LARGE SYNOPTIC SURVEY TELESCOPE

Large Synoptic Survey Telescope (LSST)

LSST Observatory Network Design

Ron Lambert, Jeff Kantor, Mike Huffer, Chip Cox, and Sandra Jacque

LSE-78

Latest Revision: January 25, 2016

Change Record

Version	Date	Description	Owner name
1	3/18/2009	Initial version as Document-7354	Ron Lambert and Jeff Kantor
1.2	9/10/2010	Added US national network information	Jeff Kantor
1.3	9/22/2010	Minor Updates	German Schumacher
1.4	6/3/2011	Summit network updates	German Schumacher
2	6/16/2011	Elevated to LSE handle and updated with latest information from MREFC update and 10-40-80 plan	Jeff Kantor
2.1	7/8/2011	Update for Directors Review and PDR; added AmLight figure	Jeff Kantor
2.2	8/5/2011	Updated Summit diagram	Jeff Kantor
2.3	2/25/2013	Adjusted for guaranteed 10 Gb/s (vs 10/40/80), updates to MIA - NCSA link, and AP move to NCSA	Ron Lambert, Jeff Kantor
2.4	6/7/13	Updated international text and diagrams. Reorganized text into sections for clarity. Numerous edits to reflect latest network baseline.	Jeff Kantor, Chip Cox
2.5	9/6/13	Updated for increase Mountain – Base network bandwidth to 200 Gb/s	Jeff Kantor
	10/4/2013	Changes approved 10/3/2013 – LCR-139	R. McKercher
2.5.1	11/1/2013	Corrected Figure 2 to 10M alerts/night	G. Dubois-Felsmann
3.0	3/9/2015	Mtn - Base and La Serena - Santiago	Ron Lambert
4.0	8/25/2015	Implementation of LCR-375; no content change since Post-INAR iteration (version-30904)	R. McKercher
4.1	12/15/2015	Mtn-Base diagrams and fiber updates.	R. Lambert

4.1.1	12/17/2015	Post-CCB-meeting edits, Figure 5 replaced.	Ron Lambert
4.1.2	12/18/2015	"architecture… may change" sentence removed; "to be updated" markings for Fig. 5 removed	Gregory Dubois-Felsmann
4.2	1/14/2016	New Figure 5 received from Ron Lambert	Ron Lambert / gpdf
4.2.1	1/25/2016	Figure 11 rescaled to fit page, no content change.	Gregory Dubois-Felsmann
	2/1/2016	Implementation of changes approved via LCR-385	R. McKercher

Table of Contents

Change Record i

Introduction 1

LSST Network Requirements and Data Flows 3

Network Requirements 3

Data Flows 4

Normal Operations – Alert Processing 4

Normal Operations –Data Release Processing 7

Sizing and Architecture 7

Reliability and Availability 9

Failure Scenarios on the Summit to Base Link 10

Failure Scenarios on the Base to Archive Links 10

Mountain Summit Infrastructure 12

Summit to Base Network 15

Base Site Network Infrastructure 21

Chile and International Network 21

Chile National Circuit 25

International Circuit 28

U.S. National Network 30

Data Access Center Network Connectivity 32

Future Variables 36

Appendix A: La Serena – Santiago Network Rings 38

ALMA (Paso de Jamas) 39

Aguas Negras 40

Chile-Argentina Ring 41

References 43

Introduction

The LSST Observatory is distributed over five sites: the Summit Site, the Base Site, the Archive Site, the Headquarters Site, and a proposed French Site. Figure 1 shows the location, features, and function of each site. Each site hosts one or more LSST Centers and Facilities, which are purpose-specific, staffed, operating entities within the LSST Observatory. Each Center/Facility has personnel and equipment housed there in order to perform its assigned functions.

While the sites are geographically distributed, they are all functionally integrated. Dedicated, AURA-owned fiber optic lines connect the Summit and Base Sites, while the others are connected through fibers on carrier-owned networks. Control functions are distributed for operational efficiency and to provide robust, reliable, safe operation.

During operation, the LSST Project will broadcast near real-time transient alerts within one minute of image collection. Digitally processed images and feature catalogs will be available to the U.S. and Chilean communities via public networks within 24 hours. Full data product releases are scheduled every year, with two planned releases with data from the first year of operations.

Figure 1: LSST Sites, Facilities, and Centers

The Mountain Summit Site is located on Cerro Pachón, Chile, at the telescope site and performs readout of the camera and crosstalk correction of raw image data. The Base Site is located at the AURA compound in La Serena, Chile, and provides a Data Access Center (DAC) for the Chilean astronomical community. The Archive Site is located at the University of Illinois National Center for Supercomputing Applications (NCSA) in Urbana-Champaign, Illinois. It provides transient alert processing, calibration data products processing, and data release processing, as well as complete data product archiving. The Archive Site houses both a DAC for the U.S. community and the computing and storage infrastructure for the Education and Public Outreach Center (EPOC). The French Site is proposed, is located at the CC-IN2P3 computing center in Lyon France, and would provide one half of the processing capacity for the annual Data Releases, subject to execution of an approved Memorandum of Agreement. The Headquarters Site is currently in Tucson, Arizona and provides overall observatory management, science operations, and EPO functions.

LSST Network Requirements and Data Flows

Network Requirements

The requirements most directly driving the LSST Observatory Network design cover the capabilities to transfer data and control information within and between the subsystems of the LSST (Telescope and Site, Camera, Data Management, Education and Public Outreach), as well as providing access to the data to scientists and non-scientist users.

The LSST has a wide range of time scales in these requirements, including:

- order of milliseconds in the Observatory Control System (OCS) and control systems of each LSST subsystem (TCS, CCS, DMCS)
- order of seconds in acquiring and processing the raw image stream from the Camera through Data Management in order to create transient alerts
- order of days in archiving the raw image stream and other metadata for subsequent use in the production of astronomical catalogs and other data release products
- order of weeks to months in producing and deploying data necessary for calibration of the Telescope and Camera
- order of years in producing and deploying Data Release of LSST processed images and catalogs.

In addition, there are large geographical distances between the various LSST sites, which imply non-trivial latency in certain data transfers.

The above requirements are documented in the LSST Observatory System Specifications (OSS) [LSE-30], and they are flowed down to the LSST subsystem requirements documents (Data Management System Requirements [LSE-61]).

There are also data transfer requirements associated with supporting user access in DACs in Chile and the United States, and in the Education and Public Outreach Center (EPOC) in the United States. In the case of Chile, these requirements are derived from Memoranda of Agreement (MOA) between LSST and AURA [AURA LSST MOA], and between AURA and Chile [AURA Chile MOA] that require providing a DAC in Chile in return for authorization to site the LSST observatory in that country.

All of these requirements and MOA drive the need to provide high-speed networking between the various LSST sites and centers, and in turn have driven the network design.

Early in the design process, we evaluated the alternative of shipping media for the less time-constrained requirements, i.e. where the full preparation, packing, shipping, unpacking, and loading time, estimated at 3 - 5 days, was shorter than the required time frame. We realized that we could effectively use shipped media for the annual deployment of Data Releases, since we have annual storage hardware refreshes being shipped in those time frames. In the event that the anticipated bandwidth between the Archive and Base Sites is not realized, or has an extremely long extended outage, this remains an option for transferring Data Releases to the Base Site. We may also be able to meet requirements by shipping media for certain calibration data, and we will evaluate that once the calibration data transfer requirements are fully defined. However, given the need to provide high-speed networks for the more stringent time-constrained requirements, it appears likely that the data release and calibration data transfer will be accomplished primarily on these networks.

LSST is investing in infrastructure that will guarantee a minimum of a 10 Gb/s, a most likely case of 40 Gb/s, and a best case of 100 Gb/s path diverse network links from La Serena to North America during Commissioning and Operations. The Santiago to Champaign network is specified as 2 separate 40 Gb/s links, one off the Pacific coast of South America and the other off the Atlantic coast. Both links independently have sufficient bandwidth to support all science requirements including transient alert production.

While this is considered the very low end of the bandwidth that will be available, it has been considered the most conservative, minimum guaranteed bandwidth for this design. A more realistic bandwidth estimate of 40 Gb/s has been adopted as the design baseline.

Data Flows

Normal Operations – Alert Processing

Terabyte-sized volumes of LSST data and data products must be moved each night and day from the Summit to the Base Center in Chile, from the Base Center to the Archive and Data Access Centers, and finally from those centers to the users when necessary, as depicted in Figure 2.

With each pointing, the Observatory Control System sends the next pointing to the Archive Center for moving object predictions. On camera readout and crosstalk correction at the Summit, the crosstalk-corrected images flow via the Summit to Base Network to the Base Site.

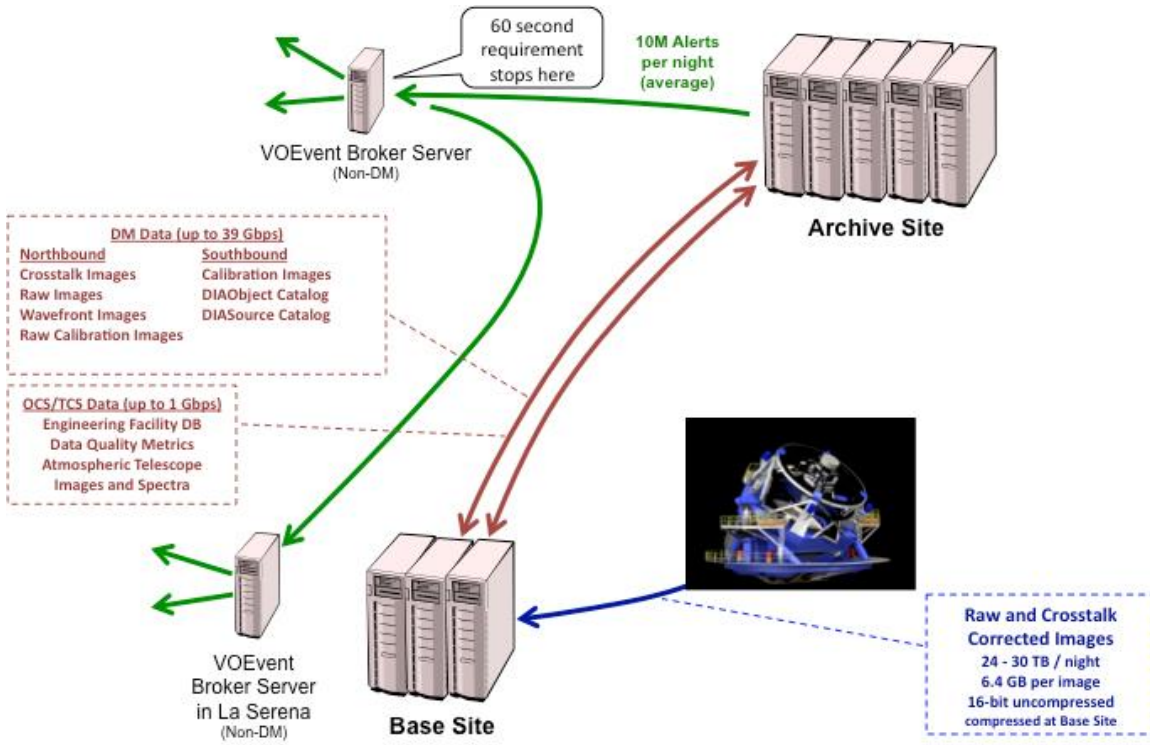

Figure 2: Nightly International Data Flows

There, the image is buffered for failure recovery purposes, and forwarded to the Archive Center, where it is processed to remove the instrumental signature, calibrate it, and subtract it with a sky template to detect difference sources. The resultant World Coordinate System (WCS) and Point Spread Function (PSF) are sent back to the Summit for use as a quality check by the OCS. For every pair of images in a visit, the difference sources are associated with known objects and solar system object position predictions.

Transient events are detected, and alerts are generated at the Archive Center. From the Archive Center, the alerts are sent to alert distributor/brokers such as Skyalert, GCN, and other VOEventNet entities via public networks. LSST will also operate a limited distributor/broker capability to allow alerts to be sent to observatories with significant follow-up capabilities. Alerts are also transferred to the Base Center, where they may be sent to additional distributor/brokers. Alert distribution/broker capabilities will be tested during LSST's Commissioning phase. The exact identification of distributor/brokers and observatories to receive alerts will be determined as LSST approaches full operations. A list of existing distributor/brokers as of the writing of this document is shown in Table 1.

Distributor/ Broker	Operator	URL
VOEvent Compliant	(see below)	http://voeventnet.caltech.edu
Astro GCN/ eStar	University of Exeter, Liverpool John Moores University	http://www.estar.org.uk/wiki/index.php/EStar:About
Skyalert	Caltech	http://skyalert.org
RAPTOR	Los Alamos National Laboratory	http://www.lanl.gov/quarterly/q_fall03/raptor_science.shtml#observatories
Gamma Ray Coordinates Network	NASA GSFC	http://gcn.gsfc.nasa.gov
Non-VOEvent Compliant		
Minor Planet Center/ Central Bureau for Astronomical Telegrams	SAO, Harvard University	http://minorplanetcenter.net/iau/mpc.html http://www.cbat.eps.harvard.edu/index.html
Astronomer's Telegram (Atel)	astronomerstelegram.org	http://www.astronomerstelegram.org

Table 1: Transient Alert Distributor/Brokers

As transient alerts are generated, the raw images are downloaded from the summit to the Base. The raw images are archived at the Base Center and also transferred to the Archive Center. Before the next night's observing starts, the entire night's raw images have arrived at the Archive Center. By the end of each period of nightly observing, the latest astronomical catalog updates are transferred from the Archive Center to the Base Center. Finally, after the nightly observing is done, Summary Data Quality Analysis statistics are generated at the Archive Center and sent to the Base Center and Headquarters.

On a continuous basis, the OCS commands, metadata, and telemetry are also downloaded to the Base Center, and portions of it are sent immediately to the Archive Center for transient alert processing. The full Engineering and Facility Database (EFD) is forwarded to the Archive Center for permanent archiving.

At various intervals ranging from days to months, Calibration Products (e.g. Master Flats, Bias, Fringe, Crosstalk Correction data) used in Alert and Data Release Processing are produced at the

Archive from raw calibration images taken and transferred from the Summit to Base to Archive. The Calibration Products are transferred from the Archive to the Base.

Normal Operations –Data Release Processing

The Data Release is transferred to the US DAC via networks internal to the Archive Center, and to the Chilean DAC via the international network (with media transfer with an annual hardware upgrade as a secondary option). There may also be opportunistic transfer via the network of sky templates, co-added images, calibration data, and catalog updates from the Archive Center to the Base Center at more frequent intervals.

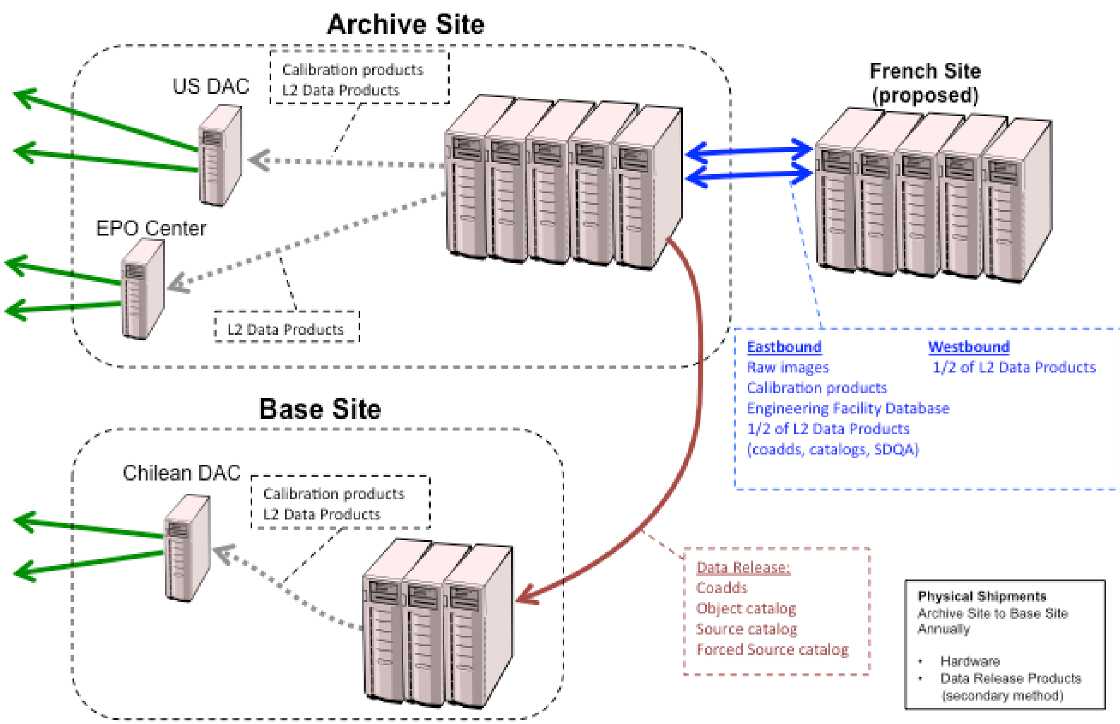

Figure 3 LSST Data Release Processing Data Flow Diagram

Sizing and Architecture

Summit to Base Link

The key driving requirements for the Summit to Base communications are the bandwidth and reliability required to transfer the crosstalk-corrected image data and associated EFD meta-data for alert processing, transfer the raw image data to the Base Center, and handle OCS command and control traffic.

LSST will install a fiber pair utilizing DWDM technology to provision two 100Gb/s Lambdas with end nodes that will be owned and operated by LSST. This capacity is divided into 100 Gb/s for image data transfer and 100 Gb/s for OCS command and control data. Because this link closely follows the path of existing CTIO networks, REUNA will operate, and LSST will manage and maintain this mission-critical link, providing the reliability and availability necessary to run the combined Summit to Base infrastructure as a single component of the LSST system.

Base to Archive Link

The minimum required peak bandwidth from the Base Center in La Serena to the Archive Center in Illinois is established by the need to transfer the crosstalk-corrected images for transient alert processing. The peak requirement for raw image data is much less, because the raw image data only need to arrive at the destination before the next night of observing starts. The peak requirement for annual data release transfer from the Archive to the Base is even less, as this can occur over the period of a month or more.

In order for the Archive Center to process transient alerts within 60 seconds of crosstalk-corrected image readout from the camera electronics, we assumed 7s of this time budget is allocated to data transfer, including compression/decompression. We assumed 2s of that budget is in the Summit to Base transfer, and the other 5s is in the Base to Archive Transfer and compression/decompression.

At 10 Gb/s, analysis and testing show that the transfer time from the Base to the Archive can be accomplished within the required time frame on tcp-based networks, even when considering packet loss and latency in the network at much higher levels than are currently being experienced. This is feasible with Virtual Local Area Network (VLAN) protocols to ensure dedicated bandwidth, use of jumbo frames, and configuration of TCP congestion control to optimize utilization in LSST traffic patterns. Additional analysis and testing is ongoing to verify that the effect of protocol translation overhead (given SONET-based links within the overall path) is minimal and does not affect the design. [Freemon2013]

Commercial carriers and research networks are the only realistic source of this capacity and reliability over this distance. LSST has partnered with NOAO, CTIO, and is coordinating with other South American observatories to leverage and expand existing NSF-funded network initiatives that have allowed increasing bandwidth capacity between Chile and the United States at the same funding level.

This primary network infrastructure is path diverse from Santiago flowing either side of South America. The Eastern side is via Sao Paolo to St. Croix to Miami to Jacksonville to Chicago to Champaign. The Western side is via Santiago to Panama to Boca Raton to Atlanta to Chicago to Champaign. LSST-available bandwidth in each side is expected to be 40 Gb/s. Provisions for "bursting" to available bandwidth (up to 100 Gb/s) are possible. Bursting in this context means the ability to go above 40 Gb/s up to the available bandwidth on the link. If no other traffic is present on the link, it can be up to the full capacity, or less if there is other traffic present. Moreover, the ultimate goal is to establish a continuous wave at layer 1 from La Serena to

Miami and Champaign over the east cost of South America link, at 100Gb/s for the sole use of LSST traffic. This is currently be worked on by Ampath and REUNA and LSST personnel.

These bandwidths presume continued investment in these networks by the NSF and Brazil. Contingency plans exist in case those investments are not sustained and only the minimum 10 Gb/s is available.

Archive to DAC/EPOC to User Links

Within the United States and Chile, LSST must provide open access to LSST data—not only to the scientific community, but also to diverse segments of non-specialist users whose curiosity may be triggered by LSST educational and public outreach programs. LSST will provide DACs in Chile and the US co-located with the Base and Archive Centers so data transfers between each site will be via local internal networks with 100 Gb/s or higher bandwidth. LSST will also provide EPO portals for educational and outreach efforts. User access to the DACs and EPO portals will be via public and Research and Education Network (REN) connections (e.g. internet and Internet2, XSEDE, ESNet), and the aggregate bandwidth will be limited only by the connectivity of the hosting and using institutions. In cases of stand-alone DACs or science centers funded outside the project, the entity developing and operating the center will be responsible for providing network connectivity to LSST Archive Center to enable data transfer.

Reliability and Availability

The above links must also be reliable to avoid data transfer bottlenecks, so path diversity and redundancy are required where economically feasible, and there must be spare capacity on the order of the normal bandwidth to "catch up" in the event that a failure or slow-down occurs.

The above designs are targeted to achieve 98% or higher mean data transfer reliability, which is especially important for transient alert processing and requires higher than 98% network availability. The historical data in Table 2 suggest that this is achievable.

Link	Years in operation (REUNA/AMLIGHT)	Cuts/outages	Longest Outage	Mean Time to Repair	Mean Availability
La Serena – Santiago	10*	multiple	4.5 hrs (27 hrs**)	4.5 hrs	99.5%
Santiago – Sao Paolo	5	16	23 hrs	5 hours	99.5%
Sao Paolo – Miami (GX)	3.5	multiple	34.5 days		99.9%
Sao Paolo – Miami (LAN)	5	multiple	22 hours		99.9%
Santiago – Panama - Miami	1 month	NA	NA	NA	NA

Table 2: Chile and International Link Reliability

Failure Scenarios on the Summit to Base Link

The only link in the network without path diversity is the Summit to Base link. A tradeoff study was conducted to assess the feasibility of path diversity. The result of this study was that the single path was sufficient to achieve the overall reliability necessary, and that a second path would not be a cost-effective increase in reliability. A second path remains a possible upscope should funds be available.

In addition, tradeoffs were conducted assessing whether to install the fiber on poles or underground, and various maintenance scenarios were analyzed. The result was to employ a combination of poles and underground for 3 segments in this link. [LSST11811]

Finally, the size of the Summit buffer is currently 2 nights of observing. In the unlikely event that an outage lasts longer than this, provisions for transferring media down the mountain have been incorporated into the system design.

Failure Scenarios on the Base to Archive Links

This analysis summarizes the tradeoffs associated with operations scenarios that involve failure of one or both of the international links. For this analysis, we use the term Pacific and Atlantic when referring to physical cables and not the word *links,* and *Science* and *Secondary* when referring to the data being carried, though those terms may be used in other sections of this document.

Atlantic Cable Failure

In normal operations, the science data is routed on the Atlantic cable. If the Atlantic cable fails all science data flow is automatically and immediately rerouted to the Pacific cable and supersedes all other network uses. This maintains operations and meets the science requirements. When the Atlantic cable is returned to service it can now be used for non-time critical data. There is no need to immediately reroute science data back to the Atlantic cable during observations.

Pacific Cable Failure

A Pacific cable failure does not affect science operations but interferes with other non-critical data transfers. This does not have an impact on operational requirements.

Dual Cable failure

A failure of both links simultaneously interrupts operations and images are not transferred to the Archive. The images will be buffered at the Base Center for up to 14 nights, which is 200% of the anticipated mean time to repair at least one of the failed links. While both links are down each night's observations will be written to physical media and shipped to the Archive. Typical shipment times will be on the order of 3-5 days.

Single Cable Repaired

The mean time to repair at least one of the cables is estimated at 7 days. As soon as one of the failed cables is repaired the normal nightly science data flow is restored. There may be several nights of buffered data en-route to the Archive on physical media at that time, but this does not affect the current science data flow. LSST has returned to normal operations and no data has been lost due to the network failures, although the Alerts have not been delivered within the Alert latency requirement of 60s for the buffered data.

Both Cables Repaired

When both links are returned to service the science flow is normal on the Atlantic cable and the Pacific cable may be placed into "catch-up" mode. Images that have not yet arrived at the Archive are transmitted on the second cable, whether or not those images are already in transit to the Archive on physical media. The entire network returns to its pre-failure usage as soon as all buffered images have been transmitted to the Archive.

With respect to shipping media, a separate analysis has been performed that assess whether the media should be spinning disk, flash, solid-state disk, or tape. Several of these options are feasible, and the current baseline is solid-state disk, but we will continue to monitor technology/cost trends. [LSST14313]

Mountain Summit Infrastructure

The Summit Computer Room will be located at the Summit Site, which is at 2700 meters elevation. There are three main elements at the Summit: the Camera Control System (CCS), the Camera Data System (CDS) and Telescope Control System (TCS/OCS) with a minimal presence of the Data Management System (DMS). Each of these systems will be developed independently and rely on the network infrastructure for interconnections. The network facilitates data flow between the elements and transfers images and other data to and from the Base Site.

We have adopted a mesh architecture with Access and Core level switch/routers, with dual trunk 10 Gb/s transfer rates between the Access and Core and 1 Gb/s and 10Gb/s at the Access level to the individual VLANs. Care will be taken to avoid oversubscription. The exception to that is for image data from the Camera system readout destined for the Base Site, which will be 10 x 10 Gb/s ports. These data will not enter the general Summit network, as the policy is to be available at the Base Center and the Archive Center in close to real time. The Camera system may also output a third complete image for real time display and other monitoring and diagnostic operations. Camera will store two days of raw image data in the event that there is a disruption in the link to the Base Site and for queries of recent data.

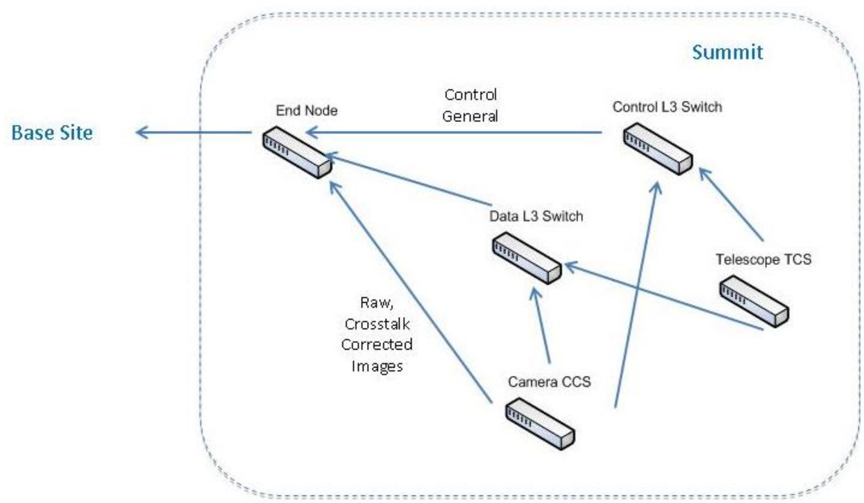

Figure 4: Control and Data Network Diagram

As indicated in Figure 4, lines will run over separate networks in order to decrease latency in the system. We will maintain separate subnet/VLANs between the different entities in order to impose levels of security with access control lists. Figure 5 shows the detailed implementation of the summit networking system.

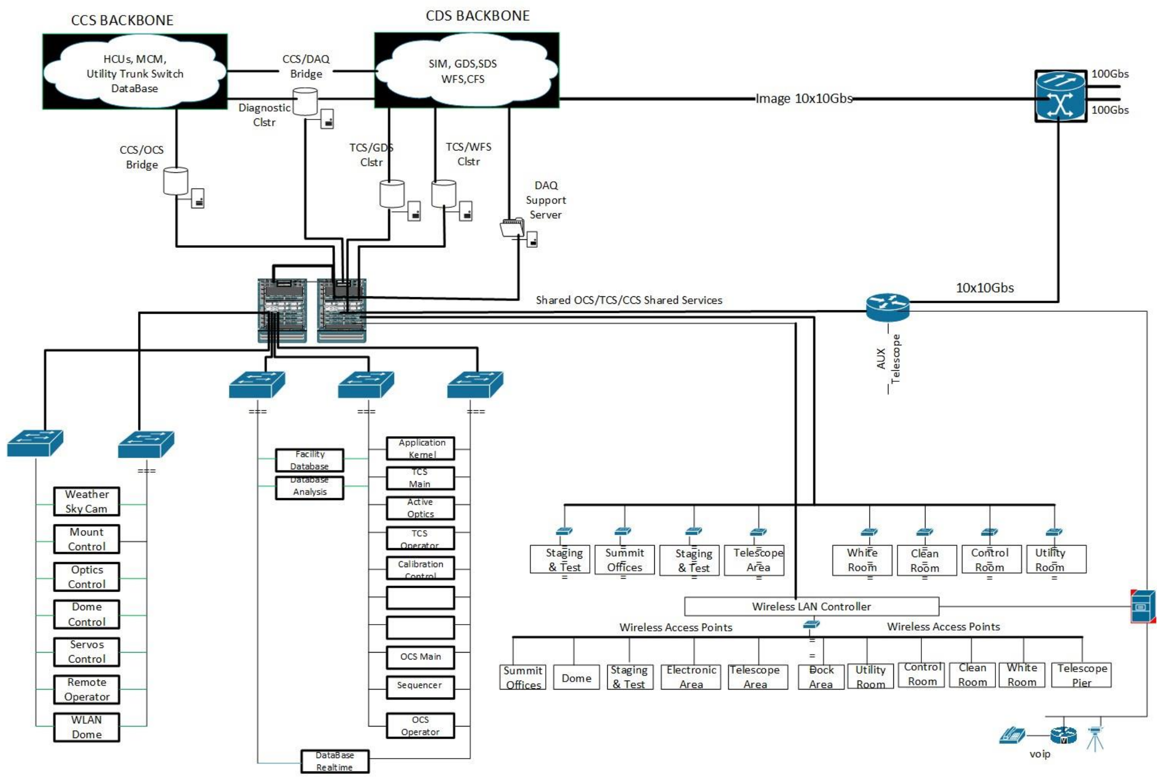

Figure 5: Summit Network Implementation Detail

A representative design using current technology is described in this paragraph. The central Core data center switches, (represented as Cisco 5548), will be redundant and serve the entire summit network and computer systems integrating the Telescope and Camera subsystems. Automatic failover will occur in the event of a switch failure. The Core and Access switches will be state of the art technology with low latency and ~192 Tb/s switching fabric. The Core switches will use dual ported uplinks in a mesh configuration and connected by fabric extenders to offer 10 Gb/s bandwidth throughout the summit. Up to 16 ports can be aggregated to form a maximum of 1.6 Tb/s channel. Interconnects will be fiber cable wherever possible and optimize transmission with SFP+ connectors. The diagram indicates in each network box where that equipment will reside, i.e. cr = computer room, tu = telescope utility area, do = dome, etc.

It is expected that there will be private wireless access throughout the summit building and a separate public subnet for visitors and guests. The Observatory will deploy voice over IP telephony as shown in the DMZ of the Firewall equipment with up to 96 handsets installed around the summit facilities. There will be two streams of 100 Gbs data transiting the Summit to Base link. In the baseline design this will be accomplished over two 100 Gb/s optical streams with DWDM end node equipment. A Level 3 image data stream will also be fed into a real time display system for instantaneous viewing of the complete image. The camera installed on the telescope is the normal mode of operation, but for maintenance it will require fibers to the computer room from the clean and white rooms.

The OCS/TCS/Camera systems will be provisioned with the two 100Gb/s link from the Summit to the Base. The Camera Image Data stream will occupy one entire 100G channel whilst CCS, TCS, OCS and all other subsystems will share the second 100G channel. Links that are considered to be a security risk will be connected through a firewall at the summit, that is, any general data that does not directly pertain to OCS/TCS/CCS/DM flows. Since the equipment will be acquired and deployed starting in the LSST construction phase and operated through LSST operations, the design baseline will change with respect to specific model numbers and suppliers without negative impact to cost or performance.

All rooms in the summit facility will be equipped with fiber and copper outlets for data and telephony. It is anticipated that an 8” cable tray inside the computer room and running the length of the facility with 4” dropdowns into each room will be adequate for the network and computer cabling. All cable runs will terminate in patch panels that are remotely configurable, and all network equipment should have IP remote power switching and monitoring. Fiber cables will run inside plastic Lightpath ducting.

The computer room networking equipment will be contained in two 42U racks. The telescope utility area will have a closet with 6-8U rack space for remote switches. The Clean Room area will have a closet for 4-6U rack space. The ground floor Utility area will have a closet for 1-2U rack space.

Summit to Base Network

The distance from the summit to the base is 55 km point to point or 90 km via the public highway and Observatory access roads. The plan below shows the routing.

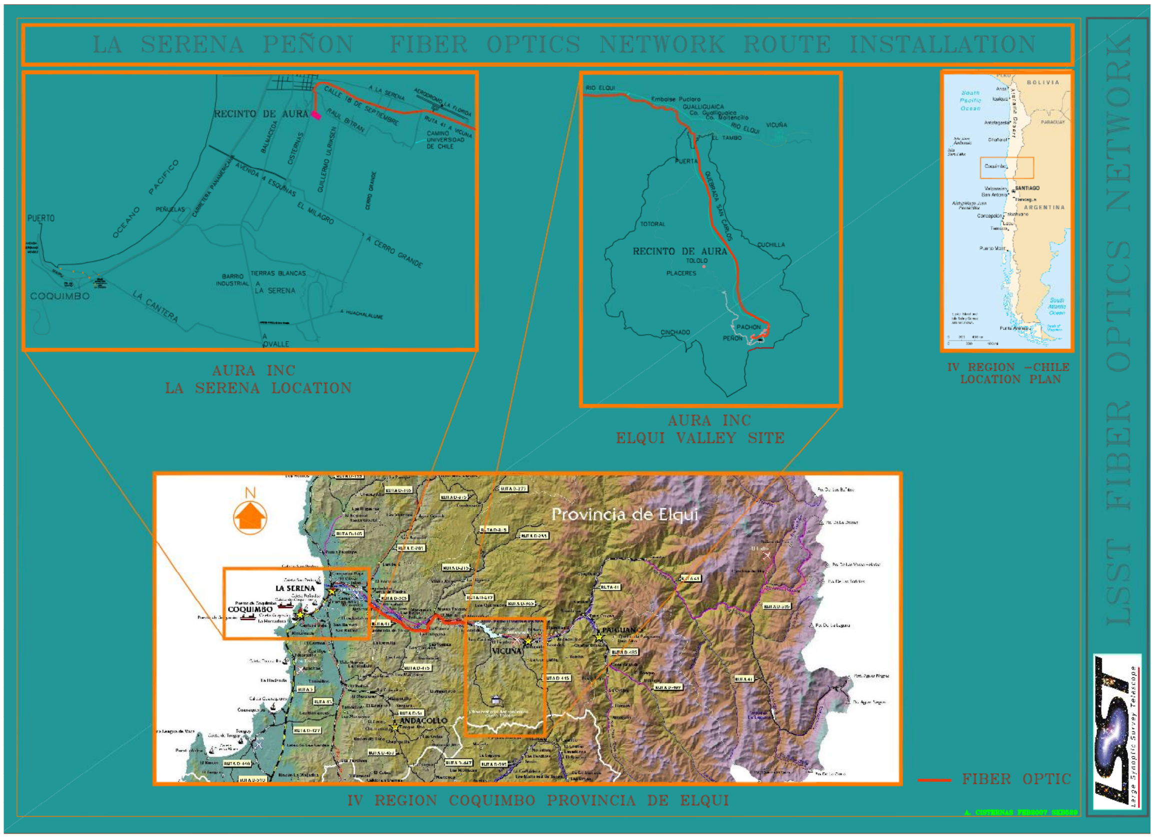

Figure 6: La Serena – Cerro Pachon Fiber Optics Network Route

The geographical diagram above shows the physical routing of the network that will be installed for LSST and AURA traffic from summits to base during the first years of LSST construction.

The Summit to Base fiber cable link will consist of two distinct portions installed by REUNA/TELEFONICA. Part A La Serena to Gate and Part B Gate to C. Pachon.

Part A is La Serena to the AURA Gatehouse in Vicuna.

As shown below in Figure 7, there will be two pairs of dark fibers provided by Telefonica on the public highway, one pair of which will carry only LSST traffic and the second pair AURA traffic. These fibers are provided by Telefonica to AURA through an 18-year Irrevocable and exclusive Use (IEU) and are part of a bundle utilized by the three major telecoms in Chile. Most of the 90Km span will be an aerial installation. On the public highway REUNA will guarantee 4-6 hour mean time to repair, and on the AURA property 6-8 hours in daylight.

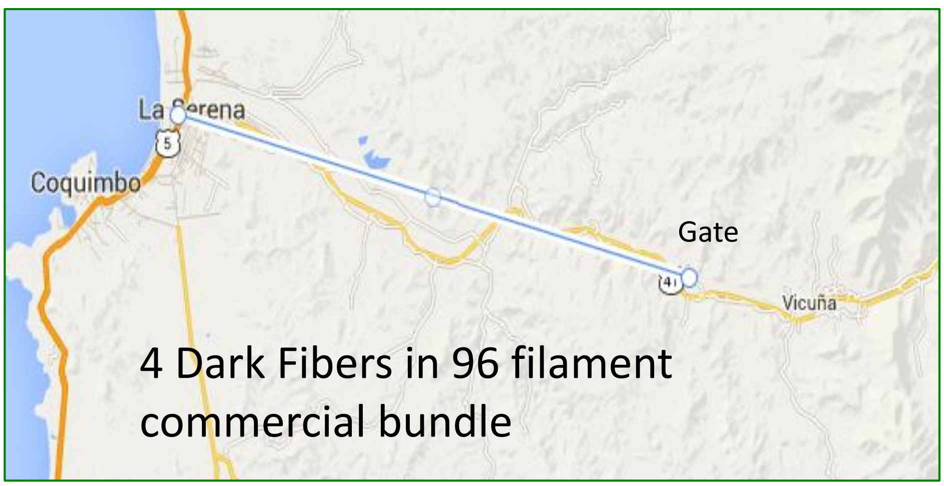

Figure 7

Part B is Gate to C. Pachon

Figure 8 shows the section on AURA property where there will be a 12-pair fiber bundle spliced into the Telefonica provided filament fibers in a manner that is consistent with our use and offer spare fibers for some redundancy. The cable in fig.8 (blue line) will be installed on the electrical posts direct to C. Pachon with a bifurcation at the water pump in San Carlos valley. At this point cable will be installed to jumper C. Tololo's traffic over to C Pachon . Fig.9

Figure 8

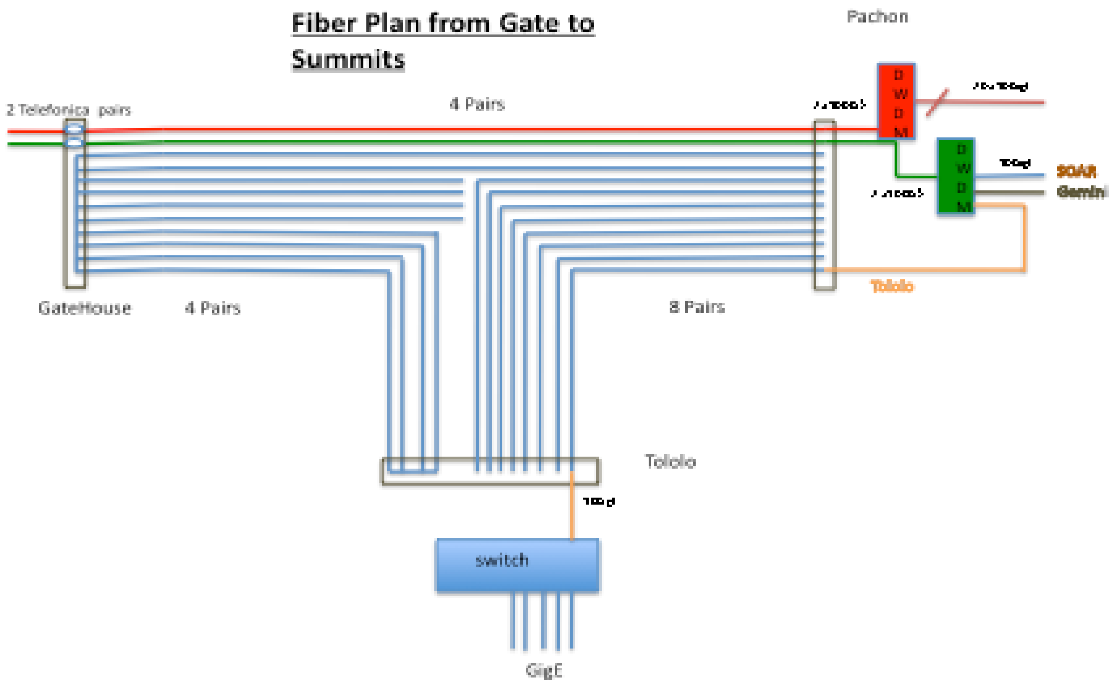

Figure 9

Equipment for Base to Summit.

Initially, equipment to provision the AURA traffic will be purchased and installed either with a 10G Lambda or 100G Lambda but whatever the case it will be upgraded to 100G prior to early commissioning for potential LSST testing. Later, the LSST equipment will be purchased and installed. The DWDM chassis will house two active sets of modules of 10x10GigE for customer input and two active OTU-4 optical modules for a total bandwidth of 200 Gb/s. On Cerro Pachon the LSST chassis will physically reside in the LSST telescope computer room and the AURA chassis in the shared infrastructure building. In La Serena the two chassis will reside in the LSST/AURA data center.

The AURA/LSST DWDM equipment will be configured such that the AURA DWDM will have the capacity to backup the LSST Camera stream in the event of a failure in the fiber. This is termed a collapsed ring which offers optical redundancy. It should be noted at this time that redundancy will not be offered for AURA traffic in the LSST circuits. The redundant link between DWDM equipment would be transported over a 100GigE interface.

Presently there are 9 vendors under consideration. An official competitive tender will be conducted early in the LSST construction phase. Our RFP states that the line of equipment being proposed will be actively supported through LSST operations and that typical active service life is approximately 15 years. All the companies would install the equipment in situ, give service and guarantee advertised bandwidths and throughput.

REUNA will provide operational support.

Typical latency from summit to base is considered to be:

CP WDM equipment (Depending on vendor)	5-65uS
Fiber	432uS
LS WDM equipment (Depending on vendor)	5-65uS
Total	440-560uS

Summary of VLAN Labels and Bandwidths

Groups of vlan labels are shown for the systems, whose values will be conserved at the base and summit. The image data will occupy one 100G link.

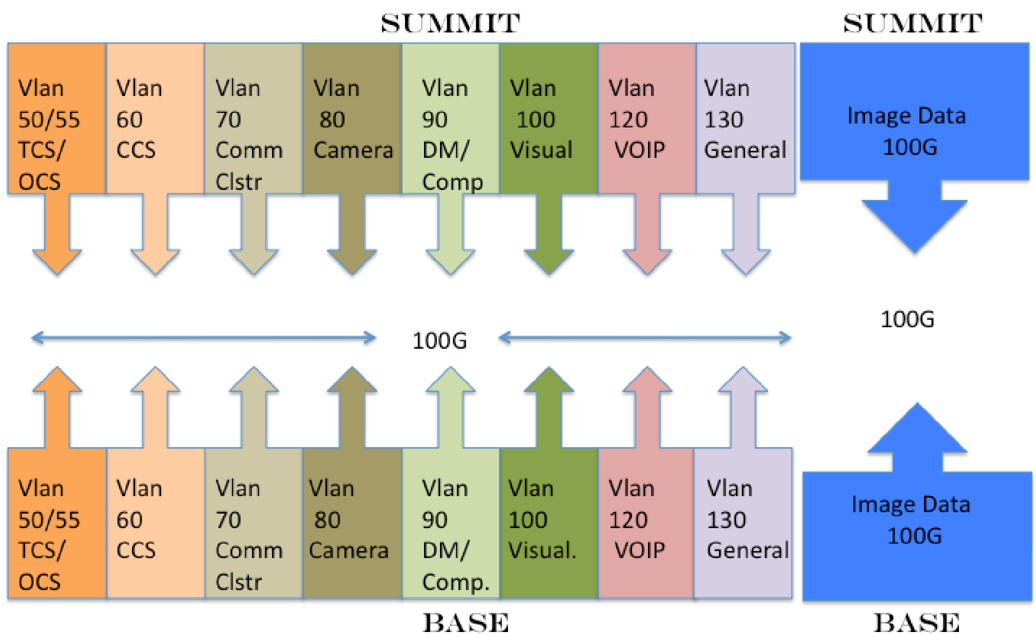

Figure 10: Vlan Labels

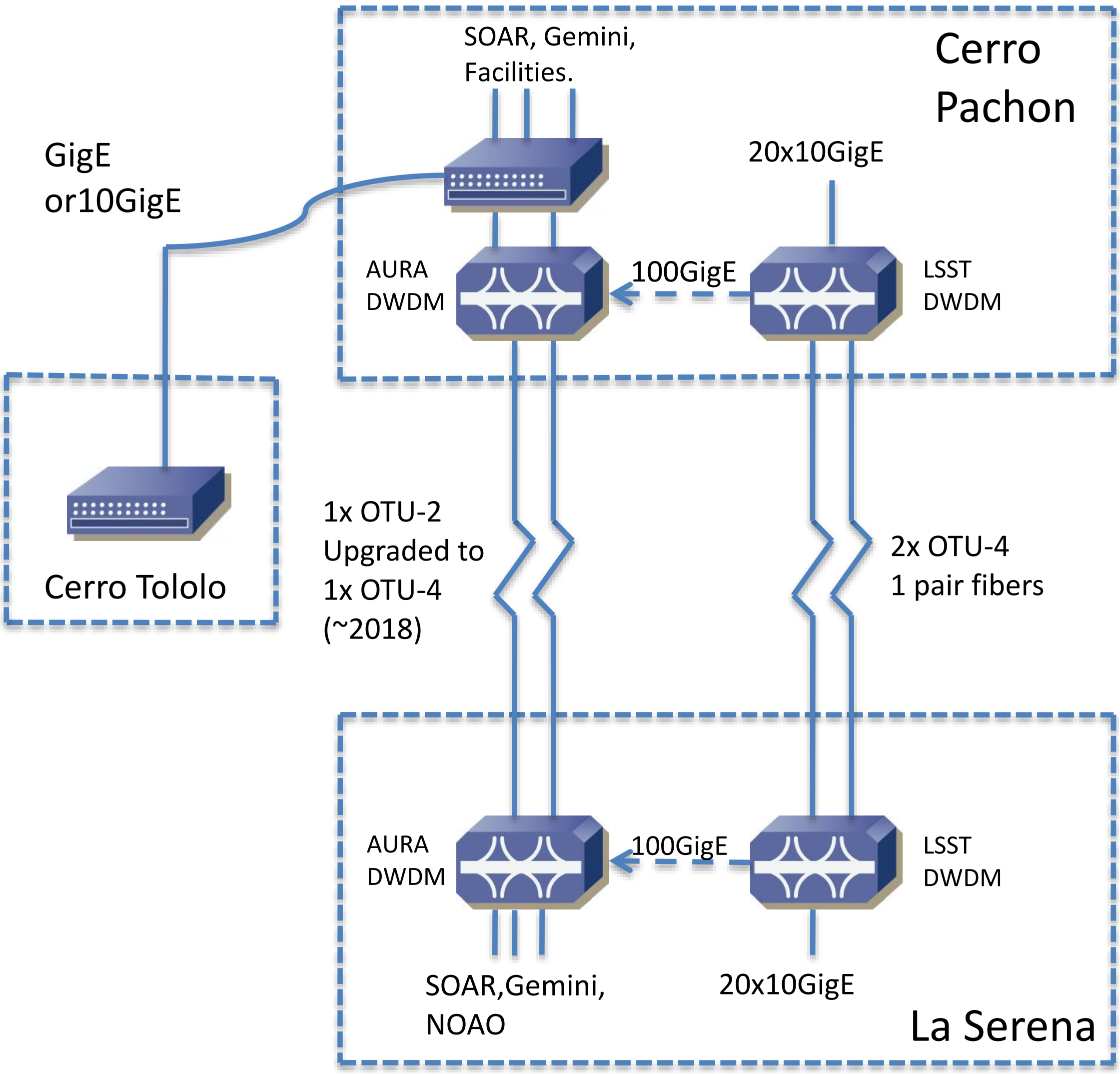

Figure 11: Mountain Summit - Base Network

Base Site Network Infrastructure

The Base Site network is situated in La Serena, Chile, at the AURA compound. Similar to the Summit Site, an Ethernet meshed core-access architecture will be developed. Major network-related functions performed at the Base are: storing a complete copy of the image data and data releases for use by the Chilean DAC, providing a temporary 14-day buffer of data in case of Base to Archive network outage, and transmission of raw and crosstalk-corrected image data to the Archive Center. Crosstalk-corrected image data is not retained at the Base but will be replicated and forwarded to NCSA. The raw image data will bifurcate in the Base Center computer room for storage in the local Base Center and for forwarding on to the Archive Center.

The bandwidth to the wide area network will be a permanent private 100 Gb/s, on which LSST will transfer 15 TB per night to the Archive Center. VLAN trunking will be employed across all switches between the Base Facility and Mountain Summit to maintain homogenous networks between Telescope, Camera, and Data Management traffic for security.

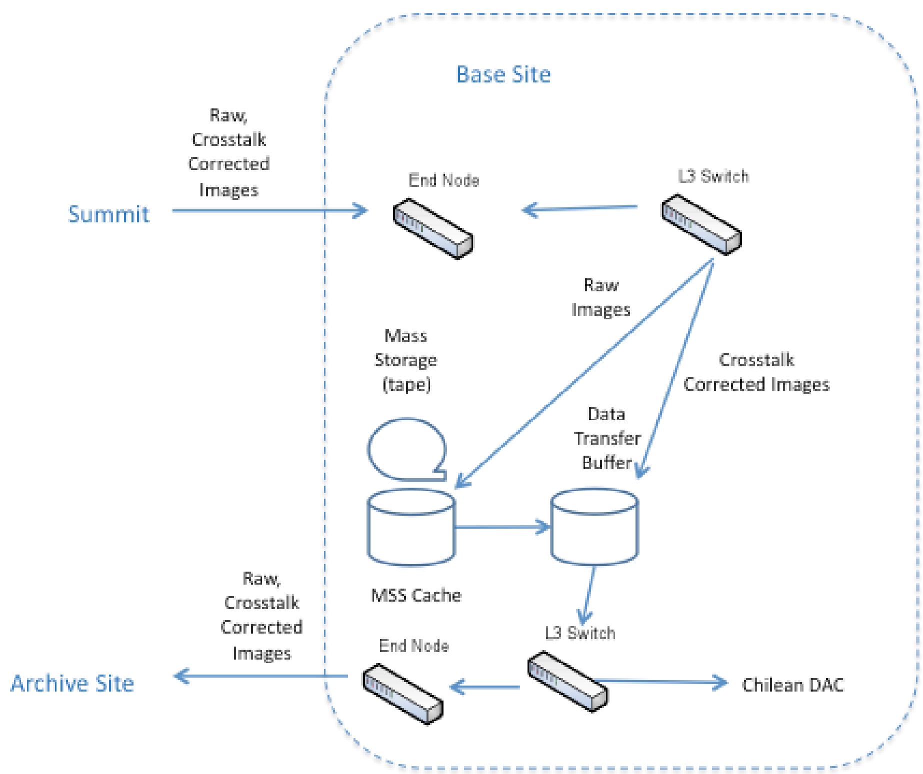

Figure 12: Base Facility Infrastructure Diagram

Central to the Base network are two virtual core switch/routers VSS, which are redundant in the event of failure. Every rack of machines have two top of rack switches configured to be redundant and connect to the core as fabric extenders.

In Figure 12 below the summit connects at the top left of the diagram showing two Lambdas of 100 Gb/s each. An optical image of 100 Gb is converted to 10x10GigE circuits and sent to 10GigE NICs in a rack of replicator servers where they are processed and forwarded to the Alert production machinery and DAC for storage. The images will also be added to a FIFO in the backup storage area in case of network outage to NCSA. The WDM rack will employ 10GigE interfaces on the customer edge so the image data is deliverd to a camera client without added delay. It is noted that NCSA will employ Infiniband as the network fabric for DAC Storage, Alert Production functions which will be copied in La Serena along with the Commissioning Cluster .

At the base there are racks for EFD, Visualization, Operations and Camera equipment. The connections shown to these functions are representations of redundant links to the core router/switches. Vlans can be trunked from sumit to base so these functions can be extended to the summit if required.

Each rack will house two top of the racks network access fabric extenders for reliability and redudancy. Each rack will be installed with PDUs with individual outlet V/A monitoring capabilities and network switchable power outlets.

In addition all services such as Wireless LAN, Email, Web, DNS, VOIP communications will be provided from the base.

It is assumed that the only time critical traffic is that of the image data from the camera to the replicators in La Serena. All other IP functions can be satisfied by regular switch and routing components at GigE or 10GigE bandwidths.

In the bottom right hand corner is shown the WDM router to REUNA/NCSA. All traffic will egress/ingress through this equipment to a gateway router/switch and Firewall. Non-scientific traffic, email,web,voip etc will be filtered at the gateway router and switched to a non-priority circuit either on LSST streams or AURA links.

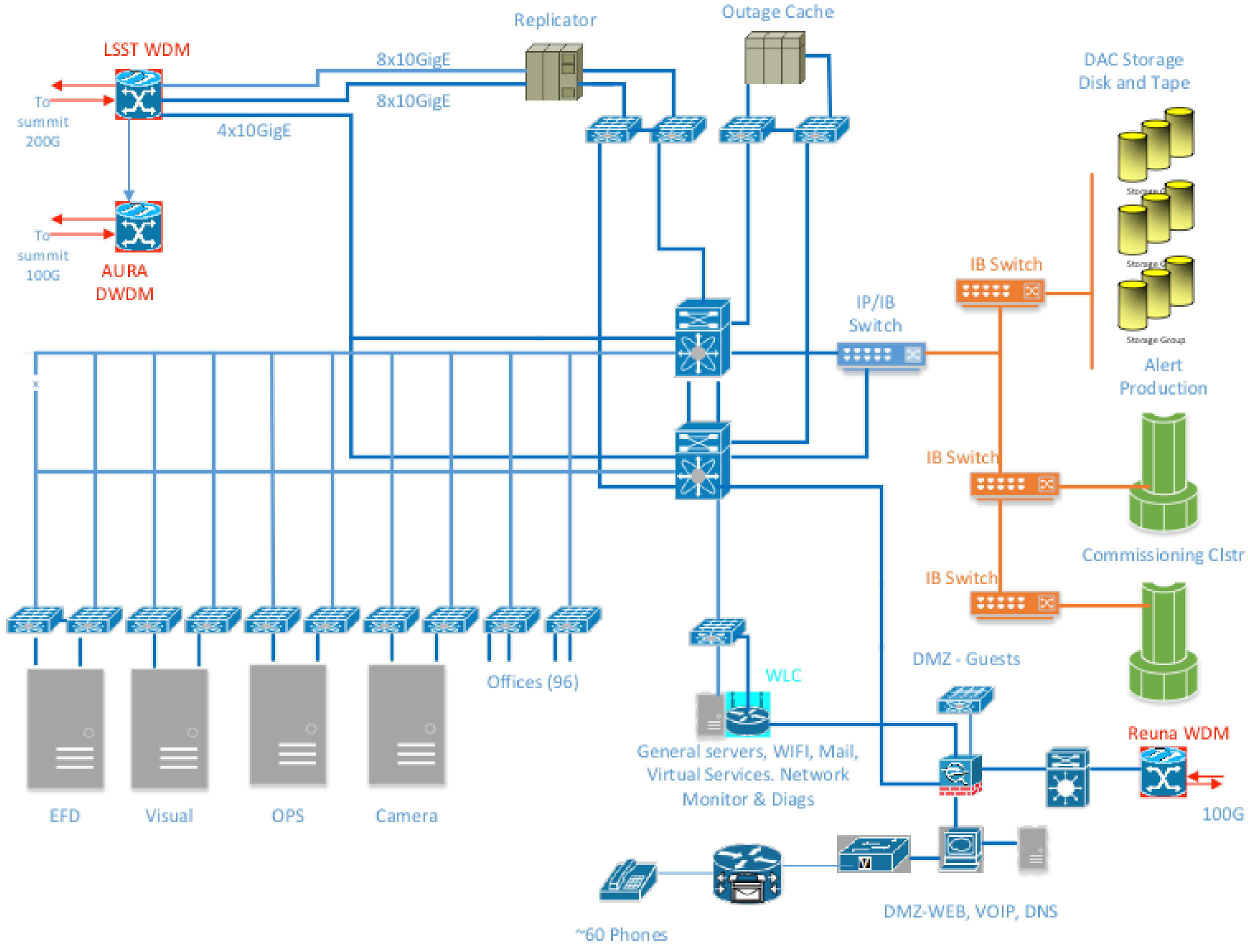

Figure 13: Base Site Network Infrastructure

Chile and International Network

Since 2003, USD $40 million (USD $12 million NSF) has been invested in U.S. – Latin America connectivity. The East coast of South America will be at 40 Gb/s by 2014, driven by a wide range of academic and research needs including high energy physics (HEP) and the Large Hadron Collider (LHC) project. Brazil (RNP) is investing internally for 100 Gb/s between Rio de Janeiro, Brasilia and São Paulo, and is formulating plans for upgrading the Brazil – U.S. links to 100 Gb/s.

This is a ring that will be provided via submarine and terrestrial networks owned by a broad consortium of network operators (including GVT, AmLight/FIU, Silica, Level3, Telecom Italia, and others) seamlessly available to LSST through a 15-year IRU and operations agreement with AmLight/FIU.

This primary network infrastructure is path diverse from Santiago up either side of South America. The Eastern side is via Santiago to Sao Paolo to St. Croix to Miami to Jacksonville to Chicago to Champaign. The Western side is via Santiago to Panama to Boca Raton to Jacksonville to Atlanta Chicago to Champaign. LSST-available guaranteed bandwidth in each side is 40 Gb/s. Provisions for bursting to available bandwidth (up to 100 Gb/s) are possible. These bandwidths presume continued investment in these networks by the NSF and Brazil. Contingency plans exist in case those investments are not sustained and a minimum of 10 Gb/s is guaranteed.

On the West coast, Santiago to Miami is currently 10 Gb/s with a link from Panama across the Caribbean to Miami. AmLight is in discussions with a number of entities to increase this capacity to 40 Gb/s. The capacity from Panama to San Diego is available on former Global Crossings, now Level3 but this known to be old cable and without significant investment ($M) not suitable for high speed links. CLARA is also working on terrestrial links to the U.S., however a legally impassable protected rain forest between Panama and Colombia will mean that other strategies will need to be developed for a potential CLARA backbone.

Two providers, REUNA, and AmLight, will be employed to reach the U.S. from La Serena and the entire circuit will utilize DWDM technology.

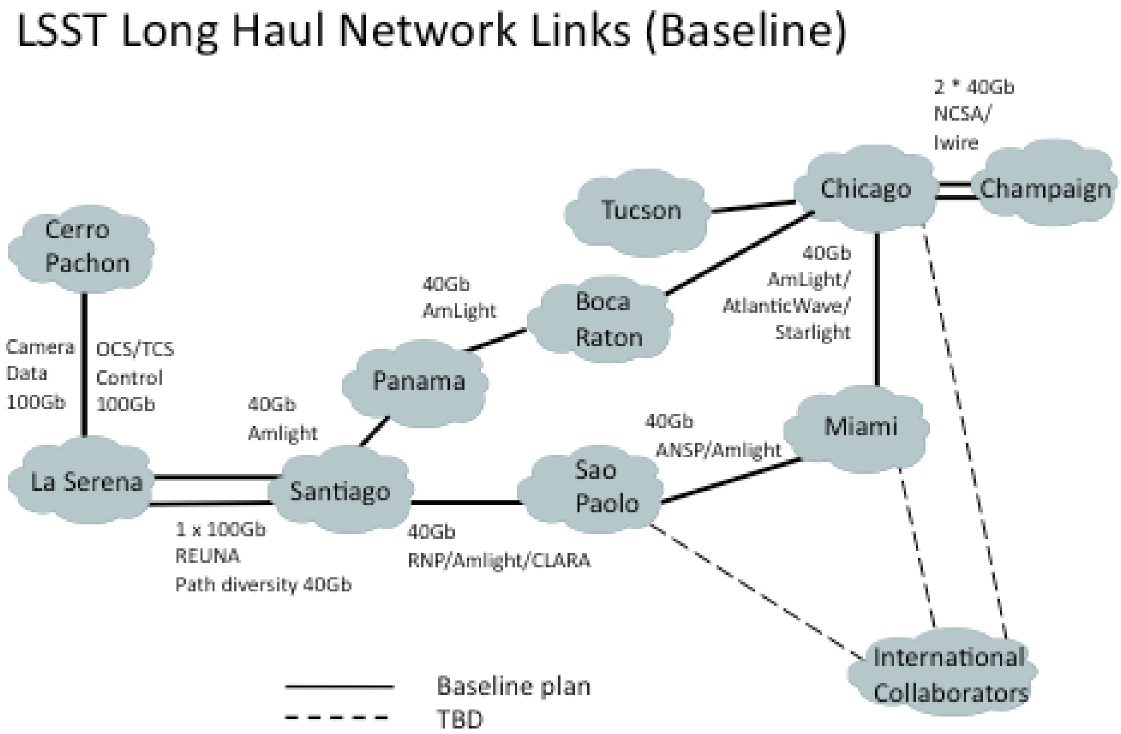

Figure 14: Chile and International Network

Chile National Circuit

Within Chile, REUNA (Chilean National Research and Education Networks) will provide bandwidth from La Serena to Santiago. This link, that provides the last mile circuit from Santiago, will transfer LSST data between the Base and Archive Sites, and will also be used extensively by the Chilean and Latin American communities to access the Chilean Data Access Center. The Latin American traffic will interact with Chile over CLARA, a South American network.

In 2009 REUNA formed a consortium with ESO named EVALSO and currently has a 10 Gb/s lambda from Santiago through La Serena to Antofagasta. LSST will provide REUNA with up front funding during the construction phase in order for them to upgrade their network and in return receive no annual service costs during the survey period. This will provide LSST with bandwidth for the transfer of all LSST data between La Serena and NCSA.

As a means of validating the AmLight and REUNA cost estimates, LSST approached commercial companies within Chile directly for bandwidth needs. The three operating telecommunications companies within Chile were approached to offer solutions for this capacity. Two offered viable technical solutions with an average monthly recurring cost of USD 50,000 with availability in 2012. Commercial letters of offer were made, insuring the creditability of the solution. The resultant costs quoted are unsustainable given the funding level, with no opportunity for access to dark fiber and the minimum bandwidth requirements of LSST.

LSST will utilize the services of REUNA, Chilean REN, for transport of data from La Serena to Santiago. AURA has had a working relationship with REUNA since 2005 in which it is an active member. Figures 14a below illustrates that there will be a dark fiber pair (blue line) in 2017 operated by REUNA with an IEU of 18 years from La Serena to Santiago which is inside a bundle of 96 fibers owned and utilized by the three major Telecoms in Chile. This fiber cable is new infrastructure and an integral part of services provided by the Telecoms to the North of Chile so in the event of failure or cut it will be repaired in the shortest possible time, i.e. 4-6 hours, by the Telecom. There will also be a redundant path to Santiago over an existing cable (red line) which will initially provision 4Gb and increase to 40Gb in time for LSST commissioning. This is a service provider and not fiber. The bandwidth will be available for LSST related traffic.

Figure 15a: La Serena to Santiago link

Figure 14b shows that on the dark fiber pair between La Serena and Santiago REUNA will provision two OTU-4 optical paths, one (100Gb) exclusively for use by LSST and the second for REUNA research and education network which extends to the north of Chile. Additionally AURA has reserved the rights for a further 9 Lambdas within the dark fiber for future use. The end equipment in La Serena will reside in the LSST Data Center and will be operated by REUNA.

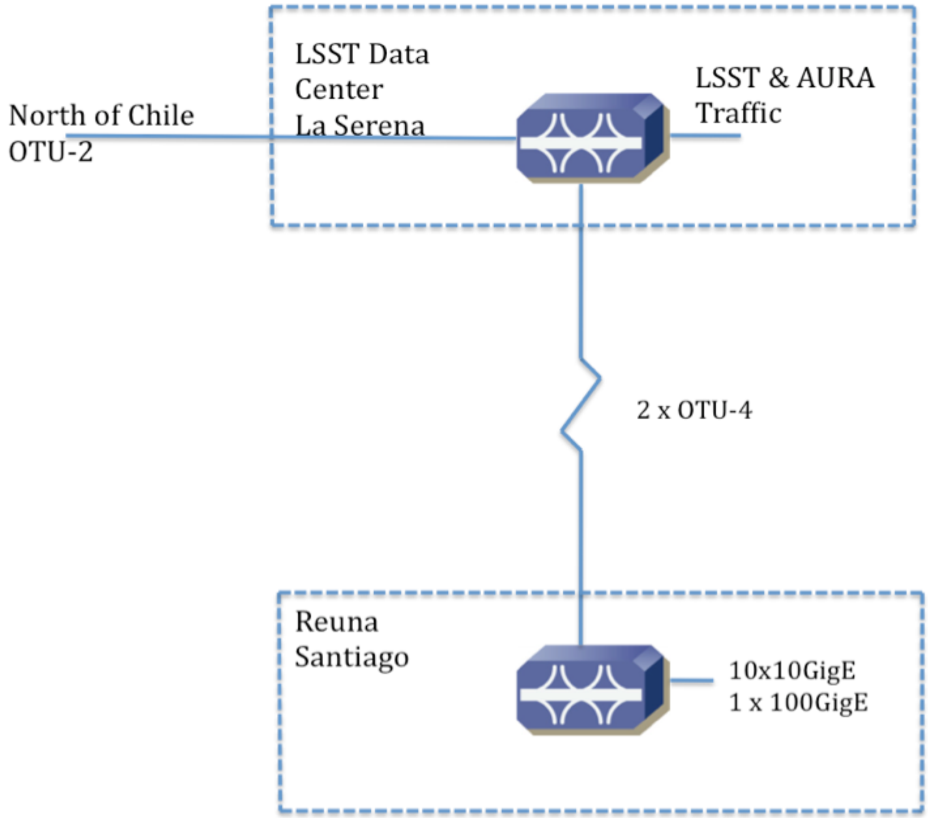

Figure 15b: La Serena to Santiago Link

International Circuit

Current carrier technology is typically capable of 100 Gb/s on a single fiber or "Lambda" but 10 Gb/s Lambda this can be considered a guaranteed minimum bandwidth available on these links during LSST operations. In fact, this exists on most of these links today.

Dedicated bandwidth can be acquired via a services lease, but that means that the carrier decides how to provision the service, rates can change from year to year, and this is typically not the most cost-effective option. An alternative is to acquire Indefeasible Right to Use (IRU) on the lambda or on the fiber. In the former case, the dedicated bandwidth is defined by whatever equipment the carrier is operating at the time. The latter permits the acquirer to put end equipment at any bandwidth on the fiber they desire up to the limits of technology, and increase it at their own discretion.

It is expected that these carrier-provided long-haul links will be upgraded to 100 Gb/s lambdas before the time of LSST operations. LSST's planned investments in long-haul networking during Construction, along with continued investments by AmLight's South American networking partners in Chile and Brazil, as well as potential involvement by commercial entities make this a virtual certainty.

The current design thus specifies 40 Gb/s as the baseline dedicated available bandwidth on these links. Further, it is our intent to negotiate IRUs on the fiber wherever possible, so that we are not tied to the carrier's equipment upgrade plans. LSST actively participates in the South American Astronomy Coordination Committee (SAACC) that leverages multi-observatory network bandwidth requirements to achieve the best possible terms with the carriers.

Negotiations have taken place such we have letters of intent from all the providers that a ring of 40 Gb/s guaranteed bandwidth will be available for use by LSST at the end of construction stage. This will also enable bursting to use available bandwidth (up to 100 Gb/s when required, for catch-up. Amlight (FIU) and Global Crossing have both quoted a 15-year IRU giving MTBF greater than 98%. This funding will be paid during the construction period to allay operational costs.

The link from Santiago to São Paulo will utilize a leased carrier line managed by AmLight. LSST will contract an IRU for a 40 Gb/s Ethernet circuit that has reliability better than 98%.

The São Paulo to Miami traffic will share the AmLight network of 40 Gb/s, provided by an NSF and Brazilian collaboration. However, LSST may increase the bandwidth to 100 Gb/s on a capacity available basis for catch-up, and negotiations with the carriers are being conducted in order to secure this option.

U.S. National Network

The U.S. national network infrastructure also has two “sides” relative to the path to Chicago. One side begins in Boca Raton at the termination of the international Pacific side network connection from the Base Site and encompasses the link from Boca Raton and on to the Archive Site via the west coast of Florida The other side begins at AmLight in Miami with the termination of the international network connection from the Base Site and encompasses the link from AmLight on the east coast of Florida to the Archive Center. Each side of the U.S. National network will support 40 Gb/s with burst capability of up to 100 Gb/s.

The national component of the link consists of the following segments:

- National Backbone from AmLight in Miami to StarLight in Chicago - data will travel via a 200 Gb/s (100 Gb/s protected) AtlanticWave link to both the FLR Jacksonville 100 Gb/s Internet2 AL2S port and (as a back up) WIX Washington D.C.100G AL2S Internet2 port, then through the matrix protected AL2S network to StarLight in Chicago
- National Backbone from Boca Raton to Atlanta to StarLight in Chicago – (Multiple options available, e.g. Internet2, NLR, ESNet, selection is TBD)
- From StarLight in Chicago to NCSA in Urbana-Champaign, IL - is described in the next section

For this network a Layer 2 path would be established from Miami to Champaign. A VLAN would be configured with dedicated bandwidth of 40 Gb/s and be configured to also allow for bursting to 100 Gb/s, dependent on the actual traffic on the link. A backup VLAN would be configured as well to provide an automatic backup path for the connection.

Figure 15 illustrates the basic network infrastructure assumed for the Archive Center. In addition to the 40 Gb/s connection to the Base Site, the Archive Center must have the following connections: 10 Gb/s to the R&E Layer 2 Network, 10 Gb/s to the R&E Layer 3 Network and 1 Gb/s to the Commodity Network. Note that except for the 40 Gb/s connection to the Base Site, the remaining connections can be shared, both internally and regionally, as long as the following requirements are met:

- Multiple simultaneous sustained 1 Gb/s single-stream data transfer to the R&E Layer 2 and Layer 3 networks.
- Headroom on the commodity network that allows reasonable bursting ability.

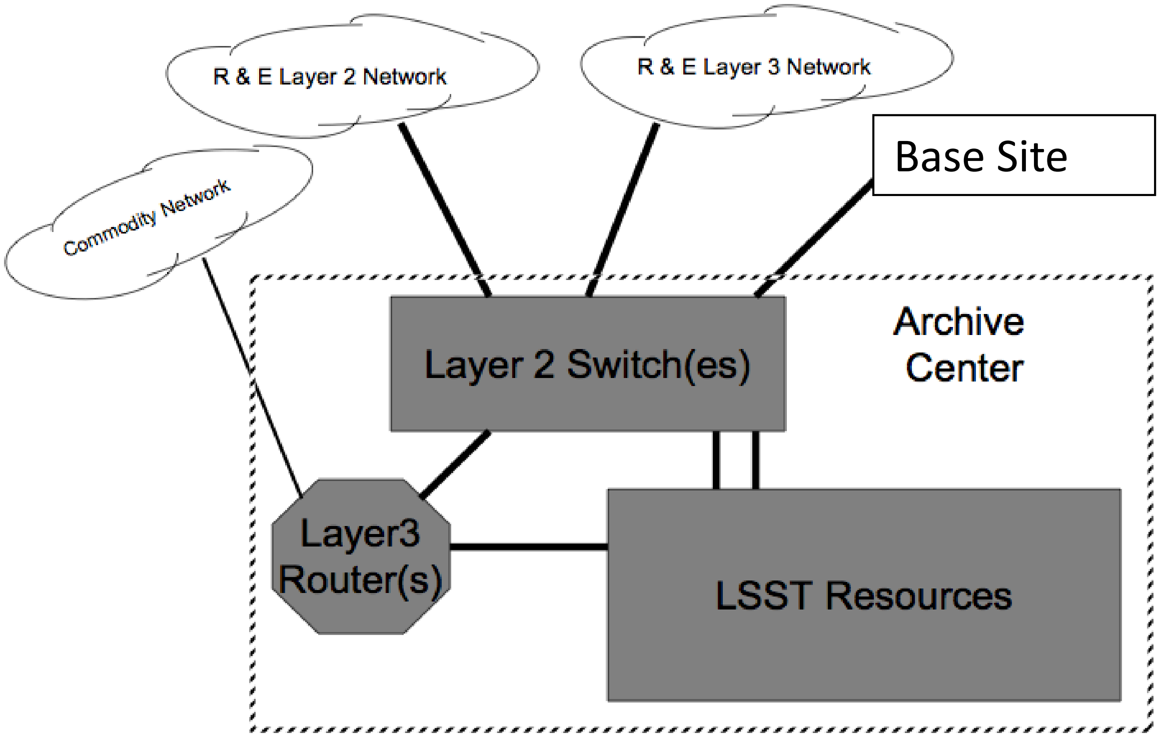

Figure 16: LSST Archive Network Infrastructure

Chicago to Urbana-Champaign Segment

LSST will utilize the institutional connectivity of NCSA on this segment. NCSA has always enjoyed a rich set of connectivity to Chicago. In 2002, as part of the Teragrid project and in collaboration with Argonne National Labs, NCSA built I-wire. A linear optical network over leased dark fiber from Urbana to Chicago in order to carry the multiple 10G circuits required for the Teragrid project. In 2010, NCSA transitioned away from running the I-wire optical network and migrated I-wire into the optical network known as ICCN, developed and managed by the University of Illinois – Urbana, Champaign. ICCN is an optically protected ring connecting Urbana-Champaign to multiple peering locations in Chicago and also to peering points in other Illinois cities such as Springfield and Peoria. See Figure 16 of the optical ICCN network for an overview of the paths and add/drop sites.

As of this writing, NCSA utilizes three 10G lambdas on the ICCN optical ring for connectivity to MREN, OmniPop, and XSEDE all located physically at 710 North Lakeshore Drive, commonly referred to as 710NLSD. Within 710NLSD are numerous peering opportunities with entities such as StarLight (for international peerings), ESnet (for peering with DOE sites), MERIT, ATT, CenturyLink, Zayo, Internet2 and many more. For a comprehensive list of networks and companies at 710NLSD visit: http://www.startap.net/starlight/NETWORKS/

In addition to the 3x10G circuits to Chicago, NCSA has 2x10G circuits to the ICCN layer3 network which are used to provide backup connectivity to the R&E networks as well as commodity connectivity to two different Internet providers. NCSA also has 5x10G circuits to the campus research network known as CARNE (Campus Advanced Research Network Environment) which has a 100G circuit to 710NLSD in Chicago terminating on Internet2 hardware and providing the Internet2 AL2S and AL3S services.

Started in Q1 of 2015, NCSA is in the middle of a project to substantially increase the WAN capacity into its facilities. The core routers and exit routers have been replaced with 100G capable hardware and the ICCN Adva optical gear is being upgraded to support 100G lambdas. The new 100G circuits will be provisioned on the ICCN optical ring in a protected capacity to mitigate the effects of fiber cuts or DWDM equipment failure. When completed, the new network topology will look like Figure 16.

NCSA will have four 100G circuits in production with three of the circuits terminating at MREN, Internet2, and CARNE. The fourth 100G circuit will be used to support specific science cases such as LSST. CARNE is also in the process of being upgraded to provide a 100G circuit to OmniPop in addition to its current 100G circuit to Internet2. NCSA and CARNE will establish bi-lateral peerings allowing for failover in the event of fiber or equipment outages.

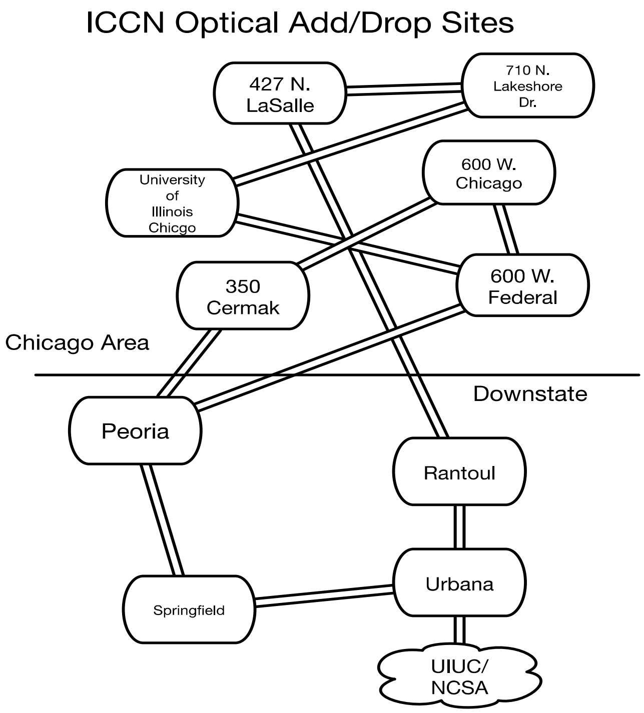

Figure 17 ICCN Optical Add/Drop Sites

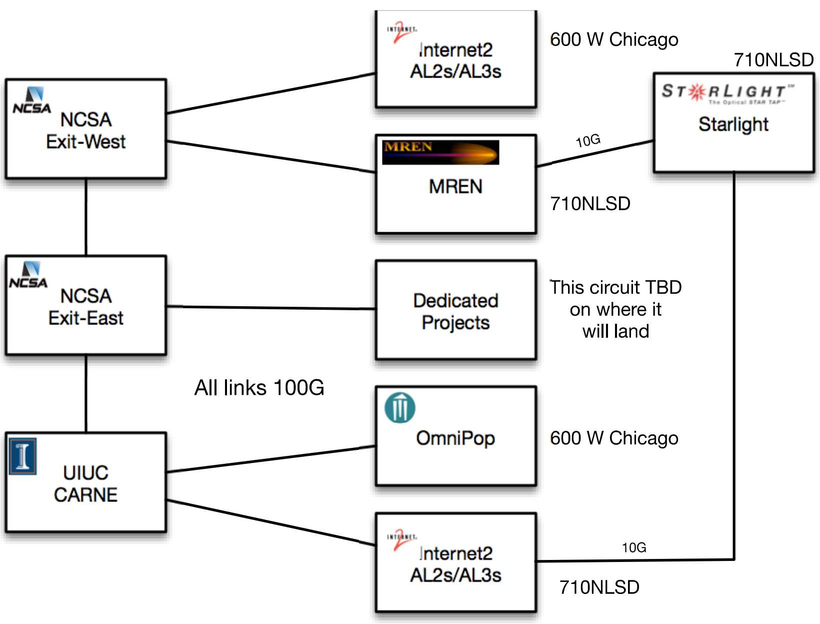

Figure 18 NCSA Network Topology by 2016

Data Access Center Network Connectivity

There are two project-funded DACs, one is in the Base Site in La Serena, Chile, and the second is in the Archive Site in Urbana-Champaign, Illinois. Figure 18 illustrates the network connectivity between the US national components of the LSST DMS, specifically between the Archive Center and the US Data Access Center (within the Archive Site) and between the US Data Access Center and the User Communities.

The Archive will transfer large data sets to both Data Access Centers (DAC). Each DAC will have the two most recent data release catalogs, the Level 1 catalog, and the as-yet unreleased next catalog data. Each DAC will have access to the full set of the raw images co-located at that site on tape, plus the current sky template and a 30-day cache of the most recent calibrated images on disk.

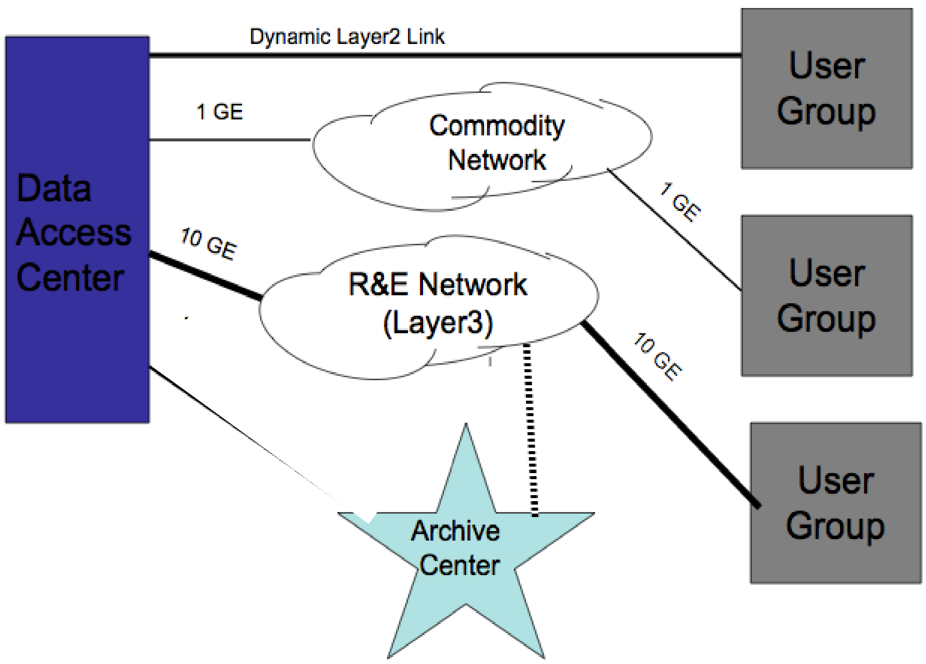

Figure 19: Data Access Center Connectivity

The Data Access Center provides data access and processing services on demand to the various LSST User Communities. Note that users do not access the Archive Center or Base Center. Instead, they access all LSST data sets via the Data Access Center. The primary data path between the DAC and institution-based User Community will be a 10 GbE to one (or more) of the Layer 3 based Research and Education networks such as the Internet2 Network, NLR PacketNet or DoE ESnet. Each DAC connection will be tuned to support, a minimum sustained data transfer rate of 1 Gb/s with supporting higher sustained data transfer rates encouraged.

User Communities not connected to one of the R&E backbones will access the DAC using the Commodity Internet.

The DAC will be able to support the dynamic allocation of Layer 2 infrastructure, such as the NLR Layer 2 Dynamic VLAN service or the I2 Network DCN service, in order to support large data transfers to an end user site. These services allow the allocation of point-to-point VLANS with non-dedicated or dedicated bandwidth for short durations (hours or days).

Future Variables

The above networking design assumes the ability to leverage, including cost effective pricing, existing research and education networks. Changes in these networks are expected over the years between now and LSST operations, and should be taken into consideration when selecting national and regional network providers and technology. Below is a summary of expected changes over the next few years.

The XSEDE backbone, a private Layer 3 network that connects associated supercomputing, visualization and data centers in the U.S., has been used as the national network infrastructure for the LSST Data Challenges. This infrastructure is funded through NSF grants and given the history of such grants, we expect repeated funding of this infrastructure as long as the research demand remains strong. The expectation is that the follow on project will provide similar infrastructure. That said, we are not depending on XSEDE for connections between the Archive Center and the Headquarters and/or science users.

National Lambda Rail (NLR) is a fiber network infrastructure based on a 25 year IRUs between NLR and Level3 that will begin to mature in December 2022. The fiber is lit with equipment that is owned and operated by NLR. NLR infrastructure supports both 10 GE and 40 GE technology, with a backbone capacity of 128 x 10 GbE and 160 x 10 GbE lambdas on the southern and northern paths, respectively.

The Internet2 network services are built upon 7-year contract for managed services with Level3. The basic service is built upon ten 10 GbE nationwide lambdas. The number of lambdas per segment can be increased based on demand. While the contract with Level3 expired in 2013, as with previous backbone transitions, Internet2 is looking at new technologies, providers and contracts. At the time of writing NLR and I2 have 100 Gb/s circuits planned for the near future.

Static and dynamic Layer 2 services have emerged and may provide a great opportunity for LSST to match specific project location and bandwidth needs cost effectively. The two services mentioned above, NLR's FrameNet based Dynamic VLAN service and Internet2's DCN are still under development. Dynamic provisioning using either service requires a specific software environment (Sherpa for NLR and Dragon or OSCARS for I2 DCN) and are not yet interoperable. The cost schedules for these services are changing as well, however both are expected to be usage based. As these services mature, commercial offerings from companies such as Level3 and Qwest are expected as well.

The network connectivity from Chile to the U.S. is being designed to maximize the committed and available bandwidth by collaborating with the AmLight (America's Lightpaths) consortium. The AmLight consortium (evolving from WHREN: Western Hemisphere Research and Education Networks) has been seed funded by the U.S. National Science Foundation's Office of Cyber-

Infrastructure (OCI-0441095 and OCI-0963053). This $0.75M annual network funding has historically been through cooperative agreements awarded to the Florida International University (FIU). Significant additional annual funding has been provided by other AmLight consortium members, including funding from ANSP (the Academic Network of Sao Paulo) of $3.8M and funding from RNP (the Federal Academic Network of Brazil) of $3.7M.

There is no reason to believe that the AmLight consortium or its successor (as AmLight is the successor of WHREN) will not continue to operate through the period of operations of LSST. However, the NSF cooperative agreements are re-competed every five years, without an OCI commitment to the program in perpetuity. Similarly, the ANSP support is renewed every five years, and the RNP support is renewed every three years.

As such the LSST support plan is to enter into a contract with FIU/AmLight to provide a minimum of 10 Gb/s path diverse capacity from Santiago to Chicago. This level would be contractually guaranteed irrespective of the evolving participation of other AmLight members. It is anticipated that through the continued consortium participation of ANSP and RNP, a bandwidth of 40 Gb/s will be available for LSST operational use.

LSST Brazil (through ANSP and RNP) could commit to continue AmLight funding for the benefit of LSST through the operational period. If LSST Brazil were to do this, the contract for bandwidth would be able to guarantee protected 100 Gb/s. The difference between 10 Gb/s and 100 Gb/s is significant, but not critical, to the science operations of LSST. It improves the reliability to transfer the image data in time to process the data for transient alerts at the Archive Center. Of course, with 10 Gb/s, LSST can also operate as in the original baseline, with transient alert processing done at the Base Site and still meeting the 60s latency requirement.

In 2011, LSST and the French particle phsysics institute IN2P3 signed a Memorandum of Intent to have IN2P3 provide computing, storage, communications infrastructure and labor to enable half of the annual Data Release processing to be performed in the CC-IN2P3 computing center in Lyon, France. This will include the provision of a high-speed network connecting CC-IN2P3 and NCSA. There are existing connections between CC-IN2P3 and Fermilab via the LHCONE network, and these connections may provide the foundation for the CC-IN2P3 NCSA link.

Appendix A: La Serena – Santiago Network Rings

Several institutions have already started initiatives to create path diversity between La Serena, Antofogasta, and Santiago Chile that LSST may leverage, as depicted in this Appendix. Leveraging these rings would further increase the reliability of the LSST data transfers and provide additional security against a cut on the traversal of the Andes, which has traditionally been the area most prone to such outages.

There are at least 2 distinct efforts currently underway to establish a ring including La Serena and Santiago, both of them involving Andean crossings and having part of the link in Argentina. There is also a ring that may connect Santiago with Buenos Aires that may be leverage by LSST. AmLight is coordinating with these efforts to leverage such a ring for LSST use, via the SAACC and direct discussions.

ALMA (Paso de Jamas)

Alma Site

This ring contains several segments.

1. ALMA OSF to Calama - New fiber install, currently 10 Gb/s, joint funded with Silica, Silica owns and operates, but ALMA has (25 year IRU) on one dark fiber. ALMA (or REUNA) operates end equipment on that dark fiber.

2. Calama to Antofogasta: Telefonica owns fiber and operates. ALMA funded, currently 2.5 Gb/s, IRU on the lambda, upgraded if carrier upgrades to 40 Gb/s.

3. Antofogasta to Santiago: EVALSO funded. Telefonica owns and operates fiber. Currently 10 Gb/s (split into 4* 2.5 Gb/s channels). IRU on lambda through 2020, upgraded if carrier upgrades to 40 Gb/s. Expect to happen in 2 to 5 years. Equipment last 5 - 7 years, and had been in use for ~2 years before this IRU started. To increase before carrier is ready, have to add

equipment or buy another lambda. REUNA currently operates end equipment to split into channels, would need separate investment on this equipment to go to 40 Gb/s. ALMA would prefer to change IRU to fiber, rather than lambda, to maximize flexibility for all users. Also discussing with REUNA to upgrade fiber. Given that La Serena is on link 3 already (CTIO donated equipment to REUNA, there until 2020), take advantage of window of opportunity in 2015 - 2017 to amend IRU to fiber. Otherwise, wait until 2020 and then have to renegotiate IRU.

At a point after congressional appropriation for LSST MREFC start, ALMA would need confirmation of LSST project funding by US, and an LSST commitment letter to provide budgeted $6.5M USD funding to REUNA from 2015 - 2020. This will allow REUNA and AURA to negotiate for IRU on fiber (vs lambda) and fund link after 2020. Our current baseline already meets this, just needs confirmation.

4. OSF to Santiago via Argentina (Paso de Jamas): Silica owns fiber. Expected within 5 years. Elastic capacity, 10 year IRU for one lambda, when carrier upgrades will upgrade to 40 Gb/s. Currently "gentleman's agreement" to only use as backup, not as another primary path, but inquiry about 100 Gb/s was considered not impossible.

Aguas Negras

There is another planned ring connecting La Serena and Santiago going up Elqui Valley and via the planned Aguas Negras tunnel into Argentina. This was presented to us by Florencio Utreras of Alice/CLARA. It requires a planned $800-900 million USD investment, apparently mostly Argentinian, that has to be made to make this tunnel a reality first. Thus the Paso de Jamas is considered more "real" as it doesn't require a tunnel, but instead simply runs along an existing pipeline for most of the way. This path would give also path diversity from both Cerro Pacon to La Serena, making it even more attractive. LSST could share installation of fiber on public road, and have an IRU on some number of fibers. Other paying customers on this fiber will likely drive a fast MTTR to meet our specifications. There is also a neutrino lab that might provide funding on this link.

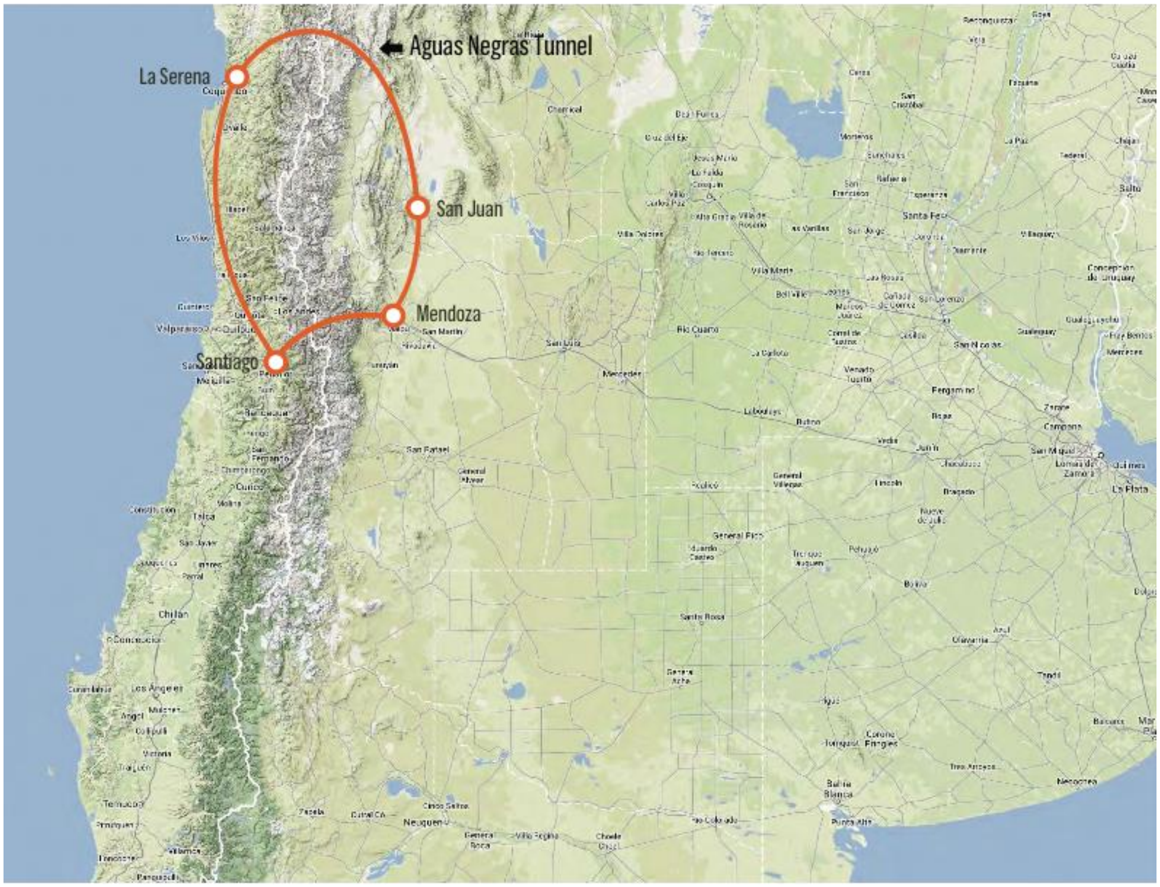

Aguas Negras Tunnel

Chile-Argentina Ring

AmLight is also in discussions with Entel and other entities about completing a third ring, and it is considered likely to be available in LSST's operations period. The solid lines are existing and owned by Silica. The ring is partially complete with ENTEL at this time; meaning that this map shows dark fiber, there is additional leased capacity that makes this work. Odds are very good that the ring will be complete because of commercial interest. There are detailed plans to complete a Malargue to Santiago link.

Chile-Argentina Ring

References

[Freemon2013] Freemon D. M.: Optimizing Throughput on Guaranteed Bandwidth WAN Networks for the Large Synoptic Survey Telescope (LSST). CoRR abs/1303.7467 (2013)

[LSST14313] Dossa, D., Freemon D.M., Kantor J., Shipping Media Analysis, rev. 2013

[LSST11811] Lambert R., Feasibility of Redundant Fiber from Mountain to Base, rev. 2013

[LSST11918] Ibarra J., Cox C., Jacque S., Lambert R., Kantor J., LSST Network Operations and Management Plan, rev. 2013

[LSE-30] LSST Observatory System Specifications (OSS)

[LSE-61] Data Management System Requirements (DMSR)

[AURA Chile MOA] Memorandum of Agreement between AURA and Chile

[AURA LSST MOA] Memoranda of Agreement between LSST and AURA