ACS Architecture Review

Meeting held 2001-05-22

Present:

G. Chiozzi	(Author)
B. Glendenning	(Chair)
B. Gustafsson	(Author)
R. Heald	
K. Morita	
J. Pisano	
M. Pokorny	
G. Raffi	
J. Schwarz	
S. Scott	

Summary Minutes

1. Agreement on agenda

The agenda was agreed to without comment.

2. Software packaging and distribution (e.g., SLALIB)

001, 038, 043, 173, 176, 190

Considerable discussion on this point occurred. The technical requirements call for astronomical libraries to be part of ACS, whereas they are silent on issues related to the "bundling" of tools. It is considered to be important for support reasons that the standard distribution use tested versions of tools and libraries. Although not an architectural issue, it was agreed to unbundled the distribution into more "tarballs". Similar discussions have also taken place in the context of SPR#5, which has some additional commentary and actions.

3. Use cases and process

002, 003, 197

It was noted that the ACS development process was not entirely consistent with the one outlined in the draft software engineering practices document, particularly with regard to use cases. There was some disagreement amongst the reviewers about whether having use cases would have been of technical assistance, however an interesting sociological comment was made that the process of creating use cases might have aided the "buy-in" processes. It is not practical to stop development now to produce use cases, however ACS is still in its infancy and we can use the experience to date while formulating the process for later iterations.

4. Property types

049, 050, 051, 055, old022

It was agreed that the following scalar types would ultimately be supported (not necessarily in ACS1.0), as well as array<scalar> and sequence<scalar>.

Long, Double, String, Boolean, Bitfield (formerly Pattern), Complex, Enum

5. Object life cycle/MACI and naming

059, 060, 064, 068, 069, 070, 187

It was agreed that the description in the document needs to be extended (to the ~2 page level), as presently many of the most important concepts are only described in the ACS MACI 1.0 paper.

6. Local configuration database copy
   071, 072

We agreed to remove the local configuration database from the architecture document. At DO initialization, a CORBA configuration database reference will be given to the object from its Activator. In this way if for performance reasons a local configuration database is determined to be necessary it can be introduced as an implementation detail. This solution also allows for alternative (e.g., from a laptop) configuration databases to be substituted at DO initialization for testing or other purposes.

7. Usage of XML
   109, 128, 129

While the use of XML for log messages appears appropriate as they are relatively complex (~6 fields), there was considerable skepticism about the utility of it for sampling files. It was agreed that this usage is largely an excuse to learn the technology and we may have to write conversion programs (e.g., to comma separated value files) to a format more useful for data analysis. This was generally, but not universally, accepted by the reviewers.

8. Smaller issues
   a. Abbreviations
      033, 115
      No discussion – leave to the authors' choice.

   b. Write-only properties
      048
      As discussed in the replies, they are not provided for.

   c. Property inheritance hierarchy
      053
      We agreed to simplify the Property inheritance hierarchy. The current description reflects unnecessarily pre ACS 1.0 implementation details.

   d. Simulations
      082
      As discussed at the Grenoble joint software meeting, ACS will provide property-level support only (turned on and off with a characteristic), and device level simulation support will be an application decision.

   e. Value retrieval mechanisms overview
      094
      As discussed in the replies some improvements organization of the text should be implemented. In particular, the comparison of the various value retrieval mechanisms will be moved to the end of the data channel section, perhaps as a separate subsection.

   f. Time formats
      164
      ISO formats only should be used "internally." It is understood that at user interfaces

more formats might be required. We could start with ISO format only and allow users to "insist" on their preferences via SPR's.

g. Root access
194
While not an architectural issue, the elided comments on this subject should remain in the "notes" document to ensure that we do not lose institutional memory of the discussions. This subject may be best carried further via SPR's.

h. Have attributes as well as characteristics & engineering support
old004
As described in the replies, this feature will be removed.

i. Remove "quality"
old015
As described in the replies, this feature will be removed.

9. Meeting summary and actions

The document is accepted, after revision to reflect the distributed replies to comments and the discussion at this meeting. The actions are only:

- ⇨ Glendenning to issue these minutes.
- ⇨ Document authors to revise the document and submit it to Glendenning for posting on the web (expected to occur in 2001-July).

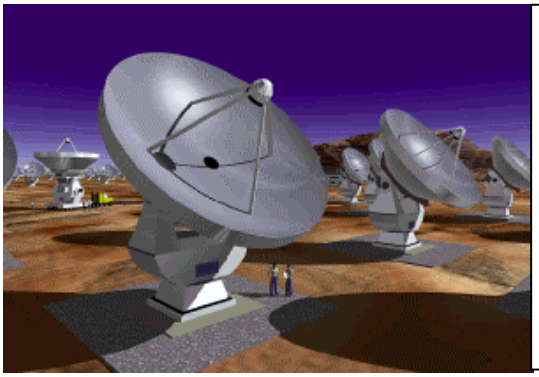

Atacama Large Millimeter Array	ALMA-SW-NNNN Revision: 2.0 2001-09-25

Notes on ACS Architecture

ALMA Common Software

G.Chiozzi, B.Gustafsson, B.Jeram
ESO

Keywords:	
Author Signature:	Date:
Approved by: Institute:	Signature: Date:
Released by: Institute:	Signature: Date:

Change Record

REVISION	DATE	AUTHOR	SECTIONS/PAGES AFFECTED	REMARKS
2.0/Prep.1	2001-04-11	G.Chiozzi et al.	All	Created. This document goes side by side with the ALMA Common Software Architecture document and has the same revision number.
2.0	2001-09-25	G.Chiozzi et al.	All	Added comments and replies for official review of ACS Architecture

Table of Contents

1 Introduction

1.1 Scope

This document collects notes to the ALMA Common Software Architecture document.

It has been decided to keep the notes separated from the main document, to avoid making it heavier to read.

1.2 Overview

The document includes notes in the following areas:

- Justification of architecture choices
- Replies to comments submitted for the previous issues of the ACS Architecture
- Considerations coming from the usage of ACS 0.0 prototype release, in particular related to the Kitt Peak 2000 test
- Considerations related to the distribution of the development in successive ACS releases, with particular reference to ACS 1.0.

2 ACS Basic Architecture

2.1 Introductory sections

2.1.1 Comments to issue 2.0/Prep 1 and replies after official review. General comments and chapter 1 included

```
---001--------------------------------------------------------------------
----

(Dirk Muders)

General comments:

There are many useful details concerning the various sections that ACS
Will cover. Sometimes the amount of detail seems to be too much at
least for the current use. The individual deliverables could be
prioritised further (extending maybe the traceabilility matrix).

However, the overall ACS structure / architecture is still not yet
clear
to me. The packaging diagram on page 11 indicates some kind of
levelling,
but it is not yet clear how ACS is deployed on the individual nodes
and what is necessary to actually "run" it.
Who controls the starting of servers and the creation of DOs on
a given node ?

The boundaries of ACS are somewhat fuzzy. On the hand it sounds like
ACS
is supposed to be the cross-platform "glue" between applications and
hardware devices. On the other hand the document mentions device
drivers and ACS applications. Isn't this outside the specs of ACS ?

A number of items like astronomical libraries or scripting languages
seem to be merely a packaging of existing software into a
distribution. While one might do that for practical purposes, I would
not
consider it part of the ACS architecture.

Quite a few subsections list either requirements or wish lists. These
Should occur in the previous document. Here one should
concentrate on real decisions concerning ACS
architecture / design.

REPLY:
Thanks for the very interesting comments, from which
we can extract very useful suggestions.

Right, this document is not a "pure" architecture
description as it should be and as described in the "SW Development
Process Methodology and tools". This is due to historical reasons.

Due to the iterative process we want to apply,
we will extend and complete the document in the next releases
and we plan already a new review with major updates and feedback
from "on the field" usage at the beginning of 2002.

On the specific comments from Dirk:
```

```
- We will see how to extend the traceability matrix.
  Any specific suggestion?
  We want anyway to keep architecture and planning separate.
  We have planning documents that details what a release
  (or, at least, the next one) will actually contain.

GCH 2001-09-07 Nothing specifically done in this document.
   We will take this suggestion into account for planning documents.

- The deployment discussion is missing and needed.
  We will add a section.
  The actual idea for deployment emerges from
  the ACS 0.0 prototype. We also had some discussion with
  Mick on the need of differentiating the deployment
  for development and run-time environment.

GCH 2001-08-28 DONE Added a deployment discussion

- We have tried to explain in several occasions what we think
  are the boundaries of ACS.
  In general we should assume that these are known from the ACS
  requirements, which define the scope of ACS.

    TO BE DISCUSSED AT THE REVIEW if this docuemnt
       needs to be improved in this respect.

GCH 2001-08-28 DONE added some explanations.
     No major restructuring done on this respect, as agreed
     at the review,

- Dirk is right in pointing out that in some cases we just list
  requirements, without describing how we will make sure they
  are satisfied.
  We will discuss the points one by one while replying to the
  respective comments from Dirk.

GCH 2001-09-11 Done.

---002-----------------------------------------------------------------
----

(Jim Pisano)

General Comments
=================

1. Documentation Level

I still have difficulty determining where this architecture document
fits into
the hierarchy of software engineering documentation. I view this
document
as an system overview document which should:

1. Identify the major packages.
2. Present an overview of the package functionality including:
  a. what are the main goals of each package
  b. definition the interaction between the packages
3. Provide a segue for more detailed designs of each package.
```

Item 1. is done on pp. 12 in the diagram.

Item 2. is not done. I noticed that it was in the ver. 1, but then removed due
to another reviewer's comment. IMO, this was a mistake.

Item 3. is done.

From what I understand, this document is part of the "Inception leads To Preliminary Design Review" step (from your SE practices doc) which is fine. But then this document then dives into many details of each package without providing a comprehensive overview. Once an overview is understood by the
reader, then it leads to a more detailed description of the packages. The approach in this architecture document makes the description Difficult to follow and leads the reader to ask why is such-and-such done? Also I find that I am unable to see what parts fit into TICS and how to fit ACS into TICS
- this should definitely be a primary goal of this document: basically what are the parts, how do they fit together and how & where would a user
implement them.

Along these lines, this document is already somewhat stale as Significant progress on the component design *and* implementation has already been made.
This complicates the delineation of high-level and package-level design as documents written after implementation always define what *has been* developed instead of what software *should* be developed.

REPLY:
As you mention, the overview was removed because of previous comments. We also agree that this was a mistake.

On the (design) details, see reply to 001.

On the problem of fitting TICS (or, better ALMA) to ACS and the other way around.
We agree that this document alone is not sufficient for this purpose. Our idea was that the examples and prototypes we have developed would have done the job, together with a strict co-operation between the TICS and ACS teams.
With respect to the "Inception leads to Preliminary Design Review", we miss the Use Cases that would have absolved the need of describing how ACS is supposed to be used and as a consequence described how TICS and ACS would have fit together.
As already said, we did not have resources to do both Use Cases and prototypes. We gave preference to the prototypes with the rationale that ACS is used by software developers that would have well understood the concepts by looking at examples of usage in the form of code and that would have learned by trying to use and discussing with the ACS team (while on the higher level, where "final users" are involved Use Cases are the only viable solution).

Unfortunately we did not have enough
interaction with the TICS team and many things remained and still are not clarified.

TO BE DISCUSSED AT THE REVIEW:
Reintroduce (improved) the overview, with a short description of the packages and of their relations.

GCH 2001-08-28 DONE

For the reasons given it is instead not proposed to add
Use Cases. At the review meeting it will have to be decided if
this is critical.
In this case we need more resources to allow this work to be done
without impacting in an unreasonable way the development milestones.

GCH 2001-08-28 This was agreed at the review. Nothing specific to be
done on this point

---

2. In reviewing the comments to revision 1 of ACS Architecture titled
"Notes
on ACS Architecture", there were many references to
ACSArchitectureNotes.doc.
As this explained some of the examples in ACS Architecture document, it
would
be useful to include as a reference.

REPLY:
Here we probably do not understand your comment.
The two documents will be available together, but we do not
feel that referencing in one document comments to the same document
is correct.

GCH 2001-08-28 Nothing specific to be done on this point

---

3. This design is "front loaded" and from the release of ACS 0.0, there
are many features which don't seem necessary but are included.
These extra features often seem to be carried over from VLT
software making the implementation difficult to understand,
learn and use. This point was brought up in the review
of the first revision, but only minor changes were made. I think a
good,
hard, honest look at the design with the intention of critically
analyzing each component for what can be trimmed or discarded
must be made. In the detailed comments, specific cases of trimming are
mentioned.

4. In following the iterative model of software development, namely one
that expands its scope with each iteration, the focus of this
document should be limited to requirements that are specific to
TICS. Especially since the test interferometer is meant to
be a prototyping test bed for all of the ALMA
software in multiple stages.

REPLY TO 3. AND 4.:
This architecture is clearly heavily influenced by the VLT experience.
There are also a number of very important differences.
We have tried to analyse these aspects and to profit as much as
possible from the lessons learned from the VLT.
If there are (as there are) features carried over from the VLT,
it is because in our experience they have been crucial for the VLT.
Some of these may not appear important at the beginning of the ALMA
development, as well as we discovered their value only after a number
of

years into the VLT development, but this is not a good reason
to ignore these lessons.
There are also many features of the VLT that we critically discarded,
because we have now learned that were done in the "wrong" way
or could have been done better.

This architecture document has much more of what could ever have
had if we would have all started from scratch, but this is because
it tries use make good use of all the experience gained with
previous projects.

Concretely there is one relevant VLT component (the Database), which
will be removed in favour of a commercial database (and files).

GCH 2001-08-28 Nothing specific to be done on this point

---003----------

(Joe Schwartz)

As I haven't been involved in the use of ACS or in discussions about
it up until now, many of my comments may simply reflect my ignorance
of the subject. Nevertheless, they may encourage you to make the
document clearer to the non-ACS expert.

In general, I am concerned that there has been very little contact
between the HLA work and the ACS. I understand "architecture" to be
something that is fairly general and shapes the software product well
beyond the next 1 or 2 releases, so the HLA and ACS groups need to
understand each other sooner rather than later. It's unfortunate that
we will once again be divided into parallel sessions on Thursday in
Grenoble, so that ACS and HLA will be passing each other like ships
in the night. Maybe we can make up for some of this on Friday.

General question: Were Use Cases ever done for ACS? They might help
the reader focus on how ACS is to be used in typical situations.
Without a context, it is hard to follow the requirements document,
AD01, as well as this architecture description. Many of the
requirements have an arbitrary flavor to them...

REPLY:
The first step is to have a more tight collaboration between TICS and
ACS. This is the reason why Joe and Dirk are invited to this review.

We assume that HLA is likely to impact some aspects of the ACS
architecture. Whatever feasible within the scope of existing
requirements
will be accomodated from release to release.
A revision of the ACS requirements when HLA needs are clearer would be
though the best starting point, in case the scope of ACS has to be
more substantially modified.

For the Use Cases see reply to comment 001.

GCH 2001-08-28 Nothing specific to be done on this point

---004----------

(Martin Pokorny)

General comments: The architecture document is quite thorough and complete; much thought has obviously been put into the document. It ought to provide a good foundation on which to develop its ideas in design and implementation. I believe that the document could be improved by a tighter focus on it subject (software architecture), and some clarification on a few issues, however.

AGREED

GCH 2001-08-28 Nothing specific to be done on this point

---005---

(Mick Brooks)

General Comments
==========

The explanation of the 3 tiered model is much better handled now.

Figure and table numbering and captions are required throughout as in the
document standard.

Although the overall architecture is becoming much clearer to me, the following issues are still of concern:

- The handling of sampling and monitoring for hardware access properties

- How to transmit data such as total power samples. This is not strictly
an ACS problem but more of an application, however some
service such as the data channel should be used for this.
- The life-cycle aspects of a DO. How and when are they created and destroyed needs to be made clearer.

ACCEPTED
We will extend these sections.
Some are clarified by the ACS 1.0 design documentation but it is worth to extend their architectural description.

GCH 2001-09-11 Done. I am not sure the work we have done
in extending these sections is sufficient, but more details will be available in the design documents.

---006---

(Ron Heald)

In these comments I mention only those areas I see as being terribly wrong and/or controversial. As such, they will be good topics for discussion in Grenoble. There are many more areas that need less significant changes, but time does not allow discussing them here.

General Comment: The purpose of this document has been muddled with the

introduction of other documents. On one side there's the "ACS Technical

Requirements", as it's been for some time. The relationship between the

"Requirements" and the current document is clear.

There are now new documents appearing that more fully elaborate the architecture and design of certain ACS sections. I'm referring to the titles "ACS Error System Architecture" and "Logging and Archiving. In "Logging and Archiving" it states "Where in conflict with the ACS Architecture document, this document shall take precedence". If this is true, what purpose does the current document's logging section serve?

Another such area is the Time System where I authored the TICS design documents "Time Distribution" and "Time System and Representation". All the material covered in the present document's "Time System" is repeated in my documents.
The present document's section is at best obsolete in many cases, and at worst, misleading and inaccurate. My documents explore the architecture and design in greater detail and are more up to date.

What purpose do you see for the present document? How do you intend to avoid duplication and conflict with the other documents? Perhaps the current document should be broken into parts, and eventually be replaced by the other documents?

REPLY:
This document is an Architecture document. As such it has to describe what the system is going to provide in order to satisfy the requirements
and how it will be structured without getting into the details of the implementation (this is the purpose of design).

Once the architecture is defined, the design has to come.
This is the purpose of the new documents that are coming out.
Design documents do not replace the architecture, but go inside into implementation details.

We think that the only practical way of proceeding is iterative and the waterfall does not work.
- We start from an architecture
- We make a design.
 - The design will make us discover new things and errors in the architecture.
 - We will then go back to align and fix this one and
 - Then back again to design with more details.
As described in the "SW development process", we need to have always both architecture and design information and to keep the documents aligned by iterating on them.

The phrase "Where in conflict with the ACS Architecture document, this document shall take precedence" is from and ACS 1.0 design document
and means that, for ACS 1.0, that document describes the actual design and that this can be in conflict with the ACS Architecture Document.

This conflict can be due to two different reasons:
- It has been discovered that the Architecture description was wrong or incomplete.
  At the next iteration (and after review) the Architecture will be modified accordingly.
- For practical reasons what described in the architecture cannot be implemented in ACS 1.0 (lack of time, lack of resources, other needed components not yet ready....).
  Than ACS 1.0 has to do in another way. This is documented and will be re-implemented or fixed in a next release.

GCH 2001-08-28 Nothing specific to be done on this point
a part from this clarification

---007----------

(Ron Heald)

General Comment:
IMO, the ACS Architecture document is NOT ready for review. It lacks coherence, cohesiveness, and polish. It does not present a clear vision. It contains many verbatim parts from the previous version along with pieces cut and pasted from comments made at that time. Several sections begin with paragraphs from the old version and end with new material; little attempt has been made to reconcile them, and the two parts are sometimes contradictory.
It's clear the document has been hastily put together to meet a deadline.

As proof of this conclusion I offer the comments received so far. More than half the paragraphs have someone who disagrees with them or requests them to be clarified. A document cannot be reviewed when there
is so many things wrong.

A couple of decisions were made at the Joint Software Meeting that effect the corresponding document sections. The biggest of these is the
decision to eliminate callbacks from ACS commands and property.get. This will result in the complete rewrite of many paragraphs in section 2.7.

I have been requested to re-write section 2.12 Time System to better align with the development I've been doing. There little sense in reviewing this section before I have completed the changes.

There is bad grammar in many places. This makes reading the document difficult, and worse, it makes several passages impossible to understand. How can a document be reviewed when its meaning is unclear? (I feel these corrections are your work, and not for me to do. I have made no attempt to correct grammar, and have guessed at the meaning of several sections.)

Given all the above, I feel that the best way to proceed at this point is to produce a new version of the document, and then commence the comment and review procedure. All the comments received so far, the Joint Software Meeting decisions, my changes to the Time System, and grammar corrections should be incorporated in producing the new version.

REPLY:
This comment comes a couple of weeks after 006 and seems to us in contrast with that. In that comment the situation did not seem so dramatic.
We do not agree neither with your analysis of the work we did nor with your opinion on the comments received.
There are many comments and they suggest a lot of improvements to the document, but none (or the sum) of them seem to us to mean that the document is terribly wrong.
At the Joint Software Meeting we have on purpose discussed major open points. The final decision on those points will be taken at the review and then the document will be modified accordingly.

We insist that you should re-write section 2.12 on the time system. This is your contribution to ACS and should be done in agreement with the rest of the document.
What we need is a one or two pages architecture description and it cannot be too different from what we already have.
At this point you are in an advanced development phase and you should have, as a consequence, already a detailed design.
We cannot believe that you do not have enough information to provide a reliable architecture description at this point.

For what concerns grammar, we can propose to write the various sections of the document in our native languages (Italian, Swedish and Slovenian).. but it would not help. It would also not improve the way we have all to work in an international collaboration, namely making some effort to be collaborative.

GCH 2001-08-28: We have tried to clarify as much as possible.
It was agreed at the review that the situation was acceptable and could have been fixed within the limits
of what decided in the review meeting.
TimeSystem section still missing from RHE.

NOT DONE GCH: thorough spell checking

---008---
(Koh-Ichiro Morita)
I have no strong comments about technical details of ACS architecture, which is described in this draft. On the other hand, I have a few comments on structure and contents in this draft.

Through whole text in the draft, logical explanations are not enough for me. These are desirable for person who is not familier with CORBA or new partner of this developments. For example, I need more explanations
about 4 layer package structures in p.11 and the logical concept of 3 tier model in p.12

Sequence of sections does not seem logical.
For example, I do not understand why second layer package "Sampling" appears just after the section of bottom layer package "Data channel".

You should write the caption for each diagram for reader to understand easily.

ACCEPTED:
- We will reintroduce the overview of packages removed from

```
   2.1 in previous issue
 - We will re-order some sections. You are right, sampling is in the
wrong
   place
 - We will add captions.

GCH 2001-08-28    Reintroduced overview of packages

---009-------------------------------------------------------------------
----
(Mick Brooks)
p4, S1.1 Para 1: This whole section sounds a little tentative given
that
much of the architecture
has already been implemented to some degree. The document doesn't just
propose an architecture any more,
it explicitly explains an architecture implementation.

REPLY:
The architecture is really still a proposal and we have implemented
essentially prototypes for demonstration.
We would prefer to strengthen this section at the next release.

GCH 2001-08-28:   Nothing to be done

---010-------------------------------------------------------------------
----
(Mick Brooks)
p4, S1.1 Para 2: This is not just a possible architecture, it is the
implemented architecture for ACS 0.0.

REPLY:
Same as 009

GCH 2001-08-28:   Nothing to be done

---011-------------------------------------------------------------------
----
(Mick Brooks)
p4, S1.1 Para 3: This paragraph is pretty hard to read. Can I suggest
for
the first sentence:
"This document provides a complete picture of the desired ACS
functionality
for the entire development phase, but
individual concepts and features will be developed incrementally over a
number of releases, according to the Software
Life Cycle described in [RD17]. For each release, a detailed plan will
be
developed, identifying the components to be
added or revised. Development priorities will be discussed with the
community of users during the planning phase
of each release."

ACCEPTED.

GCH 2001-08-28: Done

---012-------------------------------------------------------------------
----
(Mick Brooks)
```

p4, S1.1 Para 4: Again this is not a proposal any more. I think this sentence says "This
document confines itself to Common Software, including Control Software."
Control software is not part of
Common software, and if this is not what is meant, then the meaning needs
to be made clearer.

REPLY:
This is not we want to say.
What we mean is that we describe the part of the common software to be used
mainly for writing control software.
There will be most probably other need for data flow software, pipeline or
proposal preparation that we do not understand yet, because not enough work
Has been done in that areas.
We will then have to extend this document to cover also that needs.
Would this explanation be ok?

GCH 2001-08-28: Done according to this explanation

---013---
(Mick Brooks)
p4, S1.2 Para 1: You should mention that ACS will include OS builds and commercial device drivers as well as the
software written explicitly for ACS.

ACCEPTED

GCH 2001-08-28: Done

---014---

(Dirk Muders)
p4, s1.2, 1. paragraph:

I would stress right here that ACS will provide a generalized common interface
between applications and the hardware in order to facilitate the implementation
of new hard- and software.

ACCEPTED

GCH 2001-08-28: Done

---015---
(Brian Glendenning)
p.4 s1.2 paragraph 2: Reuse concerns might also be a reason to not use ACS.
For example, suppose that some existing pipeline infrastructure was adopted.

ACCEPTED
The requirement 3.1.3 actually allows exceptions and your example

would follow in this case.

GCH 2001-08-28: Done. Added explanation on exceptions

---016---
(Dirk Muders)
p4, s1.2, 6. paragraph:

... explicit decision "of" (the) ALMA (project) to use CORBA ...

ACCEPTED

GCH 2001-08-28: Done

---017---
(Ron Heald)
p5, last 3 para's of S1.2: These give the document history and should be
moved to the change record.

ACCEPTED

GCH 2001-08-28: Done

---018---
(Mick Brooks)
p5, S1.3 Para 2, Second sentence: suggest re-wording to "At both the antenna and the central control building there will
be not only Ethernet LAN connectivity but also a Field-bus[AD01 - 10.4.5
Field-bus] (the AMB) [AD01 - 10.5.12. Field-bus]
connected to various devices."

ACCEPTED

GCH 2001-09-03: Done

---019---
(Mick Brooks)
p5, S1.3 Para 2, Second bullet point: should be "ALMA Monitor and Control
Bus (AMB), a serial multi-drop master/slave bus based
on the CAN bus for communicating with simple I/O devices." The AMB term should also be referenced in the glossary.

ACCEPTED

GCH 2001-09-03: Done

---020---
(Mick Brooks)
p5, S1.3 Para 2, Bullet points 4, 7, 8 and 10: These all talk about the backbone being ATM or TBD. The current design concept
has no ATM but Gb Ethernet instead. As the italicized comment at the end
suggests, this whole section is irrelevant and could

be removed which would mean it wouldn't need to be edited every time the
network technology changes. I think that the salient
point here is that the primary transport protocol will be TCP/IP regardless
of underlying physical layers.

ACCEPTED
Will be removed as you suggest

GCH 2001-09-03: Done. Bullets removed.

---021----------
(Dirk Muders)
p6, s1.3, top:

Hasn't it been decided to use Gigabit ethernet instead of ATM ?

ACCEPTED
See 020

GCH 2001-09-03: Nothing to be done (removed). See 020.

---022----------
(Brian Glendenning)
p.6 s1.3 point 1&3: The baseline plan is to use Gb Ethernet as the antenna
network (even for the TI). For the TI 100Mb Ethernet will be used elsewhere,
with no decision yet for ALMA.

ACCEPTED
See 020

GCH 2001-09-03: Nothing to be done (removed). See 020.

---023----------
(Steve Scott)
p.6 First phrase: "suitable for a connection..." could be replaced
with a sentence that the Antenna LAN and the central building LAN
will be as much as 25KM apart, but will be interconnected by the
Antenna network (which is really a WAN).

ACCEPTED

GCH 2001-09-03: Nothing to be done (removed). See 020.

---024----------
(Jim Pisano)
p. 6 1.3 "Other Central Computers" - the correlator computer is already
mentioned in a previous bullet - "Correlator Control Computer (CCC)".

ACCEPTED

GCH 2001-09-03: Nothing to be done (removed). See 020.

---025----------
(Steve Scott)
p.6 1.3 Italicized comment at the end of section about the
irrelevance of the network architecture. The network architecture
is important, particularly the distance. Round trip light travel
time in fiber will be 0.3 - 0.4 milliseconds, long enough to be a
consideration in protocol choice.

ACCEPTED
See 020. This document does not make any choice on network protocol.

GCH 2001-09-03: Nothing to be done (removed). See 020.

---026----------
(Mick Brooks)
p6, S1.4 and 1.5: Given that the descriptions of applicable documents
and
reference documents are identical, what is the reason for
grouping them separately? A single list of references should be
sufficient,
with the requirements doc being ref no 1.

ACCEPTED.

GCH 2001-09-03: Done

---027----------
(Mick Brooks)
p6, Reference documents: Some dates are not in the ALMA standard
format:
RD1, RD2, RD12, RD29, RD30, RD36.

ACCEPTED

GCH 2001-09-03: Done

---028----------
(Mick Brooks)
p6, RD01 is not an ALMA software memo. It is an ALMA-US software memo
which
has never been reviewed and is hidden on a
different web page than the rest of the memos.

ACCEPTED
Reference will be removed/replaced

GCH 2001-09-03: Done. Reference removed.

---029----------
(Mick Brooks)
p6 RD02 is also not an ALMA Computing Memo. It is also hidden on a
different web page.

ACCEPTED
Reference will be removed/replaced

GCH 2001-09-03: Done. Reference replaced with ALMA Project Book.

---030---
(Jim Pisano)
p. 6 1.5 Some of the reference documents no longer apply, e.g., RD01, RD06,
RD09, RD14, RD15, RD20, RD21, RD23, RD24, RD26.

ACCEPTED

GCH 2001-09-03: Done. References removed and packed.

---031---
(Mick Brooks)
p8 RD36 is a draft software memo and has a newer revision than the one mentioned.

ACCEPTED

GCH 2001-09-03: Done. Latest version on the web has date 2001-08-11.

---032---
(Mick Brooks)
p8 S1.6 and S1.7: Information about the location and function of the glossary is repeated. Should be referenced only once.

ACCEPTED
We will merge sections 1.6 and 1.7

GCH 2001-09-03: Done. Sections merged.

---033---
(Ron Heald)
p8, S1.6: Several of these appear in the glossary, and should be removed
from here. Also please remove "WS" and "SW", and use "workstation" and "software" in the text instead.

TO BE DISCUSSED AT REVIEW
We think that having some abbreviations and acronyms at hand is useful for
people reading the document far from a terminal.
We will reduce the list.
What is wrong with WS and SW?

GCH 2001-09-03: Nothing to be done. At review we decided that WS and SW are acceptable abbreviations and can be used.

---034---
(Jim Pisano)
p.8 1.6 Many abbreviations & acronyms are listed but not used, e.g., CASE, CCD,
CCS, CDT, DBMS, EVH, FDR, ICD, IEEE, OCL, OO, SNMP, SRS, UIF, WAN & WS. Please
remove them.

```
ACCEPTED

GCH 2001-09-03: Done.

---035-------------------------------------------------------------------
----
(Mick Brooks)
p10 S1.7, last sentence: sentence is unnecessary because previous
sentence
says the same thing.

ACCEPTED

GCH 2001-09-03: Nothing necessary. Section removed

---036-------------------------------------------------------------------
----
(Ron Heald)
p10, S1.7: This repeats what was said in S1.6.  Remove this section
completely.

ACCEPTED

GCH 2001-09-03: Nothing necessary. Section removed
```

2.1.2 Comments to Issue:1.1/Prep.2. General comments and chapter 1 included

```
---001------------------
(BrianGlendenning)
p.0
> This is an important comment: A  discussion of the "data pipes" and
perhaps
> publisher/subscriber is missing but important as much of the system
(including
> ACS itself) will use them.

ACCEPTED
We agree that this is an important missing part.
We have up to now just defined the concept we have in mind for the
"data pipe"
and the high level architecture, but still a detailed architecture and
a
design are missing.
We will add references to documentation in books/web to the
"data pipe" and publisher/subscriber design patterns, that are our
starting points.
The actual detailed design and architecture requires more investigation
and
prototyping work, in part to be done for ACS 0.0, but mostly later.

GCH 2001-04-06: Added section on data channel

---002------------------
(JimPisano)
>
> I find this document describes a well thought-out system and a good
basis
> for a library of operations used by many devices.
>
```

> A general concern that I have is that there is a lot of functionality which
> requires implementation that may not be necessary. As a system matures these
> features may be useful, but front-loading a software project at the beginning
> can lead to steep learning curves and many delays. Most of my comments deal w/
> removing or postponing features which fall into this category.
>

(RobertLucas)
> General impressions:
> This looks very complicated. I didn't understand: are all those
> features required (it looks like there are more features than in the
> requirements document)? or do many come features for free by using
> existing
> systems?

(RonHeald)
> Important comment: The ACS is too big and too complex. It tries to do
> too much. It fulfills the requirements and continues on for another 50
> km. Instead let's start with basic functionality we know is needed, and
> add things later as experience dictates. Remember "complex is easy,
> simple is difficult".
>

AGREED.

Einstein's principle that "a model should be as simple as possible AND YET NO SIMPLER"
is very well understood.
The idea of simple first and more staff later is also well received. This is
why we have releases. ACS 0. proposes the very first things.
Agreed also to keep mechanisms simple and few in the beginning and expanding
after some experience.
It is anyway to be taken into account that what is described in this document is extensively
based on years of experience we all have in similar projects.

Please note that it is intended to review ACS TR and ACS Architecture after
the KP tests and so the software should be constantly upgraded and refined.
At this point a development plan will be developed to identify in detail
which features will be developed first and which in successive releases.
Adjustments will be done according to the feedback received.
This is stated in Requirement 11.1.4 (Releases).

GCH 2001-03-26: Added comment in Introduction.Scope section.

> The numbering of the document is very difficult to refer to. All heading

> numbers should be fully qualified, e.g., 1.3.4.7.1. I had a
difficulty
> figuring out which number to refer to in my comments.
>

ACCEPTED.
The current numbering will be fixed.

GCH 2001-04-06: Done. Document uses modified H3 and H4 styles now

> As has been mentioned, there are many grammatical issues which are
> understandable
> due to the main authors (Chiozzi & Gustaffson) have English is a
second
> language.
> But as there is an American author (sorry Fritz)for most of these
joint
> documents,
> it should be the American who should take responsibility for editing
these
> docs for grammar as they will make them must easier to read & more
concise.
>
> My only specific grammatical pet peeve is that "warranty" is a noun,
not a verb.
> The word you want to use is guaranty.
>

ACCEPTED.
We will do our best to fix grammatical issues. All suggestions are
welcome.

BGU 2001-03-26: I have checked and corrected spelling and grammar.
To be done again after editing of the document.

---003------------------
(MickBrooks)
The document seems to be coming together well and gives a great
overview. I
hope we can get down to some hard coding soon!

ACCEPTED.
Coding for ACS 0 is already started. Many components need prototyping.

GCH 2001-03-26: Nothing to be done

---004------------------
(RonHeald)
>
> Detailed Comments on ACS Architecture, Issue 1.1/Prep.2
>
> Differences with DRUC -
>
> 1. Default values are contained in the property and are set whenever
an
> initialize is given.
>

(RonHeald)
p.11 2.2.1
> Add "Attribute" for non-static data associated with a property.

TO BE DISCUSSED AT REVIEW
The difference here seems to be between compiled-in values (as you argue
for) and initilized values.
Although the distinction between Attributes (read/write) and Characteristic
(read-only) is nice in principle:
- the distinction is less unambiguous than it seems (e.g. you say that defaults are read-only, ranges read-write; it could also be argued that they should
be the other way around).
- it does not seem correct to have half of the values hard-coded and half-
 in a configuration database. This is also an arbitrary interpretation of
 the ACS TR requirement 4.2.1 (DO information in Config. DB)
- the present system is simpler, as characteristics are all R/O and
  only properties can be R/W.
To make a distinction is OK, while to advocate a different structure and behavior adds complexity for a feature of debatable value.

The proposal in 2.2.5.4 is to start simple with read-only characteristics and, if
experience shows a real need, make them read/write at a later time.

If this approach is accepted and as already requested in the past, the definition
of Attribute should be removed from the ALMA Glossary.

GCH 2001-03-26: REMAINS OPEN FOR REVIEW

> 2. When read-only properties are read they should determine if their
> value is acceptable, marginal, or outside the proper range.
>
CLARIFICATION NEEDED
When a property is read, the read returns just the value (together with a
completion code, that contains a timestamp).
What you described is in the current design handled by alarms: it is possible to
attach call-backs to properties to be triggered when the value of a property changes band
between acceptable, marginal or outside range.

Do you want also a public method to explicitly ask a property in which band the value is?
This can be done, but does it really add something useful to the alarm mechanism?

GCH 2001-03-26: NO ACTION. DO WE NEED TO ADD SOME DETAILS? OPEN FOR REVIEW?

> 3. Provide properties with some type of "engineering support". Perhaps
> when in engineering mode the property would accept input from
> engineering and ignore other input.
>
TO BE DISCUSSED AT REVIEW
The engineering access mode is meant to be a way to set a
value and to make it read-only. It is similar to changing the

permissions on a file. It is felt that some values might be changed by software, and to debug/test things, we would like to force the values.

Scenario:
  change permission to read-write
  set value
  change permission to read-only
  test, debug
  change permission to read-write

This can be implemented adding a Lock(TRUE|FALSE) method to Properties.

G.Chiozzi and other people at ESO are against such an "engineering mode", based on
the experience with the VLT system: such a concept has been foreseen, implemented and
later on removed on request of the commissioning and operation teams. It adds complexity and adds overhead to the
work of engineers, that are supposed (and suppose) to know what they are doing.

The suggestion is propose the Lock() solution in the document, but give low priority to the implementation so that we can first verify the need for it.
Implemented in this way, it does not have an impact on other components of the system.

GCH 2001-03-26: REMAINS OPEN FOR REVIEW

> 4. Explicitly say that monitoring and sampling will not interfere with
> control.
>
OK, but this is in the ACS requirements: 4.1.2 and 4.3.1

BGU 2001-03-26: I have added something like this for sampling.

> 5. Properties that deal with physical values should input/output values
> in SI units.
>
We fully support strong commitment to SI units.
It is though more a requirement to M&C. It does not make sense for ACS to check this.

A library for unit conversions and handling would be very useful and will be mentioned in
section 2.18.3.
Can anyone provide a reference to such a library? I have in mind the unit handling
capabilities of my HP48, but I have never found such a library for C,C++ or Java.

GCH 2001-04-06: Mentioned AIPS++

---005------------------
(RonHeald)
p.0
> Replace all "ALMA Common Software" with "ACS", or vice versa.
>

GCH 2001-03-27: Done

> Please don't justify the right margin as it makes some lines
difficult
> to read. (For example, p.12, line 3.)
>
GCH 2001-03-27: Template now handles that

> Please put a blank line between paragraphs to improve readability.
>
GCH 2001-03-27: Template now handles that

> Please remove (or move to an appendix or other document) all the
little
> "footnote" information; it's not part of the architecture and
clutters
> up the document.
>

GCH 2001-04-06: Done. Some notes move in the Architecture Notes
document

> Please make the paragraph numbering consistent. Sometimes it's 2.17,
1,
> 1, etc., sometimes it's 2.18, 2.18.1, etc. (The latter is easier to
> reference.)
>

GCH 2001-04-06: Done

> Please turn on the MS Word spelling and grammar checker and fix the
> problems it finds.
>
OK

GCH 2001-04-06: Done

---006------------------
(BrianGlendenning)
p.6 1.3 LAN
> I think maybe the maximum cable run is 25km, not 20km (I think this
was my
> mistake in the technical requirements).
OK

BGU 2001-03-26: Done
---007------------------
(BrianGlendenning)
p.6 1.4 Authors reversed.

OK

BGU 2001-03-26: Done
---008------------------
(MickBrooks)
p.6 Section 1.3 You might want to include the ARTC and CCC in here as
these
are new architectural elements.

```
OK

BGU 2001-03-26: Done
---009------------------
(MickBrooks)
p.6 Section 1.5 There is a new revision of RD01 which you should
reference.

OK

BGU 2001-03-26: I have not found anything like that
---010------------------
(MickBrooks)
p.7 There is a new revision of RD12 which you should reference.

OK

BGU 2001-03-26: Done
---011------------------
(MickBrooks)
p.7 Section 1.6 Perhaps you should add ACC and ARTC, although the link
to
the glossary should suffice. As you say, the network architecture is
kind of
irrelevant.

(RonHeald)
p.8 1.7
> A definition for "alarm" and "error" is needed.  A definition for
"array
> time" is needed, and standardize its use, for example, no "array
> standard time".
>

(RobertLucas)
Glossary:
- why not only refer to the ALMA glossary ? or take definitions from
there.  I don't see the need to include a glossary in *all* documents,
specifically if they do not agree.

- observation block: disagree on definition. (see ALMA Glossary).

- generally several terms still refer to optical observations

- (observing program -> nights) service observing : not foreseen for
ALMA

ACCEPTED.
The ALMA Glossary was not available at the time!
It will now be referenced here.
A "local" glossary will include only terms not part of the ALMA global
one.

GCH 2001-03-27: Cleaned up and added link to ALMA glossary
```

2.2 Overview

2.2.1 Comments to Issue 2.0/Prep2. Section 2.1

```
---037------------------------------------------------------------------
----
```

(Ron Heald)
p10, S2.1, first 2 para's: This repeats what was said in S1.2.

REPLY
Very often in books, papers and document in general, the introduction is made
up of small essential paragraphs extracted from the body of the document.
We do not see anything wrong in the fact that first para in 1.2 is similar to
2.1. Actually it is a summary of that para.

GCH 2001-09-03: Nothing to be done. Clarified

---038---

(Dirk Muders)
p11, s2.1, package diagram:

As mentioned above, I think the astro libraries are too general to be part
of ACS. If they are in ACS, then I don't think they belong that far up in
the diagram. They typically do not use any of the ACS functionality beneath
them because they are unaware of it. The libraries might instead be used by
some lower level ACS component. Again, I think they should live parallel to
ACS, not within ACS.

ACCEPTED
You are right saying that astro libraries will be unaware of lower level
ACS components in the diagram. We do not believe that any of the lover level
components in this diagram will use them.

Astro libraries should be part of ACS because will be commonly used by many
different teams for many different purposes.
This means that they will be provided as integral part of the ACS package, in
general without any development on the side of ACS.

We will move the parallel to the hierarchy of modules providing basic ACS
services.

GCH 2001-08-28 astro libs are now in layer 1

---039---

(Joe Schwartz)
p. 11 states "Each package will be described in a separate section in this chapter", but the section names are often different from the names on the package diagram, e.g., "ACS Application Framework" vs. "High level application framework". Moreover, sections 2.3 and 2.4, "Direct Value Retrieval" and "Value Retrieval by Event" are not, of course, packages at all. These two sections ought to be subsections of 2.2, "Distributed Object", I believe.

ACCEPTED
We will fix these inconsistencies

GCH 2001-09-03: Done.

---040----------
(Martin Pokorny)
p.11 The description of the levels is not clearly associated with the figure. Depending on whether you take the "first" level to be that including "CORBA" or that including "Distributed Object", it's easy to be confused. Using the description of the fourth layer, I think that I finally got it right, but it would helpful to be more specific here.

ACCEPTED

GCH 2001-08-28 Done. Text and figure aligned.

---041----------
(Martin Pokorny)
p.11 Are the libraries and scripting language really a _part_ of ACS? I have some doubts about this. I noticed that you comment on this fact in section 2.17, but I still don't understand why these should be a considered as part of ACS. It seems that you are raising issues of software engineering and system configuration with respect to these inclusions, not software architecture.

REPLY
They are part of ACS requirements.
ACS should be seen as a packaging of commonly used components, some to be developed and as much as possible are off the shelf (may be with minor
adaptations)
For libraries, see 038.
For scripting languages, it will be typically necessary to provide a small
support library to make easy to access other ACS components/services/features
from the scripting language itself.
The reasons for this approach of having a common software are actually the software engineering and system configuration issues you mention.
The common software wants to be the main tool to make easy and acceptable enforcing the SW engineering rules and the two go hand in hand.

GCH 2001-08-28 Nothing specific to be done, a part from this clarification

---042----------
(Mick Brooks)
p11 Diagram: Should have a figure number per the document standard. Some
explanation of the different colors used for packages would
be helpful. The text "Packages are allowed to use services provided by other packages on the lower layers and on the same
layer, but not on higher layers." seems to imply that DOs cannot call astrometrical library routines, for example. Is this true?

```
ACCEPTED
Captions and explanations will be added.

Having the DO at a lower level than astro libraries, means that
the libraries implementing the basic classes for DOs cannot (and
do not) use them.
But DOs at application level, like for the Control System, will
be developed in packages on top of the whole ACS. They will inherit
from the
classes provided in the DO package and will be free in their
implementation
to use astro libraries.

GCH 2001-09-11 Done. Captions added. Clarified here.

---043----------------------------------------------------------------
----
(Ron Heald)
p.11 2.1 (also p.41 2.17.3) There are two concepts of ACS: one is as a
software package used throughout ALMA, and another is as a version
control and software release system.  ACS should only be an ALMA
software package, and therefore should NOT include the FITS, astronomy,
and user interface libraries.  Nor should it include scripting
languages
or any other externally provided software.  It's fine to specify the
versions of external software ACS needs to work properly.

REJECTED
FITS, astronomy, UIF libraries, scripting languages are in ACS
requirements.
ACS is really both things.

GCH 2001-08-28    Nothing specific to be done on this point
                  Clarified at the review

---044----------------------------------------------------------------
----
(Steve Scott)
p.11 The grouping of packages into layers is subtle but very
important. It will significantly reduce the cross coupling that
makes system maintenance (and even building) difficult. But in
general the astro libs are uncoupled from eveything else and
belong on the bottom layer - they have been misplaced. The layers
are a useful construct.

AGREED

GCH 2001-08-28    astro libs are now in layer 1
```

2.2.2 Comments to Issue:1.1/Prep.2. Section 2.1

```
---038----------------------------------------------------------------
----
(Dirk Muders)
p11, s2.1, package diagram:

As mentioned above, I think the astro libraries are too general to be
part
of ACS. If they are in ACS, then I don't think they belong that far up
in
```

the diagram. They typically do not use any of the ACS functionality beneath
them because they are unaware of it. The libraries might instead be used by
some lower level ACS component. Again, I think they should live parallel to
ACS, not within ACS.

ACCEPTED
You are right saying that astro libraries will be unaware of lower level
ACS components in the diagram. We do not believe that any of the lover level
components in this diagram will use them.

Astro libraries should be part of ACS because will be commonly used by many
different teams for many different purposes.
This means that they will be provided as integral part of the ACS package, in
general without any development on the side of ACS.

We will move the parallel to the hierarchy of modules providing basic ACS
services.

GCH 2001-08-28 astro libs are now in layer 1

---039---
(Joe Schwartz)
p. 11 states "Each package will be described in a separate section in this chapter", but the section names are often different from the names on the package diagram, e.g., "ACS Application Framework" vs. "High level application framework". Moreover, sections 2.3 and 2.4, "Direct Value Retrieval" and "Value Retrieval by Event" are not, of course, packages at all. These two sections ought to be subsections of 2.2, "Distributed Object", I believe.

ACCEPTED
We will fix these inconsistencies

GCH 2001-09-03: Done.

---040---
(Martin Pokorny)
p.11 The description of the levels is not clearly associated with the figure. Depending on whether you take the "first" level to be that including "CORBA" or that including "Distributed Object", it's easy to be confused. Using the description of the fourth layer, I think that I finally got it right, but it would helpful to be more specific here.

ACCEPTED

GCH 2001-08-28 Done. Text and figure aligned.

---041---
(Martin Pokorny)
p.11 Are the libraries and scripting language really a _part_ of ACS?

I have some doubts about this. I noticed that you comment on this fact in section 2.17, but I still don't understand why these should be a considered as part of ACS. It seems that you are raising issues of software engineering and system configuration with respect to these inclusions, not software architecture.

REPLY
They are part of ACS requirements.
ACS should be seen as a packaging of commonly used components, some to be developed and as much as possible are off the shelf (may be with minor
adaptations)
For libraries, see 038.
For scripting languages, it will be typically necessary to provide a small
support library to make easy to access other ACS components/services/features
from the scripting language itself.
The reasons for this approach of having a common software are actually the software engineering and system configuration issues you mention. The common software wants to be the main tool to make easy and acceptable enforcing the SW engineering rules and the two go hand in hand.

GCH 2001-08-28 Nothing specific to be done, a part from this clarification

---042----------
(Mick Brooks)
p11 Diagram: Should have a figure number per the document standard. Some
explanation of the different colors used for packages would
be helpful. The text "Packages are allowed to use services provided by other packages on the lower layers and on the same
layer, but not on higher layers." seems to imply that DOs cannot call astrometrical library routines, for example. Is this true?

ACCEPTED
Captions and explanations will be added.

Having the DO at a lower level than astro libraries, means that the libraries implementing the basic classes for DOs cannot (and do not) use them.
But DOs at application level, like for the Control System, will be developed in packages on top of the whole ACS. They will inherit from the
classes provided in the DO package and will be free in their implementation
to use astro libraries.

GCH 2001-09-11 Done. Captions added. Clarified here.

---043----------
(Ron Heald)
p.11 2.1 (also p.41 2.17.3) There are two concepts of ACS: one is as a software package used throughout ALMA, and another is as a version control and software release system. ACS should only be an ALMA software package, and therefore should NOT include the FITS, astronomy,

```
and user interface libraries.  Nor should it include scripting
languages
or any other externally provided software.  It's fine to specify the
versions of external software ACS needs to work properly.

REJECTED
FITS, astronomy, UIF libraries, scripting languages are in ACS
requirements.
ACS is really both things.

GCH 2001-08-28    Nothing specific to be done on this point
                  Clarified at the review

---044-------------------------------------------------------------
----
(Steve Scott)
p.11 The grouping of packages into layers is subtle but very
important. It will significantly reduce the cross coupling that
makes system maintenance (and even building) difficult. But in
general the astro libs are uncoupled from eveything else and
belong on the bottom layer - they have been misplaced. The layers
are a useful construct.

AGREED

GCH 2001-08-28    astro libs are now in layer 1
```

2.2.3 Comments to Issue:1.1/Prep.2. General comments and sections up to2.1 included

```
---012------------------
(BrianGlendenning)
p.10 2.1
> "Each package provides a basic set ... that shall be used by all ALMA
> applications." It sounds like you are saying that all applications
will use
> all packages, which is presumably not so (e.g. few things would use
FITS).

ACCEPTED: It was meant so. We should say that application
shall use what is applicable to them, as it is stated in AD01 3.1.3.

BGU 2001-03-26: Done

---013------------------
(RonHeald)
p.10 2.1
> Most of the verbiage after the drawing can be removed as each point
is
> described better (and with more detail) later.  Replace it with
> something about the 4 levels and by saying each point is now
described.
>
ACCEPTED

BGU 2001-03-26: Partially cleaned up.
GCH 2001-03-27: OK now

---014------------------
(RonHeald)
p.10 2.1
```

```
> How can the Distributed Object be the core implementation for
> Distributed Objects? Maybe the sentence should read "Distributed
Object:
> core implementation that is the base of the distributed model"?
>
CLARIFICATION
This is a list of PACKAGES.
The Distributed Object PACKAGE contains the core implementation for the
Distributed Object, Property and Characteristic classes.
We will anyway re-phrase to make it more explicit.


GCH 2001-03-27: Removed. Nothing to be done
```

2.3 Distributed Object

2.3.1 Comments to Issue:2.0/Prep.1

```
---045--------------------------------------------------------------------
----
(Ron Heald)
p12, S2.2: There are some objects (such as the Time System's Epoch and
Duration) that are dynamically created and destroyed by clients.  How
will they be handled in this scheme?

REPLY
They are not DOs

GCH 2001-08-28    Clarified here. Nothing specific to be done on this
point

---046--------------------------------------------------------------------
----
(Ron Heald)
p12, S2.2.1: Is "3 tier naming for the logical model" the same as "3
tier logical model"?  If so combine this section with S2.2.2.

ACCEPTED

GCH 2001-09-07 Done.

---047--------------------------------------------------------------------
----
(Steve Scott)
p12 2.2.1: What are the principles of the three tier naming? An
explanation and example would help here.

REPLY
Sections 2.2.1.* were supposed to give the explanation and example.
We will try to discuss with you directly and find out how to make
The description clearer.

GCH 2001-09-07 We hope the changes i plemented in the document
               make the concept clearer now.

---048--------------------------------------------------------------------
----
```

(Jim Pisano)
p.13 2.2.1.2.1 I thought that it was decided that we would not have write-only
properties. You stated in the Notes document that write-only properties won't
get implemented, no use writing about vaporware. Strike 2-nd sentence.

TO BE DISCUSSED AT REVIEW
We agree.
We left the text because decisions can be taken only at the official review.

GCH 2001-09-07 Done. No write only properties any more as agreed at review

---049----------
(Steve Scott)
p.12 2.2.1.2.2: Why is the set of basic data types so limited? No boolean? It seems like this will result in replicated code to solve similar problems over and over.

REPLY
Actually we want to avoid implementing the same code for many different types
and conversion problems between similar types (like short, int and long).
On the VLT we had many more types and we had a lot of problems of this kind,
so we have decided to radically reduce them to the minimum that can cover all
needs.
Also, nowadays saving a couple of bytes using a short instead of a long usually does not pay:
- it introduces performance problems (CPUs now always works with longs
and every operation on a short requires a conversion to long)
- it introduces alignment problems in memory and when transferring structures from different platform
We have been discussing the introduction of structures of the basic types.
This would be probably useful, but before doing it we would like to have a real
Need and a proofed advantage, since it introduces a significant increase of
complexity in the handling libraries.

GCH 2001-09-07 Done. Added explanation.

---050----------
(Dirk Muders)
p12, s2.2.1.2.2 (this level of subsectioning is quite high):

If ACS deals with science data, one should add "Complex" as a type.

TO BE DISCUSSED A REVIEW
This seems a very reasonable request

GCH 2001-09-07 Done.

```
---051-----------------------------------------------------------------
----
(Jim Pisano)
p.13 2.2.1.2.2 "An array of ...". This should be changed to "A sequence
of ..."
as IDL distinguishes between a sequence and an array as the former is
variable
length & the latter is fixed length. Static length is a special case of
variable length and probably won't get used enough to warrant a special
case.

REPLY
Actually this requires more explanation.
We have two cases, both handled by ACS:
- A sequence of "long properties" (where we started from)
Is an IDL sequence of Long property types.
This means that each item in an object, an instance of the long
property.
This is handled directly by CORBA and IDL and there is no special code
to
write.
- A property of type "array of long" (requested by Brian in previous
issue)
In this case the property itself contains an array and it is one single
object.
   This is much more efficient than the sequence of properties, but has
to be
   developed as a special kind of property.

GCH 2001-09-07 Done. Clarified and extended.

---052-----------------------------------------------------------------
----
(Joe Schwartz)
p. 13, Section 2.2.1.2.3: If this is the third of the "three tiers",
then the section ought to be numbered 2.2.1.3, to correspond to the
sections on DOs and Properties.

ACCEPTED

GCH 2001-09-07 Done.

---053-----------------------------------------------------------------
----
(Joe Schwartz)
p. 14 (and elsewhere). Figures should have numbers and figure
captions. In addition, it would be nice to have some narrative for a
(to me) complex UML diagram like that for the NamedComponent
hierarchy.

(Brian Glendenning)
p.14 Figure - The Property hierarchy appears a bit deep - 5 layers to
get to
RO/RW property. I would combine Property/Typeless property and probably
P<type>/RO<type>.

TO BE DISCUSSED AT REVIEW
The named component hierarchy comes straight from the design UML model.
From the architecture point of view, it is too complex and could be
simplified to show only the architectural classes, more as it was in
the previous issue.
```

Should we go back to a simpler architectural diagram?

GCH 2001-09-10: Added narrative and made Architecture diagram lighter

---054----------
(Brian Glendenning)
p.14 Figure - Figures (here and elsewhere) should have figure numbers (and
usually captions) for referencing purposes.

ACCEPTED

GCH 2001-09-12 Done.

---055----------
(Mick Brooks)
p14, Diagram: Contains property type called pattern which was not mentioned
in Section 2.2.1.2.2

ACCEPTED
Pattern and Enum are similar, but not the same.
We will explain.

GCH 2001-09-10: Done.

---056----------
(Jim Pisano)
p.15 2.2.3 Is the 3-tier logical model the DO-Properties-Characteristic model
or the client-server w/ CORBA middleware model? I don't see how the 2-nd
sentence has anything to do w/ AD01 13.1.1. Strike it, you can still keep the
reference for the 1-st sentence.
(Mick Brooks)
p15, S2.2.3: This section repeats what was said in Sections 2.1 and 2.2.2.
Section 2.2.3 should be deleted.

ACCEPTED
2.1, 2.2.2 and 2.2.3 will be merged avoiding redundant information.

GCH 2001-09-07 Done.

---057----------
(Martin Pokorny)
p.15 s2.2.3.2 In the current design, there are not "specific subclasses" that provide access to specific implementations. The last sentence of this section should be deleted.

REPLY:
Your CAN access properties are specific subclasses

GCH 2001-09-07 Clarified. Nothing specific to be done here.

```
---058-------------------------------------------------------------------
----
(Joe Schwartz)
p. 15, 2.2.3.4, I don't understand the "generic name=value type of
interface"; can you supply an example?
(Ron Heald)
p15, S2.2.3.4: What is the "generic name=value type of interface", and
what is its purpose?

REPLY:
We will add an explanation.
This feature was correctly requested by Brian.
The interface of properties is defined by their IDL and the IDL is
the same independently from the implementation (logical, CAN...).
But specific implementations will have also specific characteristics.
For example a CANLong property has a CANID characteristic.
This means that from the property's IDL there is no way to retrieve
The CANID using CORBA calls.
We provide then a generic interface that can be used to retrieve
Any characteristic just querying by name.
This is described with more details in the ACS 1.0 IDL and design
documents.

GCH 2001-09-07 Done.

---059-------------------------------------------------------------------
----
(Steve Scott)
p.15 2.2.3.6: There is a lot hidden here. This is a very important
part of the architecture. How do you make sure that some computer
will implement the DO servant? How do you find it, start it? Etc,
etc, etc. This topic should be several pages, not a single
paragraph.

(Mick Brooks)
p15, S2.2.3.6: The discussion of the manager and activators role in a
DOs
life cycle should be expanded. I don't understand
what is being proposed here. Under what client service requests does
the
activator perform code loading and DO construction
and destruction? The diagram on the next page helps a little but some
explicit explanation would be good.

(Jim Pisano)
p.15 2.2.3.6 Since MACI appears is a crucial component of this design,
then at
a reference to the MACI document(s) should be included to clarify what
the
other terms mean in this bullet, e.g., "Activators", life cycle of
objects,
etc. Better yet, it would be useful for a summary of MACI features
similar to
what Gianluca wrote in "Answers to Questions of Fritz Stauffer about
the
benefits of MACI". BTW, the "Managers" in the various diagrams (e.g.
"Manager"

REPLY:
You are right. This section should be expanded.
```

The concepts are described with much more details in the MACI Design Document
For ACS 1.0.
Since we expect that this part will be accurately reviewed next year, we
propose to say that the section will be expanded in the next issue of the
document and to reference for the time being the MACI documentation.

GCH 2001-09-12 Done.

---060---
(Brian Glendenning)
p.15 s2.2.3.6 The usefulness of MACI is not apparent to me. At the least the
fact that it is an obstacle to preventing implementation of BACI objects in
languages other than C++ needs to be fixed.

REPLY:
As discussed in Grenoble, MACI does not prevent to make servers in languages
other than C++.
From the MACI point of view, servers just need to implement the Activator IDL
interface.
The major problem, inrelated to MACI, is that properties are currently
implemented in C++, so a non C++ server has to find a way to use the C++
implementation or to provide an implementation for the IDL interface of DOs,
properties and characteristics.
We have discussed in Grenoble what can be done for ACS 1.0 and Python and we
will do it.

GCH 2001-09-11 Added section in deployment about non-C++ servants

---061---
(Jim Pisano)
pg 15, "archiveManager" on pg. 31, "logManager", pg. 29) - are these MACI
Managers?

REPLY:
No.

GCH 2001-09-07 Clarified here. Nothing to be done in the document.

---062---
(Mick Brooks)
p16, S2.2.4: First two sentences contain too many derivatives of the word
"heirarchy". Suggest: "Distributed Objects will be
organized hierarchically. There are essentially two ways to build such a
structure:"

ACCEPTED

GCH 2001-09-11 Done.

---063---
(Mick Brooks)
p16, S2.2.4: I do not understand bullet point two, needs better explanation.

ACCEPTED
Will be expanded and an example given.
This was also discussed in Grenoble and it is hopefully clear now

GCH 2001-09-12 Done.

---064---
(Steve Scott)
p.16 2.2.4, first bullet:"A naming scheme..." is this the Name Service? If so, say so. If not, why not?

REPLY:
The OMG Naming Service specification does not define a representation of
names as string.
CORBA applications usually define a naming convention to map a simple string
Into a Naming Service "Name" (i.e. a sequence<NamedComponent>)
This is what we are going to do.
The actual naming convention is a design detail described in the MACI design
document.

GCH 2001-09-12 Done.

---065---
(Jim Pisano)
p.16. 2.2.4 Grammar - 2-nd bullet "the actual servants that instantiated ..."
remove "that".

ACCEPTED

GCH 2001-09-11 Done.

---066---
(Jim Pisano)
p.16 2.2.5 "... edit all values of properties & characteristics." The changing
of characteristics on the fly doesn't seem like a good idea for "static"
characteristics. I understood that there would be tools to change the configuration database and reload them (see item 2.2.6.3). Having only one way
to change characteristics simplifies the code.

(Mick Brooks)

p16, S2.2.5: If "all objects in the system can be reached by navigating the
hierarchy", this would seem to mean that the option in
bulletpoint 2 in S2.2.4 cannot be used.

REPLY:
The Object Explorer is an engineering user interface and not a programmatic
API.
It will be used mainly for debugging/tuning purposes.

GCH 2001-09-12 Done.

---067---
(Brian Glendenning)
p.16 s2.2.6 I don't understand what you mean by this. It sounds like you are
saying that values in the configuration database cannot be created
programatically (i.e., the database API does not have any write functions),
which would seem unnecessary.

ACCEPTED:
The phrase is not clear. Remove "independently from..." and keep the first
part as an introduction to the sub-bullets that follow.

GCH 2001-09-07 Done.

---068---
(Brian Glendenning)
p.17 s2.2.6.1 Some objects will have to be initialized in a certain order
(e.g., bottom layers before containers). How is this implemented?

REPLY:
This was discussed in Grenoble and we will add a description.
In quick summary:
- Whenever a client needs a CORBA reference for a DO, a request to MACI Manager is done
- If the object is not already instantiated, the Manager asks the Activator
to create it
- When an object contains references to contained objects, the dependency is
expressed via CORBA references and resolved through requests to the MACI
Manager.
This warranties that all objects are automatically created in the right order
and when needed.
Some objects are needed immediately at bootstrap. They are directly specified
in a MACI table and MACI instantiates them as soon as it is boostrapped.
If there is a root top-level object, just putting this object in the MACI
table will trigger a cascade instantiation of all dependent objects.

GCH 2001-09-12 Done.

---069---
(Mick Brooks)
p17, S2.2.6.1: So DO life-cycle is determined by system startup and shutdown and not by the actions of clients?

REPLY:
No. Some objects are created at bootstrap, others when needed.
See previous point, discussion in Grenoble and MACI design document

GCH 2001-09-11 Done. See previous point

---070---
(Steve Scott)
p.17 2.2.6.1: Is there any notion of sequentially setting properties? E.G. need to set the synthesizer frequency, wait for it to come alive, and then start the oscillator locking state machine.

REPLY:
If we understand what you mean, you have the DO handling the synthesizer
frequency that have to be alive before the DO that handles the oscillator
locking state machine
This is handled by the MACI Manager:
- Both DOs are in the MACI table.
- To be able to access the synthesizer frequency, the state machine needs a
  reference to the corresponding DO.
- It then resolves the reference with a request to the manager
- The manager will return it if it exists already or will ask the activator
  to instantiate it.
As a consequence, the synthesizer frequency is always already running when it
is accessed by the state machine.

GCH 2001-09-12 Done.

---071---
(Brian Glendenning)
p.17 s2.2.6.2 This seems unnecessarily complicated (e.g., having a local
copy that might be in a completely different file format). I think having a
central database (which could of course be internally federated/redundant)
is fine. Of course at startup time it should be possible to nominate a different startup database, and even change it at run time.

(Jim Pisano)
p.17 2.2.6.2 ".. This Local Configuration Database ... to make the system more
fault-tolerant." I don't see how an ASCII file which can be touched by humans

with emacs makes the system more fault-tolerant. Having an editable file just introduces opportunities for errors and out-of-synch problems which reduces fault-tolerance. There should be a mechanism which allows for the generation of temporary changes that does not require editing a text file. There are a plethora of DB's and DB editors, let's stick to them instead of text files & vi.

An approach that satisfies the TICS' requirements of warm boot and cold boot follows:

The system configuration comes from the central config. DB.

An application listens to the monitor stream and creates a snapshot DB that contains current monitor values.

When a computer reboots, the default DB and the snapshot DB are merged according to if it is a cold, warm, or hot boot. To configure the degree to which defaults are overridden, another DB can keep a list of values which are overridden from the snapshot.

To allow for user modified values, a third DB with user preferences is kept.

So, the scenario to boot a configuration is to merge the snapshot with the user preferences, and then this is merged with the default DB according to the override DB. This DB is then downloaded to the computer.

It might be that a rule is the user preferences are always included, so this would make the scenario - merge defaults and snapshot according to overrides, and then always merge user.

How would we implement this as local copies? The problem is that the snapshot would have to be made and kept locally to the computer, but the ABMs do not have file systems. But, if the requirement later is we need this for performance, then we would need the hardware and need to move a copy of the snapshot application to the local computer. The default DB and the override DB and the user preferences would be copies from the center to a local file system.

Thus there is no local DB to the LCUs. The issue of performance is an analysis one which can be evaluated with TICS.

(Ron Heald)
p.17 2.2.6.2 There's no requirement or need for a local configuration database. How does it fulfill the stated purpose of "to avoid performance problems at startup or during operation and to make the

system more fault tolerant"? It's a concept left from the VLT real-time
database and only adds complexity. Any time the same value appears in
than one database there's a big problem keeping those values
synchronized.

(Steve Scott)
p.17 2.2.6.2: The Local Configuration Database removes a lot of
the location independence of distributed objects and adds
complexity. It seems very unlikely that performance really demands
this enormous complication. Relocation of DOs would require
reconstructing the LCDB.

(Dirk Muders)
p17, s2.2.6.3:

My 0.02 Euro worth on the local vs. global config DB issue: A local
copy may
facilitate the testing but synchronization will always be an issue.
Especially
when it comes to deciding who's got the reference. Example: local
testing
shows that a parameter should be changed globally. How does one
transmit this
to the global DB AND other affected local DBs (e.g. at other antennas)
? It
would mean that the global DB masters the local one, thus inverting the
idea
of the client downloading the global DB once at start time.

TO BE DICUSSED AT REVIEW
It really seems that everybody agrees in disagreeing with us :-(
We are then probably really biased by the VLT architecture on this.

As discussed in Grenoble, the main justification for having a local
database
was to avoid that LCUs are dependent on the database on a central
workstation
and are querying this central place for data.

If we introduce a mechanism to pass to DOs a reference to the database
they
have to use, we remove the problem from the architecture and we leave
it as
a detail design. We can then easily switch from local/remote/federated
config
DB according to easiness of implementation and performance
requirements.

In this way it is also easy to switch between different databases at
startup
time.

We see 2 possibilities for implementing this:
- To pass the reference to the DB explicitly as a constructor
  parameter to DOs.
- To have the DB configured at Activator level and have the DOS
  to implicitly retrieve the reference to it from the Activator
  itself.

Solution 2 seems to provide simpler interfaces and less code.

Anyway, all DOs under control of the same Activator should anyway use the same
database.

GCH 2001-09-07 Done.

---072---
(Martin Pokorny)
p.17 s2.2.6.2 "Each Local Configuration Database is periodically and automaticallly built from the central database, kept under configuration control". I'm assuming that "automatically" means that on every occasion that the central database is changed, all local databases are also changed. You might want to emphasize this (or clarify it, in case I have the wrong interpretation).

REPLY
The idea was to have the local databases changed on restart of the local
database only and not at every change, but see previous comment.

GCH 2001-09-07 Local configuration database removed from architecture

---073---
(Brian Glendenning)
p.17 s2.2.6.4 If characteristics are static this means that they cannot be
used for, e.g., calibrations that change at run time.

REPLY:
This is correct. Calibration values that change at run time should be implemented as properties, not as characteristics.
"static" calibration values, for example measured by an engineer that are
not supposed to change for months can be characteristics.

GCH 2001-09-07 Clarified in the document.

---074---
(Jim Pisano)
p. 18 2.2.6.4 "Characteristics can change ... with context". What context?
characteristics are very specific to the DO or property. As defined in 2.2.1.2.3, " Characteristic - Static data associated with a Distributed Object
or with a Property, including meta-data such as name, description, version and
dimensions, and other data such as units, range or resolution." I'm not sure
that characteristics would change much w/ time either, but, granted, they
might. I would remove the "with context" at least.

ACCEPTED
We remove "with context"

GCH 2001-09-07 Section rewritten and (hopefully) clarified.

---075----------
(Joe Schwartz)
p. 18, 2.2.7, Maybe I just don't understand what is meant by "standard state machine", but objects in the High-Level Software will have states that will be defined by the nature of the application, e.g., a Scheduling Block might have "checked out", "awaiting verification", "ready for scheduling", "waiting on breakpoint", ... states. How will this fit into a standard state machine?

REPLY:
See TICS design.
There, a standard state machine for devices is described.
To have a common set of states and transition commands is important to be able
to build/change the global state of the system.

GCH 2001-09-07 Nothing specific in the document. Clarified here.

---076----------
(Brian Glendenning)
p.18 s2.2.8 The importance of this feature is not obvious.

(Ron Heald)
p.18 2.2.8 There's no need for the concept of "groups". When you study
this idea in detail you find its added value is small, and is outweighed
by the complexity it adds. Whenever this issue arises the example of 64
identical antennas is given. In the design of the mount software, I originally had the idea of using "group" commands to control all antennas in a "sub-array" simultaneously. I was quickly told that moving all antennas in different directions while performing the same offset measurement is more efficient, and therefore each individual antenna requires unique position commands. I think similar reasons will
be found in other control areas, and that this feature will not be used. This fact has already been realized in the configuration database
design as it provides unique entries for each device.

ACCEPTED:
As discussed in Grenoble, we propose to remove groups.
We will file a change request for the requirements.

GCH 2001-09-07 Done. Groups removed and SPR ALMASW2001079 submitted.

---077----------
(Jim Pisano)
p.18 2.2.9 Since DO's are built from the config db, this point of serialization
is unnecessary. I realize that "it's a requirement", but I understood this
requirement as being able to construct objects from a persistent database which
the config db does. Maybe we should drop this requirement?

REPLY:

The configuration database in itself contains only the configuration, not
necessarily the current state. For example, the current value of logical
properties is in principle not in the configuration database (currently it
in the ACS 0.0 config. db because we use the VLT database and we have written
the code of properties to do it).

GCH 2001-09-07 Nothing done in the document. Clarified here.

---078----------
(Jim Pisano)
p.18 2.2.10 If each DO has to be serializable, then this is overkill for DOs
associated only with the control software. This bloats the code of each DO -
use inheritance to add this capability for serializable DOs when needed.

ACCEPTED
This is what we mean. Will be specified.

GCH 2001-09-07 Clarified in text.

---079----------
(Jim Pisano)
p.18 2.2.11 Grammar - "wraps" -> "wrap"

ACCEPTED

GCH 2001-09-07 Done.

---080----------
(Martin Pokorny)
p.18 s2.2.11 Is the usage of JavaBeans part of the ACS architecture?
It seems that applications, not ACS, will determine the necessity of
JavaBeans. I suppose that the I'm still confused regarding the line
separating applications and ACS.

REPLY:
Java and UIF libraries are part of the ACS requirements.
As described with more details in 2.14, Java is the language selected for
user interfaces and ACS provides Java tools to support the development.
Part of this support is ABeans and a code generator that from the IDL of a
DO generates automatically the corresponding JavaBean.
In this way the developer has libraries that provide him direct support for
ACS concepts like DO/Property/Characteristic, Monitors, Data Channel...

GCH 2001-09-07 Clarified in text.

---081----------
(Ron Heald)

p18, S2.2.12: The ORB independence discussed is only on the client side. On the server side we're very much tied to TAO and ACE. This should be made clear.

REPLY:
This is not correct.
Also on the server side there is very little CORBA code that depends on orb
specific calls. If you provide another C++ ORB, porting should be easy.
What is true is that we are dependent on ACE, but you can use ACE with any
other C++ ORB (TAO uses ACE, but they are two separate packages and ACE does
not need TAO).

The major limitation is the usage of ACE exceptions, but this has nothing to
do with CORBA.
As discussed in Grenoble, moving from ACE to native exceptions is easy, once
we get a C++ compiler on VxWorks that supports them.
We are investigating, but we are not optimistic.

GCH 2001-09-07 Done. Added Section on servants C++ and not
C++ to deployment

---082---

(Jim Pisano)
p.15 2.2.3.2 "Specific subclasses will provide polymorphic access to specific
implementations like 'logical', 'simulated', 'CAN', 'RS232', 'Digital IO' and
so on.". Further on in section 2.2.13, you state: "The granularity of simulation mode is at Distributed Object level and not a Property level ...". On one hand you are advocating for simulated Properties, but then you
say that you will not have simulated properties. IMO, simulated properties is a
valuable tool and should be allowed, but if you're definitely not going to
allow them, then delete "simulated" from 2.2.3.2.

Ron Heald)
p18, S2.2.13: As discussed at the Joint Software Meeting, simulation can
be implemented with a simulation mode of the DO, and by replacing a DO with its simulation equivalent. Also, simulation can be done at the DO
and Property levels.

(Brian Glendenning)
p.19 s2.2.13 I prefer that it be possible to load a different object than a
single object have a simulation mode which can be thrown. This means that
simulation cannot be turned on without breaking object references, which
does not seem like a problem to me.

(Jim Pisano)

p.18 2.2.13 I don't see that each DO should have an on/off simulation switch. I
think that there should be a simulated version for a given DO so that it can be
switched in as needed via the CORBA naming service. This has 3 advantages: 1)
it makes the DO code smaller and consequently simpler by avoiding this
dual-mode mechanism of simulated/not simulated; 2) it allows for development of
the simulated device without the hardware and 3) the simulated device can
reside anywhere, on the LCU, on a workstation, or a desktop PC which is
accessible via CORBA. The disadvantage is that there are more DO's a "real" DO
which connects to hardware and a simulated DO.

(Jim Pisano)
p.18 2.2.13 Last sentence doesn't make sense. "The granularity of simulation
mode is at Distributed Object level and not at Property level to allow the
implementation in simulation mode of interrelations between properties inside
the same Distributed Object (for example, setting a value for the set-point of
a thermostat will affect after some time the actual temperature)."

One could easily envision a Property being simulated which would really test
the DO's decision logic for one or more properties narrowing the focus of
simulation (and consequently debugging). This is preferable to replacing the
entire DO's mode which could introduce all sorts of different interactions and
logic between the simulated mode and real modes of the DO. As I understand,
this capability do simulate individual Properties has been rejected in previous
discussions, but I think that it should be allowed as it seems the best way of
all 3 possible simulations scenarios, i.e., 1) simulated Property, 2) simulated
DO, 3) simulated mode switch.

(Martin Pokorny)
p.18 s2.2.13 I prefer having the granularity of simulations be at the
level of Property, in addition to Device. (In fact, in s2.2.3.2, you
mentioned simulation mode for Properties.) In diagnosing problems, it
is often useful to change only a single "variable" at a time. On the
other hand, maintaining a consistent simulation may sometimes require
device-level simulation. In other words, both levels of simulation may
be useful.

(Mick Brooks)
p18, S2.2.13: The alternative of instantiating explicit simulation DOs has
not been discussed. Perhaps it should be.

TO BE CONFIRMED AT REVIEW
The section on simulation will be re-written according to what decided

```
at the Grenoble meeting:
   - ACS provides support for simulation at property level.
     All properties that access hardware can be switched in simulation
     by setting TRUE a simulation characteristic in the configuration
     database (this is a proposal for the way to implement simulation
at
     property level and was not explicitly discussed in Genoble).
     After this, they behave like "logical properties".
     This provides basic simulation capabilities
  -  If an application wants to provide a more sophisticated level
     of simulation (for example simulating interrelations between
     the values of properties), a specific simulated device should
     be implemented in parallel to the real device.
     Switching from the real to the simulated device is handled in
     the configuration of the MACI manager, telling it to start a
different
     implementation of the same device's CORBA interface.

GCH 2001-09-07 Revwritten according to this proposal, that was accepted
at review.
```

2.3.2 Comments to Issue:1.1/Prep.2

```
---015------------------
(BrianGlendenning)
p.11 2.2 point 1.3
> I would think that quality should only be part of properties related
to M/C
> points, not necessarily those involving, e.g., data flow.

(JimPisano)
p.11 2.2.1.2.3 Property Quality
> Although this concept of data quality for a read value is listed in
the
> requirements, I don't see how useful this is. Either a value is
retrieved
> (OKAY) or not (ERROR). SUSPICIOUS can only be implemented w/ by
keeping a
> history of retrieved values, something that seems useful for the
indirectly
> retrieval mechanism as these values are retained in a DB.
>

(RonHeald)
p.11 2.2.1.2.3
> Important comment: "quality" has no use in direct value retrieval;
> either you get the value or an error.
>

TO BE DISCUSSED AT REVIEW
The concept of quality is in the requirements.
To scrap quality would anyway simplify the system and we could probably
leave without:
just always return error whenever the quality would have been not OK.

We should really discuss if we should remove the concept at all.

In any case, if we decide to keep the concept, quality is also
important
to pure soft properties, e.g. server down, LAN connection broken etc.
```

As Ron notices, it is normally with indirect readout that you might have old
obsolete values, without having an error, because the readout simply was stopped.

However there is a requirement for transparency (4.1.4) and applications
shall be the same for getting direct or indirect values. So structures must also be the same.

To have properties with/without quality adds complexity to the system. If quality does not make any since for some specific property, then it is always OK.

GCH 2001-03-26: For the time being, quality is removed for the Architecture.
   If we will decide that it is important is can be added later.
   The notes on Architecture document is used to keep track of this.
   The requirement will have to be removed as well.
   If accepted at the review, and SPR change request for the requirements will be filed.

---016------------------
(MickBrooks)
>p.11 Section 2.2 This section is pretty clear, but I'd like a little bit
>more explanation of the URI. Is it a pointer to documentation only?

(BrianGlendenning)
p.12 first point 3.1
> URI -> URL? (In the figure as well).

(BrianGlendenning)
p.12 first point 3.1
> Do we really need to have megabytes of description information in the running
> system? Why not just hold an index to it.

(RonHeald)
p.12 2.2.1.3.1
> There's no requirement for a URI. It would be bad to separate the
> documentation from the Property as it would quickly get out of sync.

(RonHeald)
p.12 drawing
> Again there's no requirement for URI()
>

REPLY
Whenever you see URI you can just read URL, although URL is actually a special case of URI
(see glossary and W3C referenced page).

The idea is to keep in the run time system only a short description of properties in the
Description characteristic and, optionally, a pointer to more extensive documentation
in XML/HTML format in the URI characteristic. There are no MB here.

The documentation is written together with the code in the source files,

then extracted to produce cross referenced and nice to read XML/HTML pages.
The model is Javadoc (doc++ for C/C++ and hidl for IDL)
The URI will contain just the reference to this generated documentation, that
can be on any web server and does not need to be kept online on the running
system (as Brian, understandably, fears).
There is no point in asking users of the system to browse with an editor in the
source code.

GCH 2001-03-27:
Brief explanation added. Extended clarification not added in order not to make heavier the text.

---017------------------
(RonHeald)
p.11 2.1
> There is no requirement for "API for run-time command syntax checking".
> What is meant by "tools for dynamical command invocation". "forth"
> should be "fourth"
>
REPLY
Yes, there is a requirement or even two for syntax checking: 6.1.3 and 6.1.5
dynamical -> interactive command invocation, as from requirement (6.1.5)

GCH 2001-03-27: Done

---018------------------
(RonHeald)
p.11 2.2
> Remove "put at" and "whole"
>

OK

BGU 2001-03-26: Done

---019------------------
(RonHeald)
p.11 2.2.1.1
> "Distributed Objects are the smallest directly addressable entity in the
> system." I would say this is the definition for Property. Maybe you
> mean directly addressable via the CORBA name service?
>
CLARIFICATION
Properties are only addressable through a DO.
If you do not get the device first, you cannot access properties.
This means that properties ARE addressable, but NOT DIRECTLY from any entity
that is not the owner DO itself.

GCH 2001-03-27: Done

```
---020------------------
(RonHeald)
p.11 2.2.1.2
> Should be "[...] that are monitored and controlled, for example,
status,
> position, velocity, [...]".  Remove unmatched ")".
>
ACCEPTED

BGU 2001-03-26: Done

---021------------------
(RonHeald)
p.11 2.2.1.2.1
> Remove write-only Property.  It was recently decided that if hardware
> does not support the get, then software would.
>
OK

BGU 2001-03-26: Done

---022------------------
(RonHeald)
p.11 2.2.1.2.2
> Remove Property types strings, enum, binary, and sequence.  Strings
> negate compile time checking.  The others are not required.
>
TO BE DISCUSSED AT REVIEW
We think that strings and sequences are important and we would not
remove them.
Enum are probably also very useful.
Binary are just suggested and experience will tell if they are
necessary or not.
They are not for first implementation but it would be wrong to
remove the TBD reference from the document.

Enum was to address the problem of bit twiddling magic constants, so a
dictionary
associating names to bit patterns was envisioned.
For example, a motor might have 0x1 = "on", 0x2="off", 0x4="forward",
0x8="reverse", 0x10="slow", 0x20="fast".  To set a motor to on,
forward, fast is
"on forward fast".
It might be that strings are really enums, so that might go.
Sequences seem that they would be useful sometime in the future.
For sake of implementation as fast as possible, sequences
and binary have low priority, strings are possibly enums,
and enums seem useful to have.


GCH 2001-03-27: Added Array. Moved in notes binary and sequence and
added structure as possible future property types

---025------------------
(BrianGlendenning)
p.12 first point 3.2
> Again, things like range and resolution have to do with M/C points,
not
> properties in general (IMO).
```

REJECTED: we do not agree.
Most quantities in the system will have range and resolution and other characteristics.
To identify a common set of characteristics is essential to keep the interfaces simple.
There are case where a characteristic is not applicable, but handle a completely free
format would make the system much more complex.

GCH 2001-03-27: Nothing changed in document

---026------------------
(BrianGlendenning)
p.12 figure
> The diagram shows that access to characteristics is through a configuration
> record. Is that right? I would guess that characteristics should come directly
> off the ACS Object. Is a characteristic (essentially) just a typed name=value
> pair? Do we really need to keep the configuration record after the ACS Object
> has been constructed?

(BrianGlendenning)
p.13 point 3.4
> This seems to indicate that there is no access to characteristics other than
> those which have been "nominated" by being made explicit member
> functions. That is, there's no access to "unknown" characteristics via a
> name=value type of interface at run time. True?

>From the interface point of view, characteristics come for sure directly off the
ACS object. No user of a DO should see the ConfigurationRecord.
The diagrams wants to show that characteristic are retrieved BY THE DO from a record
on a database, through a ConfigurationRecord class that hides
all the details of the access to the database.
The final database will be most probably XML, on ACS 0.0 it is the VLT database,
for ANKA it is their configuration database but for the DOs
this does not make any difference, since everything is hidden inside the implementation
of the ConfigurationRecord class.

In the present design (and in ANKA) only "nominated" characteristics are accessible
via explicit member functions, to keep the system simple.
On the other hand, the extension you ask for is very easy:
 - The ConfigurationRecord is now a private member of a DO
 - It already provides name=value interface at run time, to access the config database
 - Would be just enough to make it public, et voila!
If experience will show that this is necessary, we have already the solution at hand.

What follows is a general comment about configuration database that applies to

a lot of comments following.

The intention is that the configuration database is the system documentation
in one place with easy access by user interfaces.
It is the users' interface to the system description.

The device hierarchy at the antenna (class hierarchy) is easily represented by XML.
XML gives an OO view of the system.
It is a monolithic block of data, but there is structure to the data, and there is well-formedness because the DTD, document definition, enforces the organization.

There are opinions that documentation goes with the code.
But, if the documentation is useful to some potential user in the future,
then searching through the code does not make the information very easy to access.
The documentation should go as much as possible with the source code, to limit
alignment problems, but should then be extracted from there and made available
in more convenient ways (see the model on Javadoc).
This is another argument for having a central location for documentation.

Finally, do we want constant values in code, a part from "universal constants"
like PI or c?
Constant values embedded in code works, and it is rarely changed.
But, the constant values are interesting documentation, too.
The cost for implementing loading constants from a database is already paid for
because we have built the mechanisms for values that do change.
This choice can go either way, but the information is hidden and involves messing
with the code which seems to be a poor choice and sways the argument to a configuration database.

Okay, so we have a really big pile of data describing the system.
It can be accessed by the booting computers in several ways.
The proposal is to pull out just the needed characteristics and
put them in an XML database that is embedded into the running software.
It is just a text file that is parsed by a local server that knows how to
locate the information by the ACS object's name.

What kind of database should hold the central configuration file?
Is it XML or RDBMS? The database size will be
64 antennas * 2k objects/antenna = 128 k objects + all what is on top of the antennas
This is not an unreasonable size for a RDBMS to handle.
It's not clear if XML can not handle this, too.
We need to make up our experience in this and we cannot say now, without prototypes and
performance tests if we can do with just XML or not.

GCH 2001-03-27: Added interface to access characteristics name=value

---027------------------
(BrianGlendenning)
p.12 figure
> I would like the figure to indicate that M/C point, Device, and
everything
> under device are only illustrative.

ACCEPTED
It was meant so, but colors do not come out well on BW!
We will make it clear in a better way.

GCH 2001-04-06: Done

---028------------------
(JimPisano)
p.12 2.2.1.2.3 Characteristics
> URI should be optional. For most simple monitor points, the
description
> should suffice for documentation.
>
ACCEPTED
This is how it was meant.
The characteristic is always there for symmetry, but does not need to
contain
a URI. It is then always possible at any time to fill in the value of
the
characteristic in the configuration database with a pointer to newly
written
documentation.

BGU 2001-03-26: added that URI is optional

---029------------------
(JimPisano)
p.12 2.2.2... "this 3 tier logical model..." Are you referring to the
CORBA
> middle-ware logical model, or the package diagram in 2.1 or what?
>

(RonHeald)
p.12 2.2.2
> What's a "3 tier logical model"? (I know 3 tier distributed systems,
> i.e. client, business logic, and database.) Remove "because the ACS
does
> not address directly the Control System, as in the case of the
mentioned
> references", it's unclear and misleading.
>
Three tier refers to DO, property and characteristic model (3).

ACCEPTED
3 tier is: 1) DO, 2) Property and 3) Characteristic.
You all had problems with this, so the text must be made more clear.

BGU 2001-03-26: done
GCH 2001-03-26: tried to make more clear 2.2.2

---030------------------
(MickBrooks)
p.12 point 2. I assume this is the Heald Compromise?

It has been decided in March with Brian and Gianni to use Distributed Object
as a generic term that is applicable also outside the control system, while Device
would be a subclass used inside the control system.

GCH 2001-03-28: Nothing to be done

---031------------------
(RonHeald)
p.12 2.2.1.3
> A characteristic is not associated with a DO, only with a property.

REJECTED
No, DOs have characteristics.
have a name and there will be information in the configuration database such as revision, description, etc. for all ACS objects.

> There's no requirement for "resolution".

(RonHeald)
p.12 2.2.1.3.2
> Again there's no requirement for resolution.

CLARIFICATION
Right, but we get it from ANKA, where it proves very useful.
A number of applications are: automatic definition of scales in plots, number of digits to be used in the logs and monitor stream, automatic definition of thresholds for alarms and events.
Don't we need resolution at least for M&C points?

GCH 2001-03-28: Clarified

---033------------------
(RonHeald)
p.12 drawing
> Handling monitoring in the Property is TBD
>
CLARIFICATION
The property has just methods to create a monitor. The monitoring itself
will probably be in a separate monitor process, but this is implementation,
not architecture.
The interface is defined, a first implementation comes with ACS 0.0 (ANKA),
but will need to be optimized in the final system to scale up to our needs.
This will not affect interfaces, but only performances.

GCH 2001-03-28: This clarification should be enough, not to make heavier the text

---034------------------
(RonHeald)
p.12 drawing

> validate() should be added to RWProperty to check an input new value.

CLARIFICATION
validate() exists but as a private method.
The setValue() method itself calls validate, so that there is no need to explicitly call validate first. The setValue() would just fails.

Validate() uses the characteristics to perform its job.

Since there is a central description for everything in the system, then the user interfaces can validate values with the description. This is discussed in the sections about command checking.

GCH 2001-04-06: OK

---035------------------
(RonHeald)
p.12 drawing
> There's no requirement for increment() and decrement(). They have
> little use in RA.
>
ACCEPTED
Actually we get them from ANKA, where they are used in user interface to increment/decrement stepwise properties.
A typical example is the "double arrow" control or a slider:
you want to have two buttons to rise/lower a value at each press;
one uses increment() the other decrement().

They are helpful in tine tuning devices, where you want to increase/decrease the value
step by step of the minimal allowed unit.

They will be anyway dropped, is you do not consider them useful, in the final system.

GCH 2001-04-06: Left for the time being, since used by Java Abeans

---036------------------
(RonHeald)
p.12 drawing
> defaultValue() in ROProperty should be moved to RWProperty as it's only
> useful there.
>
REJECTED
The default value is the initial value upon restart/init and is used in simulation mode. If there isn't any hardware attached this is the value to return.

GCH 2001-03-28: Nothing changed in the text, to keep it lighter

---037------------------
(RonHeald)
p.12 2.2.3
> Section 2.2.4.4 (p.13) needs to be together with this.
>
ACCEPTED

The two sections are clearly strictly related.
We will rewrite the text grouping i a more logical way concepts.

GCH 2001-03-30: merged

---039------------------
(BrianGlendenning)
p.13 point 4.1
> I would say the global namespace of the running system, not just the
control
> namespace.

ACCEPTED: It was meant so

BGU 2001-03-26: Done

---040------------------
(BrianGlendenning)
p.13 point 4.2 - 4.5
> This is one of my major comments: I consider these points to be
controversial
> and would rather say that we will prototype various solutions to get
real
> experience. If you insist on a decision now, these points require
> considerably more justification (IMO).

ACCEPTED (except for Properties in the name service)
We have very clear in mind the fact the you consider these points
controversial.
Actually, the current formulation comes after many discussions and does
not push
for any decision: it clearly shows that you have the means to do what
you want,
and it just proposes guidelines.
Prototypes and experience will show what is better.

Only one point is firm: properties DO NOT go to the name server. This
is a corner stone
of the architecture

GCH 2001-04-02: text rewritten. Point open for review after TICS

---041------------------
(JimPisano)
p.13 2.2.4.2 The organization of DO's needs clarification. What do you
> mean by the physical distribution of the DO's on the server? Aren't
they
> physically distributed at the devices for the most part? References
to
> them are logically distributed at the servers. I'm confused.
>
CLARIFICATION
There will be DOs as Devices in ABM but also in the ACC etc. They are
where they have been designed to be, depending on applications.
ACS has to cope with them.
We will try to improve the text.

GCH 2001-03-30: text clarified

```
---042------------------
(RonHeald)
p.13 2.2.4.2
> Since DOs are software how can they be in a "physical" organization?
> Period after "shallow", then "Servers will probably [...]".
>
OK, physical will be removed.

BGU 2001-03-26: done

---043------------------
(RonHeald)
p.13 2.2.4.6
> It's not clear that the Object Explorer will have write capability.
Add
> "by" after "reached".  Remove "status".
>
ACCEPTED
It will be made explicit. It shall have them according to 5.1.3
(Browser)


GCH 2001-03-30: made explicit

---044------------------
(BrianGlendenning)
p.14 point 5
> I would say "for all Distributed Objects that require them". It's not
clear
> that all will.
ACCEPTED
All DOs have anyway a minimal set of characteristics, like the DO name
itself.

GCH 2001-03-30: clarified

---045------------------
(BrianGlendenning)
p.14 point 5.3
> The scheme here where the configuration DB will sometimes be on an
RDBMS and
> sometimes using XML seems unnecessarily complex to me. I'd say that
we should
> choose one, and if actual experience dictates we should have two
mechanisms so
> be it. XML seems like the better fit here. In general, the role (or
lack
> thereof) of XML is not very definite in this document (this is
another of my
> major comments).

ACCEPTED:
This whole section is clearly TBD based on experience and prototypes.
It does not try to make any decision, but just lists the reasonable
possibilities
that will have to be investigated.

Just ONE is clearly simpler and to be preferred.
```

XML looks very promising in many areas, but we have no experience yet.

For this reason we just list areas where we think it should have an important role, but it is all to be verified.

GCH 2001-04-02: RDBMS is used for the Configuration Database. XML is just used as the
transfer protocol.

---046------------------
(JimPisano)
p.14 2.2.5.1 I find the need to store so much info in one monolithic
> database, namely the Configuration Database, is contrary to OO designs.
> Information to configure a system at startup should be placed into this
> DB, but part numbers and other extraneous info should be placed in separate
> DB's. These other DB's can be referred to by the applications using the
> configuration DB. Maybe even sub-divide the configuration info into smaller
> subsets so that DO's get info from a DO configuration DB while the
> RO Properties get their info from a RO Property configuration DB and
> similarly w/ the RW Properties.
>
The usage of the configuration database does not violate OO design principles,
but is part of the implementation.
Access from clients to characteristics is always and only through the object.
How the object gets them, is a private problem of the object itself: one
implementation is the configuration database, another could be data members
with coded values. This is completely transparent to clients and warranties
that the object model is consistent.

The Config. DB is in the requirements, for configuration control reasons,
nothing to do with OO design.
Part numbers will in fact not be in the Configuration DB; this was a wrong
example.
The local copies of the DB (TBD if XML files) subdivide the monolithic DB
at run time, but are "generated" from that one.

> I'm also concerned about the size of the configuration DB w/ ~100K monitor
> points plus a "central configuration DB" & a "local configuration DB"
> that there is too much complexity and room for errors with these large
> quantities. This adds a large layer of configuration management which
> may be unnecessary or problematic.

The local DB might be a XML file downloaded at start-time. It might be complex

to have configuration management for ALMA, but it would be impossible if
all information would be scattered here and there.

GCH 2001-04-06: Clarified. Central Config DB will be RBDMS.

---047------------------
(RobertLucas)
p 14: 5.1 "every piece of information related to ALMA": this is frightening. What is the need for both XML and RDBMS ?

CLARIFICATION
The actual role of XML and RDBMS is not defined.
Hopefully we can do with just one.
XML is very interesting, but we need prototypes and experience to find out
how well it scales up. If it does not scale well enough, then the RDBMS can take care of "mass" data.
The flow XML<->RDBMS would need anyway to be seamless.

The NOTE can be misinterpreted and should be clarified.
We want to point out that there should be in the system some central database
(we could call it "repository DB") the really contains "every piece of information related to ALMA".
In many system, there are many independent and often not consistent and not
compatible databases that contain parts of description of the system.

This is bad and should be avoided, providing one central database (necessarily
and RDBMS) that can be consistently used for the different purposes.

In this context, the central configuration database described here would be
a subset (a set of tables) in the repository DB and from this it would be
possible to extract the parts used by the DOs at startup/init, most probably
in the form of XML files.

GCH 2001-03-30: text changed and parts moved in notes.

---048------------------
(RonHeald)
p.14 2.2.5
> Important comment: The configuration database should NOT have every
> member of every DO and Property, and certainly not "every piece of info
> related to ALMA". It should only have the values needed for the DO
> constructor, and then only for top-level DOs. Most attribute and
> characteristic initial values should be static values in the code, not
> in the DB.
>
REJECTED
Agreed that the Config. DB should not be confused with a repository DB (see
also previous comment)
The Config. DB however shall contain all the characteristics and

properties, which are needed for the DO constructors
for all the objects (indep. if they are or
not in the naming service).
It is wrong to split between values in code and DB, because:
- This will simply introduce two locations for values, which sometimes still
  have to be changed. Everything will end up in the Objects at run-time.
- It is against the purpose of a Configuration database, which is to keep a
  complete static image of the parameters of DO's.
- hopefully we will not want to have 64 copies of the Mount code, because
  there are different pointing model parameters. It will be better to
  down-load this values at start-up time from a configuration DB.

GCH 2001-03-30: no changes to text

> Support for temporarily modifying values comes from settable properties
> and attributes. There is no requirement (or need) for a "local" DB.

TO BE DISCUSSED AT REVIEW
See also previous discussion on attributes.
Requirement 4.2.1 says only that the Configuration DB shall be downloaded at start-time. The name local DB will be removed, as it seems
misleading.
The way to temporarily modify values is via the DO browser. You seem to object to the separation between data and code. This is intrinsic in the
concept of Configuration DB but, as discussed already, does not violate OO concepts.

GCH 2001-03-30: left open for the review

> I know that ANKA has built their system this way. It makes more sense
> in their case because they have very few devices with many copies of
> each one, while we have almost all singleton devices.

REJECTED
The VLT is also built this way and has only 4 vs. 64 antennas, which ALMA
has. They will be "almost" identical, i.e. they will have similar but not identical configuration for characteristics, that cannot be put in the code.

GCH 2001-03-30: no modifications to the text.

---049------------------
(RonHeald)
p.14 2.2.5
> There's no requirement to update DOs from a local database, only that a
> DO *may* be serializable.

CLARIFICATION
No, there is a requirement to download information at start time (4.2.1).
To re-use serialized information is a possibility on top.

Reference will be changed.

BGU 2001-03-26: done

---050------------------
(RonHeald)
p.14 2.2.5
> An RDBMS or XML, not both. Again, a DB local copy is not required.

This area is TBD, as said above. See also requirement 10.5.5
See other discussions above.

GCH 2001-04-06: Clarified

---051------------------
(RonHeald)
p.14 2.2.5.4
Important comment: Characteristics do not change enough to warrant them being in the DB. On the rare occasion that they need to be changed it can be done by editing the source file. Attributes have set() methods do this. Again, a local copy of the DB is not required.

REJECTED
There are two explicit requirements against this: 4.2.1 (Config. DB) and 14.1.9 (Dynamic configuration), which are contradicting what you say.
Your comment should be passed at the next review of the ACS TRs.
The local copy of the DB is instead a TBD domain.
See also previous discussions on configuration database and "attributes".

GCH 2001-04-06: Clarified

---052------------------
(BrianGlendenning)
p.15 points 6-7
> Delete - not part of ACS architecture, and illustration has been made in
> earlier figure.
>

(RobertLucas)
2.2 5. 7. why more distributed objects than properties ? I thought a single object would have many properties ?

ACCEPTED.
The text is confusing and will be removed just referring to [AD01 - 3.3.1. Size].

BGU 2001-03-26: Done

---053------------------
(BrianGlendenning)
p.15 point 8
> If state machines are to be standardized in ACS (as this implies), that is an
> important point that needs more elaboration - standardized state machines
> would presumably be an important architectural feature, but it is not

> mentioned very prominently here.
>
ACCEPTED
To be elaborated at a later stage.

GCH 2001-03-30: Explicitly said that TICS experience will be used

---054------------------
(BrianGlendenning)
p.15 point 11
> I consider the example to be so speculative that it may not help clarifying
> the ACS point.

CLARIFICATION:
This example was given by P.Grosbol, one of the "inventors" of the Observation
Block concept inside ESO.
It does not look so bad to us, but can be removed.

GCH 2001-04-06: Removed

---055------------------
(BrianGlendenning)
p.15 point 12
> This possibility sounds so dubious that the "wrapper" approach should maybe be
> adopted and list the standard as the possibility.
>
ACCEPTED

BGU 2001-03-26: Done

---056------------------
(BrianGlendenning)
p.15 point 14
> This is a major comment: I would much rather replace a DO with a simulation
> version, rather than put a simulation switch into each DO. And after all,
> there may be many different types of simulation. Take advantage of CORBA!

TO BE DISCUSSED AT REVIEW
Both replacing DO with simulation DO and putting properties in a DO in simulation
could (and should) be available.
They satisfy different purposes.
- Replacing the DO with a simulation instance is useful for external users: they
typically need just a limited simulation, just to test their own applications
against the interface of the DO.
- Putting the DO itself in simulation allows to test the DO itself without
having the need to access the hardware connected to its own properties. This allows
off-line testing of DOs and is an important purpose of simulation.

GCH 2001-03-30: left open for review

```
---057------------------
(JimPisano)
p.15 2.2.5.6 Please define what a high-level coordination object is or
> give an example of one.

CLARIFICATION
Examples could be the Antenna Mount system in the Control domain and
and Observation Block in the data flow domain.


GCH 2001-04-06: Text changed

---058------------------
(RonHeald)
p.15 2.2.9
> Groups are not useful in RA.
>
REJECTED
May be not, but there is a requirement on Groups.


GCH 2001-04-06: No change in text.

---059------------------
(RonHeald)
p.15 2.2.10
> There's no requirement for DO to be persistent, only serializable.
>
ACCEPTED
Right, serialization will be used to create on demand persistent images
of
certain types of objects.

 BGU 2001-03-26: done
```

2.4 Direct Value Retrieval

2.4.1 Comments to Issue:2.0/Prep.1

```
---083------------------------------------------------------------------
----
(Ron Heald)
p19, S2.3.1: Remove mention of asynchronous get() and set() methods.

REJECTED
As we agreed in Grenoble, asynchronous get() and set() have to remain
because
they are used internally.
Both are then available.

GCH 2001-09-07 Nothing to be done.

---084------------------------------------------------------------------
----
(Jim Pisano)
p.19 2.3.2 Grammar - "Value setting is normally done  ...", remove
"normally

ACCEPTED
```

```
GCH 2001-09-07 Done.

---085-------------------------------------------------------------------
----
(Jim Pisano)
p.19 2.3.3 Grammar - "CAN properties will always access directly the
hardware
on the CAN...", change to "CAN properties will always directly access
the
hardware on the CAN...".

ACCEPTED

GCH 2001-09-07 Done.

---086-------------------------------------------------------------------
----
(Martin Pokorny)
p.19 s2.3.3 Add other I/O methods, for example RS232, GPIB.

ACCEPTED.
Will be added as examples

GCH 2001-09-07 Done.
```

2.4.2 Comments to Issue:1.1/Prep.2

NONE

2.5 Value Retrieval by Event

2.5.1 Comments to Issue:2.0/Prep.1

```
---087-------------------------------------------------------------------
----
(Ron Heald)
p.19 2.4 Value Retrieval by Event, p.23 2.6 Sampling, p.30 2.9
Archiving
System  There is substantial overlap in the content of these areas.

REPLY:
These facilities target different and specific needs, but it is true
that the
boundaries overlap.
When on boundaries, users will have to decide what to use.

GCH 2001-09-07 Clarified here. Nothing specific added to the document.

---088-------------------------------------------------------------------
----
(Brian Glendenning)
p.20 s2.4.2 I assume that this does not cause any I/O to happen, but
will be
checked when the I/O does occur (e.g. if it is being monitored). Or is
it
ignored completely for, e.g., CAN properties? This should be merged w
2.4.8.
```

ACCEPTED:
Should be as you say, but it is an implementation detail (with restrictions
Imposed by performance requirements).
Will be merged with 2.4.8.

GCH 2001-09-07 Done.

---089----------
(Jim Pisano)
p.20 2.4.3 I don't see that timed or periodic events are required for direct
value retrieval. The requirement, 4.1.3, refers to Indirect value retrieval. Consequently, this item should be removed.

REPLY:
The difference are described in 2.5.3, .4 and .5
Timed and periodic events provided by value retrieval by event are triggered
directly when the event occur and are delivered directly to the registered
object via the callback mechanism in a point-to-point fashion.
On the contrary, the data channel is caching events and sending them in bunches
at low frequency using the notification service.
The value retrieval by event is then much better suited for providing timely
feedback to control applications, as described also in the TICS design document.

GCH 2001-09-07 Clarified in text.

---090----------
(Ron Heald)
p20, S2.4.3: You need to mention that the interval rate can only be one from a given set.

REPLY:
No. TICS will use just a given set of interval rates, but the mechanism itself can be used at any desired interval.

GCH 2001-09-07 Clarified here. Nothing to be done in the text.

---091----------
(Jim Pisano)
p.20 2.4.6 This is just a duplication of section 2.5 & unnecessary, remove it.

ACCEPTED
It was a forward reference

GCH 2001-09-07 Done.

---092----------
(Brian Glendenning)
p.20 s2.4.7 Delete for now as it is presently a TICS feature.

```
(Jim Pisano)
p.20 2.4.7 As state machines are to be deferred, I would mention that
this will
be deferred like other references to state machines.

REPLY:
The idea is that it is implemented by TICS and integrated later on in
ACS.
As such, we think it should remain in the architecture

GCH 2001-09-07 Clarified here. Left in the Architecture

---093-------------------------------------------------------------
----
(Martin Pokorny)
p.20 s2.4.8 Derived properties of a device based on hardware monitor
points will have similar restrictions. At some level, polling must be
introduced to provide event-on-change support; where does this happen?

(Mick Brooks)
p20, S2.4.2 and 2.4.8: I would include 2.4.8 in 2.4.2 and give some
idea on
how a client knows that a property does
not support events. For example is there a typed difference between
hardware M/C points and others?

REPLY:
Sorry, our text is not clear. We will rewrite it.
CAN-Properties will have to implement notification on change also for
CAN
monitor points although CAN monitor points does not provide a specific
support via HW or in the drivers.
This can/should be done via polling.
If there are clients registered on events on change, an ACS monitor at
the
characteristic change frequency of the monitor point is used to poll
and to
generate events in case of changes in the value.
The polling frequency determines the time resolution of the event-on-
change.

GCH 2001-09-07 Clarified.
```

2.5.2 Comments and notes to Issue:1.1/Prep.2

1. CORBA defines an Event Service, but all ORBs do not provided it. The event service has many limitations that make it interesting only for event broadcasting. The Notification Service is an evolution of the Event Service and seems much more interesting, but it has to be investigated.

2. TAO [RD07] has the event service, extended with a real time event service. GTC has measured 5ms latency in event propagation between nodes and 0.5 ms in the case of remote invocation inside the same CPU.

3. If not already defined CORBA services will be used, the data channel architecture can be used to optimize the transmission of events between remote machines in a transparent way.

```
---060------------------
(BrianGlendenning)
```

p.16 2.4 point 1
> What does the DO have to do with this - it only seems to involve the property
> itself!

CLARIFICATION
Yes, you are right, but properties are accessed via their master DO.
Text will be clarified.

BGU 2001-03-30: Text clarified

---061------------------
(BrianGlendenning)
p.16 2.4 point 2 sub-points 2-6
> This is a major point. These cannot naturally be implemented with CAN as it
> is a pure polling mechanism. Do you mean that CAN will be polling the monitor
> point at some rate to support these capabilities? Why not just have the time-
> based and on-write events? If you really mean that continuous polling has to
> be supported to provide these, does property.get() go to the polled value or
> the hardware? If not the latter, how do you get the "really" current value?
> (You imply later that this is the mechanism you intend).

(BrianGlendenning)
p.17 2.5 point 2.2
> Again, how does this interact with polling etc?
>

ACCEPTED
It is correct, that the values have to be read in a polled system, but all monitor properties will be polled anyway at some rate, maybe very slowly(this is defined in the set of characteristics as the default monitor time, that characterizes the proper change time of the property).
When the property is polled, then the change events are generated if enabled and the condition is met. Reading the property at any time should generate a change event if enabled and the condition is met.
When Mick was at ESO we had some very constructive discussions on this point and we have seen a couple of different schemes for a reasonable implementation.
We have not identified any important problem.
Some prototyping work is anyway necessary and foreseen, before a final design is decided.

BGU 2001-03-26: done

---062------------------
(BrianGlendenning)
p.16 2.4 point 7
As mentioned before, state machines if included - need considerably more elaboration. (For example, the figure on page 12 & 16 should say something about them).

ACCEPTED
This is still a very TBD area.

We think that event driven, state driven software is easier to understand and build, especially in complicated systems.

GCH 2001-04-06: Will be done after TICS, using the state machine objects defined there.

---063------------------
(RonHeald)
p.16 2.4
There's no requirement (and little use) for value retrieval by event. Certainly all these conditions are not needed. There's no requirement for the ACS to collect monitor data.

REJECTED (value retrieval by event)
ACCEPTED (all these conditions are not needed)

Yes, there is a requirement: 4.1.6 on event/data channels. This is valid for any data channels, so also for monitoring purposes, even if monitoring is done by applications. So ACS provides the data channel mechanism, does not do the monitoring for applications.

It is ACCEPTED to leave the actual set of conditions for generating events on change TDB.
ACS 0.0 will provide "any change". If the other conditions listed will have to be implemented will be let TDB.
Value retrieval by event provides the features such as time events that are used by the control system to monitor data.
Value retrieval by event is useful because the control system can create alarms on out of range hardware conditions.
Alarms can be done by creating polling loops in a number of tasks, but this doesn't scale well.
The intention is to have a common method used by all applications for generating events.
Then, the implementation would be optimized in one place once.
The alternative, polling loops in a number of tasks, is difficult to maintain, debug, and optimize.

BGU 2001-03-29: Added polling of monitor points

2.6 Indirect Value Retrieval

2.6.1 Comments to Issue:2.0/Prep.1

---094--

(Jim Pisano)
p.21 2.5.3 This does not relate to the Data Channel, but to Direct value
retrieval. I would remove it from here and add it to the previous section only
if there's any new information here that's not already discussed in the previuos section.

(Jim Pisano)
p.21 2.5.5 Same comment that this should be moved to the appropriate section.

TO BE DISCUSSED AT REVIEW:
Sections 2.5.2, 2.5.3, 2.5.4 and 2.5.5 try to explain the differences,

advantages, disadvantages and proper usage of the 3 mechanisms.
We think this explanation is important, but probably this is not the right
place where to put.
Would you suggest us where to put and eventually how to restructure it?

---095----------
(Jim Pisano)
p.21 2.5.7 The correlator will publish its data on the science data channel
(see TICS document Fig. 2 & 3). I have only seen references for ACS data
channels to support monitor & logging data. It is understood that for ALMA, the
data channels probably won't support the 3 - 30 MB/s data rates required for
first light, but for TICS which will have a maximum of ~2 MB/s, the CORBA data
channel concept is valid. This point, 2.5.7, discusses applications creating
specific data channels, will this apply to the science data channel? If so, I
would like to see more verbiage related to this specific case.

I realize that the science data has been ignored by ACS and if it continues to
be so, then I want to make sure that access to the CORBA notification service
(which is the current plan) will be readily accessible or whose interface would
be simplified by ACS (which is alluded to in point 2.5.9.2). I would not want
access to the CORBA notification service to be hindered by ACS.

I also see in section 2.17.4 more discussion of Bulk data transfer for the
image pipeline which can encompass the science data channel from the correlator.

REPLY:
ACS will not hinder access to CORBA Notification Service.
We could not do any work up to now on the issue of high data rate transfers.
This will have to done in a second stage and is all TBD.
We will most likely use the experience accumulated with TICS and integrate
that work in ACS.

---096----------
(Joe Schwartz)
p. 22, The UML isn't entirely clear to me here. Is the "push data" a
message from DataChannel to Data subscriber or is it a data flow?

REPLY:
It is a data flow.
The "Data publisher" pushes data into the "Data Channel"

```
---097-----------------------------------------------------------------
----
(Mick Brooks)
p22, S2.5.9.1: Does not seem to say anything.

REPLY:
We will remove it.

GCH 2001-09-07 Done.

---098-----------------------------------------------------------------
----
(Brian Glendenning)
p.21 s2.5.9.2 Is a simplified wrapper API really necessary? If the
wrapper
is created is it forbidden to use the actual CORBA interface?

REPLY:
The idea it to provide an easier an more convenient interface.
If it is not perceived as such, we have missed our objective and you
can use
the actual CORBA interface.

GCH 2001-09-07 Clarified in text.

---099-----------------------------------------------------------------
----
(Jim Pisano)
p.23 2.5.9.4 Grammar - replace "warranty" (a noun) with guaranty (a
verb).

ACCEPTED

GCH 2001-09-07 Done.

---100-----------------------------------------------------------------
----
(Jim Pisano)
p.23 2.5.9.5 Combine this w/ 2.5.9.4.

ACCEPTED.

GCH 2001-09-07 Done.
```

2.6.2 Comments and notes to Issue:1.1/Prep.2

1. This service is not necessary in a first implementation, but will become very important on the final system, since it allows user access to Properties without affecting the behavior of the running system.

```
------------------
(JimPisano)
p.17 2.5 Diagram - Local Configuration DB
It looks like w/ the Local Configuration DB, we've lost all sense of
"static" characteristic info as the Indirect Value Retrieval (IVR)
manager can write to the config. DB. This further supports the division
of the configuration DB.

ACCEPTED
```

Sorry, the arrow in the diagram is misleading (probably in the wrong direction, depending how you see it).
The indirect value retrieval manager extracts from the configuration database the list and configuration of the properties it has to provide for indirect retrieval, it does not write in the configuration database.
Will be better explained.

GCH 2001-04-06: Section rewritten

---065------------------
(JimPisano)
p.17 2.5.6 "Access in current applications..." Is this in ANKA or VLT?
> What is the relevance of this statement?

CLARIFICATION
We are talking of ANKA and of other systems we have seen implemented in CORBA.
VLT has indirect value retrieval, implemented with what is called "scan system",
and this has proved itself very important to reach the required system performances.

The relevance of this statement is that we need to develop software for that and
that we need to spend more time investigating if CORBA solutions already exist.

Lately we have found some example/documentation on the implementation of proxy
objects, that look promising but need more investigation and prototyping.

GCH 2001-04-06: Section rewritten

---066------------------
(RonHeald)
p.17 2.5
> What's the relation between indirect value retrieval and the monitor
> data stream?
>
Indirect value retrieval means the user gets values from a mirror image of the running system. The mirror image is a copy of the running system, but the properties get their values from the monitor data stream.
Everything else acts the same.
There are timer and change on value events on the mirrored image.

---067------------------
(JimPisano)
p.18 2.5.8 Local Distributed Objects - This needs a definition. Are they
> the subset of DO's used for IVR? What is the relationship of DO's vs.
> LDO's? How are they kept in synch? Please show this relationship in the
> diagram for 2.5.
>
ACCEPTED

This need some clarification and is actually not described in details because
not fully defined.

The starting idea is that they are fully featured DOs whose properties just mirror the values of the real DO they are mirroring.
Client code should not change if it is accessing a property directly or indirectly.

We have on the other hand seen another approach, typical of CORBA based applications, that looks very interesting and that we need to explore more in detail.

With this approach there would be no need of real local DOs, but clients would
instantiate special "tie" classes instead of the bare interface classes provided by IDL
compilers (see Advanced CORBA Programming with C++ for details on tie classes).

With this approach, the "tie" client instance would be able to access directly the remote
DO or to register to the Indirect value retrieval manager.
ANKA uses this approach to implement Java Beans (Abeans) on top of CORBA IDL interfaces.

This needs more investigation, prototyping.

GCH 2001-04-06: Section rewritten

2.7 Sampling

2.7.1 Comments to Issue:2.0/Prep.1

---101---
(Martin Pokorny)
p.23 s2.6 Does sampling use the data channel? Or is the data channel restricted to the filling of the sampling database? I must admit that I don't really understand how sampling fits in to ACS, and that it's not really an application. Not that you need to address this in the current document, but I feel that the role of sampling in the TICS design needs to be clarified.

REPLY: Sampling could also be a TICS application but we think it is a general application not only used by TICS and therefore it is better that it is a part of ACS. Sampling will also use the data channel for transporting its data. This will be expressed better in the document. This is also an ACS requirement in AD01.

BGU 04-07-2001: Added description of use of data channel

---102---
(Joe Schwartz)
p. 23, 2.6.3, "...samples are stored in a data file for later analysis" seems to me to conceal an important high-level decision: is there an archiving policy for sampled data, or is this just throwaway

stuff for test use?

REPLY: Yes I think it would be useful to archive some sampling data, but certainly not everything. Much will just be throwaway for test use. But I don't think it is up to ACS to define any policy here.

BGU 04-07-2001: Nothing to be changed in the document. Clarified in reply

---103---

(Joe Schwartz)
p. 23, 2.6.4, While AD01 refers to an "existing engineering data analysis tool", no specification or requirements document exists. Did someone decide that LabView was okay? Is there to be a LabView-specific ACS interface? Will there be an I/F to other tools? It seems to me that specification of such an interface should be part of the ACS architecture--or is LabView a part of ACS?

REPLY: LabView has been considered as an engineering tool by hardware people. ACS will not define its use. ACS will also not provide any interfaces to LabView.

BGU 04-07-2001: Nothing to be done. Clarified in reply

---104---

(Steve Scott)
p.23 2.6.4: It should be possible to plot in near real-time anywhere, not only at the operator terminal, limited only by network bandwidth. The feed should be a general stream - a file is too restrictive (a file can be the underlying implementation but it must be transparent to the user).

REPLY: Yes, this is how it is meant. The text will be improved. The file is the undelying implementation, but this file can also be used for other analysis tool and for archiving data.

BGU 04-07-2001: Changed document to better describe this mechanism

---105---

(Steve Scott)
p. 23 2.6.5: "Data sampling is..." => "Data transport is..."

ACCEPTED

BGU 04-07-2001: Done

---106---

(Brian Glendenning)
p.23 s2.6.6.3 I assume you mean that it "sends" (not samples) the data packets. You should mention that it uses the data channel.

ACCEPTED: Yes this is just a mistake.

BGU 04-07-2001: Done

---107----------
(Dirk Muders)
p23, s2.6.6.4:

This is a requirement. How does the ACS architecture / design achieve this ?

(Joe Schwartz)
p. 23, 2.6.6.4, "Sampling must not reduce the performance of the control system or the network." This sounds like a requirement to me. How will the ACS architecture ensure this?

REPLY: The data are cached on the local host and sent with lower rate to
the center to reduce network traffic. The data collection on the local host will be performed with lower priority than critical control tasks. This will be ensured during the design of the ACS and the control system.

BGU 04-07-2001: Added explanation how performance could be kept

---108----------
(Jim Pisano)
p.24 2.6.7 "The implementation could be optimized ..." It seems to me that the
sampling data channel must be implemented as a data channel and it is a natural
extension of the Data Channel concept discussed in section 2.5. I would change
this sentence to "The data channel architecture will be used to fill in the
sampling database from the sampling manager." or "The sampling manager will be
a subscriber to a sampling data channel which stores samples in a sampling
database."

ACCEPTED: Yes you are right. It is foreseen to use the datachannel for transporting sampling data. The text will be changed in this way.

BGU 04-07-2001: Added description of use of data channel

---108----------
(Joe Schwartz)
p. 24, 2.6.8, XML for sampling files: I recall a discussion in Socorro about performance problems when XML was used for log files. Have any performance tests or analyses been done to determine whether the system can handle the load when XML is used?

(Steve Scott)
p. 24 2.6.8: I don't see how XML adds any information or functionality over simple ASCII files or DBMS tables. Either a

file or table would have a simple method to extract and create a sequence of time series sample objects (time, value, maybe that is all).

TO BE DISCUSSED AT REVIEW:
XML has the advantage that it is a standard format for which you easily can find parsers and conversion tools for example to ASCII. When using an internal ASCII or binary format you always have to invent syntax and parsers.
Of course we have to prototype and verify that we get the required performance.
This will be done during the design phase.
Your point is of course valid to some extent because the structure of the sampling data consists only of the data value and a time stamp so it is not that complex.

BGU 04-07-2001: XML kept as is. Added in section 4 about prototyping and testing perfomance

---110---
(Dirk Muders)
p24, s2.6.8 - 2.6.10:

These are examples of possible implementation decisions and the usage of ACS.
One should only mention the decisions that have been taken.

REPLY: 2.6.9 and 2.6.10 will be merged with 2.6.4. 2.6.9 presents the envisaged solution which have to be prototyped to verify its suitability before a decision is taken.

BGU 04-07-2001: Done

---111---
(Jim Pisano)
p.24 2.6.9 This should be combined w/ 2.6.4. Also I would add that a Java app could be written as a subscriber to the sampling data channel (if you approve of this concept).

ACCEPTED: 2.6.9 will be merged with 2.6.4. Data channel will be used and the sampling client will most probably be a Java application.

BGU 04-07-2001: Done

---112---
(Martin Pokorny)
p.24 s2.6.9, s2.6.10 I thought that applications are not a part of ACS. This may be the wrong place to mention a plotting application, and the features that it will have.

REPLY: ACS will provide some "applications" where a plotting engine for sampling is one of them. These are general tools to be used by everybody.

BGU 04-07-2001: Nothing to be done. Clarified in reply

---113---

(Joe Schwartz)
p. 24, 2.6.10, "Multiple samples can be super-imposed..." Is this an ACS feature, or a feature of the "existing engineering data analysis tool"?

REPLY: This will be a feature of the non-existing plotting tool for sampling.

BGU 04-07-2001: Nothing to be done. Clarified in reply

---114---

(Dirk Muders)
p25, s2.6.11:

Why does the TP data need to go through ACS ? Aren't the high data rates a
problem ?

(Jim Pisano)
p.25 2.6.11 Please consider my comment to point 2.5.7 here about "...the actual
receiver data stream which has its own fiber network."

(Joe Schwartz)
p. 25, 2.6.11, "...although this is actually science data..." Please don't forget that if ACS is to be used by *all* ALMA software, it's going to have to handle science data all the time. I would hate to think that something as basic as the architecture of ACS is being prepared now without serious thought about those requirements.

(Mick Brooks)
p25, S2.6.11: Does this discussion belong here any more? Will it be implemented using the sampling engine?

ACCEPTED: Science data will also be considered by ACS. The foreseen application
for transporting science data is the data channel. We need to prototype this
to verify its performance. The total power will also use this channel. So it
will be removed from the sampling chapter.

BGU 2001-09-07: Done

```
---115-------------------------------------------------------------------
----
(Ron Heald)
p25, S2.6.11 and p39, S2.15.3: "M&C" is not in the acronyms list.  I
would prefer to use "monitor and control".


REPLY: M&C will be added. But what's wrong using acronyms?

BGU 2001-09-07: Done
```

2.7.2 Comments to Issue:1.1/Prep.2

```
---068------------------
(RonHeald)
p.18 2.6.2
> There's no time limitation on sampling.  Sampled data can simply be
> stored in a file, which is the way LabView expects input anyway.
ACCEPTED
We will remove the limitation.
The limit will be indirectly imposed by AD01 4.3.1 on sampling,
that says the network load should be under control.

BGU 2001-03-26: done

> Plotting should also be available at the antenna.

CLARIFICATION
If we are debugging problems at the antenna, then a portable computer
with the user interface and software to talk to the objects if what is
foreseen.
This allows to plot values at the antenna.
> There's no
> requirement for ACS to transport total power, holography, or optical
> camera data.
>
CLARIFICATION
Since the first version of this document Mick has pointed out total
power as a very
likely candidate.
This is explicitly written in is "Total Power data interface options"
document
dated 2000-07-07.

BGU 2001-04-02: no change of dcument

---069------------------
(BrianGlendenning)
p.19 2.6 point 8
> Important: Again, I think one of XML or RDBMS should be chosen. Here
an RDBMS
> appears to me to be more appropriate as the samples are of interest
generally
> to engineering staff who aren't apt to have tools that can deal with
XML
>
ACCEPTED
As said in the text, we are not able to make the choice now. We need
more
```

```
experience and prototyping to find which is better.
XML is a good interchange mechanism between applications and computers.
If we build with a good data description language, then the tools are
more easily reused.
We would translate from XML to RDBMS which is off the shelf.
An RDBMS has usually too much overhead for "just in time" display

BGU 2001-03-30: Nothing changed in document

---070------------------
(JimPisano)
p.19 2.6.7 "data channel architecture" needs a better description as to
how
it affects the sampling implementation & its optimizations.

ACCEPTED
This still requires some analysis and design work.


GCH 2001-04-06: Added section on data channel

---071------------------
(MickBrooks)
p.19 point 11. I am making a proposal for the physical interface for
total
power this week. The samples will somehow be buffered at the ABM.

ACCEPTED
We have seen the document and we will take your proposal into account

BGU 2001-04-02: Added reference to Total Power
```

2.8 Command Handling

2.8.1 Comments to Issue:2.0/Prep.1

```
---116-------------------------------------------------------------------
----
(Ron Heald)
p.25 2.7  The subject of "command handling" should be discarded.  The
authors define in it terms of the VLT software where it forms a command
checking and communication layer.  With the use of CORBA, "commands"
are
simply (possibly remote) object method calls.  In this environment
there's no need for an additional layer.  Command checking is done
within the method call (as it will in any case), and CORBA forms the
communication layer.

A static "CDT" mechanism will never be able to perform the dynamic
checking that is required in some cases.  Dynamic checking can only be
done within the object as you point out in 2.7.10.


(Dirk Muders)
p25, s2.7 and s2.7.1.6:

What are commands ?  Are they merely methods on DOs ?  It sounds here
like
there will be a whole system of syntactically checkable commands for
each DO.
```

This would be quite complicated. I was always under the impression that CORBA /
ACS just make a normal object accessible to everybody no matter on which
platform.

REPLY: This subject was discussed in Grenoble and the conclusion was:
Commands are DO methods.
Syntax and dynamic checking shall always be done on the server side.
Clients can do syntax checking. Generic tools like sending commands from a command line or from scripts shall always do a syntax check before sending the command.

ACS will provide a library supporting syntax check only. Here we propose
to use a command description (CDT) to describe commands and their syntax,
parameters, static ranges.

This section will be modified to reflect this.

BGU 04.07.2001: This section has been modified acccrding to the discussion in
Grenoble

---117----------
(Brian Glendenning)
p.25 s2.7.1.4 You should be able to have both clients and servers in multiple languages.

ACCEPTED: This we have agreed upon in Grenoble. The document will be modified to reflect this.

BGU 04.07.2001: Done

---118----------
(Brian Glendenning)
p.27 s2.7.1.5 IMO given AMI it is not clear that the callback mechanism is
the "standard" method.

(Brian Glendenning)
p.27 s2.7.1.6 AMI should be mentioned here.

(Joe Schwartz)
p. 25, 2.7.2, It is stated here (and in AD01) that it's the responsibility of the server application to check syntax and parameter ranges of a command. I think that while this is an important safety measure, the client should have primary responsibility for ensuring that it has satisfied the preconditions necessary for successful use of a server; after all, the server can throw some kind of "bad argument" exception, but only the client can know what it's sending. Cf., Douglass, "Real-TimeUML", who says, "The responsibility for ensuring preconditional invariants are met falls

primarily in the client's realm. That is, the user of the operation is required to guarantee that the preconditional invariants are satisfied."

(Ron Heald)
p.25 2.7.3 From this section is is not clear that having both synchronous and asynchronous method calls is only from the client view point. This section is unchanged from previous versions, and conflicts with sections 2.7.18 and 2.7.19 that discuss AMI and threads to eliminate the need for asynchronous mechanisms on the server side. Section 2.7.15 implies the server receives a callback object that is no longer true.

(Jim Pisano)
p.28 2.7.19 So now we have 3 types of asynchronous commands:
1. Callbacks ala ANKA
2. Callbacks via AMI
3. Synchronous commands on a client thread.

This is too complicated, please let's select one and go with it.

REPLY: In ACS 0.0 asynchronous commands with callbacks was the standard way.
It was agreed in Grenoble that AMI and syncronous commands are the standard mechanisms to be used. AMI will be used for C++ clients with TAO and
synchronous commands with threading for other ORBs not supporting AMI and with
Java or Python.
Using asynchronous commands time-out handling and intermediate replies are
supported by ACS. With synchronous commands the applications will have to
take care of this.

The document will be modified to reflect this way of handling commands.

BGU 04.07.2001: Done

---119----------
(Dirk Muders)
p26, s2.7.4:

Should it read "must be ignored" ? If it is not mandatory, drop the subsection.
It is then up to the application programmer what to do.

(Martin Pokorny)
p.26 s2.7.4 Is this a client application specification? If so, it needs some emphasis. If not, change "are" to "may be".

(Joe Schwartz)
p. 26, 2.7.4, "Asynchronous replies sent after a timeout are ignored by the client application." What does this have to do with ACS? Isn't this the client's problem?

ACCEPTED: Asyncronous commands are no longer used so this section will be
dropped.

BGU 04.07.2001: Done

---120---
(Jim Pisano)
p.27 2.7.9, 2.7.10, 2.7.11 It seems like there is a chance of redundant checking, namely the DO checks commands & the Property checks commands. I hope
that a Property.set() function isn't first checked at the DO & then checked at
the Property.

REPLY: When setting properties on the the Property.set() method will check the
range of the value on the server side.

BGU 04.07.2001: Nothing to be done. Clarified in reply

---121---
(Dirk Muders)
p27, s2.7.11:

This is a requirement. How is this achieved in ACS ?

REPLY: Yes it is a requirement. Perhaps it does not belong to this document
and should be removed.

BGU 04.07.2001: Done

---122---
(Martin Pokorny)
p.26 s2.7.16 Change "timeout time of can have" to "timeout value or can have".

(Dirk Muders)
p27, s2.7.16:

... timeout time "or" can have a dynamic ...

REPLY: This is only valid for asynchronous commands which is no longer the
standard way so this section can be dropped.

BGU 04.07.2001: Done

---123---
(Jim Pisano)

```
p.27 2.7.16 Grammar - "A client can use a default value for the timeout
time of
can have a dynamic evaluation based on a Callback Timeout and
Negotiation
Mechanism." Change "of" to "or". Also change "mechanist" to
"mechanism".


ACCEPTED

BGU 04.07.2001: Done

---124-------------------------------------------------------------------
----
(Steve Scott)
p. 27 2.7.16: "The command execution method..." => "The command
execution method for asynchronous commands..."


REPLY: Asynchronous commands will be dropped so this section will
disappear.

BGU 04.07.2001: Done
```

2.8.2 Comments and notes to Issue:1.1/Prep.2

1. The following example show how the XML description could look for a Device (Elevation axis) that implements just one command with a single parameter (Move(newPosition))and one read-ony Property (position). Please be aware that it is just a fictitious example, which has nothing to do with the real antenna. It just gives an example of the use of XML:

```
<Device>
  <Header Name="Elevation" Version="0.0"> <Description>Elevation
hardware.</Description>
  </Header>

  <Command>
    <Header Name="Move" Version="0.0"> <Description>Move axis to given
position</Description>
    </Header>
    <Parameter>
      <Header Name="Position" Version="0.0"> <Description>Position of
elevation axis.</Description>
      </Header>
      <Value Default="0" Format="%6d" Resolution="65536:1" Type="long"
Units="Binary Parts of a Circle">
      <Range>-32768,32767</Range>
      </Value>
    </Parameter>
  </Command>

  <ROProperty>
    <Header Name="Position" Version="0.0"> <Description>Position of
elevation axis in.</Description>
    </Header>
    <Value Default="0" Format="%6d" Resolution="65536:1" Type="long"
Units="Binary Parts of a Circle">
```

```
    <Range>-32768,32767</Range>
    <IO Address="0x00040012" Interface="CAN" Type="int32*2">
    <Conversion>linear</Conversion> </IO>
    <Time Interval="0.050" Mode="retrigger" Offset="0"/>
    </Value>
  </ROProperty>

</Device>
```

2. Differently from what we thought when we have started the investigation on CORBA, we have verified that code does not always need to be recompiled if new IDL interfaces are made available or if an IDL interface is changed, for example adding new methods[AD01 - 13.1.6. Modularity]. The ANKA group has verified that at least with Inprise's Visibroker and HP's Orbacus even if one changes interfaces, both client and server still work if they do not use the new features. It is possible to have the server exporting one version of the interface and the client use another version. If an interface changes, then re-compilation is necessary (as expected) only for those programs that use the new features of the methods that eventually changed signature. This should work also with other ORBs, since it is due to the existing CORBA IIOP protocol. On the other hand, this behavior is not explicitly supported and, should the IIOP protocol ever change, different interfaces might lead to run-time errors.

   Note: The potential need to recompile all applications when IDL interfaces are modified seems now much less critical than when we started this analysis, but still needs to be considered carefully and checked against the definition of CORBA standards and the IIOP protocol.

```
---072------------------
(RonHeald)
p.19 2.7
> Commands are equivalent to method calls, and therefore need no
special
> capabilities.

CLARIFICATION
TR's says in 6.1.6 and 6.1.8 that behavior of
commands has to be defined on top of pure methods, which are a way to
convey them.


BGU 2001-03-29: Nothing changed in document


---073------------------
(BrianGlendenning)
p.20 2.7 point 3
> Just a comment  - this means that every command (a term I'd like to
de-
> emphasize, really just a DO method call) requires two IDL  methods
plus a
> callback type per unique signature.

ACCEPTED:
In a first approximation you are right and this is what we had in mind.
On the other hand, they can be generated automatically (this is what
AMI does).

There are also clever designs that do not require two methods with two
signatures. An interesting example is described in "Abeans Programming
Tutorial", from ANKA, and presents an interesting Java implementation
```

that, based on a pure asynchronous server, allows Java applications
to access it both in a synchronous and
asynchronous way.

The final solution will be detailed later. With ACS 0.0 we get ANKA
implementation. For the final system we will have to see the evaluation
of
AMI (see other comment on AMI).

BGU 2001-03-26: Changed text. Added reference to AMI report.

---074------------------
(BrianGlendenning)
p.20 2.7 point 7
> Is the IREP generally available (e.g. I didn't think TAO had one).

CLARIFICATION
TAO has no Interface Repository, but Orbacus provides one.
This information should be available from the central database, too.

BGU 2001-03-30: Text clarified

---075------------------
(BrianGlendenning)
p.20 2.7 point 8
> Important: I think doing this properly (as opposed to a proof of
concept) is a
> *very* major job. Can we "steal" someone else's work here, e.g. IML?
If not,
> do we have the manpower to do this? The text seems to imply that
parameters
> and properties are the same, although the example separates them
properly.
>
ACCEPTED
We fully support your concerns.
Fritz and Jens Knudstrup (from ESO) are working on a proposal that
should be available soon.
The starting points have been IML/AML as well as ESO CTDs and Parameter
Files used
on the VLT.

BGU 2001-03-30: No change of document to keep it light

---076------------------
(JimPisano)
p.20 2.7.3, 2.7.14-15 In a multi-threaded environment, callbacks are
not
> necessary for (a)synchronous operation.
> A thread can be started for a given command from
> a client without blocking the client's main execution. Requiring
> callbacks on the device side complicate its software and are
unnecessary
> as all OS's implementing ACS are multi-threaded. A timeout handler is
> also mentioned which automatically assumes a different thread, so why
> not put the whole command response issue on its own thread?

ACCEPTED
What you describe is for sure one possibility, and will be implemented.

The fact that callbacks are more complex than threads is not so sure
as it seems from your comment.
Actually many people says the opposite.
For this reason both have to be supported, at least at prototype level.
Experience will tell is only one or both are necessary.
See other comments regarding callbacks and synchronous/asynchronous
commands.

BGU 2001-03-29: Added point about thread handling

---077-----------------
(JimPisano)
p.20 2.7.7 Commands should be in their own library, not the config DB.

CLARIFICATION
The text is not clear and will be rewritten.
Commands are clearly "implemented" in their own library.
To be "callable", their interface must be published in an IDL file,
used by
the clients.
Clients can use IDL compile time stubs to call remote objects.
It is also possible to build calls at run time, using the Dynamic
Invocation Interfaces without compiled stub code. This uses the
Interface
Repository to get run time access to IDL definitions.

Dynamic Invocation is necessary for generic tools.
For example, a command line tool to send commands from the shell, need
to
get the command as a string, map it via Dynamic Invocation Interface to
a CORBA method call and execute the call.

This anyway does not allow client side checking of parameters against,
for
example, range.
For generic tools like that mentioned before we believe that a check
flag
should be supported to allow local checking of parameters, without
calling
the remote object.
It would be important also to implement important user friendly
features
like command completion and context sensitive help.

The way proposed is to define commands in the configuration database,
together
with the properties of the distributed objects.
Whenever a command for a certain object is needed, it is looked up in
the
configuration database (locally), so that is is possible to know which
parameters
it needs and how they can be checked.

A generic parser for this needs to be developed, but the same parser
can be used
also on the server side to provide a standard way to check call to
commands
on the receiving side, instead of having to write the checking code
every time.

At this point there is no implementation overhead: server applications HAVE to
check parameters of incoming commands, and a checking API is very convenient.
Client applications can use the same API where it makes sense to have checking on the client side and no special development is needed.

BGU 2001-03-30: Text modified

---078------------------
(RonHeald)
p.20 2.7.2
> The checking needed for a "generic command-sending GUI" would be, for
> example, to make sure a number is entered in a number field, or a valid
> date in a date field.
>
It would mean to check range validity on static ranges, which is more (see 6.1.3 Syntax check for GUIs) and requires specific DO information, namely the valid range. What you talk about is simply type checking.

BGU 2001-03-29: Added clarification

---079------------------
(RonHeald)
p.20 2.7.8-13
> Important comment: AD0-6.1.4 says validation will occur at the
> "receiving end" = DO. AD0-6.1.5 says checking will only be done in one
> place, i.e. in the DO. There's no need for a CDT equivalent in XML.
> There's no need for a "command parameter checking library". All
> checking (except that described in the previous item) should be done
> only in the DO.

6.1.5 speaks about one file, leaving free the possibility to use an identical XML file (e.g. generated at start-up time in a local and in a remote place). The information is the same.
This is a point TBD after ACS 0., given the undefined state of the Configuration DB concept at this stage. Command checking will not be part of ACS 0.0
The difficulty here for some of us, is to think that it will be practical
to use in a tool like e.g. the Observation preparation for Phase 2 the same range checks via DOs, used for control. This is unthinkable for the experience we have with this tool. Surely though it would be desirable that the check is based rigorously on the same information, like it would result by spawning it from the same Config. DB.

BGU 2001-03-29: Added clarification

---080------------------
(BrianGlendenning)
p.21 2.7 point 12
> Choose! (We keep coming back to the role of XML)
>

We need some experience/prototyping before we can choose.

BGU 2001-04-02: no change of the document here

---081------------------
(BrianGlendenning)
p.21 2.7 point 14
> Comment: AMI makes life much simpler, as then it's purely up to the caller
> which he wants, the IDL doesn't need asynch and synch versions of the same
> method, and we don't have to have all kinds of callback types defined.
>
CORBA AMI (Asynchronous Method Invocation) will be extensively evaluated after
ACS 0.0, that is based on ANKA callbacks.
If what you say is proved true, than we will use AMI.
On the other hand, M.Plesko states that implementing code using AMI is much
more complex than using ANKA callbacks.

I (G.Chiozzi) have been reading AMI documentation in the last days and I
find the concepts interesting, but not dramatically different.

For example, with AMI the clients have both asynch and synch versions of the same
method, generated by the IDL compiler.
The major difference is that with ANKA synchronous and asynchronous behavior
is handled on the server side, while with AMI is handled on the client side
using features provided by the ORB.

As said in the document, experience and prototyping will tell.
It is also not true, as said in comment 079, that we really need asynch and
synch versions of the same method.

BGU 2001-03-29: Added note on AMI and refernence to report

---082------------------
(JimPisano)
p.21 2.7.9 Implementation details like this XML description add to the
> complexity of this document by covering too many levels. There are very
> high-level concepts which this document correctly addresses, but then
> there are these details which, since they are only know in a few cases,
> are added. This makes this document difficult to read as these details
> distract me from the overview nature of what this document should be.
> If you do want include detailed info like this, put it in an appendix
> like the comparisons to ANKA & CCS.
>

ACCEPTED

example will go into an appendix.

BGU 2001-03-29: Moved example to ACSArchitectureNotes.doc

---083------------------
(JimPisano)
p. 21 2.7.11 "the implementation of the object server..." This sentence
> does not make sense. What object server? The first sentence is all
that's
> needed.
ACCEPTED:
We mean Server Object, not object server, i.e. it is a requirement that
commands
are checked on the server side, when received. These checks can make
use of the
current object context and status and as a consequence can do more than
checks
based on the command definition alone. For example, a parameter can be
in the valid
range, but not acceptable because of the actual value of another
property
of the object.
Actually the second sentence could be a point in its own.

BGU 2001-03-29: Modified text

---084------------------
(JimPisano)
p. 21 2.7.18 "proper exceptions" C++ Exception handling and TAO
exceptions
> under VxWorks are problematic. Please specify what type of exceptions
you
> mean.
These a CORBA exceptions, that are mapped by IDL in exceptions (or
errors)
for the different languages.
We are working on understanding what works and what does not in C++
exceptions
with VxWorks.
This topic is very dynamic, since there are various versions of C++
compilers
available for VxWorks.
Each new version improves things, but we still have not identified
a set of platform/compiler/ACE-TAO build flags that works fine.

BGU 2001-03-29: changed to CORBA exceptions

---085------------------
(MickBrooks)
p.21 point 9. you say that the XML is for the azimuth axis yet all of
the
words say elevation axis. I would also like to hear Ron's ideas on this
because we will not be treating the axes separately as I understand it.
I
realize that this is merely an example, but you have picked a touchy
one.
Also, the interface is a trajectory specification (position and
velocity

pairs) rather than position alone. Maybe that is a level beneath this description?

ACCEPTED
The example will be cleaned up and moved in appendix.
We will probably use the Mount interface written by Ron.

BGU 2001-03-29: Moved example to ACSArchitectureNotes.doc

---086------------------
(RonHeald)
p.21 2.7.14
> Important comment: It is not clear whether the asynchronous commands as
> described are needed. The VLBA is a very successful instrument that has
> no such concept. The same thing can be accomplished with synchronous
> commands and multiple processes, and without the complexity of
> callbacks. A fundamental question is whether the observing script will
> be based on absolute times or durations; a question for the SSR to
> decide.
>

REJECTED
This was already discussed at the ACS TR review and 6.1.7 is in disagreement
with what you say.

It would also be wrong to base scripts on absolute time. This would prevent any
Scheduler to do its work properly, namely to launch them at a time which
in practice is never exactly the one foreseen in advance.

See also other comments related to callbacks.

BGU 2001-03-30: Synchronous/asynchronous command handling expanded

---087------------------
(BrianGlendenning)
p.22 2.7 point 20
> Important: Again, if the property.get() doesn't go directly to CAN, control
> applications that want the "most current" values will have to bypass
> Properties, which would be a shame IMO.
ACCEPTED
Direct value access (not through indirect value retrieval) from the property is
always through CAN.
The CAN interface will be designed accordingly and the text will be changed
to make this clear.

BGU 2001-03-26: changed text

2.9 Logging System

2.9.1 Comments to Issue:2.0/Prep.1

```
---125-------------------------------------------------------------------
----
(Brian Glendenning)
p.28 s2.8 IMO "logging" is a bad names because it always makes me think
of
informational text messages. ("Captain's log, stardate 1234").


REPLY: This is exactly what we mean by logging (informal text
messages).
Do you have a different interpretation of this chapter?

BGU 04.07.2001: Nothing to be done. Clarified in reply

---126-------------------------------------------------------------------
----
(Jim Pisano)
p.28 2.8.1 According to 2.5.7 "ACS itself uses the Data Channel to
provide
basic services like logging and archiving.", the use of the CORBA log
service
contradicts this, please resolve. A brief perusal of RD38 shows that
the
logging service is based on the event service (like the CORBA
notification
service). It seems that this should be used in its native format
without any
extra ACS data channel services. Also, In the rev. 1 of the ACS
Architecture
doc, it was mentioned that this logging service wasn't available in
most ORBs,
is this still true?


REPLY: The ACS logging service uses the notification service for data
transport. Also the data channel is based on the notification service.
So
where's  the contradiction?

BGU 2001-07-05: Nothing to be done. Clarified in reply

---127-------------------------------------------------------------------
----
(Jim Pisano)
p.28 2.8.3 What are "Miscellaneous log messages."? This seems like a
open-ended
bucket for any sort of unspecified junk which appears to me to lead to
an
unlimited amount of random messages that clients would spend most of
their time
filtering out. Please either specify what would get logged and/or
remove this
bullet.
```

REPLY: These are applications logging what they regard as necessary to log.
It is not up to ACS to specify what applications should log.

BGU 2001-07-05: Nothing to be done. Clarified in reply

---128---

(Steve Scott)
p. 28 2.8.4: Standard objection to XML. It has value in a complex hierarchy, but not in a flat simple sequence such as this.

REPLY:XML has the advantage that it is a standard format for which you easily
can find parsers and conversion tools for example to ASCII. When using an
internal ASCII or binary format you always have to invent syntax and parsers.
The format of a logging message is also quite complex with many fields.
Of course we have to prototype and verify that we get the required performance.
This will be done during the design phase.

BGU 2001-07-05: Prototyping and performance test of logging added to section 4

---129---

(Brian Glendenning)
p.29 Figure Where does the differing types of log messages show up? In classes? In the "schema" only? Somewhere else?

REPLY: All logs are archived in the same archive. The different log types
can be filtered out with for example the log browser.

BGU 2001-07-05: Document not changed

---130---

(Jim Pisano)
p.29 2.8.7 I don't understand this phrase: "...log entries with low priority do
not get logged,..."? What's the purpose of logging something if it doesn't get
logged?

(Jim Pisano)
p.30 2.8.15 Does this mean that any device that produces debug logging always
publishes this data, i.e., it is not enabled by a command? Does the filtering
at the consumer side somehow trigger the log producer? I was under the impression that the log producer would be commanded to start logging debug info
of a certain level and then stopped as needed. If all devices are continually
sending all debugging info, then wouldn't this easily overwhelm the logging

channel, especially w/ all the DO's in 64 antennas.

REPLY: Aplications will be written with all the logs that are appropriate,
debug, trace errors etc. Then on the server side it is possible to filter
out all lgs below a certain level before they get logged. Debug logs are
for example suppressed by default in order not to fill the log file. The
user can enable all debug logs from a server with one command.

BGU 2001-07-05: Text rephrased

---131---
(Steve Scott)
p. 29 2.8.8: Logs are critical in running the system, and the resource usage is small. Send them to the Log Data Channel as soon as possible without caching on the local machine (other than a FIFO queue).

REPLY: The purpose of a local cache was to reduce the network traffic. Whet the cache reaches a certain size or after a certain time the cache is flushed and all log messages in the cache are sent to the Log Data Channel in one block. High priority logs are not cached but are sent to the log data channel immediately.

BGU 2001-07-05: Modified text according to reply

---132---
(Steve Scott)
p.29 2.8.10: Good point. Maybe one of those "Yadayadayada... repeated 50 times" messages that you see in some syslogs.

BGU 2001-07-05: Document not changed

---133---
(Martin Pokorny)
p.30 s2.8.11, s2.8.12 Not that I have anything against someone signing up for developing applications, but is the present document the proper place for this specification?

REPLY: The log monitor is part of ACS being a general tool.
This is an ACS requirement in AD01.

BGU 2001-07-05: Nothing to be done. Clarified in reply

---134---
(Dirk Muders)
p30, s2.8.13:

This is a requirement. How is this achieved in ACS ?

REPLY: Yes, this is a requirement to the design and implementation of the logging system. Network traffic is reduced by caching the logs locally
before sending them to the data channel. The local logging process will also
have lower priority than the critcal control tasks.

BGU 2001-07-05: Added use of local cache and task priorities

---135---
(Joe Schwartz)
p. 30, 2.8.15, Are these rules concerning "no compile time debugging flags" something that ACS enforces (how?) for ALMA applications? Isn't this a policy statement rather than a piece of the architecture?

REPLY: ACS just provides the engine for this. Perhaps the text should make
that clearer.

BGU 2001-07-05: Done

---136---
(Steve Scott)
p.30 2.8.15: Very nice!! Can the debugging level be set with a CORBA command?

ACCEPTED: It is not implemented as a CORBA command now, but this is a very
good idea.

BGU 2001-07-05: Added CORBA command

---137---
(Brian Glendenning)
p.30 s2.9.1 Mention non-monitor point archiving (e.g., log messages). We
have decided that archiving of monitor points is a TICS application.

ACCEPTED: ACS will provide the engine though for collecting data and pushing
them on the data channel. This should be the emphasis of this section.

BGU 2001-07-05: Section modified

---138---
(Jim Pisano)
p.30 2.9.1 According to the diagram on pg. 31, the archiving system also
archives log messages.

REPLY: Yes, the archiving system will handle monitor data and logs. The archiving of the data will be part of TICS, but this will later be incorporated into ACS. So this acrhitecture will change according to the TICS
design.

BGU 2001-07-05: Removed everything about storing data

---139---
(Brian Glendenning)
p.30 s2.9.3 I think we should only have SI units in the interfaces of items
that are away from user interfaces (i.e. seconds rather than ms).

ACCEPTED: Good point. We will do it like that.

BGU 2001-07-05: Removed ms in text. Code will be modified after ACS 1.0

---140---
(Joe Schwartz)
pp. 30-31, I hope that I have understood the discussion of the archiving system, because the mechanism that we are planning to use in the higher-level software depends on notification that, e.g., correlator data or processed calibration data, has arrived in the archive.

REPLY: The higher-level software can subscribe to monitor data from the data channel and get notified when events occur in the system, like data has been archived.

BGU 2001-07-05: Document not changed

---141---
(Brian Glendenning)
p.31 s2.9.5 It is not important to archive the properties synchronized to
the hardware tick (the *sampling* is what is relevant).

REPLY: This should read collection of data. The text will be changed.

BGU 2001-07-05: Done

---142---
(Jim Pisano)
p.31 2.9.5 What is "the hardware tick"? Is this the 48ms array-wide timing
event?

REPLY: Yes to be clarified in the document.

BGU 2001-07-05: Done

---143---
(Martin Pokorny)
p.31 s2.9.5 Hardware monitoring is generally "synchronized" to the 48ms timing event, although derived monitor points need not be. In any case, monitor points do not normally change more often than once per

48 ms. I'm not sure what it means to synchronize archiving (as opposed to monitoring) to the 48ms tick.

REPLY: What we mean here is data collection. To be clarified in the document.

BGU 2001-07-05: Done

---144----------
(Brian Glendenning)
p.31 s2.9.6 What about other logs, e.g. text messages. Again monitor point
archiving is an application.

REPLY: Logs will also be archive centrally. Archiving will be part of TICS.
So this will be modified according to the TICS design. ACS will only deal
with data collection. Archiving will later be incorporated into ACS.

BGU 2001-07-05: Removed everything about storing data

---145----------
(Dirk Muders)
p31, s2.9.7:

Can a synchronization be forced on demand ?

REPLY: Data are sent to the data channel at least every second so there's not
much sense to force synchronization with a command.

BGU 2001-07-05: Document not changed

---146----------
(Jim Pisano)
p.31 Diagram: It doesn't seem clear which data channel is the "archiving data
channel". The archiveManager gets logging & monitor data via "data channels"
(which are separate). Is the archiveServer an example of a client subscribing
to the Archiving Data Channel? If the logMonitor & monitorDisplay objects are
clients to the Archiving Data Channel, then why do they connect to the
archiveServer, not just the Archiving Data Channel?

REPLY: Archiving will be done by TICS. So this will be modified according to
the TICS design. ACS will provide the log data channel
and the monitor data channel. Archiving will later be incorporated in ACS.

```
---147-------------------------------------------------------------------
----
(Martin Pokorny)
p.31 s2.9.9 Change "provided that the value was archived prior to that
time" to "(provided that the value has been archived)". Nobody will
expect ACS to retrieve unarchived values or values from the future!


ACCEPTED

BGU 2001-07-05: Done
```

2.9.2 Comments to Issue:1.1/Prep.2

```
---088------------------
(BrianGlendenning)
p.22 2.8
> Important: IMO the logging system is only for informational messages.
Monitor
> point archiving may use the same mechanisms that it is implemented
with, but
> it is no more logging than science data is.
ACCEPTED
TO BE DISCUSSED AT REVIEW
Nothing against your comment.
But since the mechanisms would be the same we did not want to put a
separate
section with the same contents.
If you prefer we can completely decouple the two things on the
architecture document.

GCH 2001-04-06: Section split

---089------------------
(JimPisano)
p. 22 2.8.3.1 Exceptions - do you mean errors as discussed in
2.8.3.6.2? If
> so, then change exceptions to errors.
>
ACCEPTED

BGU 2001-03-26: done

---090------------------
(BrianGlendenning)
p.23 2.8 point 5
> Again, one of RDBMS or XML should be chosen. It's also not evident
that the
> "online"  DB  only keeps a few days worth of data.

CLARIFICATION
The long-term database is the Archive DB and this has to be an RDBMS.
The short-term database contains only a few days worth of data, because
it does
not make any sense to keep older information always available online.
This would make the online database much heavier and complex.
The short-term db can be XML or an RBDMS. Experience and prototypes
will tell
and we do not have enough information to decide now.
```

The idea it that relevant parts of the short-term database will be transferred
periodically to the long-term one.
This is the architecture and this is what should be decided now.
The choice between XML/RDBMS is design/implementation and, as such, is now TBD.
This documents presents proposals, but clearly marks them just a proposals that
need more investigation.
Otherwise it would be a design document, not an Architecture.

BGU 2001-03-29: No change to document to keep it light

---091------------------
(BrianGlendenning)
p.23 2.8 point 5
> Now we need a three-level scheme! Hopefully the "pipe" mechanisms (not
> described yet) could provide this "buffering" consistently.

ACCEPTED
This is how it will be.,

BGU 2001-03-29: Note about data channel added

---092------------------
(JimPisano)
p. 23 Diagram
> Callbacks are again described here. My previous discussion of command callbacks
> apply here. Just use the OS & spawn some threads instead of implementing
> a custom callback mechanism.

See other comments related to callbacks

BGU 2001-03-29: No change to document to keep it light

---093------------------
(RonHeald)
p.23 2.8.5
> It is not clear that logs need long term storage (except things that go
> into the data archive).
>
CLARIFICATION
6.2.2 defines persistency for logs.
The exception might be debugging logs. All the rest will go into some data archive, not necessarily for astronomers, possibly just for engineers.

BGU 2001-03-29: No change to document to keep it light

---094------------------
(RonHeald)
p.24 2.8.7
> There must be COTS tools available for this?

```
>
ACCEPTED
If we use XML (or an RDBMS) as foreseen, we will most probably just
have to take
and configure some Java XML (or database table) browser.
The very basic browsing capabilities are already, for example, in MS-
Explorer 5.
Point 2.8.7 just states the features we need, nothing about the
implementation.
Point 2.8.8 suggests already to use COTS tools (that for sure will have
to be packaged and
configured according to our needs).


BGU 2001-03-29: No change to document to keep it light

---095------------------
(RonHeald)
p.24 2.8.11
> What's wrong with compile time switches for debugging?
>
CLARIFICATION
Debugging proves useful when there are problems, and problems happen
when least expected.  It makes debugging and maintaining a running
system easier.
More over, run-time debug switches are useful when debugging client on
stable
server. It is nice to get more debug output from server without the
need to
restart it.


BGU 2001-03-29: No change to document to keep it light
```

2.10 Error System

2.10.1 Comments to Issue:2.0/Prep.1

```
---148--------------------------------------------------------------------
----
(Ron Heald)
p31, S2.10: From the discussion of the error system at the Joint
Software Meeting I had the idea that the error stack was passed only
through exceptions.  However, the document seems to contradict this.
Section 2.10.1 says "errors can be propagated through the call chain
...", and section 2.10.5 says "errors are propagated via completion
codes ...".  Which is it?  IMO, it should be only through exceptions.


REPLY: The error stack is passed through exceptions. The exception is
passed back through the call chain and can be catched at any level.
When it is catched the application can add more information on the
error stack or just throw an exception again with the same content.

BGU 2001-07-05: Document modified to clarify this

---149--------------------------------------------------------------------
----
(Ron Heald)
p.31 2.10  The authors admit the usefulness of exceptions, but are
```

unwilling to abandon the VLT "error stack" mechanism, and so end up having both. Project software people have clearly stated that they do not want an error stack because of its bulkiness and complexity. It should NOT be used.

Although there are still some "startup" problems when using C++ exceptions, these are temporary and exceptions are clearly the future direction. Our biggest problem with exceptions is their use under VxWorks. This is because VxWorks uses an obsolete version of the GNU compiler. These can be fixed either with the new release of Tornado II or by upgrading the GNU compiler ourselves. CORBA exceptions are part of the standard and should be used to their fullest.

REPLY: In Grenoble it was decided that the error stack will be used. It is not necessary to add information in the stack at each level of the call chain, but only where it is useful.

BGU 2001-07-05: Document modified to clarify this

---150---

(Jim Pisano)
p.32 2.10.2 "Not all levels in the stack trace need to be propagated, but only
the ones that provide useful information." Seems very arbitrary, will there be
points in the call stack where this determination is made? E.g., A() calls B()
calls C() calls D(). D() returns an error which is propagated back through C()
and B() to A(). Does C() or B() have the option to discard error stack info?
What criteria does it use. This seems all very vague to me.

REPLY: If the error is the result of a command then the error is passed back
to the origin of the command (through exceptions). What is meant is that
not each level has to catch the exception or add additional information
to the error stack. I think the text need to express this better.

BGU 2001-07-05: Document modified to clarify this

---151---

(Joe Schwartz)
p. 33, A note on the diagram says "See 'Class Diagram Logging Architecture", so it would be nice to have such a diagram--is it the one on p. 29? If so, it needs a name.

ACCEPTED: We will add a caption to all diagrams and add proper reference here.
This comes out of the Rose model and should have been modified for this document. Here we refer to the diagram on page 29.

---152---

(Brian Glendenning)

p.33 s2.10.5 Why not just say that the baseline design is to throw an exception of a standard type. When an exception is caught is by definition
an interesting place, and the standard exception type can have information
about the catching location embedded within it before it's rethrown.

REPLY: Yes, the baseline design is to throw an exception with the error stack embedded within it. You are right we should stress this more.

BGU 2001-07-05: Document modified to clarify this

---153---
(Martin Pokorny)
p.33 s2.10.7 Another application for ACS.

REPLY: Yes, ACS will supply an error display and a log monitor to browse
through the logging archive.
This is also an ACS requirement, see AD01.

BGU 2001-07-05: Noting to be done. Clarified in reply

---154---
(Jim Pisano)
p.34 2.10.10 "This would prevent building a complete trace of the call stack,
unless strict c<oding rules are defined.". Are you advocating that these code
rules be defined or is this TBD? I would end the sentence after "call stack" or
state that all throw()s must have a default catch() thereby eliminating the
entire sentence.

REPLY: The purpose of the error system is not to supply a complete trace of
the call stack, but to provide good tracing and diagnostic of the errors.
This will of course not be better that what is designed into the applications
by those who design and implement them. ACS just provides the mechanism.
The error stack could be used to include the complete call stack, but we
would not recommend that because to many entries in it makes it more difficult to filter out the relevant ones.

BGU 2001-07-05: Document modified to clarify this

2.10.2 Comments to Issue:1.1/Prep.2

---096---
(JimPisano)
p. 24 Error System

> This is the VLT error stack system which, as I understand, is not
desired.
> Simpler alternatives should be investigated.
>

(RonHeald)
p.24 2.9
> Important comment: The described error stack comes from the VLT. At
the
> ACS technical requirements review it was agreed that this is NOT
> desired. Please consider this architecture, I'm sure there's
something
> simpler. An example of a successful system is the VLBA where
functions
> return either "Ok" or an error message string.
>

TO BE DISCUSSED AT REVIEW
At the ACS TR review, there was some confusion between errors and
alarms.
The OVRO like alarm system will be implemented.
There is a requirement on tracing errors, which supports the idea of
the error
stack. ESO experience is that the error stack system is by far superior
to a
basic error reply and we consider it an essential debugging tool, so
why
should it be dropped?

Anyway, a simple error construct will be built in the first phase.
This will be possibly re-discussed later for another release.

BGU 2001-03-30: To be discussed at review

---097------------------
(BrianGlendenning)
p.25 2.9 points 5 & 8
> I don't understand point 5, and would like to see an example of the
kind of
> thing that is envisaged for point 8.
ACCEPTED
We will clarify the points in the document and add examples.
We are going to work on that in the coming weeks

BGU 2001-03-30: Point 5 clarified, Explained error definition in
ACSArchitectureNotes.doc

---099------------------
(RonHeald)
p.25 2.9.8
> Please no separate error files (like VMS). They always become
> hopelessly out of sync. Keep error messages with the code that
> generates them.
>
CLARIFICATION
This is most probably a misunderstanding as in other comments.
Error documentation is written together with the code, but generated
'a la Javadoc' in separate help XML/HTML files.

It is not reasonable to keep error documentation online inside the code.

BGU 2001-03-30: Explained error definition in ACSArchitectureNotes.doc

2.11 Alarm System

2.11.1 Comments to Issue:2.0/Prep.1

---155---
(Martin Pokorny)
p.34 s2.11.1 How do (polled) HW properties trigger alarms?

REPLY: Polled properties can trigger alarms when their get method is called.
If it is a M&C point this will be done periodically by the system.
See also comment 93.

BGU 2001-07-05: Document not changed

---156---
(Dirk Muders)
p34, s2.11.2:

Why two levels ? This seem a little coarse. There could be a whole range of
priorities.

(Ron Heald)
p34, S2.11.2: Maybe better names are marginal and out-of-range? (I'm not sure I have the correct meaning of this section.)

(Steve Scott)
p. 34 2.11.2: Do "high and low" mean "alert and warning"? Maybe some clarification.

REPLY: From your comment and from Ron and Steve we need mor than just out-of-range alarm which is provided with ACS 1.0. I think you are right
we should provide as Ron proposes "marginal" and "out-of-range" alarms or
as Steve proposes "alert" and "warning".

BGU 2001-07-05: Section modified

---157---
(Steve Scott)
p. 34 2.11.5: An event on every transition may be just too much for a value hovering around a threshold. A settable hysteresis

time is needed.

REPLY: There is no hysteresis time foreseen. ACS provides hysteresis levels.
Our experience is that should be enough.

BGU 2001-07-05: Nothing to be done. Clarified in reply

---158---
(Jim Pisano)
p.34 2.11.5 I assume that what you mean by "change its status" is when the
alarm is set or cleared.

REPLY: Correct. To be clarified in the document.

BGU 2001-07-05: Done

---159---
(Jim Pisano)
p.34 2.11.6 1-st bullet. I'm confused. It seems like an alarm is set only when
it exceeds its minimum or maximum value (from 2.11.1) & consequently sets an
event (from 2.11.5) - I would call this edge-triggered alarms. How can there be
multiple occurrences if there is only one event per setting & clearing? Does
the alarm send an event each time it's updated and exceeds the min. or
max. value - I would call this a level-triggered alarm? If so, then the text
does not describe alarms as level-triggered, but edge-triggered.

ACCEPTED: You are right. ACS 1.0 will supply this type of "simple" alarm
system. Later alarm objects will be created when an alarm occurs which will
include the properties described in 2.11.6. It will be made clearer in this
section.

BGU 2001-07-05: Done

---160---
(Jim Pisano)
p.34 2.11.6 3-rd bullet. The client creates the alarm (2.11.3) and the property
gets a reference to the alarm. Why does the Property need to hold the status of
the user acknowledge - doesn't the alarm have this? I would think that user
acknowledges the alarm (on the client side) and since the Property has a
reference to the alarm, the acknowledgement is there.

(Martin Pokorny)
p.34 s2.11.6 Bullet item 3: What does this mean?

REPLY: This is a property of the alarm object. It is an indication if the
alarm is acknowledged or not.

BGU 2001-07-05: Text clarified

---161---
(Steve Scott)
p.34 The relationship of alarms should be used to prevent "alarm cascades". This can be done with a graph that implements the relationships.

REPLY: We'd be happy to get some proposals from you. You are supposed to be
the expert in alarms.

BGU 2001-07-05: Nothing to be done now

2.11.2 Comments to Issue:1.1/Prep.2

---098---
(BrianGlendenning)
p.25 2.10
> I understand that there a COTS alarm systems including OVRO style dependency
> graphs. Could we spend some money and save some time?
ACCEPTED
Yes we should.
If you have information send it to us.
The idea is currently to start from OVRO, because we do not know of others, but
someone (Steve?) will get the job of doing the investigation and come with a
proposal.

BGU 2001-04-02: Alarm system changed.
---100---
(BrianGlendenning)
p.26 2.10 point 2
> Do you mean that the alarm class *is* a DO or *associated* w/ a DO?

ACCEPTED
This part has to be clarified.
The problem is that it not fully clear to us as well.

In a first approximation, Alarms are just like Monitors: they are CORBA Objects
associate to properties of DOs.
If we want to be notified of alarms conditions on a property, we ask it to give
us an alarm instance and we attach a callback to that.
This is fine and it what ANKA has (and we get for ACS 0.0)

The things are more complex with the OVRO scheme.
Here we have hierarchical alarms, that are the result of the
composition
of lower level alarms and, as such, do not correspond to any already
defined
property.
Here we see two possibilities:
 - We create logical properties that handle high level alarms, and we
use
   Alarms on that as in the previous case
 - We create special Alarm DO classes whose purpose is to handle the
hierarchy
   of alarms.

This last alternative is what is at the base of 2.10.2.

Most of this is still TBD and experience/prototyping will tell.
Text will be modified to make clear that many things are still TBD.

BGU 2001-03-30: Alarm system completely rewritten

---101------------------
(JimPisano)
p. 26 2.10.3 What are the purposes of "alarm sub classes"? These
complicate
> the design by providing functionality that is undescribed, e.g.,
"general
> evaluation fault tree analysis/graph Alarms" -- totally confusing.
>
ACCEPTED
The sub-classes will be left TBD for time being.
The availability of a COTS for alarms will also be looked into, to have
a
look and feel similar to the OVRO system.

BGU 2001-03-30: Alarm system completely rewritten

---102------------------
(JimPisano)
p. 26 2.10.5 Your example of correlator blanking is not required.
Blanking
> can always be done after data sets and monitor logs are written to
the
> Data Collector's archive. It may be convenient to blank data at the
correlator
> to minimize data rates, but not necessary as you describe it. If
blanking is
> handled at the correlator, the correlator holds a queue of sub-
integrations
> that are time-tagged & can be blanked via time-tagged commands from
the
> ACC before the sub-integrations are accumulated into a single
integration
> & sent to the Data Collector.
>
ACCEPTED, example will be dropped.
This was just a cut & past of a comment I received at the previous
issue

```
of the document.

BGU 2001-03-26: done

---103------------------
(RonHeald)
p.26 2.10.3
> "Alarm sub classes" I don't understand what this is about?
>
CLARIFICATION

This will be dropped and left TBD.
Anyway, just to clarify, you have different types of alarms:

1) is binary two level alarms: it has just an OK value and an ALARM
value.
   A typical example is a limit switch. 0 is OK, 1 means alarm, switch
activated.
2) is a 5 levels alarm, i.e. you need 5 characteristics to define it.
      a1 <= x <= b1                        means OK
      a2 <= x < a1  || b1 < x <= b2        means warning
      x < a2 || b2 < x                     means error

Forget for the time being 3) (it is some complex object, that would
need much
more details and explanation. It builds a composite alarm state based
on the state
of lower level alarms.

1) and 2) are very common cases and the idea here was to provide some
base classes
that implement them.


BGU 2001-03-30: Alarm system completely rewritten
```

2.12 Time System

2.12.1 Comments to Issue:2.0/Prep.1

```
---162--------------------------------------------------------------
----
(Jim Pisano)
p.35 2.12.1 "The hardware devices (on the CAN bus) ...", should be "The
synchronized hardware devices ..." as some devices will not be
synchronized and
others (the test correlator) will not be on a CAN bus.


ACCEPTED

BGU 2001-07-05: Done

---163--------------------------------------------------------------
----
(Martin Pokorny)
p.35 s2.12.1 Hardware devices being "told to start" is rather
```

vague. This is a difficult paragraph because it tries to summarize too much too briefly. Also, it doesn't directly say anything about the Time System design, but rather something about it's requirements. (As does the following s2.12.2).

(Brian Glendenning)
p.35 s2.12.2 Delete "smart." The level of local intelligence in a device
isn't correlated with requiring precise timing.

(Mick Brooks)
p35, S2.12.1, Second sentence: Only some CAN devices will be controlled in
this fashion, not all of them as is implied here.
Examples of time controlled devices are the ACU, the nutator, the second
LO. Devices not operating in this fashion include
the compressor, the front-ends and the data transmission system.

REPLY: I think we should remove the last sentence of 2.12.1 as well as 2.12.2
as this is not really part of the architecture of ACS.

BGU 2001-07-05: Done

---164----------
(Brian Glendenning)
p.35 s2.12.4.1 We need to handle more than ISO string formats since users
like to be able to type in times in a variety of formats.

TO BE DISCUSSED AT REVIEW:
 I'm not so sure about this. If you allow all types of formats
you will introduce a confusion. Just compare the US and the European time
format. ISO format is unabigouos and should be the preferred format.

BGU 2001-07-05: Added in text that other formats can be added later for user
input

---165----------
(Jim Pisano)
p.35 2.12.4.3.3 What is the resolution of this phasing in the 2-nd bullet?

REPLY: Theoretically this would be the resolution of the time system (100 ns),
but in practice the accuracy will be limited by hardware.

BGU 2001-07-05: Nothing to be done. Clarified in reply

2.12.2 Comments to Issue:1.1/Prep.2

```
---104------------------
(JimPisano)
p. 27 2.11.1 Time tagging of commands could occur at *any* timing event
in
> the future, not just the next one. I thought that we all agreed on
time-
> tagged commanding not just-in-time commanding?
>
CLARIFICATION
Time tagging can indeed happen at any time. There will obviously be an
accuracy requirement when the tagging is done, typically at the data
source,
so that, whatever happens afterwards, the time stamp is as accurate as
available at the source computer.


BGU 2001-03-29: Clarified

---105------------------
(JimPisano)
p. 27 2.11.2 Array Standard Time - array time is the counting of 48 ms
> timing events which should be used throughout the array devices.
There has
> been a lot of discussion of array time at NRAO (see ALMA memo 298)
and I feel
> that all devices should use array time only w/ conversion to UTC at
the
> Data Collector for the data processing stages. Also UTC is needed for
LST
> conversions, but again it should remain local to that computing
system &
> not necessarily distributed around. I see no need to distribute 2
time
> systems & worry about their synchronization & accuracy.

ACCEPTED
Yes, ACS is not re-inventing the Time system, but supporting its use in
a
uniform way.  It will be array time.

BGU 2001-03-26: done

> p 28 2.12.3 The SSR has proposed an observing scripting language
which is *not*
> OO, but procedural. This conflicts w/ the first sentence.
>
ACCEPTED
ACS is dealing with a technical scripting language (req.5.2.1). The
relation
to SSR requirements for a user scripting language is still TBD.

GCH 2001-03-26: Point 2.13.3 removed

---106------------------
(MickBrooks)
p.27 Section 2.11 You should perhaps reference Larry's proposed time
system
distribution: ALMA Memo #298, Timing and Synchronization, Larry R.
D'Addario, 13 March 2000.
```

(http://www.alma.nrao.edu/memos/html-memos/alma298/memo298.pdf).

ACCEPTED

BGU 2001-03-26: done

---107------------------
(RobertLucas)
p 27:
> 2.11.2 Array time becomes absolute time one line below?

ACCEPTED:
2.11.3.2 should spell array time, not absolute time

BGU 2001-03-26: done

> 2.11.5 Anything at this time resolution level is implemented by the
> hardware (I hope !)

CLARIFICATION:
This is a cut & paste of a generic comment we got for the previous issue
of the document from Brian. This is all TDB.

BGU 2001-0329: This point has been moved to ACSArchitectureNotes.doc

---108------------------
(RonHeald)
p.27 2.11
> We need more experience in this area. Almost everything said here is
> TBD.
>
ACCEPTED
Right, you might help us on this.
We are looking at TAO Time Service.

BGU 2001-03-29: Nothing added to keep document light

2.13 Scripting Support

2.13.1 Comments to Issue:2.0/Prep.1

---166--

(Mick Brooks)
p36, S2.13: It is implied that a single scripting language will be provided
by ACS and that Tcl/Tk and Python are the "candidates"
Is this the case, or should ACS support multiple scripting languages?

REPLY: This should be decided by the ALMA sw group. We can only make proposals.
ACS will provide support for the chosen script language(s).

BGU 2001-07-06: Nothing to be done now

---167---
(Joe Schwartz)
p. 36, 2.13.2.2, It is proposed to use the ACS scripting language for Sys Admin, "replacing shell script languages". This implies throwing away the many shell scripts that come with or have been developed over the years for UNIX system admin. I think that this is not a good idea.

ACCEPTED: There is no intention to rewrite existing shell scripts. This is only valid for new scripts.

BGU 2001-07-06: Document modified

---168---
(Dirk Muders)
p36, s2.13.4:

These are all requirements.

REPLY: Yes, these are requirements for the scripting language. They are listed here to justify the choice in 2.13.6. We will put them as
subpoints to 2.13.6.

BGU 2001-07-06: Done

---169---
(Brian Glendenning)
p.36 s2.13.4.4 What does this mean?

REPLY: Yes, this is too vague. It will be removed.

BGU 2001-07-06: Done

---170---
(Joe Schwartz)
p. 37, 2.13.5.4, AD01 gives rapid prototyping as a capability that ACS is *required* to support. Has it been downgraded to "nice to have" (2.13.5)?

ACCEPTED: Correct. It will be made to a requirement.

BGU 2001-07-06: Done

---171---
(Martin Pokorny)
p.37 s2.13.6 Why does "(combat)" appear here?

REPLY: Tcl/TK need combat to talk to CORBA.

BGU 2001-07-06: Document clarified

---172---

```
(Ron Heald)
p37, S2.13.6: Once again I must point out that Glish will be part of
the
system since it comes with AIPS++ that will be used for data reduction.
I feel you should at least mention this fact.


REJECTED: It has been decided that Tcl/Tk and Python are the candidates
for
the ACS scripting language and ACS will only support these two
languages
until something else is decided.

BGU 2001-07-06: Nothing to be done
```

2.13.2 Comments to Issue:1.1/Prep.2

```
---109------------------
(BrianGlendenning)
p.28 2.12 point 3
> Important: I do not believe that it is appropriate to constrain the
observing
> scripting language at present. Delete this point and subpoint 5.3.

CLARIFICATION
We do not see how the given text constraints the observing scripting
language.
Point 2.12.3 says "should" and point 2.12.5.3 says TBD.
We think it is correct to mention ideas/principles that would help in
keeping the system simple and more maintainable, leaving any decision
after more
detailed analysis and prototyping.
If you insist, the text can be anyway removed.


BGU 2001-03-30: Removed it

---110------------------
(BrianGlendenning)
p.28 2.12  point 5.6
> It would be quite useful if you could implement server objects in the
> scripting language as well.

REJECTED
We think it is not good to program in a scripting language
There should be a clear distinction between what you do with scripts
and what you
to with programming languages.
Already you have full freedom with Java.


BGU 2001-03-30: Improved decsription

---111------------------
(BrianGlendenning)
p.28 2.12 point 5.7
> For a technical ACS scripting language I do not consider it to be a
> requirement that it be easy to use by non-programmers. Languages
often trade
```

> off ease-of-use for maintainability. I would rather have a
maintainable
> language.

(JimPisano)
p. 28 2.12.5.7 Are Tcl/Tk & JPython easy for non-programmers? Non-
programmers
> should use a GUI which totally hides the implementing script
language.
>

REJECTED
The scripting language should be used by non-programmers (operators,
system
administrators, staff astronomers making prototypes.....), otherwise
they will come
with yet another scripting language.
You already have Java and C++ as "real" languages, what is the purpose
in
having a third one!

But these non-programmers are not observing astronomers.
We are not talking here of "scripting observing".
As Jim says, that should be a GUI that totally hides the implementing
script
language.

The non-programmers are "hybrid animals", if I can use this expression.
They fill in the intermediate role between software developers (the
programmers)
and the observing astronomers. There will be a lot of them around.
They are not supposed to be Java or C++ programmers, but are the first
line
of support and will be involved in writing a lot of simple
applications.

Again: there must be a distinction between what you do with scripts
and what you
to with programming languages.
Scripts has to be used for small prototypes and small applications,
where
maintainability is not a major issue.

BGU 2001-03-30: Added note in ACSArchitectureNotes.doc

---112------------------
(RobertLucas)
p 28: 2.12.3 At which stage is the choice between Tcl and JPython
planned ? The interoperability between scripts and GUIs is not
described.

CLARIFICATION
This will probably come in one year from now.
Both scripting languages provide support for developing simple GUIs.
These can be used for simple applications and prototypes, but major
GUIs
should be developed in Java.

BGU 2001-04-02: No change in document

```
---113------------------
(RonHeald)
p.28 2.12
> Glish should be mentioned here as it will be part of the system (via
> AIPS++) in any event.
>
REJECTED
Glish is not the list of the TECHNICAL scripting languages to be
investigated given
in AD01 10.5.6.
We do not have enough resources for these two, we cannot add a third
one.

BGU 2001-03-30: No change in document
```

2.14 Graphical User Interface Tools

2.14.1 Comments to Issue:2.0/Prep.1

```
---173------------------------------------------------------------------
----
(Ron Heald)
p37, S2.14.2: ACS is not a software distribution package and should not
include GUI builders.  In fact, this is illegal since such a
distribution violates the Visual Age license agreement.  The same can
be
said for the VxWorks development package that is currently distributed
with ACS.

TO BE DISCUSSED AT REVIEW:
 We think it is practical to reduce the number of packages to install.
Furthermore if you have several products in one package they can be
configured
such that you are sure that they are compatible and will work together.
VxWorks though is not distributed with ACS. We have supported
installation
of VxWorks at some site which already have a VxWorks license. License
agreements are anyhow not part of the architecture, but will of course
have to be solved with the suppliers.

BGU 2001-07-06: Nothing to be done. Clarified at IRAM meeting

---174------------------------------------------------------------------
----
(Joe Schwartz)
p. 38, 2.14.2.2, Have prospective users been consulted concerning the
use of the ANKA Java Beans as the core palette for ACS, i.e., is this
a choice that the ALMA observers and staff scientists are expected to
accept?


ACCEPTED: A good point. We did not do it up to now. But we think it is
too
early in this phase. Using Java beans it is also simple to reuse and
implement new widgets as needed.

BGU 2001-07-06: Nothing to be done now.

---175------------------------------------------------------------------
----
```

(Joe Schwartz)
p. 38, 2.14.7, Once again I'm confused about the intentions with regard to Lab View. Its I/F is not based on Java, but rather on the G Programming Language (about which I know less than epsilon) and it's not particularly portable. Is its use to be restricted to the TICS or will it be used for ALMA engineering support throughout the project?

REPLY: ACS is not supporting LabView as such. LabView has been foreseen in
the project for some specific purposes. Perhaps we should remove all references to LabView in the document as it seems to create a lot of confusion.

BGU 2001-07-06: Done

2.14.2 Comments to Issue:1.1/Prep.2

NONE

2.15 Device Drivers

2.15.1 Comments to Issue:2.0/Prep.1

---176---
(Martin Pokorny)
p.38 s2.15 Device drivers should not be a part of ACS, in my mind (where do they fit in the package diagram?). Most of the paragraphs in this section are, in any case, not appropriate for a SW architecture document. In particular, 2.15.4 through 2.15.6 are somewhat gratuitous.

REPLY: Drivers are not necessarily part of ACS but we think it is good to packed several products together to simplify distribution and installtion. See also reply to comment 213.
2.15.4 to 2.15.6 could be removed.

BGU 2001-07-06: Done

---177---
(Brian Glendenning)
p.38 s2.15.2 Why not just state that we'll use Linux on the central computers.

ACCEPTED

BGU 2001-07-06: Document modified

---178---
(Steve Scott)
p.38 There is not much information on the LabView interfaces. Is it too early to start enumerating these?

REPLY: We will remove this section as LabView is not part of ACS.

BGU 2001-07-06: Done

---179----------
(Mick Brooks)
p39, S2.15.3: This library is not merely envisaged. It exists.

ACCEPTED: We will change the text.

BGU 2001-07-06: Done

---180----------
(Mick Brooks)
p39, S2.15.4: I do not see how this statement follows from the previous one. I suggest that the statement be removed.

(Steve Scott)
p.39 2.15.4: What does "This" refer to?

ACCEPTED: We will remove this section.

BGU 2001-07-06: Done

---181----------
(Dirk Muders)
p39, s2.15.5, 2.15.6:

What does this have to do with ACS architecture ?

ACCEPTED: They will be removed.

BGU 2001-07-06: Done

---182----------
(Mick Brooks)
p39, S2.15.5: In this case I think that what you are talking about is one
method of testing hardware CAN interfaces. I do
not believe that it has anything to do with ACS and can be safely removed.

ACCEPTED: 2.15.5 will be removed.

BGU 2001-07-06: Done

2.15.2 Comments to Issue:1.1/Prep.2

---115----------
(MickBrooks)

p.29 Section 2.14, point 5. I don't know that LabView will be fast enough to
truly emulate many of the hardware subsystems on CAN. You might mention NRAO's plans to develop an ACU simulator on a Linux PC. The CAN driver for
this is currently under development.

ACCEPTED

BGU 2001-03-26: done

2.16 High Level Application Framework

2.16.1 Comments to Issue:2.0/Prep.1

---183---
(Joe Schwartz)
p. 39, 2.16, It sounds as though the application framework will only be available towards the end of the ACS development cycle. If this is so,then most applications won't be able to use it, since their development will start before it is ready. Certainly it doesn't seem that TICS can use it. Do you expect that the dataflow software will use it? Or is it only intended for the control software? The discussion here is so general that I can't really get a feel for what it would do for me as a developer.

REPLY: I agree with you that this should come as early as possible. The minimum we should have is a DO template. We have already started with
a simple device and it should be extended with certain patterns like implementation of methods, threading, clients using AMI etc.

BGU 2001-07-09: Nothing to be done. Clarified in reply

---184---
(Steve Scott)
p.39 2.16.2: What does "This" refer to?

REPLY: This should be replaced with the framework.

BGU 2001-07-09: Done

---185---
(Brian Glendenning)
p.39 s2.16.2.1 TBD whether ACS will have a standard state machine. What does
it mean to say it will implement UML state machines?

ACCEPTED: State machines will be implemented by TICS and later integrated
into ACS.

BGU 2001-07-09: Document modified

```
---186--------------------------------------------------------------------
----
(Brian Glendenning)
p.39 s2.16.6 Delete as these are high-level analysis issues and should
not
be imposed from the bottom.


ACCEPTED

BGU 2001-07-09: Done

---187--------------------------------------------------------------------
----
(Brian Glendenning)
p.39 s2.16.7 How are complex dependencies which might only be known by
higher-level software catered for? At a minimum this needs to be
customizable. (I would eliminate this point in favor of one that states
that
ACS will provide facilities for starting DO's without complex creation
dependencies).


(Mick Brooks)
p40, S2.16.7: Again the discussion of DO life-cycles recurs, this time
in
the context of application frameworks where it
is clearly out of place. I suggest a separate section to discuss the
concept of DO life-cycle.


REPLY: We will remove this section here and expand the discussion of DO
life-cycles in paragraph 2.2
```

2.16.2 Comments to Issue:1.1/Prep.2

```
---114------------------
(BrianGlendenning)
p.29 2.15
> This is ambitious. I would emphasize the portions having to do with
> startup/shutdown and interactive services (e.g. logging).

ACCEPTED


BGU 2001-03-30: Added note on ACS start-up/shut-down, life-time
management

---116------------------
(RonHeald)
p.29 2.15
> This is too ambitious.  Please let experience tell us if we need such
a
> thing.
>
ACCEPTED
This is ambitious and, although TBD now, has to planned in advance.
We cannot ignore ESO experience, that tells that such things are
extremely useful.
```

Actually, the problem there has been to develop some of these things too late,
loosing some of the advantages. This will anyhow not all be done in one go.

BGU 2001-04-02: no change of the document

---117------------------
(BrianGlendenning)
p.30 2.15 point 5
> Delete this "advice" section for a more appropriate document.

(RobertLucas)
p 30: 2.15 It's not clear whether this belongs here? what is the status of these issues ?

ACCEPTED

BGU 2001-03-26: Section removed

---118------------------
(MickBrooks)
p.30 point 5: This decision has been taken: there will be no real time control from the center of the array to the antennas. You might want to still mention the use of ATM QOS for future unforeseen requirements. The
time-tagged discussion can now be written as resolved, but still you should
include the reasons you mention. Subpoint 7 is kind of moot considering neither Ethernet or reflective memory will run over 30km.

(RonHeald)
p.30 2.15.5
> It's now been decided ACC is a general-purpose machine and there will be
> no real-time between it and the antennas. This section can be removed.
>

ACCEPTED

BGU 2001-03-26: Section removed

2.17 System Configuration Issues

2.17.1 Comments to Issue:2.0/Prep.1

---188--

(Martin Pokorny)
p.40 s2.17.1 This sounds largely like an apology for Java; is it appropriate here? In the third paragraph, the first sentence states that Java "should mainly be used for GUIs", and the last sentence states that C++ should be used in demanding applications, and Java "everywhere else". Unless GUIs are the only non-demanding applications, I'm not sure which statement is preferred. Finally, the last paragraph reaches too far below the realm of ACS to be justified in this document.

REJECTED: I don't see your point. You've cut the first sentence too early. It
reads: "should mainly be used for GUIs and for high level applications with
no strict performance requirements". How can this contradict with the last
sentence?
The last paragraph describes how ACS is implemented to guarantee portability
which very well fits into this document.

BGU 2001-07-09: Nothing to be done

---189----------
(Jim Pisano)
P.40 2.17.3 Grammar - "warranties" -> "guaranties"

ACCEPTED

BGU 2001-07-09: Done

---190----------
(Martin Pokorny)
p.41 s2.17.3 Last sentence, first paragraph: what does it mean? Last paragraph: not appropriate for ACS architecture, is it (this is more of a system configuration, or software engineering issue)?

(Ron Heald)
p41, S2.17.3: Again, ACS is not a software distribution package. It should not include third party software as described here. It's fine to
mention what outside software (including the particular version) with which ACS has been shown to work properly.

TO BE DISCUSSED AT REVIEW:
We think it is easier to distribute and install only a few packages than having to install every individual package you need. Advantages
are:

- you have only one installation procedure instead of many probably different written with different philosophy
- you are sure to install the correct versions which are configured to work together

Again libraries are an ACS requirement, see AD01.

BGU 2001-07-09: Nothing to be done here. Clarified at IRAM meeting

---191----------
(Brian Glendenning)
p.41 s2.17.4 image pipeline -> science

ACCEPTED

BGU 2001-07-09: Done

```
---192--------------------------------------------------------------------
----
(Martin Pokorny)
p.41 s2.17.5 Again, I can't see how this section belongs in an
architecture document.


REPLY: We wanted to show what type of prototype set-ups we are using.
It does not change the architecture and could be removed.

BGU 2001-07-09: Done

---193--------------------------------------------------------------------
----
(Mick Brooks)
p41, S2.17.5: This sentence seems to have been copied from S2.15.5 and
does
not make clear whether the LabView PC
is simulating a slave node (a bad idea in my view) or simulating a
master
node (not particularly useful from the point of
view of ACS). In either case I think the sentence can be removed.


ACCEPTED: We will remove it

BGU 2001-07-09: Done
```

2.17.2 Comments to Issue:1.1/Prep.2

```
---119------------------
(RonHeald)
p.30 2.16
> As you say these are SE issues, they don't belong here.
>

ACCEPTED
This (and other sections) will be removed from this document,
but not lost until they will find their place in the proper documents.


BGU 2001-03-30: Moved to ACSArchitectureNotes.doc
```

2.18 System Management Issues

2.18.1 Comments to Issue:2.0/Prep.1

```
---194--------------------------------------------------------------------
----
(Ron Heald)
p42, S3: It should be mentioned that no root access is required for the
installation and execution of ACS.  This was an accepted comment from
the previous round that somehow was lost???


REJECTED: This is a design issue and will be handled there.
```

```
BGU 2001-07-09: Nothing to be done
```

2.18.2 Comments to Issue:1.1/Prep.2

```
---120------------------
(BrianGlendenning)
p.31  2.17
> I would like to add that no root access should be required.

ACCEPTED

BGU 2001-03-26: done
```

2.19 General Issues

2.19.1 Comments to Issue:1.1/Prep.2

```
---121------------------
(BrianGlendenning)
p.31 2.18.1
> The scripting language might fit the bill here as well.

CLARIFICATION
The scripting language should be mentioned here, but we do not want to
make
control in a scripting language!!!

We are back again to the difference between what you do with scripts
and what you
to with programming languages.
Scripts are for prototypes and for the very top level, not for
applications that
are really part of the control system.

GCH 2001-03-26: Added scripting language (Python/TCL) + ACE for
standards.

---122------------------
(BrianGlendenning)
p.31 2.18.5
> Our mission is to build ACS for ALMA.  We are all at institutes with
> telescopes that will need new software in future - there's no need to
mention
> any particular future application.

ACCEPTED
This is a remnant from the very first version and has to be deleted.

BGU 2001-03-26: done

---123------------------
(RonHeald)
p.31 2.18.1
> "Unit" should be "Unix"
>
ACCEPTED

BGU 2001-03-26: done
```

```
---124------------------
(MickBrooks)
p.32 Section 2.18.6: Again, I'm not sure that the LabView PC will be
much of
a CAN node simulator. Maybe you should mention the availability of the
CAN
prototype board (AMBSI). I'll send you one to play with real soon.
After
all, Birger has already built a CORBA DO for one of these.

ACCEPTED

BGU 2001-03-26: done

---125------------------
(RonHeald)
p.32,33
> Why are these pages partially blank?
>
CLARIFICATION
Because the next section is a chapter, and chapters start on a new
page.

GCH 2001-03-26: With new ALMA template no more page breaks

---126------------------
(RonHeald)
p.32 2.18.3
> Why would ACS contain SLA,AIPS++, etc.? Why not get them from their
> source?

CLARIFICATION
In order to package them in a coherent distribution.
We have to make sure that all sites involved in the project have the
same
versions of the packages, with the same patches (if needed) and that
they
are built with the same options and installed in the same directory
structures.
Sometimes it is also necessary to make minor changes.

GCH 2001-03-26: DONE

---127--------------------------------------------
(RonHeald)
p.32 2.18.6
> "asses" should be "assess"
>
ACCEEPTED

BGU 2001-03-26: done
```

2.20 Attributes

2.20.1 Comments to Issue:2.0/Prep.1

---195----------
(Steve Scott)
p.43 3.4: Performance is very important and there is nothing here that suggests this architecture will perform adequately. By now many simple prototypes have been done along with the KP demo, so why not give the results and extrapolate them to the full ALMA? For example, if you were running the full ALMA on today's comuputers, would the collective Data Channels take up 10% of one machine or 100% of 10 machines? This is not the time for hand waving - we should be pretty sure that this architecture will perform adequately before we are locked into it. Verification after ACS1.0 is too late - reasonable performance estimates can and should be done now.

ACCEPTED: You are right performance is important and we should add a section
on that. We intend to test performances after the release of ACS 1.0 and add
the results to the next release of this document.

BGU 2001-07-09: Extended performance section of chapter 4

---196----------
(Jim Pisano)
p.42 4.0 "ACS 1.0 will be used for the implementation of the Test Interferometer Control Software (TICS) and the [RD30] features ...". Remove
[RD30].

ACCEPTED

BGU 2001-07-09: Done

---197----------
(Ron Heald)
p43, S4, 2nd para: You may plan to use the SE process, but so far, it has not been done so far.

REPLY: We have done part of the SE process but not all due to constraint of
time and resouces. So to some extent you are right. We could have done better.

BGU 2001-07-09: Nothing to be done here. Clarification in reply

---198----------
(Mick Brooks)
p43, S4.1, First sentence: This repeats the first sentence of Section 4 and
should be removed.

ACCEPTED

BGU 2001-07-09: Done

2.20.2 Comments to Issue:1.1/Prep.2

NONE

2.21 Life cycle Aspects

2.21.1 Comments to Issue:1.1/Prep.2

NONE

2.22 Final sections

2.22.1 Comments to Issue:1.1/Prep.2

```
---128------------------
(BrianGlendenning)
p.38 6.1
> I presume that sequences of the  listed basic types will also be
provided.

ACCEPTED
In principle they should, but I fear we will not have time to do that.
We will add it to the list, with a disclaimer.

GCH 2001-03-26: Section removed.

---129------------------
(BrianGlendenning)
p.38 6.1
> Given  the emphasis that is being placed on XML I think it should be
> investigated here. This kind of application is the most natural fit
to XML. I
> would delete reference to XML in 6.7 awaiting this experience. Use an
RDBMS
> instead.

CLARIFICATION
The text in 6.1 spells: "an XML-based [configuration] database will be
investigated.
This was the plan and would match what you ask, but there are no
resources to allocate
to that, unless Fritz get some significative time allocated to the
project.

On the contrary, task 6.7 is much easier and give us a first taste on
XML.

The reason for the difference is that to implement 6.7 we just need to
WRITE XML and
we can browse logs with standard tools (MS Explorer), while to
implement a configuration
database we need both to READ/WRITE, also on the LCU. This requires
substantially more
work.

It will be one of the first things done after ACS 0.0
```

```
Please, notice also an error in 6.1: Characteristics implementation is
based
on VLT database, not on ANKA static database, since it has been easier
to
use the VLT database than to port ANKA static database from NT to
UNIX/VxWorks.

BGU 2001-03-26: changed to VLT DB

---130------------------
(MickBrooks)
p.38 Section 6:Probably doesn't belong in here, but do you have any
idea on
a possible release date for V0.0?

REPLY
First release of ACS 0.0 is foreseen for end September 2000 (joint
meeting)

GCH 2001-03-26: Section removed. ACS 0.0 released already
```

2.23 Pending comments from Issue:1.1/Prep.2

The following comment submitted for Issue:1.1/Prep.2 have been analyzed and discussed at the review meeting for Issue:2.0/Pre.1.

```
---004------------------
(RonHeald)
p.11 2.2.1
> Add "Attribute" for non-static data associated with a property.

TO BE DISCUSSED AT REVIEW
The difference here seems to be between compiled-in values (as you
argue
for) and initilised values.
Although the distinction between Attributes (read/write) and
Characteristic
(read-only) is nice in principle:
- the distinction is less unambiguous than it seems (e.g. you say that
defaults are read-only, ranges read-write; it could also be argued that
they
should
be the other way around).
- it does not seem correct to have half of the values hard-coded and
half-
 in a configuration database. This is also an arbitrary interpretation
of
 the ACS TR requirement 4.2.1 (DO information in Config. DB)
- the present system is simpler, as characteristics are all R/O and
  only properties can be R/W.
To make a distinction is OK, while to advocate a different structure
and behavior adds complexity for a feature of debatable value.

The proposal in 2.2.5.4 is to start simple with read-only
characteristics
and, if
experience shows a real need, make them read/write at a later time.

If this approach is accepted and as already requested in the past, the
definition
of Attribute should be removed from the ALMA Glossary.
```

GCH 2001-03-26: REMAINS OPEN FOR REVIEW

> 3. Provide properties with some type of "engineering support".
Perhaps
> when in engineering mode the property would accept input from
> engineering and ignore other input.
>
TO BE DISCUSSED AT REVIEW
The engineering access mode is meant to be a way to set a
value and to make it read-only. It is similar to changing the
permissions on a file. It is felt that some values might be
changed by software, and to debug/test things, we would like
to force the values.

Scenario:
  change permission to read-write
  set value
  change permission to read-only
  test, debug
  change permission to read-write

This can be implemented adding a Lock(TRUE|FALSE) method to Properties.

G.Chiozzi and other people at ESO are against such an "engineering
mode",
based on
the experience with the VLT system: such a concept has been foreseen,
implemented and
later on removed on request of the commissioning and operation teams.
It adds complexity and adds overhead to the
work of engineers, that are supposed (and suppose) to know what they
are doing.

The suggestion is propose the Lock() solution in the document, but give
low priority to the implementation so that we can first verify the need
for it.
Implemented in this way, it does not have an impact on other components
of
the system.

GCH 2001-03-26: REMAINS OPEN FOR REVIEW

BGU 2001-07-09: Nothing to be done here. Clarified at IRAM meeting

---015------------------

(RonHeald)
p.11 2.2.1.2.3
> Important comment: "quality" has no use in direct value retrieval;
> either you get the value or an error.
>

TO BE DISCUSSED AT REVIEW
The concept of quality is in the requirements.
To scrap quality would anyway simplify the system and we could probably
leave without:
just always return error whenever the quality would have been not OK.

We should really discuss if we should remove the concept at all.

In any case, if we decide to keep the concept, quality is also important
to pure soft properties, e.g. server down, LAN connection broken etc.
As Ron notices, it is normally with indirect readout that you might have old
obsolete values, without having an error, because the readout simply was stopped.

However there is a requirement for transparency (4.1.4) and applications
shall be the same for getting direct or indirect values. So structures must also be the same.

To have properties with/without quality adds complexity to the system.
If quality does not make any since for some specific property, then it is
always OK.

GCH 2001-03-26: For the time being, quality is removed for the Architecture.
   If we will decide that it is important is can be added later.
   The notes on Architecture document is used to keep track of this.
   The requirement will have to be removed as well.
   If accepted at the review, and SPR change request for the requirements
will be filed.

BGU 2001-09-07 Done

---022------------------
(RonHeald)
p.11 2.2.1.2.2
> Remove Property types strings, enum, binary, and sequence. Strings
> negate compile time checking. The others are not required.
>
TO BE DISCUSSED AT REVIEW
We think that strings and sequences are important and we would not remove them.
Enum are probably also very useful.
Binary are just suggested and experience will tell if they are necessary or not.
They are not for first implementation but it would be wrong to remove the TBD reference from the document.

Enum was to address the problem of bit twiddling magic constants, so a dictionary
associating names to bit patterns was envisioned.
For example, a motor might have 0x1 = "on", 0x2="off", 0x4="forward", 0x8="reverse", 0x10="slow", 0x20="fast". To set a motor to on, forward,
fast is
"on forward fast".
It might be that strings are really enums, so that might go.
Sequences seem that they would be useful sometime in the future.
For sake of implementation as fast as possible, sequences
and binary have low priority, strings are possibly enums,
and enums seem useful to have.

GCH 2001-03-27: Added Array. Moved in notes binary and sequence and added structure as possible future property types