EVLA Project Book, Chapter 13

Transition Plan

Frazer Owen and Jim Ulvestad
Last changed 2004-Dec-09

Revision History
2001-Sep-29: Initial Release
2003-Aug-26: Updated
2004-Nov-20: New Section added

Summary

13.1 Introduction
The current version of this chapter consists of two parts. The first part (13.2) consists of a revision of the requirements for operation during the transition phase. The second part (13.3) is the summary of the major events in our current draft of the full transition plan.

13.2 Transition Requirements
A primary requirement of the implementation of the EVLA is that the current VLA continue observing and producing forefront science during the long period of construction and phasing in of the EVLA. This requirement has been specified by NRAO Users and Visiting Committees. Although it is expected that there will be some periods when the amount of observing time is reduced, and the average number of antennas available may be less than for the nominal VLA, these periods should be minimized in order that scientific observers reap the maximum benefit from the VLA during the period between 2002 and 2009. This transition plan addresses only Phase I of the EVLA Project, as described in the proposal to the National Science Foundation that was submitted in May 2000. Extension to Phase II of the EVLA Project will be addressed at a later time, if appropriate.

Here, we summarize the scientific and operational requirements for the transition from the VLA to the EVLA, as well as the plans for the transition of a few key subsystems. This section emphasizes the requirements rather than going into great detail on the transition plan for individual subsystems. More details on both the designs and the transition plans for several subsystems can be found in earlier chapters of this Project Book.

13.2.1 Scientific & Operational Requirements
We can define a series of scientific and operational requirements that will enable the VLA to maximize its scientific return while antennas, electronics, computing, and correlator are being brought up to the standards described in the Phase I proposal and in earlier chapters. These requirements have evolved via negotiation among the various divisions of the project and the scientific staff, and are stated succinctly in the succeeding subsections.

13.2.2 Scientific Availability

13.2.2.1 Stand-alone Observing
The VLA currently performs scientific observing for an average of about 77% of the hours in a year. We require that the fraction of time that the array is available for scientific observing be maintained at 60% or greater throughout the transition period, averaged over periods on the order of six months to a year. A period of substantial or complete shutdown lasting not longer than three months would be acceptable, though not desirable. If it is found that a greater degree of VLA down-time during the transition period could reduce significantly the cost and/or timescale for

completion of the EVLA, NRAO will solicit recommendations from its advisory committees, including the EVLA External Advisory Committee, about the proper course of action.

13.2.2.2 VLBI Observing
The VLA shall remain capable of participating in scientific VLBI observations, either as a phased array or with a single antenna, with capabilities similar to those available in 2001. If necessary, it is permissible for this capability to be removed for a maximum of two observing trimesters, or approximately eight months.

13.2.2.3 Pie Town Link Observing
Scientific observing with the VLA using the fiber optic link to the Pie Town VLBA antenna has been supported through 2004. Continued use of the Pie Town link after 2004 will be subject to future negotiation and planning.

13.2.3 Antenna and Array Availability
In the text below, "VLA antennas" refers to antennas with the existing IF/LO system communicating through the waveguide. "EVLA antennas" refers to antennas with the new IF/LO system communicating through fiber optics. An "available" antenna is defined as one that supports the full range of capabilities necessary for scientific observing, including monitor/control (M/C), round-trip phase correction, fringe rotation, the full range of delay settings, and normal VLA observing modes such as reference pointing, rapid frequency selection, and fast position switching. There is an important complication in that the EVLA antennas will still have a mix of old and new front-ends for a limited period, so the new IF/LO system must be compatible with both types of front end.

13.2.3.1 General Availability
Currently, the VLA operates with a "3-antenna rule," whereby loss of more than three antennas from the array of 27 antennas either for test purposes or because of "broken" antennas is unacceptable. Loss of more than three antennas during a scientific observing program constitutes sufficient cause for an after-hours callout of personnel to fix misbehaving antennas. During the EVLA construction, it is possible that subarrays of three or four antennas may be used for tests, while individual antennas also may be undergoing extensive checkout. Therefore, the 3-antenna rule may be relaxed to a 5-antenna rule; i.e., no more than five antennas may be missing during a normal scientific observing run.

Details about the exact rules for general availability during particular time periods will be established as construction proceeds and observations with mixed arrays evolve. In general, it is desirable that VLA outages for EVLA testing should take place during less desirable observing periods, such as daytime and/or bad weather, and away from “popular” Local Sidereal Times (e.g., when the Galactic Center is up). It is also permissible to create a new class of contingency observations (e.g., surveys, monitoring experiments, and observations of strong isolated point sources), where astronomers are granted observing time that can be used during tests, with the understanding that more than five antennas may be missing. The “adjacent antenna” rule, whereby the antennas adjacent to the North arm gap in C are designated as “critical antennas,” may be relaxed.

13.2.3.2 EVLA Antennas
Installation of new systems on VLA antennas will continue through 2009. During this period, it is required that the EVLA antennas be capable of scientific operation in conjunction with the "old" VLA. Thus there must be the capability to operate these new antennas with an available monitor/control (M/C) system (either "new" or "old") and to feed their data into the old (or the new, depending on final delivery schedule) correlator in such a way that the identity of an antenna as "VLA" or "EVLA" is transparent to the end scientific user.

13.2.3.3 VLA Antennas
VLA antennas must be operable by either the old or the new M/C system, with an equivalent range of monitor points and information to that which was available to the on-duty telescope operator before the transition.

13.2.4 Array Operations and Documentation

13.2.4.1 Routine Operations

During the transition period, the VLA, including both VLA and EVLA antennas, must be operable routinely by a single telescope operator, for a full range of normal VLA functions. This includes (but is not limited to) such activities as array calibration (e.g., delay, pointing, and baseline runs), holography, and all significant scientific observing modes of the VLA. Control and data output from ancillary equipment such as the Atmospheric Phase Interferometer, water vapor radiometers, ionospheric monitors, weather stations and other such devices must have at least the same availability as before the transition began. At times it may be necessary to have one operator controlling routine operations and a second operator controlling EVLA tests.

13.2.4.2 EVLA Tests

Computer-controlled tests that make use of the operable M/C system and operational VLA capabilities shall be carried out by the on-duty telescope operator. Special tests of EVLA systems will, in general, require the active participation of EVLA software and/or electronics personnel.

13.2.4.3 EVLA Documentation

The Array Operations Division shall be supplied with sufficient documentation and training so that they can operate VLA and EVLA antennas simultaneously for scientific observing, with fractional downtimes consistent with those achieved in 2001. This downtime refers to unscheduled downtime of antennas or the array, and is not to be confused with the fractional availability discussed previously.

13.2.5 Data Management and Data Analysis

13.2.5.1 Scheduling

Scientific users of the EVLA must continue to be able to create schedule files routinely for the EVLA when it contains both VLA and EVLA antennas, and a mix of VLA and EVLA receivers. Telescope parameters, available frequency ranges, and similar evolving items must be kept up to date in the software, with appropriate selections available and warnings issued about available resources. Requirements on dynamic scheduling are not yet determined.

13.2.5.2 Data Calibration

VLA antennas will continue to operate with their 19.2-Hz calibration-switching cycle. Comparable amplitude calibration information shall be supplied on-line for both types of antennas, implying that the EVLA antennas must be able to “mimic” the performance of a VLA antenna with regard to items such as noise-diode switching, Walsh functions, and fringe rotation. The Monitor/Control system must be capable of applying this information correctly from an array with mixed types of antennas.

13.2.5.3 Data Archive

The data archive shall be capable of coping with data sets that are combinations of VLA and EVLA antennas, and shall include all appropriate ancillary information to decipher the antenna properties necessary to analysis of data that are retrieved from the archive at some (considerably) later date.

13.2.6 VLA Infrastructure

EVLA activities during the transition period shall be coordinated with activities required to maintain the VLA infrastructure (e.g., track repairs and antenna painting). Systems such as the waveguide shall be maintained as required, to support scientific observing and equipment health.

13.2.7 Personnel Safety

Current VLA requirements for personnel safety (e.g., use of hard-hats and other personal protective equipment) shall be maintained during the transition period.

13.2.8 EVLA Test Antennas

There is a general desire to have two EVLA test antennas. This is subject to the "5-antenna rule" discussed above in Section 13.2.3.1. If the EVLA Project decides that three or more test antennas are needed, the 5-antenna rule and related parts of the transition plan will be opened up for further discussion.

13.2.9 Electronics Plan

13.2.9.1 EVLA Antennas with VLA Correlator

The signal from an EVLA antenna will be input to the old VLA correlator by taking a wideband digital IF from the EVLA antenna and reconstituting an analog IF using a narrow-band, multi-bit digitizer in the VLA control building.

13.2.9.2 New/Old IF/LO Systems

The IF/LO/FO systems in VLA antennas will be modified in the antenna barn during the period from 2004 through 2009, with approximately four to six antennas modified per year, depending on the funding rate. These antennas will enter the barn as "VLA antennas" and leave the barn as "EVLA antennas" (see definition in Section 13.2.2).

Fringe rotation on the EVLA second LO will have sufficient range and accuracy to provide the fringe-rotation rates required for the old correlator, at the longest baselines and at frequencies up to 50 GHz.

13.2.10 Monitor & Control Plan

A crucial part of the transition plan is the need for the new M/C system to provide sufficient monitor and control capability simultaneously for both VLA and EVLA antennas. A first step in this regard is the complete replacement of the current Modcomps, the correlator controller, and the 9-track tape recording of real-time data. Details are presented in the M/C Chapter.

13.2.11 Correlator Plan

The new WIDAR correlator will be supplied by the DRAO, under funding by the Canadian National Research Council. The general plan is to install the new correlator in the tape/computer/electronics room located immediately to the east of the VLA Control Room. By doing so, it will be possible to run both the old and the new correlators contemporaneously (i.e., the old correlator will not have to be removed in order to install the new correlator). This should enable considerable testing in place, and may prevent a long downtime that would be required if one correlator were to directly replace the other. Note that there is no requirement that the two correlators be fed in parallel (i.e., with data acquired at exactly the same time and split into two simultaneous paths). However, it should be possible to switch between the correlators on a time scale of a few minutes or less, so that comparisons can be made under similar weather conditions. See the Correlator Chapter for detailed schedule for the correlator development and delivery.

13.2.12 Interim Operations Plan

Interim operations for the EVLA will depend critically on the status of the M/C and Data Management developments, as well as any decisions to be made about reduced observing and increased test time. It is expected that considerably more test time (and less scientific observing) will be made available during daylight working hours. In the early years of the EVLA project, this time will be dominated by tests of the new M/C system and the EVLA antennas (including LO/IF/FO capabilities). In the last several years, tests involving the WIDAR correlator are likely to dominate .

13.3 Transition Schedule

13.3.1 Assumptions and Requirements

The transition plan schedule is based first on the hardware milestones. It assumes the accelerated funding schedule for EVLA. If these milestones slip in time, or if the accelerated funding does not occur, then the other parts of the transition plan will slip in a corresponding way. Second, the transition plan depends on the availability of the necessary software at the time it is needed. Third, the plan is based on the necessary personnel being available throughout the plan. Fourth, even if the first three requirements are met on time, the plan must be a living document since we continue to learn more details about EVLA, as we build it. Fifth, this plan only addresses EVLA phase I. When and if EVLA II is funded we will need to add it into the plan. We have a general idea of how to do that but not when.

11.3.2 Hardware Milestones

The number of available antennas and the availability of the WIDAR Correlator drive the transition plan. Some key milestones are given below:

- Dec 2005: Five EVLA antennas available (for use in VLA or as EVLA subarray)
- Apr 2006: Begin Prototype WIDAR Correlator on-the-sky testing
- Dec 2006: Ten EVLA antennas available
- Jan 2007: First part of Final WIDAR Correlator available at the VLA (earliest possible date)
- Dec 2007: Sixteen EVLA antennas available
- Jul 2008: Full WIDAR Correlator Installed
- Dec 2008: Twenty-two EVLA antennas available
- Oct 2009: Twenty-seven EVLA antennas available
- Dec 2009: Twenty-eight EVLA antenna completed
- Dec 2012: Final EVLA Receiver Installed

This plan assumes the accelerated funding plan will continue. It seems prudent to make our plans based on the fastest timescale possible for the hardware. However, we need to remember that any changes in this hardware plan will require changes in the rest of the transition plan.

13.3.3 Science Goals

Based on the accelerated hardware plan in 13.3.2, we can lay out a set of goals for the scientific commissioning and first science with the EVLA.

- Apr 2006: Begin Testing of Prototype 4 station Correlator
- Jan 2007: Begin to offer outside users the new frequencies on EVLA antennas with the old correlator
- Mar 2007: Begin Science Commissioning of WIDAR
- May 2008: "Shared-Risk" Observing begins
- Oct 2009: Regular Observing with WIDAR and 27 antennas begins

In April 2006, we expect to have the 4-station prototype EVLA Correlator at the VLA. This system will only have a small fraction of the channels and bandwidth of the final system. It is intended to test the WIDAR concept. However, it also will allow us to start doing tests with data like we will have with the final system. We also will be able to start exercising some of the post-processing software. We have no plans to offer this system as a user facility.

Before the WIDAR correlator is available we may be able to offer the user community some new capabilities by using the new EVLA receivers and tuning range with the old VLA correlator. We first must have time and the personnel to have our staff test, characterize and document the performance of the new receivers. We also need to document how to tell the transition monitor and control system to use the new capability. Finally, there needs to be enough antennas to make science observations useful. Presently the most likely time to begin offering this capability is the beginning of 2007, or when ten EVLA antennas are available.

At the beginning of 2007, the current WIDAR plan is to install the first boards of the final correlator at the VLA. These boards will be from a limited production run in 2006, planned to test the WIDAR concept with the final board design. Extensive tests with these boards will take place in Penticton, before they arrive at the VLA site. Such a system will provide 8 GHz of bandwidth (out of 16) for eight antennas. When this system is installed and working, we plan to begin our serious commissioning efforts. The goals will be to test the hardware, and the post-processing software, and to define and document a limited number of observing modes which could be offered to a broader group of astronomers for initial science. We think this will take about one year.

In mid-2008, we should have the full WIDAR correlator. At this time (or perhaps slightly earlier), we can begin shared-risk observing. We need to define exactly what shared-risk means. However, we expect it will start with a limited number of outside astronomers coming for extended periods to the VLA, to do projects and to work with our staff, to bring the EVLA into a state in which it can be used by the broader community. We expect this stage also to last about one year.

Sometime in mid-2009, we expect to begin regular observing. We can expect only a limited number of modes still to be debugged with the correlator at this time. In October 2009, we should have 27 EVLA antennas, although not all the receivers will be complete until the end of 2012. We should have our e2e system, including dynamic scheduling of the array, by this time. There will still be a lot to smooth out but we hope to have engaged the entire user community by this time.

Starting in 2009 there will likely be a period comparable to the early days after dedication of the VLA in 1980. We will still be adding receivers and the number of observing modes available should increase for some time. There likely will be some limitations on data rates during this period as well. Currently, the planning suggests that for many EVLA experiments, the data volume and computing hardware required may be too much for most of our users to handle. Also we will still be learning how to use EVLA and the outside user will need to keep up with the latest wisdom on data reduction, as well as contributing to our learning process. Thus, we may need to go back to most calibration and imaging being done at the AOC. Initially, this probably means that users of the EVLA will have to come to the AOC, to reduce their data. As we get used to the processing and as computing and data-transport advances, we may be able to return mostly to remote data reduction either by having our staff do the processing and/or have astronomers at remote sites control the imaging computers at the AOC. Of course, it is likely that some institutions will have their own computing facilities from very early on. We need more planning in this area and we likely will need to feel our way along as hardware and software develop.

All of this plan assumes that the schedules will be met and the resources made available at the times they are needed. Even then, the job to be done is challenging but it seems possible if all goes well. This plan also requires a lot of resources which are not automatically available. Certainly software development is key. Right now it appears that we do not have quite enough software programmers to carry out even the minimal version of the plan. We also will need some additional scientific staff, starting in 2006, to reach the minimal operation levels we have described in 2009. If we want to do more, we need even more people.

13.3.4 Computing Plan

As this version of the transition plan is being written, a comprehensive computing plan for EVLA, including all aspects of post processing, e2e, and the needed computing hardware is still being developed. The internal process to achieve this goal is well underway with very useful, positive discussions taking place at the scientist/programmer level. However, a complete plan still needs to finished and approved by management and our advisory committees.

13.3.4.1 Post-processing

Besides the hardware, the necessary software also needs to be available for the science plan to proceed. The most important thing we need is software and computing hardware which can be used for commissioning the EVLA + WIDAR correlator, starting in 2007. This requires a new imaging algorithm which can handle the wide-band, wide-field case for continuum. Right now we have no way to image wide-band data. Without such an algorithm we will be unable to take advantage of the increased sensitivity that the EVLA hardware can deliver. There are some good ideas about how to handle the problem. However, we do not have even a simulation yet. This is the most important outstanding problem for EVLA and needs to have top priority.

In 2007, we also need a basic post-processing system which will allow at least the in-house users to begin exercising the WIDAR correlator. In 2006, the plan is to use AIPS for the prototype correlator. The development work for the wide-field, wide-band imaging is taking place in AIPS++ and so that package will likely be the first place where such imaging is possible. On the other hand, auto-flagging development is currently taking place in AIPS. Given the strengths of the developers working in both packages, it seems wise at this time to keep both paths open. Unlike in the past, both groups are working with one another and algorithms can be developed in one package and then, if desired, ported to the other package. As we move closer to the EVLA commissioning and later to the telescope being available to the entire community, we may pick one or the other package to emphasize. We may also want to combine the two systems under a single framework and user interface so that from the user's point-of-view AIPS and AIPS++ are one package. These ideas are being discussed and developed but for the time being we plan to make use of both packages for EVLA.

13.3.4.2 e2e

In addition to the post-processing, an e2e plan does not exist which clearly fits in to the budget. In this case the problem is having enough personnel. A basic plan has been developed as to how to reach the minimal e2e plan described in the phase I proposal. This requires some additional programmers very soon. As a fall-back, we can operate the EVLA in the same manner as the VLA with some enhancements.

However, there is also a desire to provide services at the level that have been planned for ALMA. This takes significantly more personnel in the ongoing operations budget after 2010, perhaps eight more employees. This requires significant new operating money in excess of the 2000-level budget.

13.3.4.3 Hardware

Besides the software, it seems likely that a computing cluster will be needed to process the large volume of EVLA data, at least in the larger configurations. A large memory may be as important as the multiple processors. This may be needed sometime in 2008 when we start to have the potential to do the science promised by EVLA. There is money in the budget for this but a detailed plan for how and when this money would be spent needs to be developed. Such a computing system may also change the nature of the operation in the AOC.

In order to process the large A-array datasets, the computing cluster may need to sit next to the archive (or a least one copy of the archive). The most practical configuration is probably that both sit in the AOC. If the real-time processing goals of the EVLA are to be met, this probably requires a wide bandwidth fiber connection between the AOC and the VLA site. The most likely way to achieve this is for NRAO to lease a one-way, dark fiber connection with the VLA site and equip it with the necessary electronics to provide the needed bandwidth. Our internal fiber group certainly knows how to produce the needed electronics. We do need to complete discussions already underway with Western New Mexico Telephone to obtain the dark fiber for a reasonable cost.

A backup hard disk system would also record a copy of the data at the site in the same format as the final archive. Thus, if a failure occurred in the fiber or if an unusual experiment required a greater data-rate then the site disk copy could be physically transported to the AOC for processing.

In this model, the archive/computing cluster then would be the heart of the EVLA imaging system. In 2008 or 2009, it may be possible to purchase a standard, out-of-the-box cluster computer which will handle the EVLA phase I problem. However, it may be that the required system will be more like a correlator than a single computer, i.e. one designed with our special purpose needs in mind, possibly with FPGA processors added to the individual CPUs to speed up the most time-consuming steps in the image making. In 2005 and 2006 we need to converge on what sort of system we really need to handle the EVLA reduction problem and how we will make use of it.

13.3.5 Operations

13.3.5.1 Telescope Operations

During the transition, operations must change from regular VLA operations in 2004, to Interim operations using both old VLA and new EVLA antennas with the old correlator, to EVLA and/or Interim VLA operations in 2007, to full EVLA operations in 2009. This will require a great deal of interaction between the software and operations personnel.

During the 2004-2009 period, the plan is to continue operations activities at the VLA site as we have done for almost thirty years now. However, in 2010 when the changes have settled down the plan is to move the telescope operators to Socorro, except during a 40 hour work week when the day crews are working at the VLA site. This will require higher bandwidth fiber communications with the site and a backup at the site if those communications should fail. This should allow the number of operators to be reduced somewhat due to having VLA/VLBA/VLBA correlator operations in the same location. It may also be possible to locate more of the computer hardware for VLA Monitor and Control at the AOC, which will make maintenance easier and cheaper. These options will continue to be reviewed as the project progresses.

13.3.5.2 AOC Operations

Besides eventually having EVLA operations in the AOC, the nature of what goes on in the AOC will likely change. For many projects the uv data volume is likely too great to ship to the individual user. The calibration and imaging techniques will also be new and evolving with time. Thus in the early days of the scientific use of the EVLA, the astronomer will probably need to come to the AOC to reduce data or possibly control the processing from a remote site of the Internet. The ultimate goal of EVLA is to be able to process the data into images in a semiautomatic pipeline and then let the astronomer retrieve the images from the archive over the Internet. However, doing this will require a learning process that in the early days will likely need to take place at the AOC.

Thus, in the 2008-2012 time-frame, the AOC may need to become the main processing center for reducing EVLA data to images, either with the astronomer in the loop or automatically depending on the problem and how far in the learning process we have advanced. This means that we need to review how provide to the infrastructure and office space to allow for a significant increase in visitors especially in the 2008-2009 time-frame.

As we move toward the longer baselines of EVLA II and other telescopes such as SKA, we probably need to use the special purpose hardware to understand how to deal with the larger demands of these systems. Thus part of EVLA I will be to define the computing systems to tackle the next problem.

13.3.6 Scientific Staffing

The original concept in the 2000 proposal for EVLA included the contribution of the functional support time from five VLA scientific staff a year to the project. Since EVLA was going to be the main development project for the

AOC, this seemed like a contribution we could afford. In practice, since 2000, the size of the AOC staff working on VLA/EVLA has decreased, due to a number of causes, and no replacements have been added. FY05 is the minimum so far. This pattern is not necessarily bad since the EVLA work to date has only needed a few scientists. However, starting in 2006, we begin four years when we need more than five contributed scientists working on bringing up EVLA. Thus, we need to add personnel starting in 2006, to make up the deficit from the earlier years. We need scientists to help commission the hardware, to advise the software developers and test their results and to document the EVLA and advise new users on how to use it. During 2009-2012 we should be able to ramp down the scientific staff to the levels of 2000, if no new demands on their time are made. However, if we want to expand our support services, such as providing good images for most projects, we will need more astronomers to support this activity.

We are about 2.5 FTE's of functional contributed effort short in 2006, to keep the project on the schedule proposed above. Since our high-level scientific staff have either 25 or 50 per cent of their time for their own research, we need four or five people to make up the deficit. We probably need to make up this difference with new hires of either applied postdocs, project appointments or permanent astronomical scientific staff. We will probably need to add one more scientific staff member in each year in 2007 and 2008, before we start to ramp down in 2009. Without the necessary scientific staff support of the EVLA, just as the need for hardware and software, the project is unlikely to reach its goals on time.