

OVERALL DATA PROCESSING ARCHITECTURE PDR  
18 - 19 July 2002  
Gustaaf van Moorsel  
10/09/2002

CONTENTS

=====  
A - External Reviewers' Summary

- 1 - Summary
- 2 - Specific Recommendations and responses

B - Other Questions

- 1 - Overview of DM deliverables
- 2 - EVLA priorities
- 3 - e2e architecture
- 4 - Archive
- 5 - Pipeline
- 6 - Telescope Scheduling
- 7 - Proposal Submission/Management, Calibrator Toolkit
- 8 - Correlator Backend Interface
- 9 - M&C Interfaces
- 10 - Post-processing
- 11 - Panel Discussion

A - External Reviewers' Summary

=====  
1 - Summary

-----  
The e2e approach is vital to maximize scientific capability and productivity of the EVLA. Ease of use will encourage new users to do relatively complex projects and allow experts to get the very best from this powerful and flexible instrument. At the same time, it will make it possible for NRAO to manage an instrument which will have more users, more projects, and many new modes of operation without a significant increase in the number of operations staff.

The people involved in the e2e project have an excellent set of computing skills. The addition of Doug Tody and Lindsay Davis to this group in September should further increase the capabilities of the group.

While the e2e project management approach (spiral design) has the

advantage of focusing on just the lap ahead, a strategy for the whole race must also be provided. Since the spiral approach focuses on providing deliverables from a few toolkits in each cycle, there is some risk that insufficient attention will be paid to the design of the communications infrastructure. Communication is the essence of e2e, so this is important.

The current management of the project does not provide for any means for direct agreement between the EVLA project staff (customers) and e2e developers as to what is to be delivered and on what time scale. We recommend a slightly more formal structure in which the observatory staff come to a signed agreement as to what is to be delivered in a development cycle, and another sign-off should be executed upon acceptance of deliverables by the observatory staff at the end of the development cycle. This approach would help avoid any misunderstandings over whether current products have been accepted, and what will be delivered in the next cycle.

The overall administrative and scientific requirements appear to have been correctly identified and the breakdown into key software toolkits looks reasonable. However, the spiral design approach has tended to delay or make more difficult the tasks of defining the requirements and the priorities of deliverables for the final system and the number of human resources required to complete them. Comments and debate from the floor during the meeting imply that there are several important issues yet to be decided and also that some important requirements had fallen between the cracks, for example the real-time display of visibility data. Technical co-operation between the M&C and e2e groups, and the continued discussion of scientific requirements with the VLA/EVLA scientific staff will be important.

However, the detailed scientific requirements have yet to be clearly expressed and there was a lack of focus on the new capabilities and demands of the EVLA, in terms of new and diverse observing modes and experiment types. Consequently we urge that a high priority be given to the tasks of completing user requirements (we note that this must be a largely a responsibility of the user groups, working in conjunction with the e2e group). After user requirements have been completed, priorities can be assigned, and a spiral development can begin on the highest priority items.

The unwillingness to adopt any formal design practices came as a surprise in a project of this magnitude. Some more formal documentation such as use-cases and data-flow diagrams, derived in consultation with the scientists, engineers and operators would help to clarify the

requirements and inform the design process in several areas. While tools such as class diagrams and sequence diagrams may not be appropriate in some areas (e.g data analysis) there are many areas (proposal submission, observation scheduling) where these could be used to great effect. In several areas the designs presented were very sketchy, and not up to the level expected for a PDR, presumably this is a consequence of the spiral approach.

Along with clear requirements, there are a number of critical interfaces which must be agreed in the near future. These include the interface between the e2e system and the M&C system, and between the correlator and the archive.

The increased performance of the EVLA will clearly require considerable work in the development of new algorithms to handle, for instance, RFI mitigation and time variable primary beams. Although the e2e group has identified a number of such areas where considerable work may be required, the e2e group (as well as the aips++ group because of increasing system support requirements) may not have the time to tackle these difficult data processing problems. The aips++ group has considerable expertise in this area but would still benefit from collaboration with scientists from NRAO and other institutes working in this field. A possibility would be for NRAO to devote some of its postdoctoral positions to young researchers in the area of image processing who would bring a new perspective to such problems.

## 2 - Specific Recommendations and responses

-----

1. Continue the process of refining scientific requirements, including specific, significant input from the ELVA scientific staff. Consolidate all ELVA requirements into a single document, so that inconsistencies and omissions can more easily be identified. The e2e group agrees strongly with this recommendation. From the e2e side, Dale Frail will be the prime point of contact for this work.

2. The key interfaces between M&C, e2e and the correlator need to be defined in detail. The e2e group states this is a high priority for them. However, from the e2e side, it is unlikely to be resolved until the end of the second development cycle.

3. Define formal mechanisms to be used for acceptance of current deliverables and agreement on future deliverables for each e2e development cycle. The e2e group welcomes this as an excellent suggestion to be pursued immediately, using the Calibration Source

Toolkit as a test case. Dale Frail will be responsible for negotiating acceptance mechanisms.

4. Develop detailed Use Cases for a range of EVLA observing modes. While these may not explore the whole parameter space, it should be possible to demonstrate in detail how the relevant toolkit in the e2e suite would handle that specific observing mode. The e2e group is convinced that Use Cases are a helpful mechanism for uncovering oversights. It plans to diagram use cases for each toolkit, and may include some for the entire package if found useful at that level.

5. In light of the above recommendations we feel that it would be useful for another PDR to be held at the end of the next development cycle. The e2e group agrees that a second PDR would be useful.

## B - Other Questions

---

### 1 - Overview of DM deliverables

---

It was suggested that a 9-month cycle is too long, and would need to be shortened to one or two months. The e2e group, however, is very pleased with 9 months as the basic cycle time for requirements-planning-analysis-design-implement-test, and thinks it would be a mistake to go faster. The e2e group also points out that there is a faster cycle time for iteration (about a month) once a toolkit is in user testing.

Some surprise was expressed to e2e's claim that the design for proposal submission and management toolkits is 'complete'. The e2e group responds that it is complete in the sense that working from our current specifications, they have performed the necessary analysis and design steps. The design is documented in UNL in the project book.

To the question whether the toolkit will contain telescope simulation, the e2e group responds that this is indeed a strong requirement from many points of view, and that they plan to test the entire scheme using a simulated telescope.

### 2 - EVLA priorities

---

To the question what priority should be assigned to post processing and to those deliverables listed as TBD the response was that prior-

ities will be assigned based on the scientific requirements document to be put together by November, 2002.

A further question concerned possible descoping, as in the original proposal there was no money budgeted for any E2e related activity. The e2e group agrees, and mentions Observation Preparation as an example of where descoping might occur.

### 3 - e2e architecture

-----

Some doubt was expressed about the desirability to archive the EVLA data in AIPS++ measurement sets (MS), as this was something the previous DMSWG recommended fairly strongly against. There are already a number of proprietary data formats in astronomy tied to particular hardware and software, such as the VLA archive format, Mark 3 format, etc. Since there already is an international standard, FITS, as the fundamental storage format, it is suggested to switch the format to FITS or else give a cogent reason why a format associated with a single software system is better.

The e2e group points out that in no sense is the MS a proprietary format - the format is published and controlled by a group including NRAO. FITS specifies only the encoding of the data, not the form of the data itself. A specific format such as UVFITS and FITS-IDI is necessary to store a complex data structure. Neither of these two formats map well onto the MS (they contain less information and have different ways of representing e.g. frequency setups) and so the transformation from UVFITS to MS (say) would be expensive. It would be possible to write binary FITS tables in a format that is readily convertible to the MS (actually AIPS++ uses this as its own internal archive format) but one still incurs the cost of conversion from FITS to AIPS++ tables. The e2e group believes that it is easier to simply write the Measurement Set directly. In practice, the CBE will use AIPS++ libraries to write the MS so the work required is relatively small.

It was recommended it would be advantageous to let prospective observers know where their proposal is in the queue, since adding this possibility of allowing external viewing of the queue took lots of resources at the JCMT. The e2e group is grateful for this suggestion and warning.

The e2e group was asked whether AIPS++ based archival data storage can easily be replaced by other means. Its response is that it can be

replaced but that the work is not trivial. One motivation for using AIPS++ Tables for the archive was to save time. The question of the use of AIPS++ tables will be revisited once a running archive is available, and it is possible that these tables will have to be converted to a relational database. This would require a binding from Glish to a relational database and some reworking of the archive glish code. So far, the use of AIPS++ Tables and Glish has been very helpful and cost-effective, compared to the conventional database approach.

#### 4 - Archive

-----

When asked, the e2e group confirmed that observing schedules and observe files will be stored with the data. There is a spot in the measurement set to store the observe file.

Asked whether there was any scientific involvement with or feedback to the archive, the e2e group responded that the system is barely ready and that there has been not much time to get feedback, especially since there are not many data to retrieve yet.

#### 5 - Pipeline

-----

The question how to decide when to flag data was raised, and it was suggested that flagging may be better tested on simulated data. The e2e group points out this is an area of ongoing research, and that some cases are straightforward, and others are not. Flagging could be based on feedback from a calibrator solution, or based on statistics, such as medians etc, and a simulator would be able to simulate simple types of RFI. The e2e group agrees that a current project should be chosen to test the pipeline.

#### 6 - Telescope Scheduling

-----

Asked for more clarity on formal definitions of sessions, observing blocks, observations, etc, the e2e group responded that such clarity only comes from the design process it is engaged in currently.

There was discussion about 'robotic' versus 'human' telescope operation, and at least one of the reviewers was of the opinion that a telescope cannot be operated in a robotic fashion and that there is a role for an intelligent human. To which several expressed their

doubts that an operator would make a better decision than the system. Also, as for scientific supervision, it was doubted we could afford to have a scientist on call 24/7. The e2e group claims this is an operational issue that must be resolved via the e2e operational model.

A reviewer inquired into how block scheduling affects the possibility of calibration across projects. Others countered that calibration across projects would make the system too complex, and that it would be very unlikely to be useful in vast majority of cases anyway.

## 7 - Proposal Submission/Management, Calibrator Toolkit

---

One reviewer voiced the opinion that a fully automated system for proposal submission and management is inappropriate since many of the functions must be done by people: referee selection, proposal prioritization, time allocation, etc. The e2e group replies that this comment may stem from a misunderstanding of the UML notation that was used to display the use cases and that indeed there are humans in the above mentioned roles. The same reviewer expressed doubt about the usefulness of Java, since everything can be accomplished just as well with HTML and Perl scripting. According to the e2e group, that might (barely) be true for a single subsystem, but for a complete package like e2e, it would be foolhardy to build a complete system without taking advantage of a powerful programming language like Java.

Another question was whether we could draw from experience at other observatories. The e2e group agrees that this is indeed the case: similar capabilities are used throughout the web. Consequently there are many tools (e.g. Enterprise Java Beans) in wide-spread use that can be adopted.

A similar point was raised about the re-use of proposal-handling techniques. Since much of this, including user registration and database handling, is also common to journals dealing with submission of scientific papers, it was suggested to survey methods used by AAS publications (at least) such as ApJ and AJ to see how they handle paper submissions. The e2e group agrees that a survey of mechanisms for paper submissions is an excellent idea.

Asked about the usefulness of the calibrator toolkit for resolved sources, the e2e group replies that the toolkit will have visibility curves available. In addition, the group plans to link this toolkit to tools that script the observations.

## 8 - Correlator Backend Interface

---

Since operators and engineers may want to look at data close to realtime, the question was raised whether this possibility will be provided. Indeed, accessing some or all of the data stream from the Backend in mid-flight has been discussed in the past and it would be possible to 'tap' off a copy of a part of the data stream such as for one or a few baselines. This could be accommodated within the current design of the Backend, but will be limited by technology and budget to a small part of the total data stream. Obtaining data from all baselines so as to be able to do a quick and dirty image would be much more difficult given current and projected near term technological capabilities. Neither has been provided as a requirement by the Project Scientist.

A further, widely shared concern was the time scale on which access to data can be accommodated. In response it was stated that the original output requirement was a maximum sustained rate of 25 Mbytes/sec. A requirement for a short duration 'bursting' capability is currently being quantified and will be added. The ultimate rate and duration of the burst will be limited by technology (CPU speed, bus and network bandwidth, storage volumes and access speed) and by the amount of budget that can be applied.

## 9 - M&C Interfaces

---

Questioned about the existence of a spec on the volume of monitor data the M&C group refers to a document by George Peck, entitled EVLA Memo 34, EVLA Monitor and Control System, Monitor and Control Points, dated 3/12/2002. Based on numbers in that document, we estimate a maximum of 450 Kbytes/sec converging on the control building. This volume of data should not prove to be problematic.

In many ways, a more interesting figure is the anticipated number of packets/sec as this figure gives one a better feel for the ability of nodes on the network to handle the traffic. The above number of 450 Kbytes/sec translates into 360 packets/second, a trivially small number. More recent estimates suggest a higher monitor data rate of at most 6 Mbytes/second, which is 6% utilization of the Gbit network in the control building. This still is a very manageable figure.

## 10 - Post-processing

---



Parallelizing is going to be important in obtaining the necessary through-put. One question was why we don't do 10 - 20 projects in parallel instead of parallelizing the code, to which the e2e group responded that this would adversely affect the turnaround time for individual projects. Asked whether, seen the orders of magnitude increase in data volume, simulations with typical data volumes have been done or are planned, the e2e group responds that it has a strong interest in this issue. The group has recently purchased an extra 1 TB of disk storage for tests of this type.

On the question of establishing post-processing requirements, e2e points out that though there have been analyses of the EVLA post-processing requirements by the EVLA science team, this effort is not yet complete, and requires additional detail and prioritization of needs. Of the total estimated effort in this area (several FTE-months), approximately one third has been completed.

On a possible conflict of interest between VLA specific work and EVLA applications the e2e group comments that the needs of the VLA and EVLA overlap in many areas, and that task scheduling will take place in such a manner as to minimize any adverse impact on the VLA. The availability of scientific requirements is very important to the overall success of the software development effort for EVLA, in order to ensure that needs are addressed in the correct priority order.

The question was raised whether the group has the resources and the necessary skill sets to tackle challenging problems such as RFI. The e2e group responds that although it is fortunate to have very good people in the project, the shortage of people working in this area in general remains a strong concern. This is not helped by the fact that there is no clear career path at NRAO for those who want to do research in image processing.

A related concern is the availability of staff scientists to work on e.g. RFI. Ulvestad points out that Socorro Ops is in a constant balancing act between new developments and supporting current users; he would have to see specific requirements before agreeing to let scientists get involved.

## 11 - Panel Discussion

-----

Points raised by external reviewers are covered in their report.  
Other reviewers:

Sahr endorses an additional e2e design review at a later time, just as for Monitor & Control. He also believes archiving M&C data at an early stage is possible. Once the device interface for a MIB has been defined, the M&C group can proceed to implement this device interface on the CMP, giving a (VLA) monitor data stream that uses the EVLA device interface. This data stream can then be used as the data source for development of a prototype for a monitor data archive for the 1st EVLA test antenna. The device interface definition should be completed by mid-Oct 2002. A date for its implementation on the CMP is difficult to predict at this time, but it should be completed well in advance of the test antenna. We should have a prototype of a monitor data archive up & running in advance of the test antenna, which, in turn, should allow early implementation of a monitor data archive for test antenna data. By the time outfitting and testing of the test antenna is completed, we should have a very good handle on the issue of archiving EVLA M&C data.

Perley sees no problem with the scope of postprocessing, but would like to see prioritization, with which the e2e group agrees. He is concerned about progress on the correlator interface as the new correlator will be far more complicated than the old one. He also stresses the importance of a simulator. The e2e group agrees that a simulator is necessary and that from the development needs, it is required early on. The group believes that a suitable tool can be built on top of the AIPS++ simulator. Perley also points out the severity of the R&D problem: A  $10^6$  dynamic range will be needed, which nobody has ever done yet. Progress in these areas is needed. Perley also stressed the need to address the user confidence problem in view of the connection AIPS++ - e2e. The e2e group is of the opinion that the only answer is to push ahead with deployment of AIPS++ for current NRAO processing.

Clark's main concern are the interfaces. Lots of these are undefined and will have to be defined in short order. The e2e group is concerned about this as well. Discussions between M&C and e2e must continue with high priority. Clark also sees the need to settle more details for the interface to the data archive at an early stage. The e2e group claims that the archive, including API's, continues to be a high priority during the second development cycle. The M&C group response is that there are a number of options for the physical interface including channel attached and networked. The logical interface depends in part on the selection of the physical interface. It is very likely that data will be stored in the form of AIPS++ Measurement sets. It is less clear yet where (on the Backend or in the archive) and by whom (by e2e or the Backend) these Measurement Sets will be

created and this will not be settled until we have had a chance to test and evaluate the various hardware options. This will be done when enough Backend software and hardware is present to conduct the necessary tests.

Clark also asks whose responsibility is reference pointing and auto-phasing. The e2e group responds that this is not totally clear to them either but a working model is that a fast calibration detached from the pipeline will be required. Asked whose responsibility it is to communicate between the scheduling tool and the reduction package, the e2e group responds that the current thought is that it is the responsibility of e2e, but this is likely to become clearer as the e2e-M&C interface is discussed in more detail.

Napier points out that the EVLA should be pushing for its own priorities that may be different from e2e/AIPS++ priorities; these may be needed before proposal handling or dynamic scheduling. The e2e group and the M&C group both agree that this reinforces the need for the EVLA to develop its own requirements and priorities. The M&C group adds, though, that the current focus of the EVLA M&C effort is preparation for the test antenna. Given the overriding importance of the test antenna, the group is of the opinion it would be a mistake to divert people, at this time, to the task of defining the EVLA M&C priorities for the e2e/AIPS++ effort.

Napier also mentions the short-term requirement to get M&C data in and out of the archive, which the e2e claims to have as one of their targets. Finally, Napier stresses the need for NRAO to attract postdocs for algorithm and technical development.