve archive access
  - On-line access to archives of contemporary and historical images, surveys, catalogs, etc.
  - Technical and scientific data mining via web and NVO

To reach these goals, initiated End-to-end Project in July 2001

e2e requirements and scope

- Extensive discussion of *first pass* scientific requirements with Scientific Working Group
  - Captured in e2e project book:

    http://www.nrao.edu/e2e/documents/e2eprojectbook.doc
  - Proceeding on basis of current requirements
  - Description of workflow from proposal to observing script
    - Converted to high level architecture and data flow
- Refine scientific requirements at end of phase 1 (July 2002)
- Commit to design and scope at end of phase 2 (April 2003)
  - First e2e advisory group meeting ~ April 2003
- Spending ~ 15% of budget on planning
  - Good way to mitigate against risk

e2e development

- Current staff
  - John Benson, Tim Cornwell, Boyd Waters, Honglin Ye
  - One pipeline developer being hired (funded by ALMA), another later
  - One NVO developer being hired (part of large NSF-funded collaboration)
- Use spiral development model
  - Develop in 9 month phases
    - Get requirements, plan, design, implement, test
    - Review requirements, plan, design, implement, test…..
  - Five year development plan consisting of 7 phases
  - Add new staff incrementally

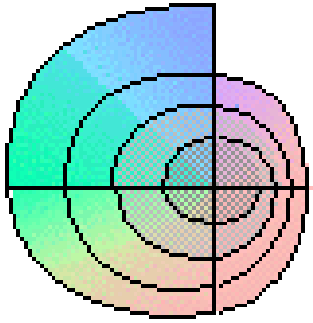

- Three important principles
  1. Keep it simple
  2. Reuse as much as possible
  3. Deliver new capabilities soon and often

e2e Architectural Diagrams

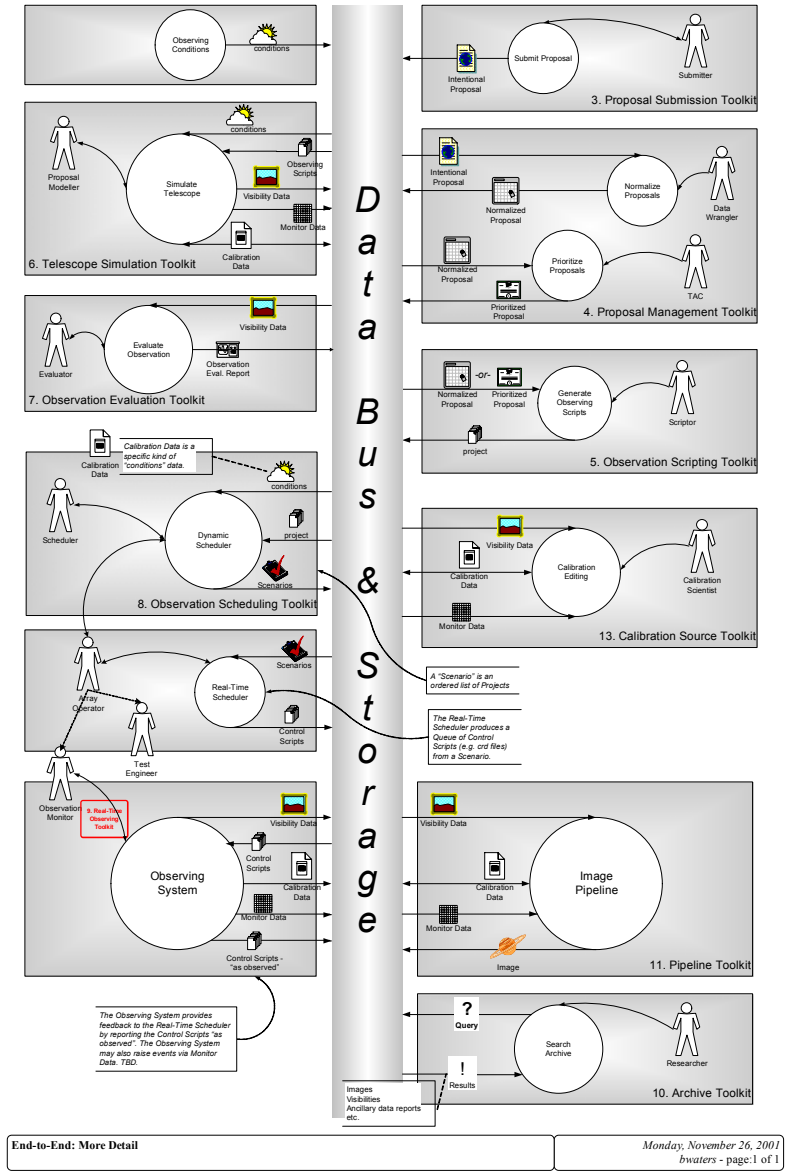

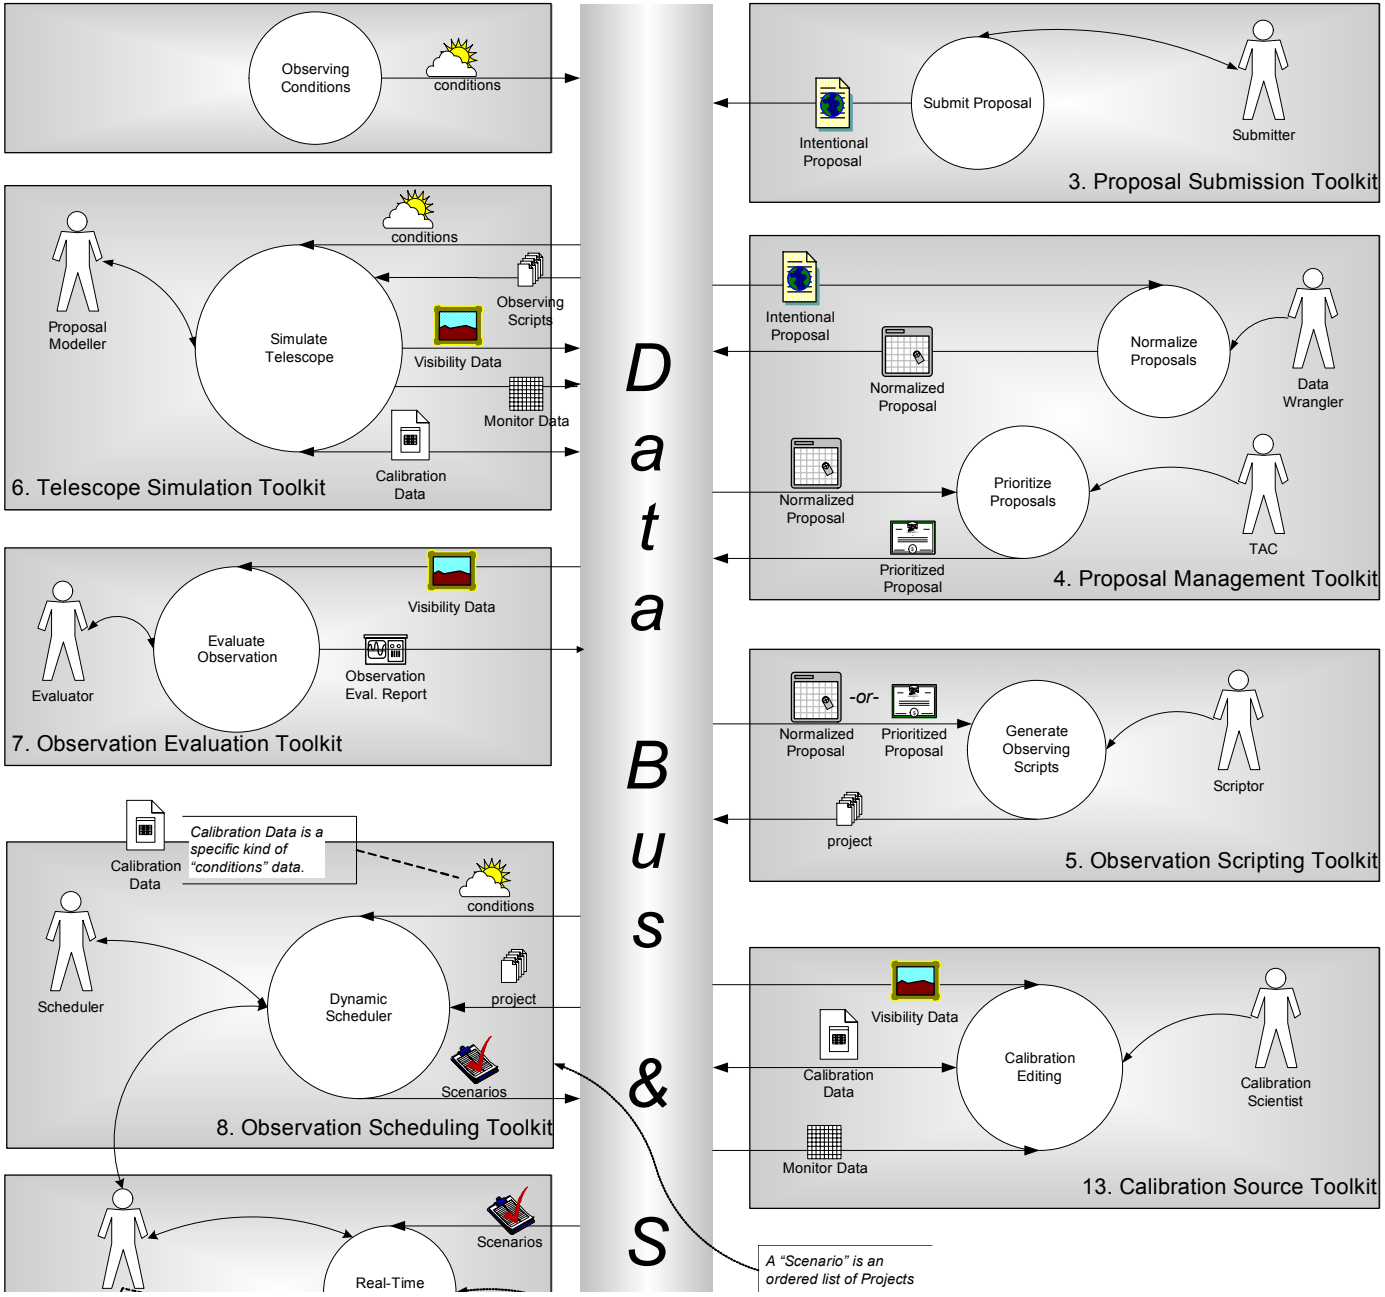
Observing Conditions
conditions
Data Bus & S
Submit Proposal
Intentional Proposal
Submitter
3. Proposal Submission Toolkit
conditions
Proposal Modeller
Simulate Telescope
Observing Scripts
Visibility Data
Monitor Data
Calibration Data
6. Telescope Simulation Toolkit
Intentional Proposal
Normalized Proposal
Normalize Proposals
Data Wrangler
Normalized Proposal
Prioritize Proposals
TAC
Prioritized Proposal
4. Proposal Management Toolkit
Visibility Data
Evaluator
Evaluate Observation
Observation Eval. Report
7. Observation Evaluation Toolkit
Normalized Proposal
-or-
Prioritized Proposal
Generate Observing Scripts
Scriptor
project
5. Observation Scripting Toolkit
Calibration Data
Calibration Data is a specific kind of "conditions" data.
conditions
Scheduler
Dynamic Scheduler
project
Scenarios
8. Observation Scheduling Toolkit
Visibility Data
Calibration Data
Calibration Editing
Calibration Scientist
Monitor Data
13. Calibration Source Toolkit
Scenarios
Real-Time
A "Scenario" is an ordered list of Projects

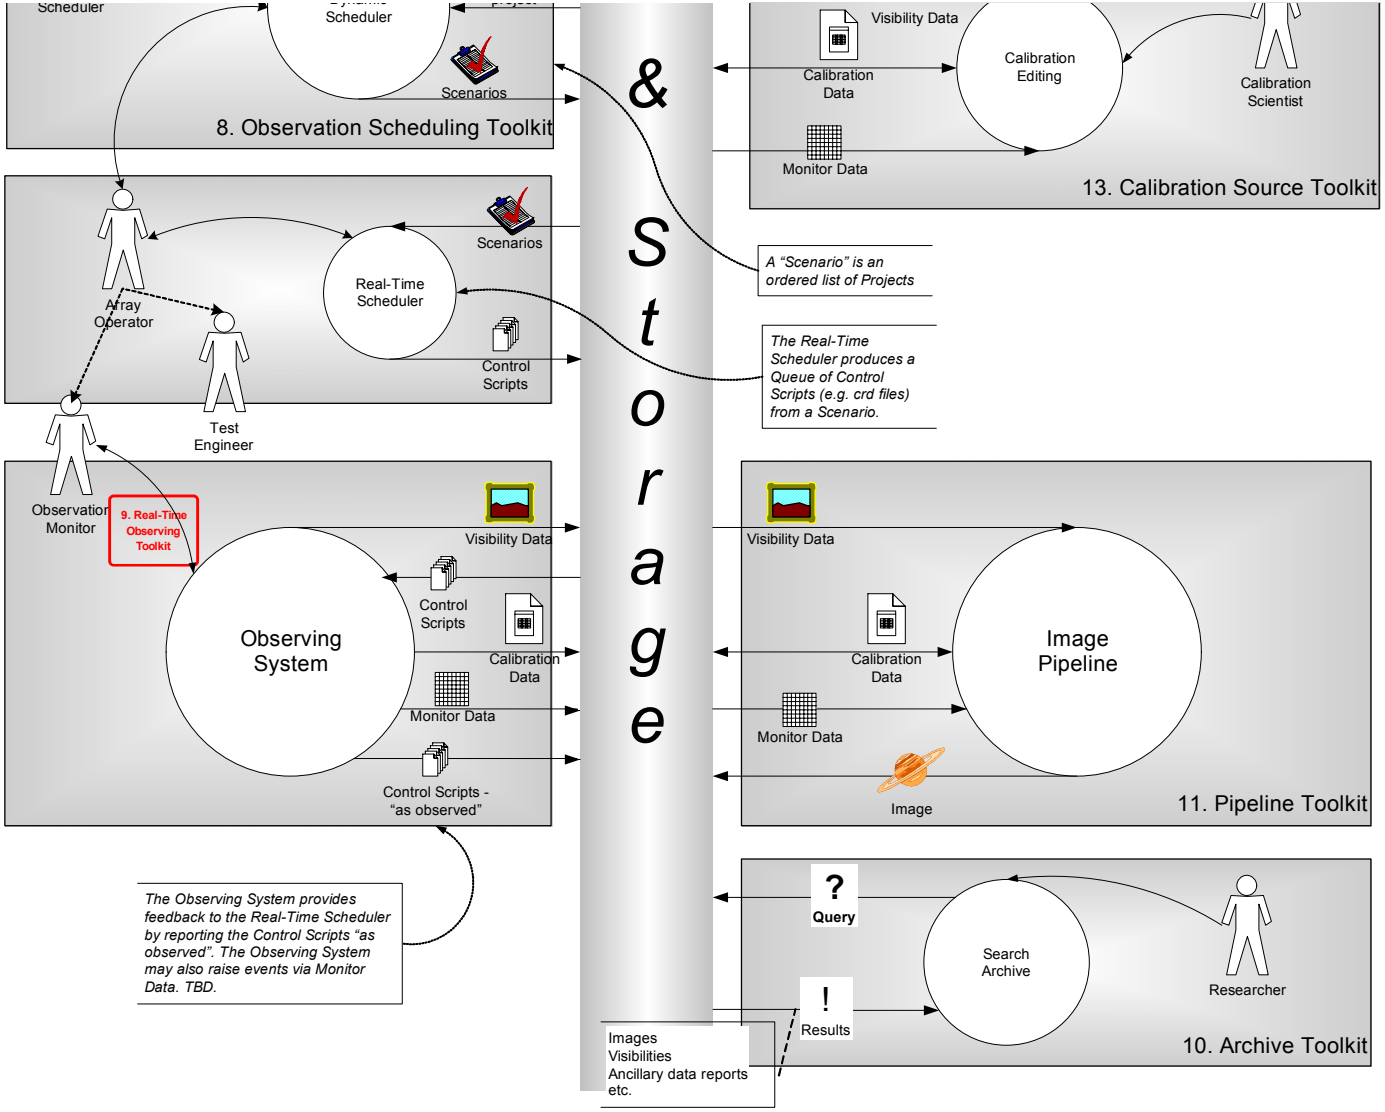

End-to-End: More Detail

Monday, November 26, 2001
bwaters - page:1 of 1

Overall e2e architecture

Data flow

Package	How?	Priority	Status
Operational Model	*Document*	*High*	*First version*
Proposal Submission Toolkit	*Web form or Java-based tool*	*Medium*	*Investigation*
Proposal Management Toolkit	*Java-based tools plus database*	*Medium*	*Investigation*
Telescope Simulation Toolkit	*AIPS++ tools*	*High*	*Deferred*
Observation Evaluation Toolkit	*AIPS++ tools*	*Medium*	*Deferred*
Observation Scripting Toolkit	*GBT Observe, GUI editor*	*High*	*Investigation*
Real Time Observing Toolkit	*Java, AIPS++ tools*	*Low*	*Deferred*
Observation Scheduling Toolkit	*OMS + local adaptations*	*Low*	*Investigations*
Archive Toolkit	*AIPS++ tables + AIPS++ tools*	*High*	*Prototyping*
Pipeline Toolkit	*Production rule software, AIPS++ tools*	*High*	*Prototyping*
Pipeline heuristics	*Glish scripts as production rules*	*High*	*Prototyping*
Calibration source toolkit	*Ingres db + Java*	*High*	*In development*

Pipeline and archive development

- VLA prototype archive and pipeline facility:
  - Under development as part of the first e2e prototype
  - Archive is to be deployed on a 2 TB Storage Area Network
  - Prototype pipeline running on 4-processor Linux IBM x370 system
    - Scope is restricted to small number of VLA observing modes (simple continuum and spectral line)
  - Prototype will test pipeline design, implementation and performance issues on a short time-scale
  - Close collaboration with AOC Computing Division
- Initial archive contents
  - Fill VLA export tapes to disk
    - Plan to offer http/ftp service later this year
  - Tests of GBT and VLBA archiving
  - Mirrors for NVSS, FIRST, *etc.*
  - Results of VLBA service reduction

Archive toolkit

- Prototype only
- Use AIPS++ tools for archiving
  - Assume disk-based storage only
  - Mirrored elsewhere
  - Store telescope data in original formats
  - Fill to AIPS++ MeasurementSets as needed
  - Extract meta-data from AIPS++ MeasurementSets
  - Form archive catalog files in AIPS++ Tables
  - Also designed image archive database for results, surveys, *etc.*
  - Query catalogs
    - From web pages and web services for data retrieval
    - From AIPS++ for data loading and pipeline processing
- Partially funded by 2-yr grant from NSF/ITR program
  - Goal is to allow other observatories to use archive software

VLA Archive tables in AIPS++

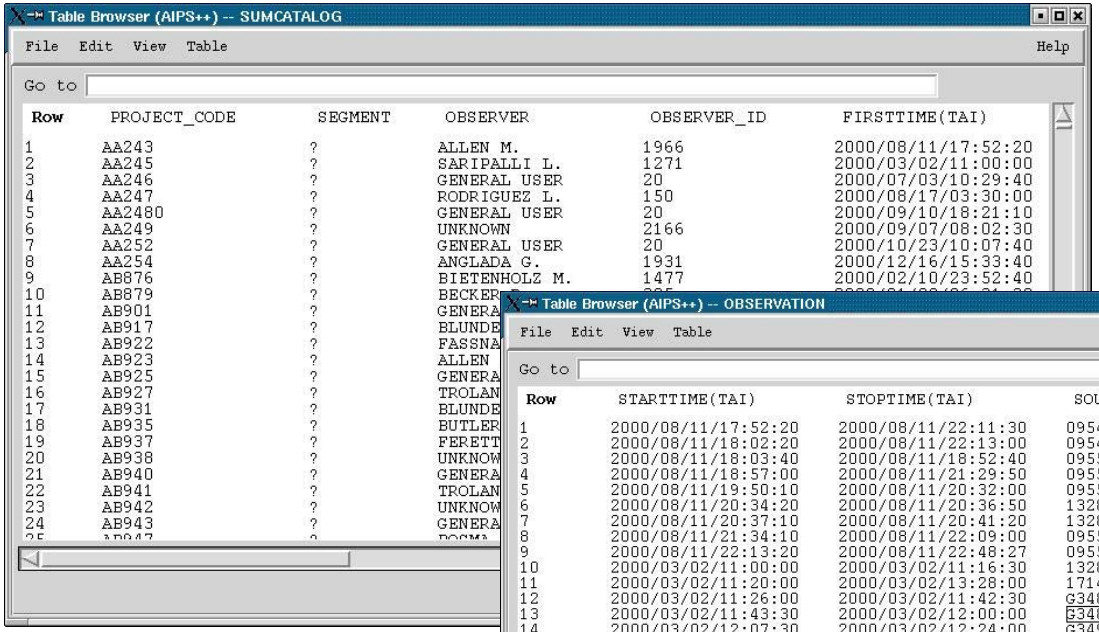

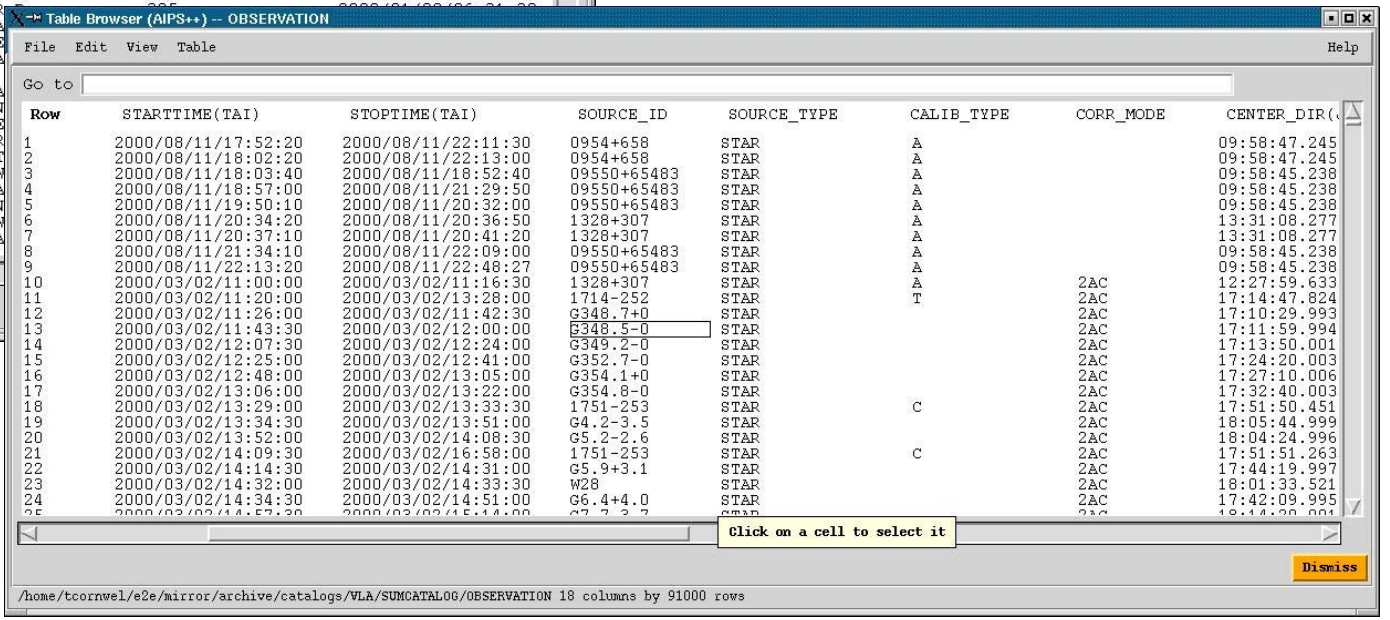

- Filled from existing archive database
- Plan to refill from original VLA tapes as they are loaded to disk

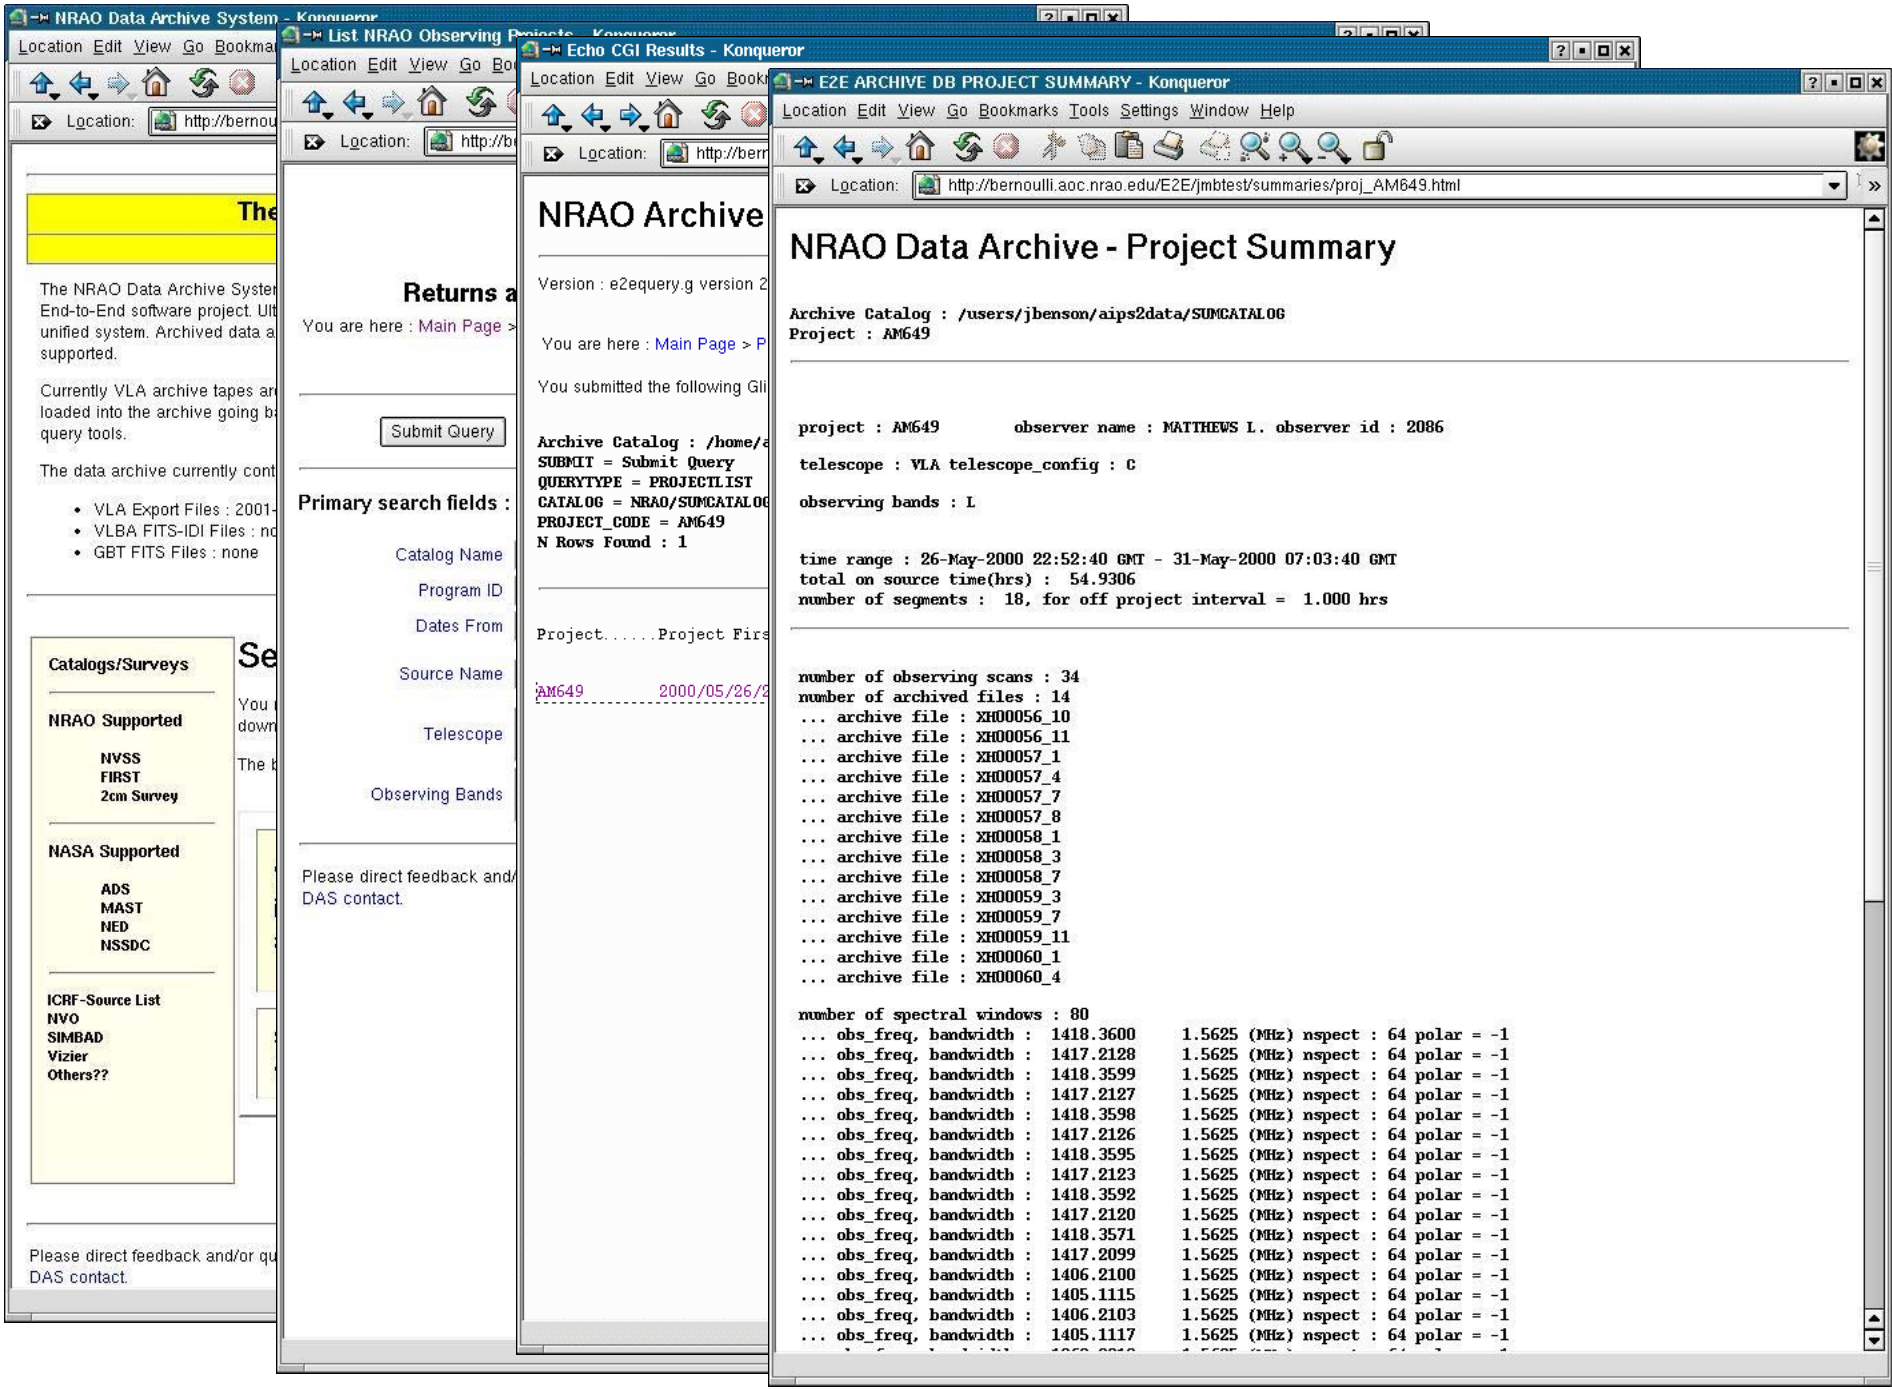
NRAO Data Archive System - Konqueror
Location Edit View Go Bookmarks
Location: http://bernou
The
The NRAO Data Archive Syste
End-to-End software project. Ult
unified system. Archived data a
supported.
Currently VLA archive tapes ar
loaded into the archive going b
query tools.
The data archive currently cont
VLA Export Files : 2001-
VLBA FITS-IDI Files : nc
GBT FITS Files : none
Catalogs/Surveys
NRAO Supported
NVSS
FIRST
2cm Survey
NASA Supported
ADS
MAST
NED
NSSDC
ICRF-Source List
NVO
SIMBAD
Vizier
Others??
Please direct feedback and/or qu
DAS contact.
List NRAO Observing Projects - Konqueror
Location Edit View Go Bo
Location: http://b
Returns a
You are here : Main Page >
Submit Query
Primary search fields :
Catalog Name
Program ID
Dates From
Source Name
Telescope
Observing Bands
Please direct feedback and/
DAS contact.
Echo CGI Results - Konqueror
Location Edit View Go Bookm
Location: http://ber
NRAO Archive
Version : e2equery.g version 2
You are here : Main Page > P
You submitted the following Gli
Archive Catalog : /home/a
SUBMIT = Submit Query
QUERYTYPE = PROJECTLIST
CATALOG = NRAO/SUMCATALOG
PROJECT_CODE = AM649
N Rows Found : 1
Project.....Project Firs
AM649 2000/05/26/2
E2E ARCHIVE DB PROJECT SUMMARY - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location: http://bernoulli.aoc.nrao.edu/E2E/jmbtest/summaries/proj_AM649.html
NRAO Data Archive - Project Summary
Archive Catalog : /users/jbenson/aips2data/SUMCATALOG
Project : AM649
project : AM649 observer name : MATTHEWS L. observer id : 2086
telescope : VLA telescope_config : C
observing bands : L
time range : 26-May-2000 22:52:40 GMT - 31-May-2000 07:03:40 GMT
total on source time(hrs) : 54.9306
number of segments : 18, for off project interval = 1.000 hrs
number of observing scans : 34
number of archived files : 14
... archive file : XH00056_10
... archive file : XH00056_11
... archive file : XH00057_1
... archive file : XH00057_4
... archive file : XH00057_7
... archive file : XH00057_8
... archive file : XH00058_1
... archive file : XH00058_3
... archive file : XH00058_7
... archive file : XH00059_3
... archive file : XH00059_7
... archive file : XH00059_11
... archive file : XH00060_1
... archive file : XH00060_4
number of spectral windows : 80
... obs_freq, bandwidth : 1418.3600 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1417.2128 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1418.3599 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1417.2127 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1418.3598 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1417.2126 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1418.3595 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1417.2123 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1418.3592 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1417.2120 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1418.3571 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1417.2099 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1406.2100 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1405.1115 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1406.2103 1.5625 (MHz) nspect : 64 polar = -1
... obs_freq, bandwidth : 1405.1117 1.5625 (MHz) nspect : 64 polar = -1

Pipeline toolkit

- Prototype only
- Implemented entirely in AIPS++
  - Standard tools for filling, editing, calibration, imaging
  - Use production rules software (make) to capture dependencies and processing rules
  - Rules for processing ("heuristics") expressed as glish/AIPS++ script fragments
    - Encapsulated in makefiles
  - Prototype framework exists and works end-to-end
    - Reads VLA archive tapes (from disk) and produces HTML pages
    - *No human interaction needed*
    - ~ 1000 lines of Glish code for entire framework
  - Working on two examples currently
    - VLA A-configuration, 8 GHz continuum including self-calibration
    - VLA D-configuration, HI synthesis
  - Expect further evolution of framework as experience accumulates
  - Large amount of work in development of heuristics for various situations

Prototype pipeline software design

include
“e2epipeline.g”

myp := e2epipeline()

myp.project(‘AB973’)

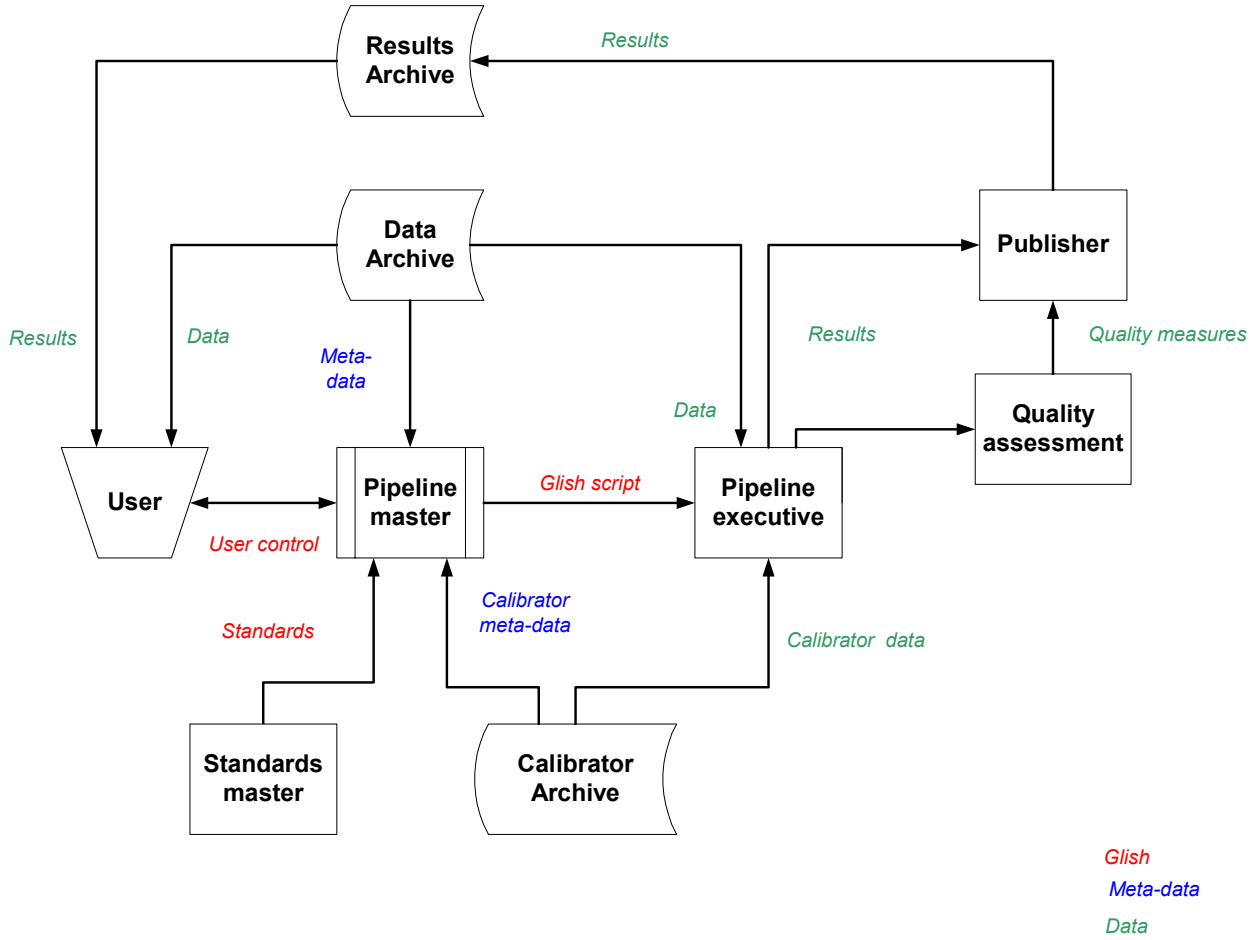

e2e/AIPS++ pipeline results

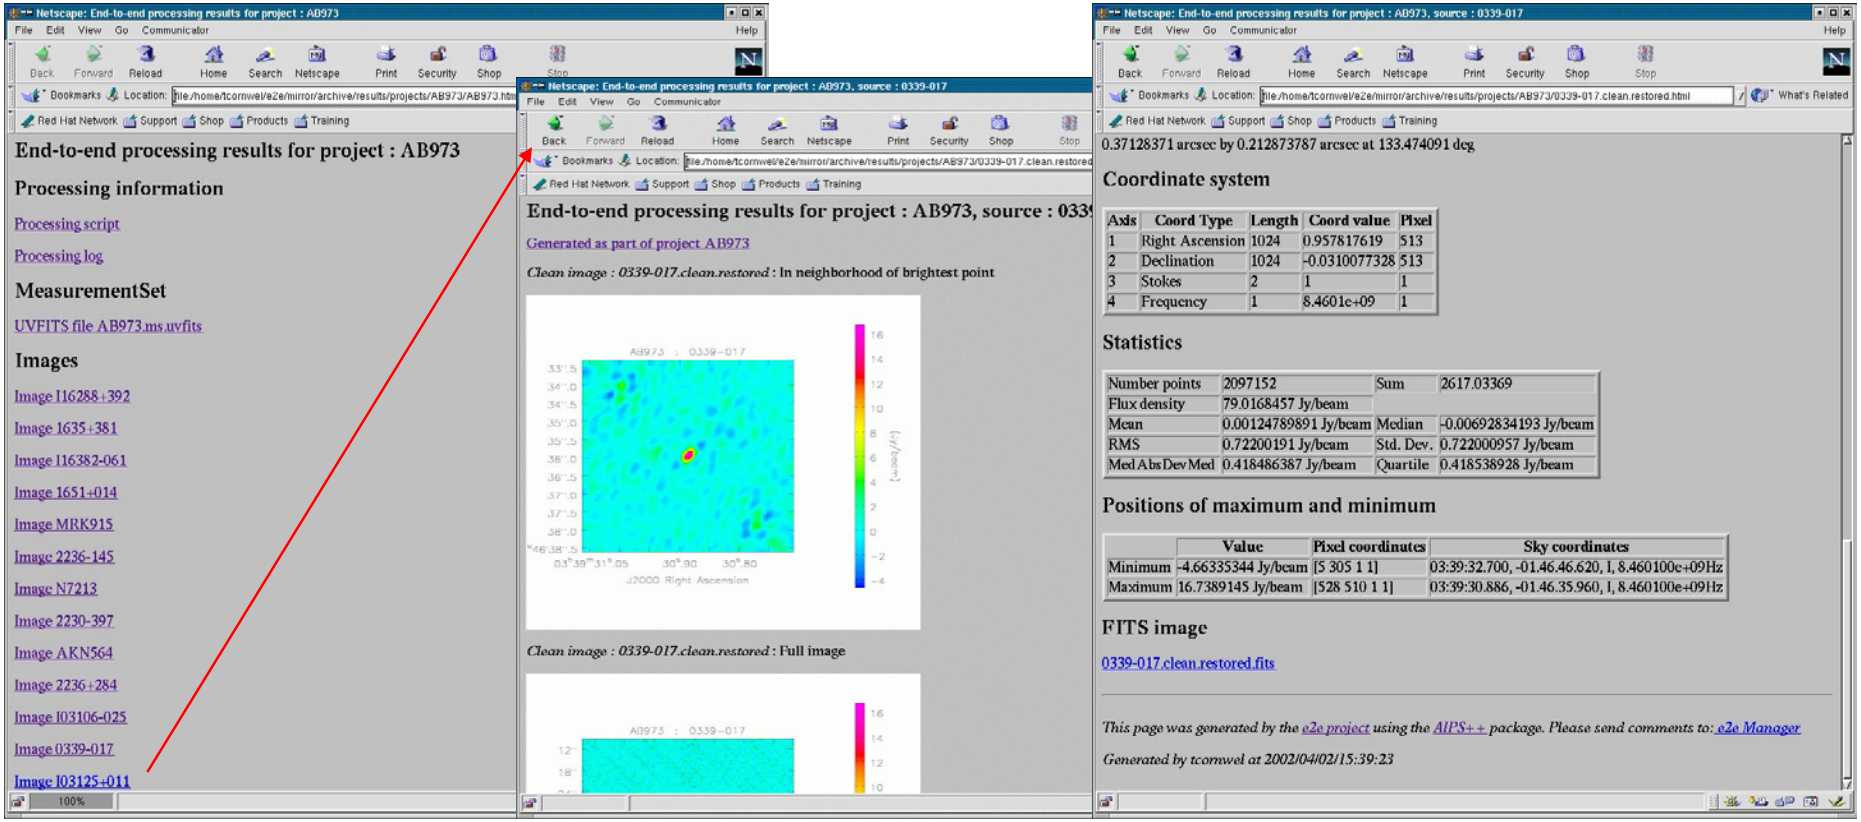

Calibrator Source Toolkit

- Aids selection of calibrator sources
- VLA/VLBA combined database
- Developed very rapidly in Java
- Now undergoing internal user testing
- Expect to deploy outside NRAO domain within a few months
- Make catalog available in AIPS++ for processing
- Also make available as an NVO cone search
  - more later

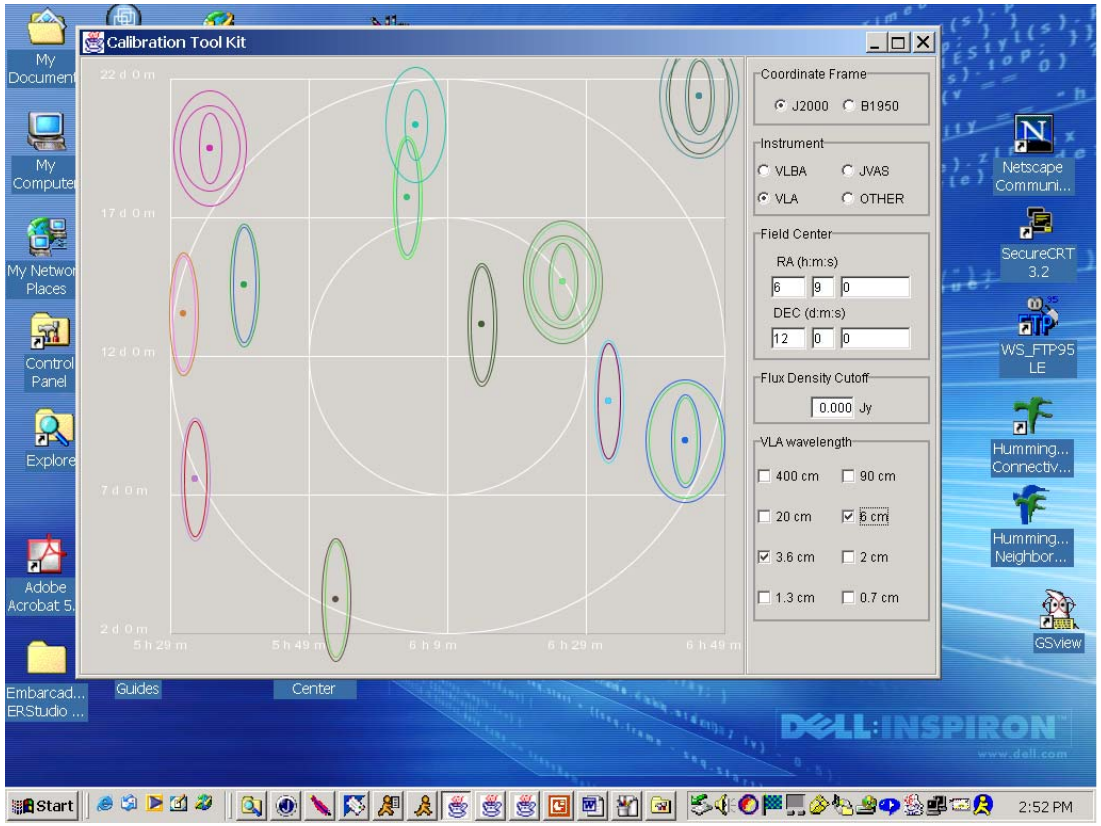

Calibrator Source Toolkit

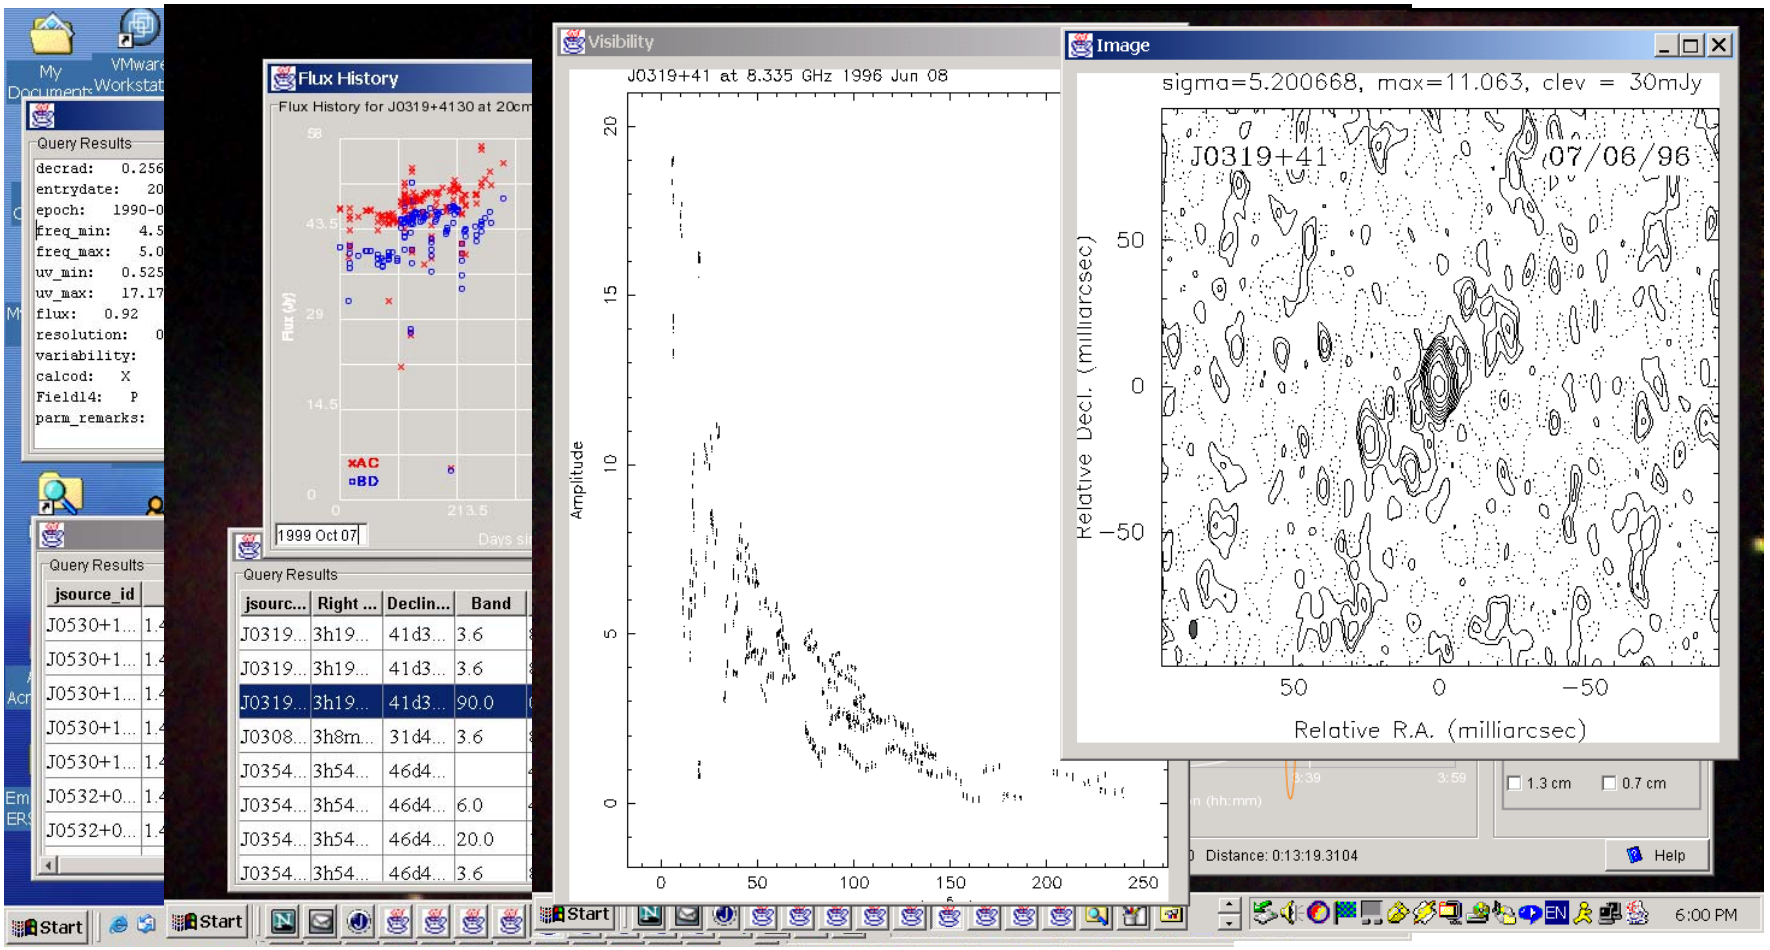

Telescopes and projects

- e2e will be retrofitted to all NRAO telescopes (GBT, VLA, VLBA)
- VLA
  - Putting archive on-line now, working towards pipeline processing
- EVLA
  - Sub-contracted to deliver entire e2e system for EVLA (for 18 FTE-years)
  - Close interaction with EVLA project team at all levels
- VLBA
  - Will start moving archive to disk after VLA archive
  - VLBA pipeline processing once AIPS++ can handle it
- GBT
  - Designing archive facility for deployment in GBT early 2003
  - Watching re-engineering of observing script generation
- ALMA
  - Sub-contracted to develop pipeline (framework only) and post-processing
  - Start development July 2002
  - ALMA has own equivalent to all parts of e2e
  - Trying for reuse if possible (*e.g.* Observation Scripting GUI from ALMA)

e2e timescales

- Customer requirements
  - EVLA PDR process in 2002, Working M&C by early 2004, Shared risk science 2007
  - ALMA development, Phase II starts this year, runs to 2006
  - GBT archive facility by end of proprietary period (early 2003)
  - NSF funding for archive work Sept 2001 – Sept 2003
  - Project book (*http://www.nrao.edu/e2e*) contains scientific requirements as currently understood
- First cycle of development (ends in mid 2002)
  - Move VLA archive to disk
  - Prototype VLA archive and pipeline software and facility
  - Improve support for VLA/VLBA calibrator database
- Second cycle of development (ends in Q2 2003)
  - GBT archive facility
  - Thorough testing of archive and pipeline for VLA
  - Development of prototype observation scripting and scheduling
  - First advisory committee meeting
- End of overall generic development (2006)
  - Working archives, pipelines, ancillary software for VLA, VLBA, GBT
  - First generation for EVLA, ALMA
- Move onto EVLA and ALMA specific development (2006+)

From NRAO to the National Virtual Observatory

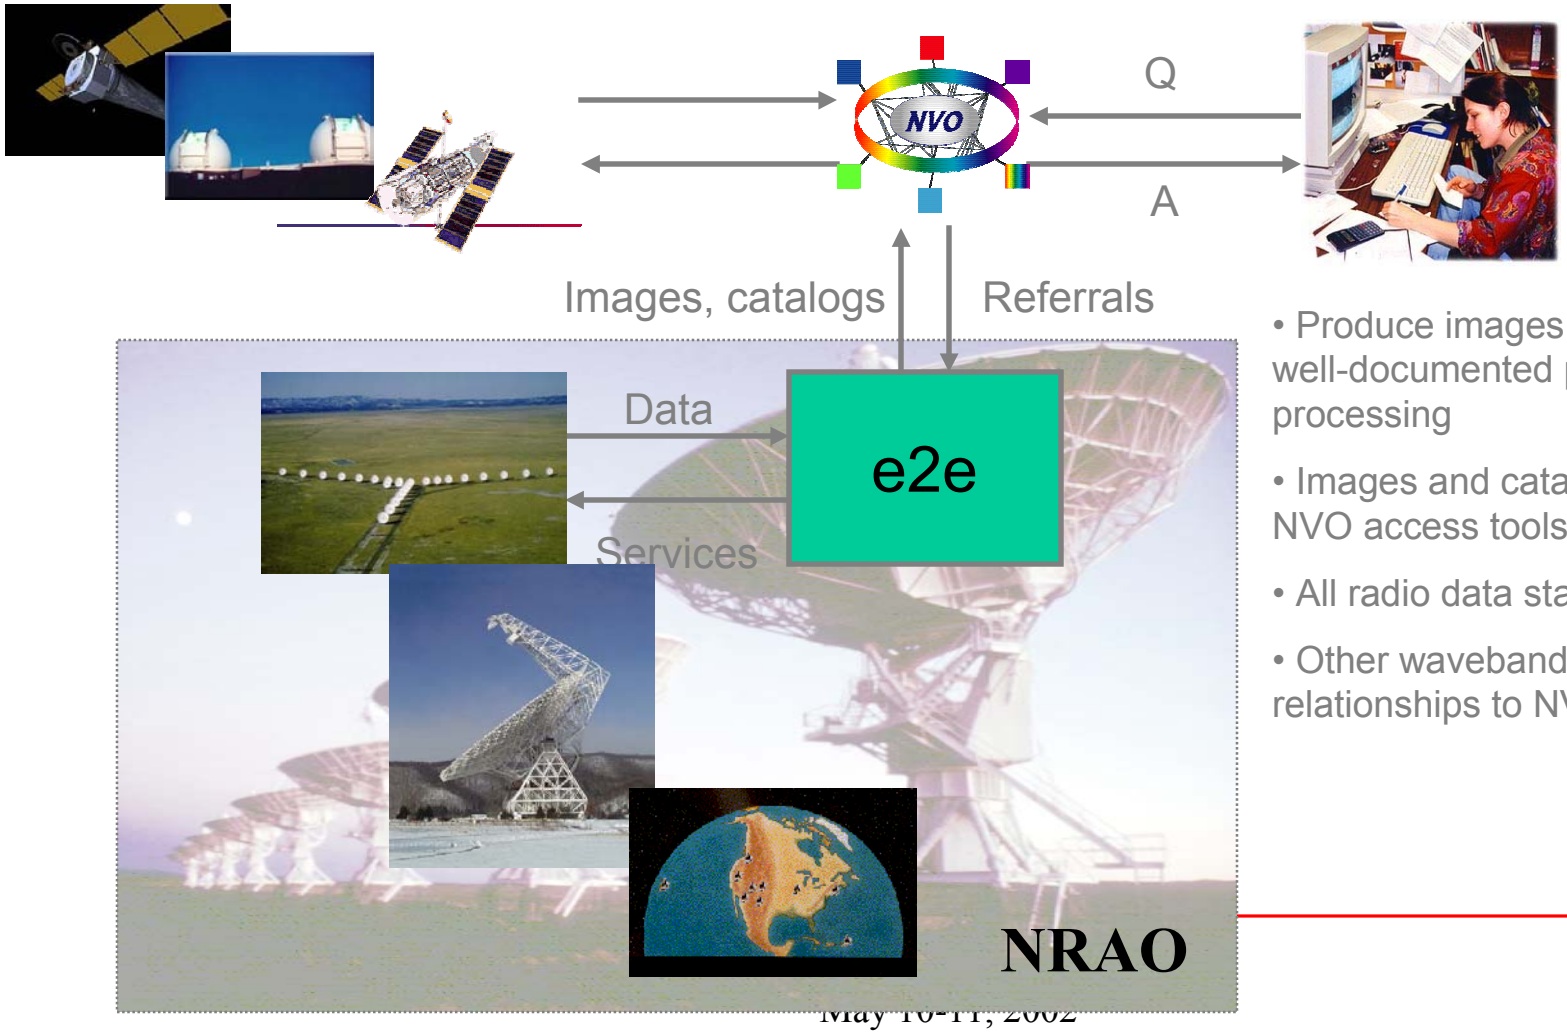

- Produce images and catalogs from well-documented pipeline processing
- Images and catalogs available via NVO access tools
- All radio data stays within NRAO
- Other wavebands have similar relationships to NVO

e2e resources

- ALMA numbers estimated by ALMA computing management
  - Seem to be in line with other ground based projects but considerably less than space based
- e2e numbers based upon straw man designs, reuse
- e2e scope will be adjusted to fit resources (~ 65 FTE-years)
- Neither constitute a detailed bottom-up derivation of resources from requirements

Effort (FTE-years)	*ALMA*	*e2e*
Proposal Handling Software	14	10
Scheduling Software	8	15
Pipeline	12	10
Data Archive	12	20
Post Processing Software	11	10
Total	57	65

Scale of EVLA data processing

- Peak data rate out of correlator backend ~ 25 MB/s
- Total data volume for Peak 8-hr observation ~ 700GB
- Floating point operations per float ~ 100 - 10000
- Peak compute rate ~ 5Tflop
- Average/Peak computing load ~ 0.1
- Average compute rate ~ 0.5Tflop
- Turnaround for 8-hr peak observation ~ 40 minutes
- Average/Peak data volume ~ 0.1
- Data for Average 8-hr observation ~ 70GB
- Data for Average 1-yr ~ 80TB

Detailed analysis

- Analyze processing in terms of FFT and Gridding costs
- Find scaling laws for various types of processing
- Express in terms of 450MHz Pentium III with Ultra-SCSI disk
- Use Moore's Law to scale to *e.g.* 2009
  - Performance/cost doubles every 18 months
- Many more details in EVLA Memo 24

Observation	# pol	FOV arcsec	Cellsize arcsec	Pointings	Facets	Pixels	
L-band full primary beam (2D)	4	7200	0.3	1	256	24000	
L-band full primary beam (3D)	4	7200	0.3	1	1	24000	
RRI Mosaic of SGRA West	2	200	0.2	64	1	1000	
HI of nearby galaxy	2	600	0.5	1	1	1200	
Observation Observation	Data rate Mb/s	Total data GB	Image Mpixel	Visibilities Mvis	Minor cycles	single d	m
L-band full primary beam (2D)	1.87	80.87	576	10108.80	10	28.50	35
L-band full primary beam (3D)	1.87	80.87	9216	10108.80	10	130.48	1
RRI Mosaic of SGRA West	0.58	16.56	128	2070.28	100	19.97	2
HI of nearby galaxy	0.65	56.58	1679.04	7072.12	10	38.30	1

Scale of processing

- Assume Moore's Law holds to 2009
  - Moore himself believes this……..
- Scale:
  - Desktop (2009) will handle many projects
  - Larger projects require ~ 10 – 20 processor parallel machine at NRAO
    - ~ $100K - $200K (2009)
  - Archive ~ 50TB per year
    - ~ $50K - $100K (2009)
- Comparable to scale of processing for ALMA

EVLA-specific post processing

- Mostly well-understood and in place
  - AIPS++ package: can reduce VLA data end-to-end
- EVLA-specific areas requiring more development
  - New modes of processing (next slide)
  - Very large data volumes
    - Automated flagging schemes
- Performance
  - Ensure that AIPS++ is efficient and fast enough (compare to AIPS)
    - AIPS++/AIPS speed ratio ~ 1 +1/-0.5 (with some outliers!)
  - Develop parallelized applications (*e.g.* imaging, calibration)
    - Well in progress in collaboration with NCSA
  - Develop location independent computing (a.k.a. Grid computing)
    - *e.g.* transparent access to archive and pipelines from remote locations

EVLA hard processing problems

Fast-slew mosaicing	~10ms data sampling rate. Remove sliding primary beam.
Full bandwidth synthesis	Deconvolve wide bandwidths while accounting for spectral index, polarization, rotation measures, opacity, *etc.*
Full-beam high-fidelity polarization imaging	Correction of time- and angle-dependent beam polarization.
High fidelity imaging	Image and deconvolve at ~ 10^7. Currently about ~ 100 away from this in best possible cases.
Wide-angle full-beam imaging	Huge images, fast data sampling rates, many imaging facets to accommodate non-coplanar baselines
Wide-angle full-beam imaging	Huge images, fast data sampling rates, many imaging facets to accommodate non-coplanar baselines
RFI mitigation	Removal of RFI post-correlation – requires high data rates

Performance tests

Test: 1.4 GHz P4 Linux RH 7.1	*AIPS++ (sec)*	*AIPS (sec)*	*Ratio (aips++/aips)*
512x512,1000,NA Stokes I	37	16	2.3
512x512,5000,NA Stokes I	41	38	1.1
1024x1024,5000,UN, Stokes I	53	37	1.4
2048x2048,5000,UN, Stokes I	146	61	2.4
512x512,5000,NAStokes IQUV	66	164	0.4
Sort TB->BT	11	34	0.3
UVFITS->local format	(402) 54	(70) (SunOS) 7	(5.7) 7.7

AXAF deep field VLA 5h, 1.3 million visibilities

(Fomalont)

Parallelized wide-field VLA imaging

VLA observations of the Coma cluster (test data courtesy Perley *et al.*)

225 imaging facets, 32 processors, speed-up factor ~20 to a net 10 hours elapsed time

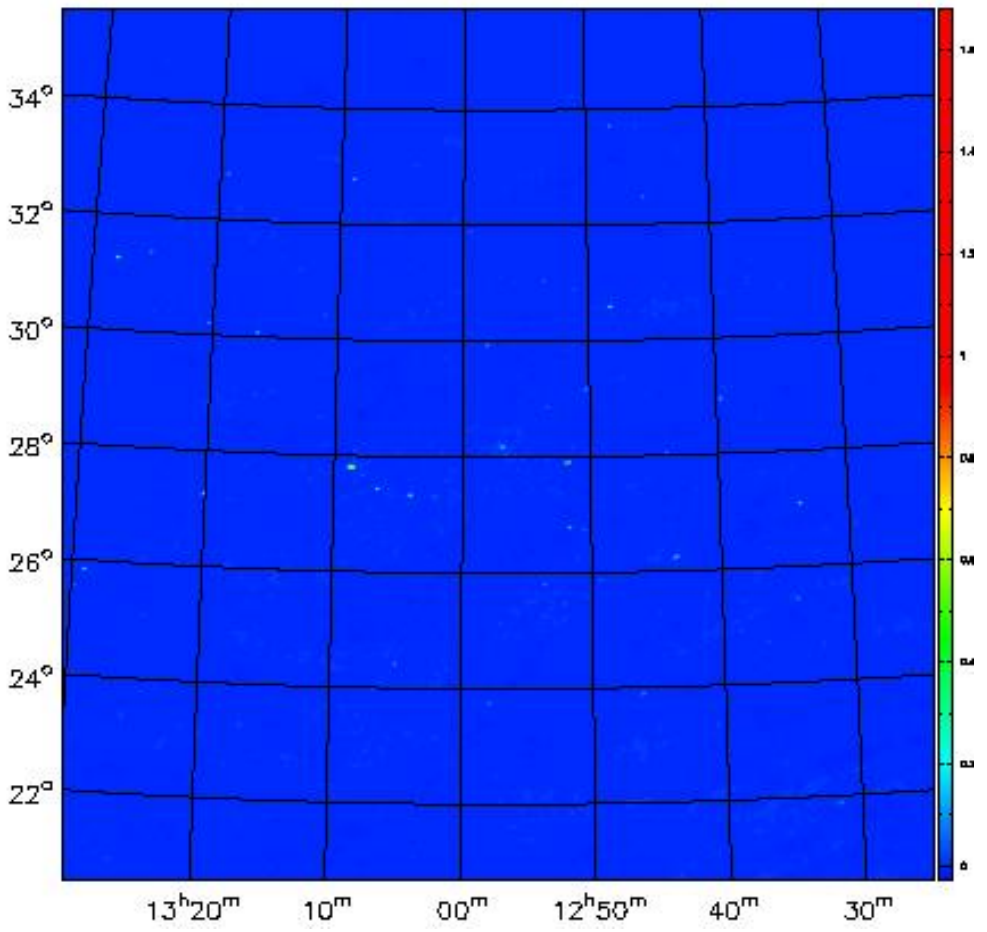

Questions

1. Are resources for e2e adequate?
2. Are there other difficult processing problems?