Australia

Expert Panel on the Provision of Data Transport

Assessment of the Response to the RFI from Australia/New Zealand

November 4th, 2011.

Please note that the Expert Panel on the Provision of Data Transport, following the instructions of the SKA Project Development Office, did not assess the *alternative* proposal presented in the response to the RFI.

I. From Receptors to the Data Processor

I.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

For each channel proposed, the implementation as a 10Gbit/s service based on standard commercial components is certainly technically feasible, and is a service which would be found today in many network service providers' infrastructures. What is unprecedented is the scale and density of such services as the individual fibres and wavelength signals converge on the central data processor. Installation, commissioning and maintenance of such a fibre infrastructure goes many orders of magnitude beyond the operational experience of a typical National Research and Education Network (NREN) which bases its services on a footprint built from a single fibre pair.

Future expansion from systems based on 40 wavelengths per fibre to 80 waves/fibre is mentioned. This presumes that all relevant optics would be installed with this in mind (50GHz channel spacing rather than 100GHz), though this is not actually stated.

Remote sites:
10Gbit/s data services across a typical NREN infrastructure should be easily achievable, and in this case the service density for the 25 sites required should be possible within normal operations and planning for an NREN provided that the small number of additional fibres required are available on the relevant routes. Current eVLBI work has demonstrated that it is possible to correlate telescope signals on a global scale at data rates of 1-2Gbits/s, and the increase to 10Gbit/s should not present any problems from a data transport perspective.

II.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

Contributions to the proposal come from credible sources with existing experience of operating radio astronomy facilities and service provision over large-scale national optical networks (AARNet, the Australian NREN).

I.3. Cost Estimates

We are not aware of any direct cost comparison material in the context of normal telecommunications operation. As already mentioned, the SKA project is on a vast scale in comparison to NREN experience in this domain. However, the local costs achieved through the work on ASKAP should be a credible baseline for the larger project, and this proposal uses these.

The RFI (section 6.1.1) specifies that fibre termination require a compliant end-to-end measurement of attenuation and reflectance. The proposal does not make reference to this and it is not made clear whether this is included – it could potentially take longer than making a splice. Testing should be straightforward on the shorter lengths, but for the longer distances more specialised test equipment may be required to deal with spans up to 180km.

There is a **substantial risk** in the assumptions made in the cost estimates for connecting the remote sites. For the commercial carrier solution chosen, network access and wavelength use is assumed to be zero cost. Without further information on the form of the commercial commitments envisaged this is identified as a substantial risk, at least in the medium to long term. For the AARNet solution, this assumption may be more believable as it might be seen as appropriate use for an NREN network, but even in this case the longer term situation could be different considering the very long lifetime of the project during which the network funders/partners will have to commit to zero cost.

I.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The proposed implementation appears to offer the capability required (with some scope for future expansion if required as transmission technology evolves over the lifetime of the project)

I.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

The RFI (section 6.1.1) specifies that fibre termination require a compliant end-to-end measurement of attenuation and reflectance. For the longer distances more specialised test equipment may be required. The proposal does not make reference to this. See I.3. Costs.

I.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed due to lack of information.

II. From the Data Processor to the Super-Computer Centre

The response mentions that the Data processor will be located at the Murchison Radio-astronomy Observatory (MRO). The nearest major population city is Geraldton and that the Pawsey Supercomputer Centre will be constructed on CSIRO owned land in Bentley which is located about six kilometres from Perth CBD (Central Business District).

II.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

The response assumes a DWDM system which transmits 25 Tbps over each pair of fibre. To fulfill the requirement of 400 Tbps, 16 fibre pairs will be used.

There is already a fibre trunk between the data processor site and the super-computer site, which offers sufficient fibre capacity. (p.130)

II.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

The proposal is based on the knowledge gained from 20-years' experience in deploying and operating national and international data transport networks. They have accumulated experience in rolling-out of optical fibre networks to support the radio astronomy infrastructure, including ASKAP, at the MRO. They also provided the national connectivity to the long baseline array in eastern Australia and the facility in Warkworth town, New Zealand. (p.123)

II.3. Cost Estimates

The capital cost for the optical fibre from MRO to Geraldton (CSIRO) and from Geraldton to Perth (AARNet) is not included in the estimate. The reason is explained as follows in the response (p.133): *"Under existing CSIRO arrangements, fibre is already available for the passive fibre component of the system. There is no additional capital expenditure required."*

The operational cost for the fibre is not included. (p.137)

Since the cost estimates depend critically on the arrangement with AARNet and CSIRO, the SKA Data Transport Expert Panel has asked additional questions to clarify the relationship between the proposer and AARnet/CSIRO: *" Please detail the agreement you have with AARNet for the connections (correlator to computer centre and the long baselines from the remote stations to the core), in particular for the long-term management, maintenance/replacement costs of the managed bandwidth".*

The answer explains the current agreements and possible future plans which seem reasonable. The costs of active transmission equipment is excluded also in this case (connecting the Data Processor to the (remote) super-computer centre). Since this cost will be substantial there is a significant **risk** that the cost evaluation presented cannot be used directly in the comparative evaluation of the two responses.

II.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The RFI describes a model for SKA. The data link from the data processor to the super-computer building will consist of a fibre optic cable capable of carrying at least 400 Tb/s.

Considering that they have an 8 Tbps link at present, they will be able to satisfy the requirement

II.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

There are no significant gaps between the RFI and the assessed response.

II.6. Sequencing

Does the plan facilitate a smooth rollout?

Yes

They already have the basic data infrastructure. They have made an 8 Tbps link from the data center to the super-computer centre which is available for early deployment.

The existing data infrastructure is shown on p.127 in the response.

- From Murchison (MRO) to Geraldton, CSIRO has 400 km of 48 core fibre and associated huts and facilities.
- From Geraldton to Perth, Australian Government has 450 km of 36 core fibre and access to associated huts and facilities.
- From MRO to Pawsey Supercomputing Centre in Perth network, AARNet and CSIRO have 8 Tbps capable, 900 km DWDM transmission system. This link is already available for early deployment.

III. From the Super-Computer Centre to Data Centres in Other Parts of the World

III.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

Data Transport between the super-computer centre and other centres in the world will be carried over standard TCP/IP and 10G or higher standard Ethernet. Therefore interoperability problems are unlikely to occur and technical challenges can be restricted to capacity and resilience issues.

Both capacity and resilience are tackled by the tier approach for data management between remote sites. This model has been proved successfully e.g. by the current distribution model for the LHC experiments at CERN.

However, there is, today, a certain mismatch of international connectivity between Perth-Asia and Sydney-US. The link Perth-Asia is currently limited to 622 Mbps. In case of outage between Sydney-US or Perth-Sydney there is no sufficient international connectivity to deliver SKA data to other centres outside Australia.

Since link capacity to Asia is bound to improve substantially over the coming years, there are no major concerns about the feasibility of the proposed solution.

III.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

The response was compiled by AARNet which is the provider/operator of the national R&E network in Australia since many years.

The proposed solution is based on current knowledge and experience and includes no evidence of limited credibility.

III.3. Cost Estimates

Cost Estimates for Transatlantic Connectivity are at hand already and were extrapolated in the response from current price indications where necessary. In case of future lack sufficient capacity between Perth-Asia or Sydney-US costs may increase to a rather unpredictable amount. Additionally the costs presented in the response may rely on partial funding by the Australian government. To avoid excessive budgetary **risk** to the

SKA project this funding has to be guaranteed over the lifetime of SKA to ensure stable costs.

The **risk** to the cost indications presume no major changes with the current situation, this will have to be confirmed at some stage. The current cost estimates sound reasonable given no unforeseen shortage of capacity on the international links and no decrease of national funding in the future.

III.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

International connectivity between AARnet and other NRENs is in place since several years already. There are no indications that additional operational and management issues will arise with SKA.

At present, the international traffic between Perth and Asia is limited to 622 Mbps. In case of outages on the link between Sydney and the US sufficient capacity can not be provided. This **risk** should be addressed before full operation of the SKA is started.

The number of remote sites wanting data from the SKA instrument may vary significantly over the lifetime of SKA. An adequate tier infrastructure of remote data centres should be implemented from the very beginning of the project to avoid the necessity of continuous adaptation of international connectivity out of Australia.

III.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

No significant technical gaps have been identified in the response. The main difficulty lies in the long term forecast about the funding situation within AARnet and reliable cost indications for international connectivity.

On the budgetary side please see the remarks about insufficient capacity in case of outages on the Sydney-US side above and also the fact that AARNet as a national NREN is dependent on national funding and cooperation with national universities (which are shareholders of AARnet). With the important contribution of AARNet to the project this could present a long-term **risk** for the sustainability of management and operation.

III.6. Sequencing

Does the plan facilitate a smooth rollout?

Yes

Capacity assumptions for SKA phase 1 and phase 2 are in line with the RFI. During both phases no additional sequencing was proposed and does not seem necessary. In case of an urgent need for more capacity immediate upgrades seem feasible within a few weeks in both phases at least on the Sydney-US side.

Resilience is provided in phase 1 already and no further improvements during phase 2 are planned.

IV. Monitor and Control (M&C) Services

The response to the RFI did not include a description of the M&C system. This is not considered a gap since the RFI could be interpreted as not requesting it.

IV.1. Feasibility

Is the proposed solution a logically possible proposition?

Cannot be assessed in depth due to insufficient information.

The response did not include a description of the M&C system. The assessment is based on the elements provided at short notice in reply to questions from the panel.

While control and management for real-time and non real-time systems and large network and computing centres is a well-developed area, an analysis of the specific challenges of the SKA project would have been welcome. It is an area of **risk for the project** since it is not evident that the integration of the M&C of the various components of SKA is straightforward and does not raise specific issues.

The response details only the fibre topology for M&C. A pair of fibre is dedicated to each of the receivers, with at least 10 Gbps of bandwidth available for monitor and control and site data service redundancy. In addition the proposal states (pdf pag.33): " For monitor & control each remote station will have a dedicated 40 Gbps wavelength running 4 x 10Gbps services to each adjacent station in a 'daisy-chain'. The topology will form a cascaded ring with both ends connected to the central M&C facility to provide a redundancy capability."

IV.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

The response is based on information from entities with large expertise in managing large networks and participations in smaller, similar projects. It is expected that the respondents can provide a solution.

IV.3. Cost Estimates

The response excludes the cost of the active equipment and reports that (pdf p.115) Monitor & Control Equipment / fibre - inner circle Item are included.

In the M&C system a major expense will be the creation of the management and control software. This expense can be quite large, in money and time, and its maintenance and upgrade costs should be added to the operational expenses. Not including it represent s a **risk** of significantly underestimating the overall cost and time scale.

The cost estimates for the buildings are part of operation and also not easily accountable just to M&C.

IV.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

Australia's Information and Communications Technology Research Centre of Excellence NICTA (National ICT Australia Ltd) has been active since a long-time on M&C for computers and networks.

IV.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

Information on the control requested by RFI section "3.4 RFI mitigation, lightning protection and other control and safety systems" has not been found in the response to the RFI.

The logic and the software development to integrate all the components can be higher than expected. As an example the logic and data flow topologies for network and power system M&C should be added, specifying the possible point of contact and integration. If the systems are not integrated, the M&C will require separate operation, increasing human effort and making more complex the reaction to failures.
The amount of information produced by all the physical components of the infrastructure (including temperature, voltages, alarms) will be very large. A specific effort on algorithms and rules to analyse and present the information is a non-trivial effort and should be assessed. A tentative estimate for the service in production should be added.

The use of passive and active monitoring is not detailed. The use of on-demand active probes can be useful to debug e.g. degradation on fibres and misbehaviours in equipment and may require dedicated hardware and systems.

It is plausible that M&C system could benefit from T&S signals but this may imply specific topological choices.

IV.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed due to lack of information.

A commented proposal should be added on the topic. It is assumed by the panel that M&C is partially implemented before the start of the deployment of the equipment and in parallel to it, as a fundamental tool to debug the installation.

V. Timing/Synchronization Services

The response to the RFI did not include a description of the T&S system. This is not considered a gap since the RFI could be interpreted as not requesting it. The assessment is based on the elements provided at short notice in reply to questions from the panel.

V.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

The solution proposed for the central site of the SKA is based on round trip phase correction systems. Based on cited experience it is considered a feasible solution, which can be developed for the distribution of timing and synchronization to receptors within 180 km of the core. The estimated cost is 700.000,00 Euro for the active equipment and about 200.000,00 Euros for two H-masers.

Two potential methods are listed for providing timing and synchronization to the remote SKA array-stations outside 180km radius: distribution using an optical fibre, the use of distributed H-Maser clocks. The first method is still in the early experimentation phase, while the second is considered more mature, albeit more expensive probably. The feasibility is accepted also for remote stations and a decision between the alternatives can be taken only as a result of detailed cost benefit analysis on the capital and ongoing cost of operating multiple clocks vs. a distributed timing system or a hybrid model.

V.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

The proponents have a known track of expertise and can leverage the research in ongoing projects in T&S to engineer or implement a solution for the SKA timing and synchronization system.

V.3. Cost Estimates

Given the short time to produce an estimate in response to a question from the panel, the quoted costs should be taken as an indication subject to change.

For the core the estimate is €700.000 for the active equipment and about 200.000,00 Euros for two H-masers. It is assumed that cost of the two fibres for T&S is accounted for in the general proposal.

There is no indication for the cost of serving the remote components. This is not considered a gap in the response, as it was not requested and further analysis is required.

V.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The proponents are considered capable of implementing the core solution.

V.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

The panel expresses its appreciation of the fact that the AUS/NZ team provided a memorandum at short notice elaborating on Timing and Synchronization.

For the core design there are no evident gaps in the brief description proposed. The experience detailed in operating large-scale optical topologies is considered an important reference.

It seems that the T&S fibres will converge physically in the same aggregation point of the data and M&C, this may increase the **risk** of losing all channels at the same time, reducing the amount of information and action capability from the control room.

In the longer run the inter-operation of the T&S system with the M&C system is recommended and the preferred topology should be made explicit (e.g. single or multiple point of contact). T&S requirements for the M&C system should be elaborated.

V.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed due to lack of information.

The information is not present or it is not complete, only when the final design has been made by SPDO and technical guidance can be given can this be assessed in depth.