EVLA Monitor and Control Transition System Critical Design Review Report

December 5-6, 2006

Stephen Scott, Mark Clark, Hanno Holties, Jeff Kern

Summary

The committee was very impressed with the state of the project, both from the excellent presentations and from the fact that upgraded EVLA antennas have been integrated into scientific operations. The task of maintaining a working array during the upgrade to the EVLA is technically complex and demanding, and the transition team seems to be on the path to success. While the committee was challenged to find issues worth documenting, we include our recommendations and observations in the spirit of expediting the identification of problems earlier than normal. Some of these issues are obvious or were presented during the review, but we have documented them here for completeness.

The Charge

The committee was asked to address the following items:

1. Are the design requirements for the M&C transition system complete?
2. Will the architecture and design of the transition system satisfy the requirements?

In regards to the first item, the committee felt that the design requirements were complete. On the second item, the committee felt that the architecture and design would satisfy the requirements, however there were four main areas that would benefit from further attention: access control, alert system, UDP protocol, and the VCI. Explicit recommendations are made on these areas later in the report.

Recommendations

Access Control

We recommend that further design work be done on access control to the operational interfaces and the Monitor and Control systems. The requirements contain sufficient detail but there was no clear path for the implementation. Although all EVLA M&C systems are within a subnetwork that can only be accessed from outside through VPN, this provides only security for misuse from outside the AOC. It does not provide security for, intended or mistaken, misuse from anybody that has access to the subnetwork. A possible reason that the design has been deferred may be the feeling that access control could impinge upon the flexibility needed during the transition. A design would not need to be implemented immediately to its full extent nor would the final controls need to be enabled. However an early design would give the opportunity to develop the necessary hooks into the infrastructure during the earlier implementation stages. As is stated in the EVLA Array Operations Software Requirements: An attempt to add security to the system after the fact would likely prove to be a costly and difficult task.

Alert System

The alert system needs a thorough design effort. It presently seems to be an extension of the monitor system, but it needs an independent design that has scope beyond the MIB framework and encompasses the entire software system. This work is likely to be significant in scope.

UDP Protocol

The innovative use of UDP as a primary communication protocol should be placed within a design that is robust to packet loss. We recommend a thorough analysis of all communication paths and recovery mechanisms, and the effects of any communication failures on scientific data and array operations. The protocol chosen is unusual for a monitor and control application, and repeated questions on the design will arise. Overall effort will be minimized if a white paper (perhaps an EVLA memo) is produced that outlines the design with an analysis of traffic (average and peak), available bandwidth, buffers, the impact of dropped packets, etc. This recommendation should not be misconstrued as questioning the efficacy of the chosen design.

VCI

The design for the Virtual Correlator Interface seems stale and it is time for renewed effort on it. The correlator is a pacing item for the schedule, so the VCI should be as ready as possible to optimize the integration.

Further Recommendations

The project is aware that Jython does not appear to be a good bet at this juncture. We would encourage the project to consider Python, the scripting language of two other NRAO instruments, the EVLA/ALMA data reduction framework and numerous applications across astronomy.

Contributions from outside the group are significant, which shows flexibility in management at the AOC. We recommend that continued care be taken in managing contributions from outside the group requiring maintenance over the long term.

Similarly, as human resources are limited within the project, single points of failure are a risk that needs to be managed carefully. In particular, it is recommended to give sufficient attention to documentation and knowledge transfer.

The implementation of code reviews is applauded, and their costs acknowledged. The project is encouraged to continue them at some level. As mentioned in the presentation, the reviews provide an excellent way to share knowledge of critical components across the team.

We recommend increasing the use of unit tests, and that they be run automatically, perhaps as part of the nightly build. A coverage metric should be added as well. Unit tests for the interfaces that will be used in both the transition and final system will prove especially valuable as they will provide an anchor as the system evolves to the final system.

We recommend that a units field be associated with the data in the monitor database. Additionally, the project would benefit from some standardization on units for similar datatypes, both in code and the database.

While we were impressed with the GUIs that were shown in the review, we caution against expending unnecessary effort on GUIs that have a limited lifetime. In particular, GUIs that are likely to only be used in the early stages of testing can lack sophistication. We endorse the decision to descope of the Plotting and Analysis Tools.

Two options were presented for the path for the production of the final scripts by the Observation Preparation Tool. We recommend that the option that directly produces scripts in the final scripting language be chosen. The other option, using intermediate observe files, might have the undesired effect of requiring the system to accept observe files indefinitely and will suffer from limitations in the observe files.

We recommend more organization and infrastructure for MIB firmware uploading and the administration thereof. This infrastructure should assume that the deployed hardware revisions will continue to be heterogeneous.

The project would benefit from a clear statement as to the flow of system meta-data, that is what meta-data is retained and what needs to be updated across subscans and scans, specifically from the scripting language to the Executor and on to iDCAF.

The monitor database should track the mapping between modules and slots, e.g. using either a hardware serial number or, in its absence, the MIB MAC addresses.

We observed that the relationship with the user community appears solid. The project is encouraged to manage the expectations of the user community by frequent communications regarding the capabilities of the transition system.

The length of the command queue in the MIBs and CMIBs should be examined. The action to be taken when the queue overflows should be explicit.

The current correlator schedule decreases the amount of time to fully integrate the correlator before the end of the project. The arrival of the correlator is likely to begin a time of very intense activity that continues until the end of the project. This implies that an attempt should be made to complete other aspects of the project before the correlator arrives.

Observations

A significant amount of effort has already gone into the transition system and given the successful use of upgraded antennas in the array, the bottom line is that it works.

The committee appreciated the effort expended by the team on the high quality presentations and design documents.

The team is enthusiastic, committed, and capable.

The management team seems engaged and has a clear path forward.

The collaboration between the WIDAR team and the AOC team seemed very good.