Visible & Infrared Survey Telescope for Astronomy

Project Office

Tel: +44 (0)131 668 8411
Fax: +44 (0)131 668 8412
Email: vista@roe.ac.uk,
WWW: http://www.roe.ac.uk/atc/vista

Document Title:	**VISTA Computer Hardware Architectural Design**
Document Number:	**VIS-TRE-ATC-00150-0002**
Issue:	**2**
Date:	**2 October 2001**

Document Prepared By:	**J M Stewart Software Engineer**	**Signature and Date:**	JM Stewart 2/10/01
Document Approved By:	**S Craig Project Engineer**	**Signature and Date:**	S Craig 4/10/01
Document Released By:	**A McPherson Project Manager**	**Signature and Date:**	4/10/01

Document Reviewed By:	**W Sutherland Project Scientist**	**Signature and Date:**	W.J. Sutherland 4/10/01

Change Record

Issue	Date	Section(s) Affected	Description of Change/Change Request Reference/Remarks
1	24 June 2001		Original presented for Close Out Review as VIS-SPE-ATC-00150-0002
2	26 Sept. 2001	5.2, 5.3, 5.5, 7	Detector controller details, data rates updated, fast tip/tilt not required and Test Camera System not required.

Table of Contents

1 Introduction

1.1 Purpose

This document describes the architecture of the VISTA computer system that will be used at the ESO's Cerro Paranal Observatory on the NTT Peak next to the VLT on Cerro Paranal. The document is intended to place in context the individual systems for the benefit of those specifying and implementing those systems.

In its current form the document corresponds to the instrument conceptual design. It will be updated to add details and clarifications as the project progresses to later design phases and implementation.

1.2 Scope

The document describes systems and networks that will be deployed at the Paranal site of the VISTA Telescope. It does not describe systems that will be used elsewhere, e.g. in the UK's archiving and data pipelining centres or at ESO HQ in Garching.

1.3 Applicable Documents

AD01 *VISTA Science Requirements*, VIS-SPE-VSC-00000-0001, V2.0, 26 October 2000.

AD02 *VISTA Operational Concepts*, VIS-SPE-VSC-00000-0002, V1.0, 28 March 2001.

AD03 *VISTA Technical Specification*, VIS-SPE-ATC-00000-0003, Issue 2, 26 Sept. 2001.

AD04 *LCU Common Software User Manual*, VLT-MAN-SBI-17210-0001, Issue 3.5, 20 October 1999.

AD05 *Central Control Software User Manual*, VLT-MAN-ESO-17210-0619, Issue 1.9, 30 October 1998.

AD06 *VLT Electronics Design Specification*, VLT-SPE-ESO-10000-0015, Issue 5, 6 March 2001.

1.4 Reference Documents

RD01 *VISTA Instrument Software Requirements*, VIS-SPE-ATC-00150-0003, Issue 2, Sepember 2001.

RD02 *Final Layout of VLT Control LANs*, VLT-SPE-ESO-17120-1355, V1.2, 12 Jan. 1999.

RD03 *VLT White Book*, http://www.eso.org/outreach/info-events/ut1/fl/whitebook/ .

RD04 *VLT Common Software*, VLT-MAN-ESO-17200-0888, Issue 1.0, 17 August 1995.

RD05 *Telescope Control System User Manual*, VLT-MAN-ESO-17230-0942, Issue 1.6.

RD06 *CCS-LCU Motor Control Module Part 1, Application Programmatic & Command Interface User Manual*, VLT-MAN-ESO-17210-0600, Issue 1.7, 2 October 1998.

RD07 *Configuration of VLT Standard VME Boards*, VLT-MAN-ESO-17210-1358, Issue 1.2, 10 November 1999.

RD08 *VISTA Software Architectural Design*, VIS-SPE-ATC-00150-0001, Issue 2, September 2001.

RD09 *Designing ATM Networks*, Chap 8 of *Internetwork Design Guide*, Cisco Systems Inc., http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/

RD10 *HP ATM 622 and 155 Mbps Adapters for HP 2000 Entreprise Seervers and Workstations*, Hewlett Packard Performance Reports A5483A, A5513A and A5515A.

RD11 *IRACE DCS User Manual*, VLT-MAN-ESO-14100-1878, Issue 1.2, 3 April 2000.

RD12 *FIERA CCD Controller Software User Manual*, VLT-MAN-ESO-13640-1388, Issue 1.3, 17 February 2001.

1.5 Abbreviations and Acronyms

ATM	Asynchronous Transfer Mode
CB	Control Building
CLIP	Classic IP (over ATM)
DHS	Data Handling Server
ESO	European Southern Observatory
FIERA	Fast Imager Electronic Readout Assembly
FPRD	Functional and Performance Requirements Document

IRACE	Infrared Array Control Electronics
LAN	Local Area Network
LANE	LAN Emulation
LCU	Local Control Unit (a VME/VxWorks system, the equivalent of an EPICS IOC)
M1	Primary mirror
M2	Secondary Mirror
OCDD	Operational Concepts Definition Document
OLAS	On-line Archive System (? need to check)
SRD	Science Requirements Document
TA	Telescope Area
TBC	To be confirmed
TBD	To be defined
TCP/IP	Terminal Control Protocol/ Internet Protocol
TCS	Telescope Control System
VLT	Very Large Telescope
VME	A widely used computer bus (acronym is unimportant)
WS	Unix Work Station (or Server)

2 Background

2.1 Relationship to VLT

VISTA will be sited close to the ESO VLT and will be operated from the VLT Control Room by ESO staff. Various site services, provided by ESO, will be used, e.g.

- archiving
- local area network
- wide area network

- astronomical site monitor
- sysadmin
- maintenance

This means that the VISTA computer architecture must be closely based on the VLT architecture and must interact with it efficiently and reliably. These constraints are stated more formally in the SRD (AD01) and the TS (AD03).

The VLT computer system is described in many documents, including overviews, requirements documents, functional specifications and user manuals made publicly available from ESO's FTP site ftp://ftp.eso.org/. The VLT White Book (RD03) provides an introductory overview of the VLT control system (Chap. 7) and the data flow software (Chap. 10). Control software runs on VxWorks/VME systems called Local Control Units (LCUs) and Unix (HP-UX) workstations (WSs). A typical control system uses software on both platforms, the LCU for real-time and hardware interaction, the WS for a scanned status database and interfacing to other systems.

VLT control software relevant to VISTA includes the VLT Common Software (RD04) that provides the infrastructure for real applications. This common software includes components that run on the LCU (AD04) and on the WS (AD05). Other VLT control software relevant to VISTA includes actual applications, e.g. the TCS (RD05) and hardware subsystem control. The other body of relevant VLT application software is the data flow system, which covers all aspects of data handling, e.g. archiving, processing and display, and also proposal preparation and scheduling.

2.2 Summary

Computer and network hardware will be used at Paranal for control, data acquisition, data assessment, data storage and transport, observation control and system monitoring. This document describes the hardware architecture and should, for many purposes, be read together with the document describing the software architecture (RD08). There is intended to be a clear relationship between the software, the hardware and potential workpackages (which are described with the software).

The network topology mirrors the modularity of the system, in terms of function, software and WBS breakdown. It is convenient, therefore, in this document to group the computer hardware by LAN (RD02).

3 Types of Computer

Two basic types of computer are used in the VLT (and VISTA) systems. Real time systems and systems that directly control hardware run on VME based computers running the VxWorks real-time operating system. Such systems are called Local Control Units (LCUs). Other systems are based on Hewlett Packard workstations and servers running the HP-UX flavour of Unix. These are called "workstations", whether they act as workstations, or servers.

3.1 Workstations

Unix hardware will comply with ESO standards i.e. Hewlett Packard workstations and servers running the HP-UX flavour of Unix. (Note that in ESO terminology the word "workstation" includes both workstations and servers.) Although ESO software can be made available on other platforms, e.g. Sun Solaris and Linux, nearly all operational systems at Paranal use HP-UX workstations. The software is most fully debugged on HP-UX and ESO hold spares of defined HP workstation configurations. VISTA will therefore use HP-UX workstations, even though potential developers may be more familiar with other flavours of Unix.

In order to support the economic provision of spares, ESO define particular models and configurations of workstations (RD02). VISTA will use these configurations wherever possible. Procurement of Unix equipment is likely to be through ESO in order to:

- ensure compatibility
- ensure software (operating system, CCS, etc.) are properly configured
- simplify software licencing
- benefit from ESO negotiated discounts
- ensure that VISTA conforms with latest standards, e.g. before these are documented

Currently the models being purchased by ESO for control applications are

- "big" configuration: J6000, up to 4 CPUs, 0/1/2 screens
- "small" configuration: B2000, 1 CPU, 1/2 screens

As of January 1999 (RD02), the maximum memory in an ESO supported control workstation was 1 GB.

Since the workstations model themselves and ESO preferences will continue to evolve quite rapidly, consultations with ESO will take place before any purchases are made. In order to benefit from the ever decreasing cost/performance ratio, decisions to purchase will not be made prematurely.

In conformance with ESO's practice for the VLT, all VISTA workstations will reside in the Control Building, where relevant connected via one or more LANs to LCUs in the Vista Telescope Enclosure.

3.2 LCUs

All VME based LCUs will conform to ESO standards (AD06). Both 68040 and PowerPC CPUs are supported by ESO, but VISTA will standardise on PowerPC (68040 is essentially obsolete).

VME interfaces will also conform where possible, (hopefully everywhere), to types supported by ESO. This is important for fault diagnosis and spares provision. Supported VME interfaces are listed in (AD06), though other interfaces are currently used or specified for use on the VLT. Whether or not an interface is specified in the current documented ESO standard, the VISTA project will check with ESO before making purchases.

The operating system used on LCUs is Wind River's VxWorks. Development is carried out on (HP) Unix workstations, which require a development licence. Installation and configuration of VxWorks development environment will comprise part of that mentioned in Section 3.

All LCUs will be located in the VISTA Telescope Enclosure. Each LCU's software process will be controlled by a workstation in the Control Building connected via a LAN. Some LCUs may have connections to more than one LAN, but will effectively belong to an environment associated with one specific workstation.

4 Network

4.1 Topology

The various LCUs all reside within the VISTA Telescope Area (TA), each LCU being in an appropriate physical location for the hardware with which it is associated. Associated with each LCU is a Unix workstation residing in the Control Building. A single Unix workstation can serve several LCUs, so long as they share a common LAN. It is important to note that most LANs extend between the VISTA Telescope Area and the Control Building, so that although the LCUs and their associated workstation are physically distant, logically they are part of the same system.

Networks at the Paranal site are described in RD02. Both Ethernet and ATM are used, depending on bandwidth requirements. Where latency or point to point bandwidth is important, special purpose Ethernet subnets are defined. Because Ethernet switch technology has essentially replaced hub technology, latency is now less of an issue, since data collisions can be avoided in many situations. However, for the moment, the network topologies shown below will assume that these separate networks do still exist.

The network architecture of the VLT (RD02) provides a useful starting point for understanding the computer systems of both the VLT and VISTA. The description below interprets this network architecture in the context of VISTA.

All LCUs are located in the VISTA Telescope Area and all workstations are located in the VLT Control Building, typically in the Control Room. Ethernet is generally used to connect nodes that are physically close, logically related and of moderate bandwidth requirements. ATM is used elsewhere, e.g. to connect different buildings and to connect to high bandwidth data sources or stores. TCP/IP and LANE protocols are used. Fibre optic is used throughout, both for Ethernet and ATM, except where some hardware does not allow this. RD02 provides a much fuller description.

Extensive use is made of ATM. This provides higher bandwidth than Ethernet (155 or 622 Mbps as opposed to 10 or 100 Mbps) and can operate over the distance between the VISTA TA and the Control Building. Additionally it allows bandwidth to be managed i.e. one can ensure that a connection between two nodes is guaranteed bandwidth no matter what other nodes are doing.

ATM and Ethernet based nodes are connected by switches running the LANE protocol (LAN Emulation). A TM switches connect ATM segments. Where there is no requirement to connect to Ethernet nodes, the ATM segments can run Classic IP (CLIP). Routers connect the ATM backbone to the outside world.

Switches and routers can be configured to filter the traffic they allow to pass through. This ensures, e.g., that malfunctioning equipment generating bursts of spurious network traffic have only a local effect. Switch and router configuration also prevents unauthorised access, whether well intentioned, e.g. from an astronomer on an Off-line Workstation, or malicious, e.g. a hacker on the global Internet. In this sense they act as a firewall.

VISTA will use several LANs, specific to itself:

- IR Instrument LAN - mechanism control and data acquisition for the IR camera
- Visible Instrument LAN - mechanism control and data acquisition for the visible camera
- Telescope Control LAN - telescope commands and status
- Guide/Acquisition LAN - bulk data associated with wavefront sensing
- User Station LAN - VISTA workstations in the Control Building

In addition, VISTA will communicate with systems on certain site-wide LANs:

- General Services LAN - e.g. off-line VISTA workstation and Archive System
- Astronomical Site Monitoring LAN - site and sky data

4.2 Equipment

Network equipment will be of the standard types specified by ESO. As documented in RD02 this includes:

- Cisco Catalyst 5500 ATM/Ethernet switch
- Cisco 2503 router
- Baystack Ethernet hubs

Nowadays Ethernet switches are used rather than hubs, allowing greater performance and greatly reducing the non-deterministic characteristics of shared Ethernet. Network equipment models evolve quite rapidly and prices continue to fall. For VISTA, first of all the network topology, including location of switches etc., will be agreed with ESO. Specification and purchase of actual equipment will be made at a later date. Choices of equipment and the network protocols (Ethernet 10/100/1000 Mbps, ATM 155/622 Mbps, CLIP, LANE etc.) have minimal effect on the software.

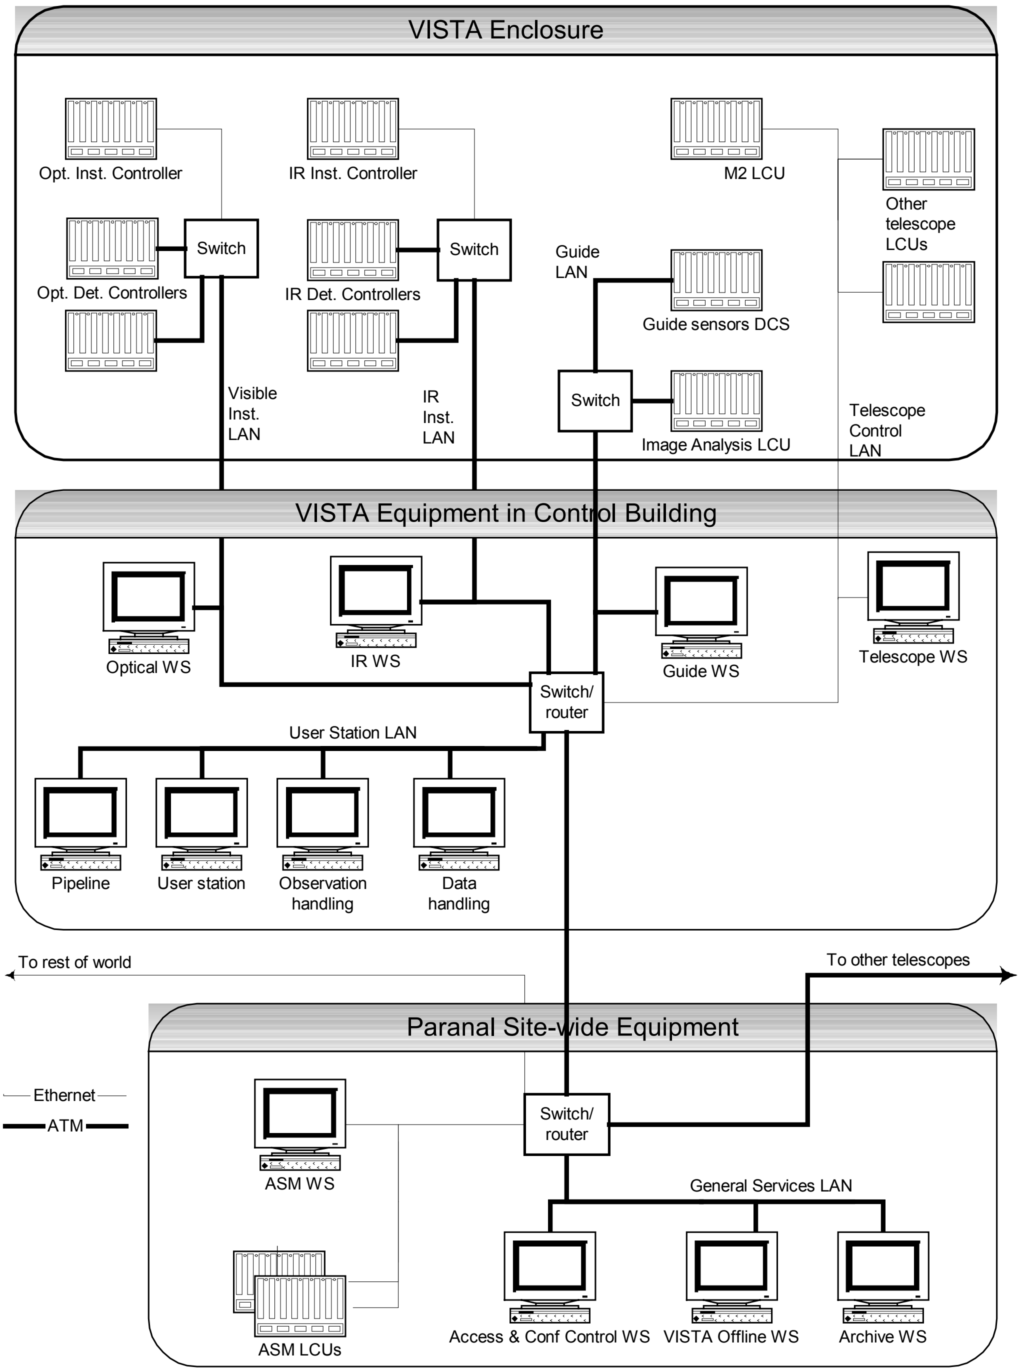

Figure 1 VISTA computer and network configuration.

5 VISTA LANs

5.1 Telescope Control LAN

The Telescope Control LAN is used for commands with associated low volume data from the TCS to its subsystems. Status returns, callbacks, etc. flow in the opposite direction. All LCUs are in the VISTA TA and the TCS workstation is in the Control Building. One or more X terminals will be in the VISTA TA for test and maintenance purposes.

Table 1 Devices expected to be present on the Telescope Control LAN.

Node	Type of device			Make/ Model	Other LAN
	WS	LCU	Other		
TCS	X			HP	
Altitude		X		PowerPC	
Azimuth		X		PowerPC	
Rotator		X		PowerPC	
M1		X		PowerPC	
M2		X		PowerPC	
Enclosure		X		PowerPC	
Bearing		X		PowerPC	

5.2 Infrared Camera Instrument LAN

Each camera has its own LAN so that it can operate without other systems effecting it and vice versa. The IR Camera LAN accommodates all systems associated with the IR Camera including mechanism control LCUs, science detector controllers and controllers associated with the guide, focus and WFS CCDs.

All software involved in the operation of the IR Camera runs on systems on this LAN, including quick look, monitoring and any on-line data reduction that may be required. However certain system that might be expected to reside on it, do not. Data handling and pipeline processing systems reside on the User LAN - it is not economic or worthwhile to provide such systems for each instruments, although clearly the pipeline processing software will be somewhat different for each camera. The Off-line Workstation resides on the General Services LAN. This ensures that the observing process cannot be accidentally affected by off-line processing, web browsing, e-mail, etc., which are only allowed on the General Services LAN.

Table 2 Devices expected to be present on the IR Camera LAN.

Node	Type of device			Make/ Model	Other LAN(s)
	WS	LCU	Other		
IR Camera WS	X			HP	
IR ICS		X		PowerPC	
IR DCS		X		IRACE	

In all likelihood, ESO's IRACE detector controller (RD11) will be used for the 16 IR detector arrays. An IRACE system comprises front-end electronics and one or more back-end "Number Cruncher" workstations, both considered internal components of the DCS. To handle the number of detectors and the data rate, using the currently defined IRACE components, a system with 4 "Number Crunchers" will be needed. The IR Instrument LAN (ATM) links the IRACE systems to the IR Instrument WS. To handle the 4 Number Crunchers, either the workstation could have a separate ATM link to each, or an ATM switch could link the 4 ATM fibres from the DCSs and the 1 ATM link from the IR Instrument WS. These options will be investigated.

5.3 Visible Camera Instrument LAN

The Visible Camera LAN (Table 3) is exactly equivalent to the IR Camera LAN.

Table 3 Devices expected to be present on the Visible Camera LAN.

Node	Type of device			Make/ Model	Other LAN(s)
	WS	LCU	Other		
Visible Camera WS	X			HP	
Visible ICS		X		PowerPC	
Visible DCS		X		4 off FIERA	

ESO's FIERA detector controller (RD12) is likely be used for the CCDs. Each FIERA can handle 16 CCDs and so 4 controllers will be needed (3 if the number of CCDs were reduced to 48). FIERA's hardware comprises front-end electronics and a Sparc based LCU, which communicates with the Visible Instrument WS via the ATM-based Visible Instrument LAN. The 4 ATM links from the FIERAs can be connected to the Instrument WS either via a switch, requiring just one ATM interface on the WS, or directly, requiring 4 ATM interfaces on the workstation, plus a fifth for transferring data from the IWS. The feasibility and relative merits of these two possibilities will be investigated taking into account ESO standards and developments in the rapidly evolving workstation market.

5.4 Guide LAN

The VLT design uses a separate LAN for transferring bulk raw data acquired for guiding and wavefront sensing. Images from the sensors are transferred to an image analysis LCU, which calculates wavefront coefficients, e.g. Zernikes, to pass on to the telescope system.

Table 4 Devices that may be present on a Guide/Acquisition LAN.

Node	Type of device			Make/ Model	Other LAN(s)
	WS	LCU	Other		
Guide WS	X			HP	
Guide camera LCU		X		FIERA	
Image analysis LCU		X		PowerPC	

VISTA's design is rather different from the VLT's in that the guide, focus and wavefront sensing systems are part of the instrument rather than the telescope (RD08). Like the VLT, a Guide Workstation controls these functions and is, like the LCUs, on a Guide/Acquisition LAN.

The VLT network configuration as documented (RD02) shows the Guide Camera LCU and Image Analysis LCUs connected to the Telescope Control LAN as well as the Guide LAN. However it appears that this connection is no longer present and these LCUs are solely connected to the Guide/Acquisition LAN. For VISTA, these LCUs will be controlled from the OS running in the Instrument Workstation and the physical connection can be through the switch/router that connects the Instrument LANs and the Guide/Acquisition LANs.

It is assumed that the same Guide Camera Detector Controller will be used with the optical sensors mounted in both instruments' focal planes. This simplifies the observing configuration and potentially requires less hardware. It has the disadvantage that sensors cannot be operated when the instrument is off the telescope unless additional hardware is obtained. The alternative is to have two separate Guide LANs with their own equipment in the same way there are two separate Instrument LANs. This assumption will be reviewed.

5.5 Rapid Guiding LAN - Not Required

The Rapid Guiding LAN is used in the VLT configuration for sending tip/tilt demands (field stabilisation) to M2 in a deterministic manner. VISTA's guider will not control M2 but will send demands to the TCS for correction of the altitude and azimuth axes, which have a lower bandwidth and therefore no difficult latency requirements.. VISTA will therefore not have a Rapid Guiding LAN.

5.6 User Station LAN

The VISTA User Station LAN hosts workstations that perform operations that are not specific to an instrument (Table 5). The User workstation is a two screen workstation that acts as a central control station for processes running on other workstations, e.g. acting as an X terminal.

The Observation Handling workstation runs high level software such as P2PP and the Scheduler. This software comprises instances of standard ESO software running on the other ESO telescopes, accessing VISTA configuration data files. This workstation may not be required if its processes are run on another workstation, such as the Data Handling workstation.

The Data Handling Workstation (actually a server) also runs instances of standard ESO software, e.g. the On-line Archive Client to store data and header information. For the current generation of VLT instruments, all data are written to the Archive Workstation (on the site-wide General Services LAN) from the Data Handling Workstation on each User Station LAN. In the era of VISTA and VST, this may present bandwidth problems and a better approach may be to constrain VISTA data to VISTA LANs. Thus the VISTA Data Handling Workstation on the VISTA User LAN would maintain the mass storage and write the transport media. This different approach will need to be discussed and agreed with ESO.

A large volume of data must be kept on-line (Section 7) and so this system is likely to include a RAID disk. This system also writes to the media used for data transport. Currently this would be performed by a DVD jukebox, but ESO are currently evaluating replacement technology.

The Pipeline Workstation is responsible for running off-line data reduction software, primarily for quality assessment to be made available during the day following a night's observation. This software will be VISTA specific and will be somewhat different for IR and visible data. This system will access data stored on the RAID array attached to the Data Handling Workstation. A RAID array with dual Fibre Channel interfaces may be appropriate to remove any potential bottlenecks from traffic on the User Station LAN and data transfer to the General Services LAN. However the baseline design is to copy data as currently done for VLT instruments.

Table 5 Devices expected to be present on the VISTA User Station LAN.

Node	Type of device			Make/ Model	Other LAN(s)
	WS	LCU	Other		
User WS	X			HP "small", 2 screens	
Observation Handling	X			HP "small" (may not be required)	
Data Handling	X			HP "big" + RAID	
Pipeline	X			HP "big" (or possibly Beowulf)	

6 Non-VISTA LANs

6.1 General Services LAN

The General Services LAN hosts various equipment that performs a Paranal site-wide rôle rather than a telescope rôle. Examples include

- access and configuration control
- Archive Workstation - copy of data, database of files, writes DVDs
- time server
- site quality control

In addition this LAN hosts one VISTA specific workstation (Table 6). An off-line workstation is made available for use by, e.g., a visiting astronomer, such that there is no risk of affecting the observing process. This workstation does not have access to services on the other LANs and so the VISTA off-line workstation can be used for functions not allowed elsewhere, e.g. web browsing and e-mail as well as off-line data reduction.

This offline data reduction, typically performed by a visiting astronomer, cannot feed back reduced data to the observing process or archive. Although it might be possible to run some software here to search for particular objects, a more effective approach would be to install appropriate software on the Pipeline Workstation.

Table 6 VISTA device that is expected to be present on the General Services LAN.

Node	Type of device			Make/ Model	Other LAN(s)
	WS	LCU	Other		
Off-line WS	X			HP	

For the current generation of VLT instruments, all data are written to the Archive Workstation from the Data Handling Workstation. As mentioned in Section 5.6, this approach has disadvantages when massive amounts of data are involved and it is proposed here to use the Archive Workstation solely for maintaining the database of files and observations, not handling the data themselves. This different approach will need to be discussed further and agreed with ESO.

6.2 Astronomical Site Monitor LAN

The Astronomical Site Monitor LAN hosts systems that provide site and astronomical data to the telescopes at Paranal. There will be no VISTA equipment on this LAN, although VISTA software will read information from it (via the TCS).

6.3 Backbone LAN

The Backbone LAN connects telescope LANs, including VISTA's, to each other (not relevant to VISTA) and to the outside world. There will be no VISTA equipment on this LAN.

7 Data Storage and Transport

7.1 Requirements

VISTA will generate high data rates and volumes due to the large number of detector arrays. The system at Paranal must be able to cope with the worst case, rather than the representative case. Worst case parameters, shown in Table 7, assumes that VISTA operates an entire night of 14 hours, generating and storing data the whole time at the peak rate.

Table 7 Worst case data rates and volumes for the two VISTA cameras.

	IR Camera	**Visible Camera**
No. of arrays	16	50
Format	2048x2048	4608 x 2048
Exposure volume	134 MB (2 byte raw data) 268 MB (4 byte coadded data)	944 MB
Time to store an exposure	5s	40s
Data rate to disk during storage	53.8 MB/s (4 byte data)	23.3 MB/s
Max. exposure rate	1 per 10s (2 byte data)	1 per 60s
Data rate to disk that must be sustainable	13.4 MB/s	15.7 MB/s
Max observing duration	14 hours	14 hours
Maximum data/night	676 GB	793 GB

Although the basic parameters are quite different, the IR and Visible Camera generate rather similar worst case data rates and volumes. The maximum datarate to disk of 53.8 MB/s is defined by the need to store coadded IR data to disk in 5s. The maximum data volume per night of 800 GB depends little on which camera is operating. Data volumes from a typical night (Table 1) are important for sizing the storage that spans several nights.

Table 8 Typical data rates and volumes for the two VISTA cameras. Figures are taken from the OCDD AD02 Section A.3.

	IR Camera	Visible Camera
No. of arrays	16	50
Format	2048x2048	4608 x 2048
Exposure volume	134 MB	944 MB
Exposures/night	3000	200
Data/night	400 GB	190 GB
30 night's data	12 TB	6 TB

7.2 Network Capacity

These figures are clearly demanding and in excess of those anticipated in the VLT design (RD02). Nevertheless the VLT design appears able to cope with these figures. ATM is used to transfer data from the detector controllers to instrument workstations and from the instrument workstations into the data handling workstations. Currently the VLT uses 155 Mbps ATM and this can cope with VISTA's requirements with multiple interfaces in the Instrument Workstations. 622 Mbps is well tried technology and its adoption would decrease the number of fibres and interfaces required. So too would adoption of Gigabit Ethernet. Although introduction of more modern network media to the VLT might be predicted and would benefit VISTA, the VISTA design does not depend on it.

Since both the Visible and IR Instrument LANs will require effectively 4 DCSs each (FIERA or IRACE), the question arises how to connect their ATM links to the Instrument WS. Each DCS could be attached via a switch to a single ATM interface on the Instrument WS. Alternatively the Instrument WS could have one ATM interface for each DCS. (This is technically feasible, e.g. an HP N4000 server with 4 622 Mbps ATM interfaces can support a unidirectional data rate of 2.1 Gbps at a CPU utilisation of 60% (RD10). These alternatives will be investigated.

7.3 On-Line Data Storage

It is a requirement (AD03) to be able to store 3 typical night's data on-line and to store data from one "worst case" night. This results in a requirement for 1.2 TB of magnetic disk storage. Although perhaps daunting a few years ago, such capacity is now fairly routine. For example the Hewlett Packard Surestore fc80 disk array can store up to 4.4 TB in a single 2 metre rack. Disks, power supplies, etc. are hot swappable and there is no single point of failure. Fibre Channel interfaces, whether or not in a SAN architecture, can provide data transfer rates of up to 170 MB/s.

7.4 Data Transport

Data will be transported from Paranal using the "diplomatic bag" i.e. the physical transport of data media on a weekly basis. This is likely still to be the case in 2004 (see 7.6).

Currently ESO use DVD as the transport media. These robust media have a capacity of 9.1 GB (or 5.2 GB) and so a week's worth of typical nights would require 310 media. This is far from convenient, but is at least possible in principle. Data rates are a problem, since this technology is inherently slow (compared to magnetic disks). To achieve a write rate of 12 MB/s (worst case night producing 944 MB written to media over 18 hours), could require 5 4-drive systems operating in parallel.

Alternative media include:

Magnetic disks

The removable disk drives used in the RAID systems store up to 73GB each (HP Surestore). One week's data would then require 40 disks. Using magnetic disks for this purpose is currently being investigated by ESO.

LTO tapes

Linear Tape Open (LTO) technology is being promoted by several major vendors and combines the advantages of several current tape technologies. A single tape stores 100 GB and this is forecast (by HP) to increase to 800 GB by 2006. Media costs are low (~$50). The current data transfer rate is 6 MB/s, so currently 2 drives working in parallel would be required to keep up with the data rate. By 2005 however one higher speed drive would be adequate. The cost of a 10 slot autoloader (710 GB capacity) is currently $8k (HP Surestore 1/9 Untrium HV dsk).

Optical Disks

Fluorescent Multilayer Disks are an example of new technology, competing in the archival data storage market, a very active area because of business needs. One CD-ROM size disk is expected to store 100 GB of data.

Data transport is a potential problem in that the technology and procedures currently employed by ESO (DVD) do not scale well to the data volumes generated by VISTA, though they are feasible. VISTA will work with ESO to ensure that whatever new system is deployed at Paranal and Garching is appropriate for VISTA.

7.5 Near-Line Data Storage

Near-line data storage is storage that is not immediately on-line, like a magnetic disk, but can be brought on line automatically and in a timely way, e.g. an optical or tape juke box. VISTA has a requirement to store 30 nights' worth of data near-line (AD02). However VISTA's requirements at Paranal are minor compared to the archive requirements for the entire project (10-15 years' worth of data) and to the archival storage requirements of UK and ESO wide field astronomy.

For Paranal, ESO use the same technology for a local off-line archive as for data transport. (The data are copied to DVDs, one of which is sent to Garching and the other kept at Paranal.) This is likely to continue to be the case with whatever new technology is chosen for data transport. An alternative for VISTA is to meet the requirement by storing 30 nights' worth of data on-line, e.g. on a RAID array, rather than near line.

7.6 Electronic Data Transfer

For completeness, the possibility of transferring VISTA data electronically must be mentioned. ESO currently use a single hop satellite link between Garching and Paranal to provide a data rate of 1-2 Mbps at not insignificant cost. To transmit a typical night's worth of IR data over 24 hours, would require a datalink with a capacity of about 40 Mbps.

Although it may be feasible to consider such a data rate between Chile and Europe by 2005, it is likely to be very expensive. Furthermore there is no great advantage in transmitting data electronically, since surveys by their nature are long term projects. Rapid identification of sources, e.g. near Earth asteroids or supernovae, can in principle be performed by running an appropriate data pipeline at Paranal and transmitting electronically only the data of interest, e.g. the objects' parameters and small images. (Such data processing is not part of the current project however.) Indeed the problem of transferring massive data electronically is one of the drivers behind current efforts in e-science and grids.

7.7 Data Compression

The figures above also take no account of data compression. VISTA data will be compressible, e.g. by embedded data compression in the data storage products, but the degree of compression is most unlikely to be more than a factor of 2 and may be significantly less. Although it may affect the final costs, it will not affect the basic decisions that ESO and VISTA will have to make over the next few years.

8 Computer Inventory

Based on the systems described above it is possible to compile an initial inventory of the computer systems that VISTA will use (Table 9). These exclude any equipment permanently assigned to a Test Camera.

Exact specifications of these computer systems have still to be determined. Most of the LCUs will be based on a standard configuration, e.g. CPU, memory, network interface and time card (RD07). Each LCU will have VME interfaces dependent on the hardware being controlled. Such hardware should be chosen as to maximise reuse of existing VLT software. For example when choosing motors, encoders and limit switch configuration, RD06 provides a list of supported configurations. Some LCUs may be significantly be more complex, e.g. those associated with the detector controllers.

Workstations will all be standard HP systems running HP-UX. Those involved with pipeline processing will need to be the most powerful in order to handle the high date rates and computational load. (The use of Beowulf clusters here is attractive, but they are not currently supported on Paranal.)

A network equipment inventory, e.g. switches, hubs and routers, will be determined at a later date.

Workstation	LCU	Model (see notes)	VISTA LANs					General Services LAN
			Tel	IR	Vis.	Guide	User Station	
TCS		HP (2) 2 screen	X					
	M1	PowerPC	X					
	M2	PowerPC	X					
	Altitude	PowerPC	X					
	Azimuth	PowerPC	X					
	Rotator	PowerPC	X					
	Enclosure	PowerPC	X					
	Hydrostatic bearing	PowerPC	X					
Visible Camera		HP(3)			X			
	Visible Inst. Control	PowerPC			X			
	Visible Det. Controllers	FIERA (4 off?)			X			
IR Camera		HP(3)		X				
	IR Inst. Control	PowerPC		X				
	IR Det. Controllers	IRACE (4 off?)		X				
Guide		HP(3) 2 screen				X		
	Guide camera LCU	FIERA (TBC)				X		
	Image analysis	PowerPC				X		
Pipeline		HP (5)					X	
Obs. handling		HP(2)					X	
Data handling		HP (4)					X	
User WS		HP(2) 2 screen					X	
Off-line WS		HP(2)						X

Table 9 Inventory of VISTA workstations and LCUs and their LAN connections.

Notes

1. PowerPC indicates an LCU comprising a VME crate, PowerPC CPU, Time Card and VME interfaces.
2. A "small" configuration HP workstation, e.g. B2000 with 1 CPU.
3. A "big" configuration HP workstation/server, e.g. J6000 with 2 CPUs.
4. Data handling workstation will be a multi-CPU server with RAID array and DVD jukebox (or ESO replacement)
5. Pipeline processing may be implemented on a high performance HP workstation/server or, e.g., a Beowulf cluster depending on developments within the ESO community.

___oOo___