Rubin Observatory

Security of Rubin Observatory data

William O'Mullane

2020-05-20

1 Introduction

There is a perceived risk that LSST images may be snooped during transmission from the telescope. This concerns the images used for alert processing which are transmitted within a minute and provide a valuable resource for identifying moving objects including satellites. Though our processing will ignore satellites other parties may be interested in this information.

In an email during September 2019 Steve Kahn indicated that all data should be encrypted. This would include transfers to Europe.

1.1 Baseline

The Rubin Observatory construction project has been built with academic level security in mind. The Information classification policy LPM-122 classifies data as "User Protected". The DM Information Security plan **?** states the "majority of network traffic will not require confidentiality"

Since our astronomy data is research data with no intrinsic value no great efforts have been made to secure the data nor the network it is traveling on. The network has been designed, and now largely implemented, for high throughput not for high security.

The baseline is for encryption of controls but not data i.e. authentication is encrypted, data transmission is not.

It would be extremely useful if we agreed on the security rating of LSST data (or subsets of it) as per NIST (Stine et al., 2008). Naively one would assume the security objective would be *Availability*, the potential impact would be *low* for confidentiality, availability and integrity. Hence the Security Category (SC) would be {low,low,low} in NIST terms.

2 Which data is sensitive ?

In communications thus far and in the security summit held on 6^{th} April 2020 all data has been considered.

We believe the vast majority of the 20TB of nightly images are not of a sensitive nature. It would be useful to understand if this is so, especially if delaying this data some period of time is sufficient to desensitize it. If this is the case do we need to encrypt all data ?

If there are patches of the sky which are sensitive can they be avoided at certain times ? e.g. Geo synchronous orbit stripe. Are there other rough orbits we should delay data from ?

Or is it really the entire northern sky?

3 Network security

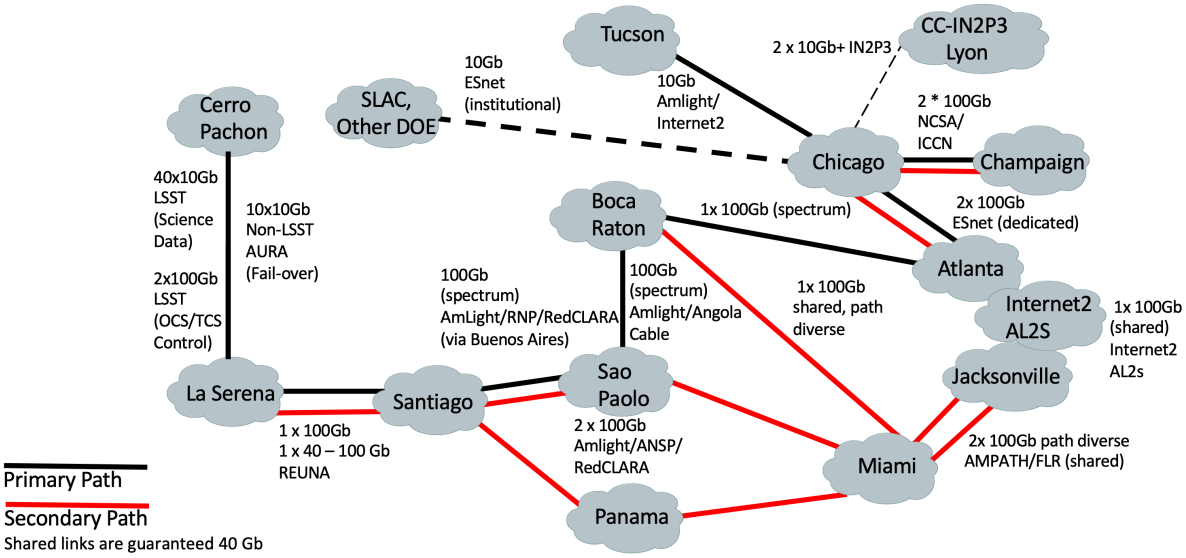

FIGURE 1: Rubin Observatory Network topology for FY22, the primary route is on dedicated lines, the back route is on shared equipment. Securing this beyond the level of the provider will be close to impossible.

On the mountain and in the base facility LSST networks and computers are in rooms requiring ID card access and the compounds have 24/7 security staff. We could increase security but it would be costly and probably not very popular with staff i.e it would require raising the

security level of the entire facility to something like FSL[1] 2 .

The transfers from Summit to Base and Base to NCSA go over private networks that are not part of the public Internet. We have our own dedicated fibers running from the mountain to the base [LSE-78], intercepting transfers would require physical access to the fibers - tapping those would probably disrupt our network at least temporarily. For an added security we could also monitor fiber attenuation. The network is also actively monitored with Bro[2] system. Hence we would at least known something was up. Though technically feasible by perhaps bribing or coercing staff somewhere this seems unlikely - a physical tap should also be noticed during fiber inspections. The image as transferred over this line is also in an obscure format (though could be put together with some effort).

International transfers to IN2P3 are still being worked out at this time. This may be done on shared commercial carrier. One assumes this is relatively secure though not perhaps as secure as our dedicated lines. ESNet could be used for this transfer as well if it did not originate at NCSA. There seems no particular reason to do this transfer from NCSA as opposed to directly from Chile or landing it temporarily at some ESNet endpoint. One may consider the open storage network (Núñez Corrales et al., 2018) for this also though its not currently the baseline.

4 Transmission security

It has been indicated, through DOE, that holding data for some time before releasing to the collaboration is an acceptable approach to securing the images - this can be ensured by limiting the LSST personnel who have access to the files during the night. We would transmit images for alert processing in near realtime to NCSA. We have suggested a six hour delay - there has been no indication of how long the delay needs to be - is one hour sufficient to prevent detection of satellite maneuvers ?

Currently alert images from the base are to be transmitted using BBCP[3]. The control channel is encrypted but the data channel is open - this is considered secure in the scientific world (Hanushevsky et al., 2001). Hence theoretically packets snooped on their way to Illinois could

[1] https://emilms.fema.gov/IS1172/groups/85.html
[2] https://www.corelight.com/about-bro/how-bro-works/
[3] BBCP is an alternative to Gridftp when transferring large amounts of data

be extracted and processed.

Our preference would be to do this in software using TLS based protocols such as HTTPS or FTPS to replace BBCP. This may cause some performance hit. We could encrypt the data channel. This would cost us in compute but could be similar to compression and not push up the number of cores needed for transfer. We would also have some software modification and setup there are projects like WireGuard[4] intending to do this at the kernel level but they are not suitable for Rubin use yet. There are also examples of hardware solutions for this which would probably be affordable, INRIA in Muehlberghuber et al. (2013) have build a FPGA based AESStandards (2001) implementation which gives 100Gbit/s network rates with μs additional latency. Table 3 of Muehlberghuber et al. (2013) lists the hardware which looks modest enough - we work with INRIA already. There is no price given in the document, but considering the cost of components and a contract to put it together one *might* consider it to be under $1M, but we would need to get a quote. This is probably more complex than we might want to get into.

We could encrypt the images themselves before transmission or even on the FPGA of the camera data acquisition (DAQ) system. The DAQ is complex construction prone to delays so we would rather not jeopardize the project by introducing changes there. Encrypting the files outside would mean more CPU and I/O - it would add a delay in the alerts processing - it may take as much as 20% of the alerts time budget (1 minute). This would not be very welcome in the science community. This would require some effort on our side to add encryption to the processing chain but this should be order some hundreds $K depending on how *secure* it is required to be.

4.1 Keeping it in Chile

The original baseline was to do alert processing in Chile in the base facility. Some years ago support for computing in Chile was sub optimal but has improved significantly - there is space in the base facility to hold machines for this. This would make the base data center the prime target of any data attack so we may need to review security there. However if we agree the base/summit are secure we could avoid on the wire snooping by doing alert processing in the base facility. We would also then have to hold images there for the appointed embargo time before releasing them to the community. We would need to do some work on the cost

[4] https://www.wireguard.com/

impacts of this but they should not be insurmountable.

5 High level summary

Though the reason behind the call for security and the level required remain totally unclear NSF and DOE asked for a brief summary of possibilities. One should also consider which data we are talking about see Section 2. Items here are not costed but an indication is given in terms of low (possibly within cost), moderate (some $100Ks), high (>$1M) Any change should be properly costed.

Security idea	Rough cost level
Delay/degrade image info. If the precise position of small objects is the driver then not providing accurate shutter times would make precise positions inaccessible	Low to Moderate.
Delay some or all images. Depending on how secure this needs to be and where it is done the cost scales.	Low to Moderate.
Do alerts in Chile. If we want to control image access for a longer period we could consider alert production in Chile. The hardware budget would remain the same but we may require extra support in Chile.	Low to Moderate.
Encrypt all images. This would have to be done on the summit or in la Serena before hitting the long haul network. Possibly no new hardware needed but a change in software.	Can be Low
Transmission Layer Encryption. Network encryption would probably require new hardware.	Moderate
Physically securing the network. This will be next to impossible and would probably require new network agreements. This however would be the only way to ensure no packet snooping.	Very high
Avoid certain time coordinates. This would require changing the scheduler to provide more constraints on pointing.	Moderate

A References

Núñez Corrales, S., Cragin, M., White (Wonders), A., et al., 2018, doi:10.13140/RG.2.2.31543.78249

Hanushevsky, A., Trunov, A., Cottrell, L., 2001, In: In Proc. of the 2001 Int. Conf. on Computing in High Energy and Nuclear Physics (CHEP 2001), Beijng, URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.132.2288&rep=rep1

[LSE-78], Lambert, R., Kantor, J., Huffer, M., et al., 2017, *LSST Observatory Network Design*, LSE-78, URL https://ls.st/LSE-78

Muehlberghuber, M., Keller, C., Gürkaynak, F.K., Felber, N., 2013, In: Burg, A., Coşkun, A., Guthaus, M., Katkoori, S., Reis, R. (eds.) VLSI-SoC: From Algorithms to Circuits and System-on-Chip Design, 1–20, Springer Berlin Heidelberg, Berlin, Heidelberg

[LPM-122], Petravick, D., 2015, *LSST Information Classification Policy*, LPM-122, URL https://ls.st/LPM-122

Standards, F.I.P., 2001, Announcing the advanced encryption standard (aes), URL http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.197.pdf

Stine, K., Kissel, R., Barker, W.C., Fahlsing, J., Gulick, J., 2008, INFORMATION SECURITY, 31, URL https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-60v1r1.pdf

B Acronyms used in this document

Acronym	Description
AES	Advanced Encryption Service
BBCP	BaBar Copy Program
CPU	Central Processing Unit
DAQ	Data Acquisition System
DM	Data Management
DMTN	DM Technical Note
DOE	Department of Energy
ESNet	Energy Sciences Network
FPGA	Field-Programmable Gate Array
FSL	Facility Security Level
FTPS	File Transfer Protocol Secure
HTTPS	Hypertext Transfer Protocol Secure
IN2P3	Institut National de Physique Nucléaire et de Physique des Particules
INRIA	French National Institute for computer science and applied mathematics

LDM	LSST Data Management (Document Handle)
LPM	LSST Project Management (Document Handle)
LSE	LSST Systems Engineering (Document Handle)
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
NCSA	National Center for Supercomputing Applications
NIST	National Institute of Standards and Technology (USA)
NSF	National Science Foundation
SC	Science Collaboration
TLS	Transport Layer Security