

SYSTEM REQUIREMENTS SPECIFICATION

EVLA Correlator Backend

Project Document: A25251N0000

Revision 2.1

Tom Morgan, Dec 20, 2002

National Radio Astronomy Observatory

Array Operations Center

P.O. Box O

Socorro, NM 87801-0387

Revision History

Date	Version	Description	Author
26-Feb-2002	1.0	Initial draft	Tom Morgan
05-Mar-2002	1.1	Revised section 3.2.2	Tom Morgan
10-May-2002	2.0	Clean-up of section 1, revisions to section 2 and 3	Tom Morgan
20-Dec-2002	2.1	General clean-ups and minor revisions to all sections. Added and quantified requirements imposed on external systems.	Tom Morgan

Table of Contents

1	<i>Introduction</i>	1
1.1	Purpose	1
1.2	Scope	1
1.3	Definitions, Acronyms, and Abbreviations	2
1.3.1	Definitions	2
1.3.2	Acronyms	2
1.4	References	2
1.5	Overview	3
2	<i>Overall description</i>	4
2.1	Product perspective	4
2.2	Product functionality	6
2.2.1	Data Input	6
2.2.2	Data Processing	6
2.2.3	Data Output	6
2.2.4	Monitoring	6
2.2.5	Recovery	6
2.2.6	Control	6
2.3	User characteristics	7
2.3.1	Array Operator.....	7
2.3.2	Engineers and Technicians	7
2.3.3	Astronomer/Scientist	7
2.3.4	Software Developer	7
2.3.5	Web User	7
2.4	Constraints	7
2.4.1	Criticality of the Application	7
2.4.2	Computer Hardware Limitations	7
2.4.3	Communications Limitations.....	8
2.4.4	Reliability	8
2.5	Assumptions	8
2.5.1	Incoming Data Stream	8
2.5.2	Auxiliary Data	9
2.5.3	Outgoing Data Stream	9
3	<i>Specific requirements</i>	10
3.1	External Interface Requirements	10
3.1.1	Correlator to Backend Interface.....	10
3.1.2	Backend to End-to-End Interface	11
3.1.3	Backend to/from Monitor and Control Interface	10
3.2	Functional Requirements	12
3.2.1	Information and data flows	12
3.2.2	Process Descriptions.....	12
3.2.3	Data Construct Specifications.....	14
3.3	Performance Requirements	15
3.3.1	General	15
3.3.2	Hardware	15
3.3.3	Software.....	16

3.4	Reliability/Availability	16
3.5	Serviceability	17
3.6	Maintainability	17
3.7	Scalability	17
3.8	Security	18
3.9	Installation and Upgrades	19
3.10	Documentation	19

1 Introduction

1.1 Purpose

The primary goal of this document is to provide a complete and accurate list of requirements for the EVLA Correlator Backend System. Upon completion, the document will act as a binding contract between developers and users and will provide a common point of reference for system expectations.

The primary audience of this document includes, but is not limited to, project leaders, the designers and developers of the system and the end user. The document may also be of interest to EVLA project scientists and engineers or as a reference for individuals involved in similar projects with similar requirements.

The requirements contained in this document are numbered based on the section/subsection in which they appear.

Note: Text found between “<” and “>” indicates questions or comments to myself and/or readers. And In most cases, the phrase “The user” can be replaced with “An authorized user”.

1.2 Scope

The Correlator Backend System lies between the Correlator and the End-to-End (e2e) System. It is the primary component of the real-time astronomical data processing capability (the processing pipeline) of the EVLA. Its primary responsibility is to perform basic data assembly, formatting and processing services and to support the desire for real-time inspection of the astronomical data stream.

The major functions the Correlator Backend System must perform are as follows:

- Receive data from the Correlator in real-time.
- Assemble time-series from the Correlator lag output.
- Perform Fourier Transforms of the assembled time series.
- Perform a limited number of additional processes upon user request.
- Deliver suitably formatted results to the End-to-End System.

This document will define those requirements that must be fulfilled by the Correlator Backend System, plus requirements imposed on other systems by the Backend System.

1.3 Definitions, Acronyms, and Abbreviations

1.3.1 Definitions

Administrator – An individual with unrestricted access to all aspects of the system.

Auxiliary Data – All other (non-astronomical) data.

Data – Astronomical observational data.

Lag Frame – The basic unit of data output from the WIDAR Correlator. For a complete definition of its contents at the bit level see NRC-EVLA Memo #014, Refined EVLA WIDAR Correlator Architecture page 70.

Lag Set – A complete, properly ordered series of lag values that can be submitted to the Fourier Transform function. The lag frames received from the Correlator will contain up to 128 lag values, so lag sets longer than 128 values will span multiple lag frames and require proper ordering and assembly into complete lag sets.

Metadata – All data about the astronomical data.

NaN – Literally, “Not a Number”. For floating point data types, a bit string that does not translate into a valid floating point number.

Non-real-time – Offline operations with data input from some external storage device or generated internal (e.g., for testing).

Processing Pipeline – The series of BE functions performed on the astronomical data, i.e., that set of functions that the data passes directly through.

Processor – A physical computation device (hardware).

Process – A data processing procedure (software).

Real-time – Online operations with active astronomical data streaming from the Correlator.

1.3.2 Acronyms

AOC – Array Operations Center

CMIB – Correlator Monitor Interface Board

CMCS – Correlator Monitor and Control System

COTS – Commodity-off-the-shelf (i.e., generic products)

e2e – End-to-End System (archive)

M&C – Monitor and Control System

EVLA – The VLA Expansion Project

RFI – Radio Frequency Interference

SyRS – Refers to the *System Requirements* document.

SRS – Refers to the *Software Requirements Specification* document.

1.4 References

- 1) ANSI/IEEE Std 830-1998, IEEE Recommended Practice for Software Requirements Specifications
- 2) ANSI/IEEE Std 1233-1996, IEEE Guide for Developing System Requirements Specifications
- 3) EVLA Memo No. 15, Scientific Requirements for the EVLA Real-Time System
- 4) EVLA Project Book
- 5) EVLA System Requirements (SyRS)
- 6) EVLA Architecture and Design
- 7) The Very Large Array Observing Log (J. Nieri, February 1994)
- 8) Refined EVLA WIDAR Correlator Architecture, NRC-EVLA Memo# 014, Brent Carlson, Oct. 2, 2001.
- 9) EVLA Correlator Monitor and Control System, Test Software, and Backend Software Requirements and Design Concepts, NRC-EVLA Memo # 015, Brent Carlson, Jan. 23, 2002.

1.5 Overview

The remainder of this document contains a more detailed description of the Correlator Backend System as well as the requirements necessary to design and build the system. Section 2 provides a general description of the Correlator Backend System. Section 3 details the requirements of the product and the requirements impose on external systems. Section 3 is the core of this document.

The format of the document follows that outlined in the IEEE STD 830 document, IEEE Recommended Practice for Software Requirements Specifications.

2 Overall Description

2.1 Product Perspective

The EVLA Correlator Backend System will be designed and implemented as a real-time data processing system. The system is expected to be implemented on a distributed memory cluster of network connected processors. Computers in the system will be commodity-off-the-shelf (COTS) products running the same operating system and message passing middleware. They will communicate with one another and the Monitor and Control System over a network. Data input to the system from the Correlator and output from it to the end-to-end System will be over high-speed networks. The networks connecting the internal processors, the Correlator and the e2e are part of the BE System. A high level conceptual diagram of the system and its interfaces to external systems is shown in Figure 1. BE Management functions will run on one of the cluster processors with one or more shadow processors standing by in case of failure. The remaining processors will be running data processing functions, or be standing by to pickup data processing in the event of a failed processor.

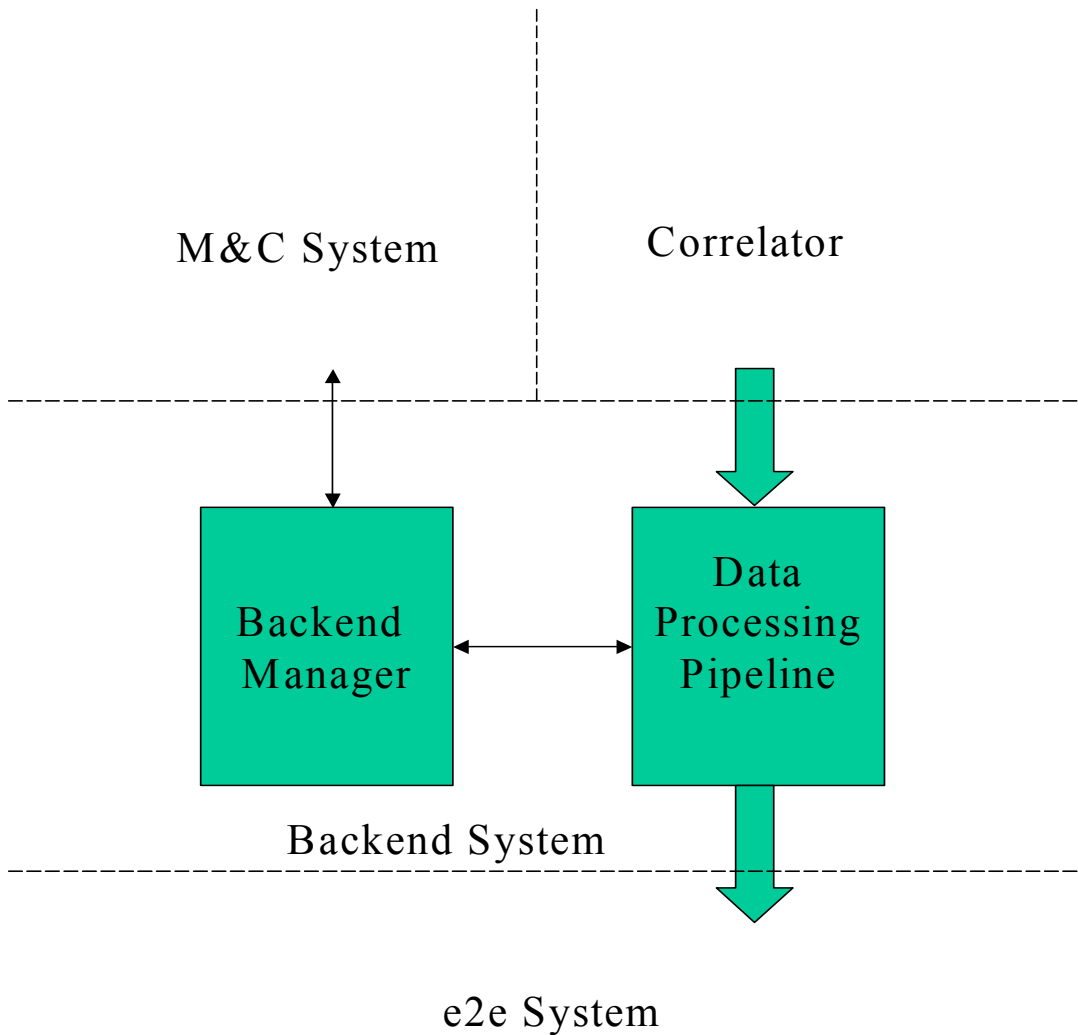


Figure 1. High level block diagram of the two major components of the Correlator Backend System with interfaces to external systems. Thin arrows represent monitor, control and non-correlator lag frame data flows. The broad arrows represent Correlator lag products flowing into the Backend and processed spectra flowing into e2e.

2.2 Product Functionality

2.2.1 Data Input

Correlator lag data will be received directly from the Correlator Baseline Boards in the form of Lag Frames. The lag frames contain correlation lag values and all auxiliary parameters needed to assemble the lags into complete lag sets (properly ordered time series). It is currently assumed that all observational modes yielding correlator results that are transmitted to the Backend will be in the form of lag frames.

Additional auxiliary data and meta-data needed for processing prior to output to the e2e System will arrive via the Monitor and Control System, whether produced by the Correlator or some other part of the EVLA System.

The BE will receive and act upon status requests and control commands originating in or via the M&C System.

2.2.2 Data Processing

The Correlator lag frames will be assembled into time ordered series, normalized, and when necessary time stamp adjusted. The time series will be Fourier Transformed and integrated over a user determined period of time. User selectable time and/or frequency domain processes may be applied to the data before and/or after Fourier Transform. Prior to output, the end results will be formatted to meet the internal needs of the e2e.

2.2.3 Data Output

Formatted spectra will be transferred to the end-to-end System. Pertinent meta-data available to the BE System will be associated with the output. The fundamental unit of output is a sub-band cross-power spectrum produced by the Correlator. No “stitching” operations that combine spectra from different sub-bands will be performed.

The BE will produce a variety of error, warning, status and other reports and messages that will be transferred to M&C for final disposition.

2.2.4 Monitoring

The Correlator Backend System will conduct a number of self-monitoring activities on application and system software as well as hardware systems to detect system failure and out of spec conditions.

2.2.5 Recovery

The ability to attempt recovery from failure and out of spec performance conditions will be built into the system.

2.2.6 Control

The system will provide control and auxiliary parameters to internal input, output, processing, monitor, recovery, and other functions and receive status and performance data from them. It will also communicate with the external Monitor and Control System.

2.3 User characteristics

All use of the Correlator Backend System will be indirect via the Monitor and Control System. The BE system will not directly produce user interface screens.

2.3.1 Array Operator

The primary contact with array operations will be via status and error messages channeled through the Monitor and Control System.

2.3.2 Engineers and Technicians

The ability of the Backend System to achieve and maintain real-time processing will be vitally dependent upon reliable operation and rapid diagnosis and repair of faults in the hardware and software systems. These individuals will be responsible for performing corrective and preventive maintenance along with periodic performance tests and upgrades. Engineers and technicians will need tools to inspect individual devices from remote locations.

2.3.3 Astronomer/Scientist

These individuals are primarily interested in the science that is obtained from the instrument. Their main interaction will be to provide the basic parameters of the observation and to select and provide parameters for any additional data processing beyond the Fourier transforms.

2.3.4 Software Developer

These individuals are responsible for developing the software and will interact with the system to ensure that it is functioning properly. They will need access to error and status data on an interactive basis and need to conduct debugging as well as recompilation and linking. The software developer requires remote access to the system so that troubleshooting can be accomplished away from the EVLA and during non-working hours.

2.3.5 Web User

A few authorized individuals may be allowed access to parts of the system that are usually considered restricted.

2.4 Constraints

2.4.1 Criticality of the Application

The Correlator Backend System is a critical component in the Astronomical data path. If it is unavailable, incoming astronomical data will be lost.

2.4.2 Computer Hardware Limitations

The ultimate throughput capability of the real-time data processing pipeline of the Backend System will be constrained by the computational performance limits of available computer hardware and the practical ability to configure and maintain large numbers of processors.

2.4.3 Computer Software Limitations

The ultimate throughput capability of the real-time data processing pipeline of the Backend System will be constrained by the efficiency of supporting software systems, data processing code and our ability to configure and tune them for maximum performance.

2.4.4 Communications Limitations

The ability to realize and maintain real-time operations is critically dependent upon the performance levels of available network systems.

2.4.5 Processing Limitations

Data processing operations performed shall be reversible. That is, uncorrected input data to a processing operation must be recoverable from the processing output. This will hold only up to the point another data processing operation is performed. Data integration is not considered to be a data processing operation. It is fundamentally irreversible once done, and no intermediate results will be retained to even partially backout an integration.

2.4.6 Reliability

The ability to maintain real-time operations over realistic extended periods of time is dependent on the mean time to failure of the hardware and software components of the computing and communications systems.

2.5 Assumptions

2.5.1 Incoming Data Stream

It is assumed that the Correlator will deliver suitably formatted network data packets to the input network of the Backend System. Lag frames will not necessarily arrive in Lag Set order. All lag frames for the same baseline will be directed to the same Backend processor. It is further assumed that the number of lags per Lag Set will always be a power of two.

The initial (for earliest science data in 2007) aggregate data volume delivered from the Correlator to the Correlator Backend shall not exceed 1.6 Gbytes/sec. This represents a 0.1 sec dump rate of 160 baseline boards of 64 correlator chips each. Each chip will produce 1024 complex lags or 2048 lag values of 8 bytes each. (160 boards) (64 chips/board) (2048 lags/chip) (8 bytes/lag) (10 /sec) = 1.6 Gbytes/sec.

An ultimate upper bound on Correlator output rate and aggregate volume is set by the fastest Correlator dump rate. This rate is currently anticipated to be 320 microseconds or one third of a millisecond. This is 300 times faster than the initial rate and will produce 300 times the aggregate data volume or about 500 Gbytes/sec.

2.5.2 Auxiliary Data

It is assumed that all auxiliary data needed for processing and formatting operations will be provided directly by the correlator or indirectly by the Monitor and Control System in a timely manner. Much of this data will originate from the Station Board CMIBS.

2.5.3 Outgoing Data Stream

The initial (for earliest science data in 2007) aggregate data volume delivered from the Correlator Backend to the e2e archive shall not exceed 25 Mbytes/sec. The present VLA system produces about 0.25 Tbytes of archived data per year, which is roughly 10% of its total capability if it were run constantly at the highest rate. Applying the same 10% factor to the EVLA will result in an archive filling rate of 75 Tbytes per year or 300 times the current VLA rate. It is assumed that the e2e System will be capable of accepting these output data rates and volumes.

Visibility data from different baselines could be processed by different Backend processors. Final assembly of all visibility data is expected to be performed by the e2e system.

3 Specific Requirements

3.1 External Interface Requirements (Including Requirements imposed on External Systems)

3.1.1 Correlator to Backend Interface

Req. ID	Description
3.1.1.1	Lag Frames – <i>The BE shall receive LTA or Speed Dump Lag Frames from the Correlator.</i> For a detailed description of the two dump formats see Reference 8, pages 69 to 71. This will most likely be in the form of one or more frames per UDP/IP packet
3.1.1.2	Transfers – <i>The transfer shall take place in such a manner that at a minimum all data needed to perform any Fourier Transform shows-up on a single processor.</i> It is preferred that all data for the same baseline be directed to a single BE IP address.
3.1.1.3	Interface Medium – <i>The interface network shall be capable of a sustained aggregate transfer rate of 1.6 Gbytes/sec.</i> It is expected that at minimum a 1 Gbit/sec switched fiber optic network will be required.
3.1.1.4	Lag Frame Delivery – <i>A lag frame shall be considered to be lost if in is not received within two Correlator dump periods of its expected arrival time.</i> Note: It is currently expected that each UDP/IP packet will deliver one lag frame.

3.1.2 Backend to/from Monitor and Control Interface

Req. ID	Description
3.1.2.1	External Auxiliary Data (e.g., Quantizer Power Measurement Data - State Counts) - <i>The BE shall receive, via M&C, auxiliary data needed for real time processing such as quantizer power levels from the digitizers at the antennas and the requantizers inside the Correlator. This data shall be delivered no later than 5.0 seconds and preferably within 200 ms after the arrival on the BE of the lag data to which it applies. (see 3.3.2.4 Memory)</i>
3.1.2.2	Observational Mode – <i>The BE shall receive, via M&C, data and parameters specific to the current EVLA Observational Mode needed for processing the Correlator Lag values. New or changes to existing Observational Mode data shall be provided to the BE at least 60 seconds prior to the time it is to take effect.</i> This time will be needed to assure that all necessary parameters have been propagated to the DP processes and all data structure changes have been made.

3.1.2.3	Observation Parameters – <i>The observational Mode parameters shall include at a minimum: the number of subbands, polarization products, basebands, subarrays, lags per lag set, and antennae per subarray; the start and end time of the observation, the Correlator dump rate, the lag sample interval, the BE integration interval, the frequency ranges of the subbands, and the frequency ranges of the basebands.</i>
3.1.2.4	Optional Processes Parameters - <i>The BE shall receive, via M&C, parameters specific to any optional processes that will be applied to the lag sets before and/or after FFT. It is expected that this information will accompany the Observation Mode information.</i>
3.1.2.5	Meta-data – <i>The BE shall receive, via M&C, all meta-data necessary to format BE results for delivery to the e2e.</i>
3.1.2.6	Operational Status and Control – <i>The BE shall provide operational status data to and receive control data from the M&C System. This includes Lag Frame destination addresses and address changes.</i>
3.1.2.7	Error and Warning – <i>The BE shall provide error and warning reports to M&C as operating conditions warrant.</i>
3.1.2.8	Debug/Test Messages – <i>The BE shall provide several optionally selectable levels of printed messages detailing operational parameters at critical locations in the system.</i>
3.1.2.9	Interface Medium – <i>The interface network shall be Ethernet (IEEE 802.3 compliant) of 100 Mbits/sec or better data rate.</i>
3.1.2.10	Availability – <i>All required M&C services shall be available no later than one year prior to first science. Earliest first science is currently anticipated in Q2, 2007.</i>

3.1.3 Backend to e2e Interface

Req. ID	Description
3.1.3.1	Formatted Output – <i>The BE shall deliver formatted final results to the e2e System. The BE shall produce all data needed by the e2e System for archiving and further processing. The output is currently expected to be in a form compatible with AIPS++ Measurement Sets.</i>
3.1.3.2	Interface Medium – <i>Interface communication media shall be capable of sustaining a minimum aggregate transfer rate of 25 Mbytes /sec.</i>
3.1.3.3	Reliability – <i>The e2e Archive shall not be offline for more than 12 continuous observing hours or a total of 18 observing hours during any 48 hour period. These times are based upon planned local BE node storage capacity. (see 3.3.2.5 Storage)</i>
3.1.3.4	Availability - <i>All required e2e services shall be available no later than one year prior to first science. Earliest first science is currently anticipated in Q2, 2007.</i>

3.2 Functional Requirements

3.2.1 Information and data flows

Req. ID	Description
3.2.1.1	Monitor and Control System – <i>The BE shall acknowledge receipt of all data received from M&C.</i>
3.2.1.2	Correlator System – <i>The BE shall notify M&C of any detected interruptions of data delivery from the Correlator.</i>
3.2.1.3	e2e – <i>The BE shall verify successful delivery of output to the e2e.</i>
3.2.1.4	Internal Data – <i>The BE shall guarantee safe delivery of all internal messages.</i>
3.2.1.5	Lag Frames – <i>The BE shall be able to handle lag frames of less than 128 values.</i>
3.2.1.6	Lag Sets - <i>The BE shall be able to handle lag sets up to a maximum size of 262,144 values (256K lags).</i>
3.2.1.7	Correlator dump rate – <i>The BE shall be able to handle different dump rates for different basebands. It is understood that the dump rates will always vary harmonically.</i>

3.2.2 Process Descriptions

Req. ID	Description
3.2.2.1	Data Receive – <i>The BE shall receive incoming data packets from the Correlator to Backend network interface. This network is a part of the BE System.</i>
3.2.2.2	Verify Receive – <i>The BE shall verify the successful receipt of incoming data from the Correlator. This includes checking for receive errors and determining that all expected data was received, accumulation of error statistics and comparison against tolerances, and reporting of all out of tolerance conditions.</i>
3.2.2.3	Input Data Management - <i>The BE shall store input data records in a memory buffer and track buffer locations of all input data until data processing is complete. Report any buffer overflow conditions.</i>
3.2.2.4	Processing Management – <i>The BE shall respond to incoming correlator mode changes, user optional processing sequence and/or parameter changes, and other external inputs that affect the data processing pipeline. Update internal parameter tables and synchronize data processing pipeline with new operational conditions.</i>
3.2.2.5	Time Series Assembly – <i>The BE shall assemble the received input data into continuous time series (lag sets).</i>
3.2.2.6	Data Integrity Verification – <i>The BE shall ensure that time series data is correctly ordered and contains valid data values along its entire extent. Compare against tolerances and report all out of tolerance conditions.</i>
3.2.2.7	Data Invalid – <i>The BE shall replace all invalid data with zero values.</i>
3.2.2.8	Data Invalid Count – <i>The BE shall keep track of data invalids.</i>
3.2.2.9	Normalization – <i>The BE shall be able to apply normalizations based on reported (from the Correlator via the lag frames) data invalid counts.</i>

3.2.2.10	Coarse Quantization Correction - <i>The BE shall be able to apply corrections based on state count and/or quantizer power measurement data. This is the VanVleck correction</i>
3.2.2.11	Time Stamp Adjustment – <i>The BE shall be able to make time stamp adjustments as required by the observational mode and correlator output parameters. This may arise when recirculation is used.</i>
3.2.2.12	Windowing – <i>The BE shall be able to perform windowing operations prior and subsequent to Fourier Transform. This will be needed for narrow band RFI mitigation. Post Fourier Transform windowing will be applied as a convolution.</i>
3.2.2.13	Time Domain Processing – <i>The BE shall be able to apply user selected time domain processes. These processes should be constructed to be chainable (output of any time domain process can be piped to input of any other, including replica of self and Fourier Transform) and repeatable in the chain. No Optional time domain processes have as yet been proposed.</i>
3.2.2.14	Fourier Transform Processing – <i>The BE shall be able to Fourier Transform the lag set time series. A power-of-two complex-to-complex Fast Fourier Transform with retention of all output positive and negative frequencies will be used. This process must be able to accept as input the output of any of the time domain processes.</i>
3.2.2.15	Frequency Domain Processing – <i>The BE shall be able to apply user selected frequency domain processes. These processes should be constructed to be chainable (output of Fourier Transform and any frequency domain process can be piped to input of any frequency domain process including replica of self) and repeatable in the chain. No frequency domain processes have as yet been proposed.</i>
3.2.2.16	Integration – <i>The BE shall be able to sum the frequency domain, spectral results. The amount (time duration) of summation will be controlled by an observational mode parameter obtained via M&C. The BE shall keep track of the number of samples/dumps integrated in each spectral channel. The summation will occur after all optional frequency domain processing, or if none, after the Fourier Transform. Integration for long periods of time is what will throttle the output of the Correlator to a rate manageable by the e2e.</i>
3.2.2.17	Output Formatting – <i>The BE shall combine the finished spectra with meta- and auxiliary data to form suitably formatted output data sets. AIPS++ Measurement Sets or compatible subsets and/or fragments thereof are the expected entities.</i>
3.2.2.18	Output Data Management – <i>The BE shall store formatted output data records in a memory buffer with backup disk buffering. Store data ready for transmission to the e2e System until successful transfer has occurred. Report any errors and buffer overflow conditions that occur.</i>
3.2.2.19	Data Send – <i>The BE shall send output data to the e2e System.</i>
3.2.2.20	Send Verify – <i>The BE shall verify that all sent data was successfully received. Report all errors.</i>
3.2.2.21	Monitor I/O Performance – <i>The BE shall monitor data transfer rates from the Correlator and to the e2e. Accumulate data transfer statistics and compare against tolerances. Report all out of tolerance conditions.</i>
3.2.2.22	Monitor Compute Performance – <i>The BE shall monitor the overall data processing rate. Compare against tolerances and report all out of tolerance conditions.</i>
3.2.2.23	Monitor Compute Errors – <i>The BE shall trap, flag and repair inf's, NaN's, underflows, overflows and other computation errors. Accumulate computation error statistics and compare against tolerances. Report all out of tolerance conditions.</i>

3.2.2.24	Monitor Processes – <i>The BE shall periodically or upon request check PID's and assure that all started tasks are alive and running. Report missing, stopped, defunct and other damaged processes.</i>
3.2.2.25	Monitor Processors – <i>The BE shall periodically or upon request check Backend physical processors and assure that all needed processors are alive and responding. Report all crashed, stopped, or unresponsive processors.</i>
3.2.2.26	Monitor Networks – <i>The BE shall periodically or upon request check all Backend internal networks and assure that all communication connections are intact and functioning. Report all non-functioning components.</i>
3.2.2.27	Start Process – <i>The BE shall be able to initiate a processing task on any Backend processor.</i>
3.2.2.28	Stop Process – <i>The BE shall be able to signal a kill for any Backend process.</i>
3.2.2.29	Alter Priority – <i>The BE shall be able to alter the priority of any of the BE tasks.</i>
3.2.2.30	Reboot Processor – <i>The BE shall be able to initiate a reboot of any Backend physical processor.</i>
3.2.2.31	Reboot network – <i>The BE shall be able to initiate a reboot of any internal network.</i>
3.2.2.32	Offload – <i>The BE shall be able to redistribute internal workload among its processors. This may involve change of destination IP address(es) for the Correlator network.</i>
3.2.2.33	General – <i>BE processes shall not violate archive data requirements. All processes shall be reversible; the raw unconverted input always being recoverable from the output.</i>

3.2.3 Data Construct Specifications

Req. ID	Description
3.2.3.1	Input Data Queue – a memory buffer of lag frames. Data entry status queue to track each record in the buffer. The lag frames will contain all information necessary to properly assemble complete lag sets.
3.2.3.2	Output Data Queue – a memory buffer plus backup disk storage of all processed spectra. These will be suitably formatted prior to transfer to the e2e. Data entry status queue to track each record in the buffer.
3.2.3.3	Processing Parameters – names, position(s) in sequence, and adjustable parameters for all fixed and user selectable processing pipeline applications.
3.2.3.4	Processing flags – a table of flags needed to identify various internal conditions relating to error response and processing state.
3.2.3.5	Metadata – All internally and externally generated data about the processed time series and spectra including invalid data flags, processes applied, coordinates, etc.
3.2.3.6	Error Report – error number (translatable into text error message), error source, error rates (as applicable), and time stamp.
3.2.3.7	Warning Report – warning number (translatable into text warning message), warning source, warning rates (as applicable), and time stamp.
3.2.3.8	Failure Report – internal system component (e.g., disk drive, processor, process, and network) failure number (translatable into text error message) and time stamp.

3.2.3.9	Recovery Report – internal system component (process, processor, network) recovery action result.
3.2.3.10	Status Report – internal system component (process, processor, network) functional state.

3.3 Performance Requirements

3.3.1 General

Req. ID	Description
3.3.1.1	Data Integrity – <i>the Backend System shall maintain input data fidelity and dynamic range across all processing, manipulation and I/O functions.</i>
3.3.1.2	Error Handling – <i>the system shall be capable of flagging and marking corrupted data segments and proceeding without interruption or effect on other data.</i> This includes, but is not limited to, partial data, zero data, underflows, overflows, infinities, and NaN's whether obtained on input or arising during processing.

3.3.2 Hardware

Req. ID	Description
3.3.2.1	Input – <i>The BE System shall be capable of accepting an aggregate data input stream from the Correlator of a minimum of 1.6 Gbytes/sec.</i> This must be done simultaneously with the output stream, but not necessarily over the same interconnects. This is an initial deployment specification and will be increased over time.
3.3.2.2	Output – <i>The BE System shall be capable of delivering an aggregate output data stream to the e2e System of a minimum of 25 MBytes/sec.</i> This includes resends and simultaneous transfer of data stored due to a previous e2e connection outage. This must be done simultaneously with the output stream, but not necessarily over the same interconnects. This is an initial deployment specification and will be increased over time.
3.3.2.3	CPU – <i>The total processor capability of the BE System shall be (combination of numbers of processors and individual processor speed) sufficient to accomplish all processing tasks while avoiding loss or delay on the input and output data streams.</i>
3.3.2.4	Memory – <i>The BE System shall have sufficient memory with sufficient access speed to buffer 60 seconds of input data and accomplish all processing tasks while avoiding loss or delay on the input and output data streams.</i> At 26.5 Mbytes/sec of input per BE node and a total of 62 data processing nodes (for an aggregate of 1.6 Gbytes/sec), this implies a minimum of 1.5 Gbytes of memory per BE processing node.
3.3.2.5	Storage – <i>The BE System shall have sufficient disk storage with sufficient access speed to meet short duration Correlator bursting demands plus a standby reserve to hold at least 12 hours of output data.</i> At 25 Mbytes/sec, this implies a minimum of 40 Gbytes of disk storage per BE processing node.

3.3.3 Software

Req. ID	Description
3.3.3.1	Applications – <i>all math/science application software shall take optimal advantage of all language, compiler, and system computational features and resources to reduce run times to the minimum practical level. They shall be coded in such a manner as to minimize the possibility of floating point exceptions during processing.</i>
3.3.3.2	Management – <i>all management software functions shall take optimal advantage of all language, compiler and system features and resources to reduce overheads to the minimum practical level.</i>
3.3.3.2	I/O – <i>all input and output, and storage and retrieval operations shall take optimal advantage of all system resources to reduce overhead and latency to the minimal practical level.</i>
3.3.3.4	Processing – <i>all data processing functions shall be chainable (outputs pipeable to inputs) and repeatable in the processing pipeline in cases where this makes computational sense.</i>
3.3.3.5	General - <i>Operating system, message passing and other middle-ware, and programming language(s) used shall follow industry standards and be commonly available and widely used. Availability of source code for the OS will be very important.</i>

3.4 Reliability/Availability

Req. ID	Description
3.4.1	Auto-correction – <i>the Backend System shall be self-monitoring. It will be capable of detecting, reporting on and automatically taking action to remedy or lessen the impact of, at a minimum, the following types of abnormal conditions: processor hardware failure, operating system hangs or crashes, computational performance below minimum specifications, computational error rates above maximum specification, internal communications failures, and external (with the Correlator and e2e) communications disruptions.</i>
3.4.2	Software – <i>the software part of the system shall be able to perform without total system restart due to internal failure between system maintenance windows.</i>
3.4.3	Hardware – <i>the hardware part of the system shall be able to perform indefinitely without complete loss of service, except in the event of total failure of primary and backup power.</i>
3.4.4	Network – <i>the internal BE network(s) and the networks connecting the BE to the Correlator, M&C, and e2e Systems shall be considered to be out of spec if packet loss exceeds 0.1%.</i>
3.4.4	Correlator mode changes – <i>the system shall be capable of responding in a loss-less manner to I/O and processing changes arising from Correlator mode changes.</i>
3.4.5	Loss of e2e – <i>the system shall continue to operate in a loss-less manner in the event of a temporary loss of availability of the e2e System.</i>
3.4.6	Loss of Correlator – <i>the system shall be able to complete processing of all onboard data, deliver the results to the e2e and maintain availability for immediate resumption of operations once Correlator access is restored.</i>

3.4.7	Loss of M&C – <i>the system shall continue to operate during the absence of the M&C System until the first encounter of unavailable critical auxiliary data. The system will cache a predetermined amount of correlator data after the first encounter of unavailable critical data and complete all requested operations on cached data if the unavailable critical data is ultimately obtained.</i>
3.4.8	Standby Mode – <i>the system shall be able to sit at idle and resume operations with minimal delay.</i>

3.5 Serviceability

Req. ID	Description
3.5.1	Hardware Accessibility – <i>all system processing and interconnect hardware shall be readily accessible for maintenance, repair, replacement and/or reconfiguration.</i>
3.5.2	Software Accessibility – <i>all systems and application source code shall be available to or on the systems that execute it.</i>
3.5.3	Debugging – <i>all software application modules shall be debuggable.</i>
3.5.4	Processes – <i>all software processes shall be killable, restartable, debuggable and testable without affecting normal operations.</i>

3.6 Maintainability

Req. ID	Description
3.6.1	Software tools – <i>software tools and pre-built applications that do not have source code available shall come with a complete diagnostic package and customer support.</i>
3.6.2	Operating Systems – <i>operating system software shall either have source code available or come with sufficient diagnostics and customer support.</i>

3.7 Scalability

Req. ID	Description
3.7.1	Hardware – <i>I/O, communications, and processing hardware shall be easily expandable, reconfigurable, augmentable and replaceable to meet increasing data transfer and processing demands imposed by EVLA science, Correlator changes, and availability of new technology.</i>
3.7.2	Transparency – <i>3.7.1, above, shall be accomplished in manner that is transparent to processing, communications and I/O software functions with the possible exception of recompilation of executables.</i>
3.7.3	Seamlessness – <i>3.7.1, above, shall be accomplished in a manner that is seamless, in that it does not affect hardware modules or software functionality that it meets at interfaces.</i>
3.7.4	Performance – <i>the Backend system shall be scaleable to an extent limited only by hardware technology and budget constraints. An ultimately upper bound is set by the Correlator upper limit of three Gbytes per second per Correlator output channel (baseline board) in real-time.</i>

3.8 Security

The Backend System needs a robust security mechanism in place so that unauthorized users are not allowed access. Authorized users are expected to be restricted to software and hardware development, testing, maintenance and operations personnel.

All users of the Backend System must be uniquely identified. This could be done via a username and associated password scheme that would authenticate and authorize the user access to the system and, if applicable, grant the user access to restricted or controlled parts of the system. If a user cannot be identified, they will not be given access. In order to monitor all past access to the system, all attempts to access the system should be logged.

Users' needs and expectations from the system will be different. Systems operations should be given unrestricted access to all aspects of the system and should have the authority to grant and revoke privileges on a per-user basis. Development, testing and maintenance personnel, on the other hand, require access to some parts of the system, but not all, indicating that an access level is needed that allows privileges to be granted on a per-user and what-do-you-need-to-do basis.

Req. ID	Description
3.8.1	<i>All users of the system shall login using some form of unique identification. (e.g., username and password)</i>
3.8.2	<i>All login attempts shall be done in a secure manner. (e.g., encrypted passwords)</i>
3.8.3	<i>A system administrator shall have unrestricted access to all aspects of the system.</i>
3.8.4	<i>Each user shall have a set of system access properties that defines the user's privileges within the system. (e.g., the subsystems a user may control or system tools the user may access).</i>
3.8.5	<i>The administrator shall have the ability to create and add a new user to the system.</i>
3.8.6	<i>The administrator shall have the ability to remove a user from the system.</i>
3.8.7	<i>The administrator shall have the ability to edit a user's system access properties.</i>
3.8.8	<i>The administrator shall have the ability to block all access to the system for all users or selectively by user. (All blocked users with active sessions shall automatically be logged off.)</i>

3.9 Installation and Upgrades

Req. ID	Description
3.9.1	Operations Activities – <i>the system shall continue operations, although not necessarily at full capacity, on all unaffected resources during partial shutdowns for maintenance, repair and/or upgrade.</i>
3.9.2	Test Mode – <i>the system shall be able to handle non-real-time operations in a transparent fashion (i.e., as if real-time). Note: non-real-time refers to input data from a source other than the Correlator.</i>
3.9.3	Replaceability – <i>modular design principles shall be employed to the maximum extent possible. Maximal practical use of available “hot-swappable” devices and components shall be made.</i>

3.10 Documentation

Req. ID	Description
3.10.1	Hardware – <i>complete and comprehensible hardware systems specifications and configuration information shall be readily available.</i>
3.10.2	Software Coding Practices– <i>software system and application code shall be well documented and written in a generally familiar language or languages (preferably not more than two). Software shall be written in a style that is easily readable and using practices that allow for minimal confusion.</i>