



Vera C. Rubin Observatory
Rubin Observatory Operations


The LSST Camera commissioning

LSST Camera Team, Pierre Antilogus, John Banovetz,
Alexandre Boucaud, James Chiang, Claire Juramy-Gilles,
Sean Patrick MacBride, Stuart Marshall, Theo Schutt,
John Gregg Thayer, Yongqiang (Brian) Qiu, and Yousuke Utsumi

CTN-007

Latest Revision: 2026-09-10

DRAFT



Abstract

This technote describes what we found during the LSST Camera commissioning.

Draft

Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased.	Utsumi

Document source location: <https://github.com/lsst/ctn-007>

Draft

Contents

1 Introduction	1
2 Camera operations	1
2.1 Recovery of CCDs	1
2.1.1 R30/S12: dead-looking segment and suspected short	1
2.1.2 R30/S12: running two of the three CCDs on Reb1	6
2.1.3 R20/Reb2/S21: faulty RD voltage read-back	7
2.2 Cryo system	11
2.2.1 Brief introduction to the refrigeration system	12
2.2.2 Architecture and Thermal Design of the Cryostat	14
2.2.3 Architecture of the Cryogenic Systems	17
2.2.4 Cryogenic Mixed-Refrigerant System	22
2.2.5 Brief introduction to the temperature sensitivity issue of the Cryo refrigeration system	28
2.2.6 Detail of the cryo-temp sensitivity issue	29
2.2.7 Solution to temperature sensitivity issue	30
2.3 CCS development	36
2.4 Firmware development	38
2.5 Observation control	42
2.6 Filter Exchange System Commissioning	42
2.6.1 Hardware: problems of youth & fixes	42
2.6.2 Slow Control Upgrade	43
2.6.3 Performances	44
2.7 Shutter	46
2.7.1 Shutter Timing issue	46
2.7.2 Shutter blade trajectory fitting	54
2.8 Guider star finder	54
2.9 Historical telemetry of infrastructure cut events: electricity, glycol, dynalene and air	54

3 Camera optimization	54
3.1 Sequencer changes	54
3.1.1 Long-range correlations	54
3.1.2 3 s sequencer file development	56
3.1.3 Natural frequency	59
3.2 Re-implementation of "clear" operation during the idle time	85
3.2.1 Introduction	85
3.2.2 Background/Context for the sequencer, DAQ and CCS	85
3.2.3 IdleClearing	86
3.2.4 Science and Guider Details	86
3.3 Electro-Optical Testing	88
3.3.1 The eo_pipe Software Package	88
4 Camera on-sky performance	90
4.1 CCD Sensitivity to High Flux	90
4.1.1 Test Description	90
4.1.2 Data	91
4.1.3 Method	91
4.1.4 Results	91
4.1.5 Conclusion	92
4.2 Pinhole repeatability tests	93
4.3 Phosphorescence	95
4.3.1 The difference between phosphorescence and persistence	95
4.3.2 Previous characterization of phosphorescence	96
4.3.3 Phosphorescence in u-band exposures	96
4.3.4 Phosphorescence in g-band exposures	97
4.3.5 Conclusion	97
4.4 Persistence	98
5 Acknowledgements	100

6 References	100
7 Acronyms	101

Draft

The LSST Camera commissioning

1 Introduction

2 Camera operations

2.1 Recovery of CCDs

2.1.1 R30/S12: dead-looking segment and suspected short

R30/Reb1 could not be powered on. During powerCCDsOn the OD current rose steeply as the output drain (OD) of the S12 slot was raised, and the over-current protection tripped within about 100 ms, ending the sequence before the CCDs reached their operating state; three such trips were recorded (§2.1.1.2). R30/Reb1 and the CCDs belonging to it were left off.

REB5 has no per-segment OD control, so the fault could only be isolated at the level of the CCD: S12 is powered partially, with its OD setpoint reduced to 0.15 (measured ODV $\simeq 2$ V), while S10 and S11 on the same REB operate normally (§2.1.2). The configuration has been in effect since 2025-06-25, at the cost of one CCD out of the three on Reb1.

Only with the REB powered in that state could the segment currents of S12 be read again. They are all low, as expected with OD off, and Seg17 sits at the noise level. The REB measures each segment's current on the low side of its current source, so a segment whose output node is held near 0 V reads $\simeq 0$ mA. The discussion of these data points to a short between the output node of the affected segment (OS, the source of the output transistor) and FS, which is at or near ground; a short involving OD was also raised as plausible, given the large OD current seen at turn-on. An ITL sensor showed the same signature about five years ago and was concluded to have a short.¹

2.1.1.1 Location and pin arrangement. R30 is an e2v raft; Reb1 reads S10, S11 and S12 (Fig. 1). The read-out section of the CCD250 is organised in a repeating group; in CCS names

¹Summary on dead looking channels, LSSTCam Confluence, <https://confluence.slac.stanford.edu/spaces/LSSTCAM/pages/240279059/Summary+on+dead+looking+channels>.




FIGURE 1: LSSTCam focal plane layout, after Plazas Malagón et al. (CTN-001, Fig. 1). The large red box marks R30 = RTM-012 and the small box indicates the sensor S12.




FIGURE 2: Pin arrangement of the CCD250 read-out section in CCS names, drawn after the pin order given in the CCD250 Design Report (e2v technologies, 2009). The names are explained in the text.

the pin order is S1 S2 S3 FS OD1 OS1 RD1 clamp1 RG1 OG1 BS GD FS OD2 OS2 RD2 clamp2 RG2 OG2, preceded by the parallel clocks P1–P4 (Fig. 2). Here P1–P4 are the parallel (image) clocks and S1–S3 the serial clocks; OD is the output drain, OS the output amplifier source, RD the reset transistor drain and RG the reset transistor gate; OG is a DC bias at the end of the serial register, ahead of the sense node and the reset transistor; FS and BS are the front and back substrate and GD the guard diode. The clamp gate (*Reset 2* in Fig. 4) is connected to P4.

2.1.1.2 Three over-current trips. Three attempts to power on R30/Reb1 failed with an OD over-current condition (Fig. 3, Table 1). The log of the first reads

```
[2025-04-07T20:51:27.238+0000] SEVERE: , R30/Reb1/Bias2 OD2 Overcurrent measured on
ODI: 0.213 > 0.082 (odIMax (A))
```

followed by `powerCCDsOn()` FAILED for R30/Reb1 and, 0.8 s later, an automatic `powerCCDsOff()`. In each case the DAC was set back to zero immediately, returning OD to 0 V within 100–140 ms of the step. Bias voltages are read every 5.3 ms and the parent currents every 33 ms, so each trip is documented by only a few samples and the voltage at which the current peaked is loosely constrained.

TABLE 1: The three failed power-on attempts. ODV is `focal-plane/R30/Reb1/S12/ODV` and ODI is `focal-plane/R30/Reb1/ODI`, the REB-level total. The threshold `odIMax` is 82 mA. Δt is the interval between the step and the sample that exceeded it.

Event	Time (UTC)	OD step	last ODV	peak ODI	Δt
1	2025-04-07T20:51:27	0 → 15 V	14.94 V	213.1 mA	~ 140 ms
2	2025-04-07T20:54:32	15 → 22.3 V	22.27 V	220.6 mA	~ 100 ms
3	2025-06-03T19:09:08	15 → 22.3 V	22.3 V	> 83 mA	~ 100 ms

Taking the excess current of the two April events to flow through a resistive path, $R_{\text{short}} = V/\Delta I$ gives 125 and 208 Ω (Marshall, 2025b). Each April trip is accompanied by a rise of the cryostat pressure from 1.31×10^{-7} to 1.47×10^{-7} Torr, decaying over a few minutes, and by a ~ 0.25 °C excursion of `focal-plane/R30/Reb1/Temp9`, indicating localised power dissipation at the fault site.

2.1.1.3 Interpretation. A shorted segment reading zero current follows from where the currents are measured (Fig. 4). Each segment current is monitored on the low side of that




FIGURE 3: Trending of the first power-on attempt (Marshall, 2025b). Top: the REB parent currents, with ODI (green) leaving its baseline as the bias sequence proceeds and reaching 213 mA. Middle: the clock rails. Bottom: the CCD bias voltages; the last trace to rise, to ≈ 15 V, is the OD of S12, and all biases return to zero at 51:28 when the protection has fired and powerCCDsOff runs.




FIGURE 4: The e2v output stage and the first stage of the REB circuit, drawn by C. Lage (Lage, 2025). The two on-chip source followers (*Output1*, *Output2*) are coupled through $C1$, with the clamp *Reset 2* on that node driven by $P4$; *Output2* drives the OS pin into the current source $J1$ on the REB, whose return to ground through $R3$ and $R2$ carries the per-segment current monitor. The suspected short is from the OS node to ground.

segment's current source, its return to ground, whereas the OD current is measured only as a total, for the CCD's OD supply as a whole. If the output node of a segment is held at 0 V by a short to ground, the current source cannot force its current and the segment monitor reads 0 mA. An open connection reads the same, since a current source cannot impose a current into an unterminated node either.

Bias2/OD2 in the log message is the OD supply of S12, and its current is measured for that supply as a whole, so a short anywhere in the CCD raises it: the over-current localises the fault to S12 but does not identify the segment. It is the zero reading on Seg17 that points to that segment. It was also noted in the discussion that the shorted node need not be at ground potential but could sit a couple of diode drops above it; the current source still needs some positive compliance voltage, so the effect would be the same, through a diode rather than through a resistance.

The ≈ 2 V that ODV still reads with the OD DAC at essentially zero has its own explanation, and is not evidence about the CCD. The current source on the REB is not turned off in that state: its gate DAC (csGate) passes ≈ 1 mA at a setting of 0, while shutting the 2N4416 JFET off would take about -3 V on its gate, which the REB cannot supply. That current flows through the on-chip output transistor and the two $1\text{ k}\Omega$ resistors in the JFET source, developing ≈ 2 V at that node, which pulls up OS and with it OD; the same floor is measured on a healthy CCD at UC Davis (Snyder, 2025). The question raised at the time was whether current could flow from the REB circuit into the OS side of the CCD and across the short to ground; at ≈ 1 mA and ≈ 2 V it was judged small enough to proceed.

2.1.2 R30/S12: running two of the three CCDs on Reb1

2.1.2.1 Concept. S12 is powered only partially, leaving its ODV at ≈ 0 (Marshall, 2025b). The OD setpoint is 0.15, for which ODV measures ≈ 2 V and decreases further as the back-bias is applied. RD and GD stay nominal, so the CCD remains protected when hvbias is turned on, while the OD current stays far below odIMax . S10 and S11 are powered normally; S12 produces no signal.

2.1.2.2 Firmware and CCS changes. The REB firmware protects against writing GD, OD or RD below hard-coded thresholds and prevents closing the back-bias switch while any of them is below its threshold, and the CCS power-on and power-state methods both reject an OD

of ≈ 0 V. Firmware version 0x31395011 (Table 9) makes the bias thresholds readable by CCS (registers 0x401100–0x401124) and adds a per-sensor sequencer override (0x400010–0x400012); a dedicated instance of it for R30/Reb1, identified by R30_Reb1 in the design name and checked before loading to EEPROM, carries the OD threshold set to zero. On the CCS side the power-on and power-state tests were changed to accept ODV ≈ 0 , and for the S12 bias block odMin was lowered from 21.30 to 0.0 while odTo1, odZeroErr, gdTo1 and odiAmin were widened; odIMax was left at 0.082.

2.1.2.3 Validation and behaviour on the summit. The configuration was exercised beforehand on a single e2v CCD at the UC Davis beam simulator, which found ODV settling at 2.01 V with the back-bias off, ODV, ODI and the per-segment currents all decreasing as the back-bias is raised, no anomaly in the supply voltage or current during bias-image acquisition, and bias images resembling those of a powered-off CCD (Snyder, 2025). The same behaviour was seen on the summit during the back-bias ramp of 2025-06-25: as rebpower/R30/Reb1/hvbias/VbefSwch was stepped from 0 to 50 V, focal-plane/R30/Reb1/ODI fell from 66.4 to 55.3 mA and .../S12/Seg15/I from 0.83 to 0.29 mA, .../S12/Seg17/I stayed at the noise level, and the cryostat pressure was unchanged.

2.1.2.4 Operational consequences. The configuration recovers two of the three CCDs on Reb1 and keeps the raft protected during hvbias turn-on, at the price of S12. Two points follow from how it is achieved: the special firmware must remain installed on this REB for the hvbias threshold test to pass, and the relaxed parameters weaken, for this REB, the tests meant to catch this class of fault; both should be tracked in configuration management. S12 is not masked in the acquisition configuration, so its frames are read out and written even though they carry no signal. The granularity is set by the electronics rather than by the fault: OD can only be lowered for a whole CCD, so a single shorted segment costs a sensor.

2.1.3 R20/Reb2/S21: faulty RD voltage read-back

R20/Reb2 could not be powered on from the first days of April 2025 until 2025-04-24. Neither the CCDs nor the biases were at fault: the RD voltage read-back of Bias1, the bias block of S21, had stuck at ≈ 54.2 V. That value lies in neither the ON nor the OFF window of the bias level test, so the CCD power state of the board could not be resolved, setCCDsPowerState() failed and the sequence was aborted (§2.1.3.1). All three CCDs read by that REB, S20, S21 and

S22, were therefore off; they were the only CCDs in the focal plane not powered on, and not exercised by the CCD shorts test, during the first week of April.

The read-back is not a physically attainable voltage (§2.1.3.2), and the fault is understood as an open connection between the RD sample point and the monitoring ADC, upstream of the analogue multiplexer (§2.1.3.3). The mitigation is a software exception for this one channel (Marshall, 2025a): its RD read-back is excluded from the shorts test, from the range check during power-on and from the ON/OFF bias count, but only while the reading stays above 53 V (§2.1.3.4). With that code in place the REB passed the shorts test and powered on at the first attempt on 2025-04-24, and S21 delivers nominal images (§2.1.3.5). The cost is that S21 is operated without an RD voltage monitor.

2.1.3.1 The failed power-on. `setCCDsPowerState()` evaluates eight tests on a REB: firmware bias protection, the board-level ODI, CLKLI and CLKHI currents, the bias levels of each of the three bias blocks, and the clock rails. On R20/Reb2 seven of them passed. The attempt of 2025-04-08 reads

```
[2025-04-08T14:15:35.088+0000] SEVERE: R20/Reb2/Bias1 RD1 voltage = 54.21 NOT in
range for ON|OFF
[2025-04-08T14:15:35.089+0000] SEVERE: R20/Reb2/Bias1: getBiasDacsPowerState() found
0(3) ON(OFF) of biases
[2025-04-08T14:15:35.089+0000] SEVERE: R20/Reb2: CCDs Bias voltage test FAILED
[2025-04-08T14:15:35.089+0000] SEVERE: Test 6: FAILED Bias level test for R20/Reb2/Bias1
[2025-04-08T14:15:35.090+0000] SEVERE: R20/Reb2: setCCDsPowerState() FAILED: 7:0:0
OFF:ON:DELTA of 8 tests passed
[2025-04-08T14:15:35.090+0000] SEVERE: R20/Reb2 CCDsPowerState updated to FAULT
```

followed by an automatic `powerCCDsOff()`. The OD, GD and OG read-backs of that block were all in the OFF range, at -0.03 , -0.03 and -0.01 V, and Bias0 and Bias2 reported all four of their biases OFF. `getBiasDacsPowerState()` required all four biases of a block to agree before it would return OFF or ON, so with the RD read-back outside both windows only three of four could ever be counted and the state was unresolvable by construction, not merely out of tolerance. The board stayed off throughout the period: focal-plane/R20/Reb2/ODI never exceeded 8.8 mA.

2.1.3.2 The read-back is not a real voltage. The channel returned only two distinct values, 54.18071 and 54.20757 V. Read as adjacent codes of the 12-bit slow ADC, the step of

$$54.20757 - 54.18071 = 0.02686 \text{ V} \quad (1)$$

is one LSB, which implies a full-scale span of $0.02686 \times 4096 = 110.0 \text{ V}$, or $\pm 55 \text{ V}$. That is the range of this channel. The REB driver scales the 12-bit conversions as

```
VOLT_12_SCALE_R5 = 10.0 / 4095,    // +-5V ADC input range
VOLT_12_SCALE_M5 = VOLT_12_SCALE_R5 * 3,
VOLT_12_SCALE_H5 = VOLT_12_SCALE_R5 * 11,
```

and CHAN_RD_2 uses VOLT_12_SCALE_H5. The ADC input range is $\pm 5 \text{ V}$, so one code is $10/4095 = 2.442 \text{ mV}$ at the ADC input; RD reaches that input through a divide-by-eleven network, so in units of RD volts one code is $2.442 \times 11 = 26.862 \text{ mV}$, matching Eq. (1) to five digits. The channel is therefore pinned near the top of its range with about one count of noise, which is what a floating ADC input looks like.

RD itself cannot reach that voltage. It is generated by a DAC output of at most 5 V multiplied by five, so it is limited to $\approx 25 \text{ V}$, as verified in the REB board acceptance tests; even a failure driving the output to its amplifier supply would give the regulated 35 V , not 55 V . A reading of 54.2 V thus cannot represent RD, and the bias itself was never suspect.

2.1.3.3 Where the fault is. RD is monitored through the REB slow ADC, which multiplexes of order seventy parameters. Following the board schematic, VDD_RD is generated on p. 32, divided on p. 31, enters an eight-channel analogue multiplexer on p. 8, is buffered on p. 65 and digitised on p. 64. Other quantities on the same multiplexer continued to read correctly, which places the fault upstream of it, in the path specific to this channel. Since the failure is most likely an open connection rather than a failed component, the behaviour may be temperature dependent.

Three questions were left open. Whether the break is before the voltage divider or after it was not established, the second case being the less damaging. The raw ADC code was not recorded, and a value of identically zero, as opposed to one moving by a few bits, would distinguish a divider tied to ground from other failure modes; the corresponding channels

of the other two sensors on the board provide the comparison. No review of the affected section of the board by an electronics engineer is documented.

2.1.3.4 Software mitigation. The decision was to make an exception for this channel in the CCD turn-on, turn-off and status procedures and to ignore its value (Marshall, 2025a). The exception is keyed on the bias path R20/Reb2/Bias1 together with the ADC identifier REBDevice.ADC_RD_1, and is applied in three places in BiasControl.java (Turri, 2025):

- testShorts() no longer drives and tests the RD channel of that block. The GD and OG channels are still tested. The skip is logged, so that it is visible in the CCS log.
- checkAdc(), which compares each read-back against its setpoint during power-on, skips the comparison for that channel *only while the reading exceeds 53 V*. The exception is therefore self-disarming: should the read-back ever recover, the normal check applies again without a code change.
- getBiasDacsPowerState() counts three biases instead of four for that block, under the same 53 V condition, through a new variable nbiases that replaces the literal 4 in the OFF and ON conditions. This is the change that removes the FAULT.

The work was tracked as LSSTCCSRAFTS-777, raised on 2025-04-19, reopened on 2025-05-05 as not yet merged, and resolved on 2025-05-07 with the merge of commit 829a8f2, released in org-lsst-ccs-subsystem-rafts 3.0.10. The code was exercised on the summit two weeks before that merge, as described next. The same commit also carries changes to the OD ramp and to the sampling of the slow currents during power-on that are unrelated to this mitigation.

2.1.3.5 Verification and recovery. On 2025-04-24 the REB initialised cleanly with the patched CCS: firmware 0x3139500e, serial number 418254201, sequencer idle, six ASPICs loaded, back bias enabled, CCDsPowerState OFF. The added log message is not printed on a successful initialisation, so at that point the log did not by itself confirm that the exception was in the running code.

A CCD shorts test was run on this REB alone before any attempt to power on:

```
[2025-04-24T19:57:38.742+0000] INFO: R20/Reb2 lowering sequencer clocks for shorts  
test: 0x3bc --> 0xc  
[2025-04-24T19:57:40.249+0000] INFO: R20/Reb2/Bias1: skipping testShorts() on RD  
voltage, jira:LSSTCCSRAFTS-777  
[2025-04-24T19:57:44.258+0000] INFO: R20/Reb2 restoring sequencer clocks after shorts  
test: 0xc --> 0x3bc  
[2025-04-24T19:57:44.556+0000] INFO: testCCDSHORTS() finished successfully
```

which confirmed both that the exception was active and that it behaves as intended. In the trending of the test the block shows eight bias excursions instead of the usual nine, three GD, two RD and three OG; the missing one is S21/RD, whose read-back stays at ≈ 55 V and dominates the vertical scale unless that trace is excluded.

powerCCDsOn was then invoked once, with the agreement that work would stop on any failure, and succeeded at the first attempt. focal-plane/R20/Reb2/ODI read 91.3 mA and the per-CCD currents and voltages were nominal. The faulty read-back moved slightly upward as the biases came on, so the broken channel is not entirely decoupled from the ADC. Clears and images followed, with stray light visible in the frames. A quick statistics pass on the first image of S21, MC_C_20250424_000016_R20_S21, gives read noise of 3.44–3.70 ADU across the sixteen segments, serial CTE of 0.99997 or better and parallel CTE of unity to the precision quoted, with no dead or anomalous segment.

2.1.3.6 Operational consequences. Three CCDs were recovered by a change of a few tens of lines, and no hardware intervention was needed or attempted. Two points follow from how it is done. First, S21 runs without an RD voltage monitor: a genuine RD fault on that sensor would no longer be caught by the bias level test, and the ON/OFF state of that block is decided by three biases rather than four. Second, the exception is hard-coded against a bias path string and a threshold rather than being configuration, so it must be carried forward through releases of org-lsst-ccs-subsystem-rafts and would silently stop protecting the board if the path naming changed. Both should be tracked in configuration management. The hardware fault is unrepaired and only partly characterised, and the initialisation logging that would positively confirm that the exception is active was left to be revised.

2.2 Cryo system

2.2.1 Brief introduction to the refrigeration system

The refrigeration system for the LSST camera (LSSTCam) is comprised of a cryogenic or “Cryo” system to cool the focal plane CCDs, and an independent “Cold” or Pumped-Coolant System (PCS) system to remove process heat from the imaging sensor’s readout-electronics boards (REBs) which are contained within the same vacuum vessel. The Cryo system operates at approximately -130°C , removing 500 W, and the Cold system operates at approximately -50°C , removing 1100W. The two systems act independently, save for the radiative coupling of the two distinct evaporators within the cryostat vacuum vessel, and through weak radiative coupling and conduction paths in each Raft Tower Module (RTM). In addition to performance requirements on the evaporator temperature for nominal heat loads, its temporal stability, and spatial uniformity, there are several operational requirements that address the systems’ performance under various extreme process load conditions, and a strict environmental interface (heat and vibration) to the telescope’s dome enclosure, such as to avoid any impact to imaging.

The refrigeration system must meet the requirements shown in Table 2 which represent the interface between the telescope and LSSTCam. These ensure that while operating, the refrigeration system does not impact the quality of the telescope images, neither by influencing the dome environment, nor by creating mechanical vibrations that could feed through to the focal plane. The refrigeration system must also operate in the anticipated temperature and humidity ranges at the observatory summit. These environmental requirements are shown in Table 3.

Finally, Tables 4 and 5 show the performance requirements for the Cold and Cryo systems, respectively. These include the temperature requirements for nominal and extreme process heat loads, as well as the temperature stability, uniformity and controllability of the two evaporators

The cooling requirements for the LSST Camera are rather unique amongst astronomical instruments, driven by the large focal plane area with 189 distinct CCDs, the proximate location of the front-end electronics, and the parallel digitization and readout with large power dissipation, all occupying the same vacuum space. In addition, the cooling system must not influence the dome environment of the observatory itself. Heretofore, smaller but similar systems could use conventional coolants such as LN₂. However, to provide adequate refrigeration capacity with such coolants would have required either a large refillable vacuum insulated reservoir

TABLE 2: Summit Facility Interface between Telescope and Camera

Requirement	Specification	Notes
Refrigerant transfer line skin temperature	Within $+4/-15^{\circ}\text{C}$ of ambient dome's temperature	Applies to all refrigerant line surfaces exposed to dome's ambient environment. Insulation is optional wherever possible.
Compressor cabinet outer skin temperature	Within $+1/-3^{\circ}\text{C}$ of dome's ambient temperature	Foam insulated cabinet walls, and coolant used to control the cabinet interior temperature

TABLE 3: Operating Dome Environment Requirements

Requirement	Specification	Notes
Ambient working temperature range	-5°C to 30°C	System should deliver the required refrigeration capacity in the required evaporator temperature range.
Ambient turn-on temperature range	-10°C to 30°C	System should be able to turn-on and operate, but the refrigeration capacity may be reduced.
Ambient survival temperature range	-15°C to -10°C & 30°C to 40°C	System should survive such temperature extremes. The system does not need to be running in this condition.
Ambient working humidity range	15% to 60%	System should deliver the required refrigeration capacity in the required evaporator temperature range.
Ambient survival humidity range	10% to 15% & 60% to 90%	System should survive these humidity extremes but does not need to be running in this condition. If the refrigeration system is running, the refrigeration capacity may be reduced.

directly behind LSSTCam, or a set of flexible, large cross section vacuum insulated transfer lines running up to the camera from a reservoir and pump at the base. Both solutions would require similar transfer lines to return Cold boiled off N₂ vapor, the combination of which proves highly inconvenient. Such would also be the case with more modern cryocoolers (e.g., Polycold) that locally produce Cold refrigerant mixes, but which then must be transported to the evaporator and returned to the cryocooler.

Instead of these cumbersome conventional solutions, the design that has been selected incorporates two distinct but classic vapor compression refrigeration systems (“cold” and “cryo”). The Cryo refrigeration system consists of 6 individual circuits, which utilizes the auto-cascade refrigeration technology. The refrigerants are compressed and cooled back to ambient dome temperature (by compressors and condensers remote from LSSTCam), and then transported at high pressure through the dome, and up to the camera. There they are expanded to produce the required cooling of the evaporators. At the rear of the camera each system has counterflow heat exchangers which pre-cool these incoming refrigerants before expansion, utilizing the Cold returning refrigerant. This refrigerant then exits the camera, now close to the dome temperature and at a low pressure and returns back to the remote compressors.

The Cold system utilises conventional cascade refrigeration technology to cool the heat transfer fluid to -50 C at the PCS chiller cabinet, the heat transfer fluid is pumped through a vacuum jacketed line to LSSTCam, now at -48 C, to cool down the electronics. The heat transfer fluid then exits the camera, now at about -44 C, returns to the remote PCS chiller cabinet via a vacuum jacketed line.

The transfer line system carries the refrigerants to LSSTCam and returns them to the compressors – a total path length of ~60 m with a maximum elevation change of ~16 m.

2.2.2 Architecture and Thermal Design of the Cryostat

The cryostat is designed to be a vacuum insulating vessel ($<10^{-7}$ Torr) containing the Cold and Cryo systems’ evaporators, the “coldplate” and the “cryoplate”. Figure 5 shows a cross section of the cryostat with the principal thermal regions and typical operating temperatures. Refrigerants enter the cryostat, pass through the evaporators and return through 4.5” vacuum conflat flanges on the rear annular feedthrough plate. The refrigerant tubing itself is continuous and joint-free within the cryostat vacuum to eliminate any potential for a leaking. The refrigeration system components outside the cryostat (in the Utility Trunk) have an in-

TABLE 4: Cold Refrigeration System Performance Requirements

Requirement	Specification	Notes
Temperature range for operation	-40 °C to -30 °C when the load ranges from 754 W to 843 W	Nominal load is 800 W at -35 °C
Maximum survival load	1056 W	The system does not run away.
Minimum survival load	490 W	The system does not run away. Trim heat and/or compressor operating frequency and/or hot gas bypass can be used when running the system with the minimum survival load.
Temperature uniformity	± 1.5 °C	Temperature across the Cold plate
Temperature controllability	± 2 °C when process load changes by 89 W	The controller should adjust the heat load of the trim heaters to compensate for the change of the process load when turning on/off electronics.
Temperature drift	± 1 °C at constant load	Duration not specified.

TABLE 5: Cryo Refrigeration System Performance Requirements

Requirement	Specification	Notes
Cryo plate temperature range	-130 °C to -125 °C with load from 361 W to 498 W	Any temperature from -130 °C to -125 °C is accepted when the load is in this range. The nominal temperature and load are -127.5 °C @ 429.5 W.
Cryo plate maximum survival load	520 W	Cryo system should not run away.
Cryo plate minimum survival load	219 W	Use trim heat to increase the minimum survival load. Cryo system should not run away.
Cryo plate temperature uniformity	± 1.5 °C	Temperature across the Cryo plate
Cryo plate temperature controllability	± 1.5 °C when process load changes 136 W	The controller should adjust the heat load from the trim heaters to compensate for the change in load from when turning on/off electronics.
Cryo plate temperature drift	± 0.5 °C at constant load	Duration not specified.

dependent insulating vacuum, referred to as the heat exchanger (“HeX”) vacuum). Ceramic insulators are used throughout the system to isolate the electrical ground of the evaporators from the rest of the cryostat. Figure 6 shows the detailed routing and makeup of the system at the interface between the cryostat and HeX vacuum system, including the mechanism for isolating grounds.

The cryoplate layout is shown in Figure 7. It comprises a single machined and polished copper plate which is brazed to a stainless-steel plate to rigidize it, with cutouts for each of the Raft Tower Modules (RTMs). There are 21 “bays” where the science RTMs are attached, and 4 corner bays for the corner RTMs containing guider CCDs and wavefront sensors. The six stainless steel refrigerant lines – one for each Cryo circuit - run continuously through the top and bottoms of the six copper ribs, having been brazed in place. Three lines are on the top, and three on the bottom, and the refrigerant flow direction is organized to be counterflow. The rafts are bolted down to the copper on the front side of the orthogonal ribs.

2.2.3 Architecture of the Cryogenic Systems

The physical layout of each Cryo system is shown schematically in Figure 8. With the exception of the transfer line lengths, a largely identical system supports LSSTCam in the clean room Maintenance Facility at the Vera C. Rubin Observatory.

All of the ground-based components reside in two temperature-controlled housings [“Cryo cabinets”] attached to the underside of the Telescope Mount Assembly (TMA). The transfer lines to and from LSSTCam (a liquid supply, a vapor supply, and a suction-return for each circuit) originate and return to these cabinets. They are a mix of orbitally welded - electropolished 316-L stainless steel tubing and flexible refrigerant lines. The counterflow HeX systems reside in a separate vacuum insulated housing within the camera’s Utility Trunk, just behind the cryostat.

To achieve the total cooling capacity required, each Cryo system is duplicated six times, with each of these circuits providing ~85 to 110 W of cooling capacity around the nominal operating temperature of -130 C in order to attain the required ~500W at the cryoplate. The circuits run completely in parallel through the cryoplate (see Figure 7).

Each of the two Cryo cabinets supports three of the six circuits. The compressor system for each circuit is itself modularized and easy to replace/exchange for servicing. Each compressor




FIGURE 5: Cross section of the cryostat with its principal thermal regions and typical design temperatures during operation (in red).




FIGURE 6: Detailed routing of the refrigeration system at the interface between the cryostat and the heat exchanger vacuum system.




FIGURE 7: Cryoplate (a) Routing of the six refrigerant circuits in the copper cryoplate (b) with stainless steel stiffening frame attached.




FIGURE 8: The physical layout of each Cryo system is shown schematically in Figure 8. With the exception of the transfer line lengths, a largely identical system supports LSSTCam operation in the Clean Room Maintenance Facility at the Vera C. Rubin Observatory.

module contains a commercial hermetically sealed compressor that has been modified to add a water cooling jacket and an oil-reservoir with both visual and electronic oil level gauges.[1] The compressor is followed by an oil separator, a refrigerant condenser (or "after-cooler"), and a refrigerant phase separator. A small sight glass in the liquid supply line indicating both flow and moisture content is incorporated after the phase separator. Three diaphragm valves on the low- and high-pressure sides isolate the module from the rest of the circuit. A small heated "surge tank" on the low pressure side of each compressor module is used during charging and cooldown to increase the volume of the system. Each Cryo circuit is completed with a set of replaceable filter-driers (and isolation valves) on the suction return line as well as the liquid and vapor phase supply lines.

Each Cryo cabinet is insulated and has a system of three fans and three heat exchangers to match the external (dome) temperature. A mix of 35% V/V Ethylene Glycol and water solution (EGW) is provided at fixed supply temperature to each Cryo cabinet to remove heat from the compressor, the condenser, and the phase separator, as well as the cabinet's own heat exchanger. The composition of the coolant mixture used for the Maintenance Facility is the same (see Table 6).

TABLE 6: Fixed temperature Glycol Solution for Cryo Cabinets

Location	Coolant Composition	Coolant Supply Temperature (°C)
TMA	35% Ethylene Glycol + 65% water	9 ± 3
Maintenance Facility	35% Ethylene Glycol + 65% water	9 ± 3

Within each cabinet, there is readout for an extensive set of instrumentation including thermocouples measuring temperatures and transducers measuring refrigerant pressures throughout the system. The compressor power, oil levels, environmental temperatures (cabinet air and cooling water), and valve-states are also measured. All signal conditioning and digitization is done internally and communicated via an Ethernet connection to the Camera Control System (CCS). Finally, an independent hardware protection system uses a distinct and independent set of sensors and a programmable logic controller (PLC) within the cabinet, to shut the system down if any of the operating parameters move out of their safe allowed range.

A similar system of instrumentation and protection exists in the PCS chiller cabinets (described in the next section) and in the Utility Trunk for both the HeX systems, the cryostat and HeX vacuums, and for the cryoplate and coldplate evaporators themselves. Conditions measured

in these regions, even though they reside outside the cabinets, can also initiate a safety shut-down of the compressors through the integrated protection system.

In addition to the aforementioned hardware, a system has been built to mix and test the custom multi-component refrigerant deployed in the Cryo systems.

2.2.4 Cryogenic Mixed-Refrigerant System

Overview

As noted previously, the Cryo refrigeration system is comprised of two temperature-controlled cabinets on the ground, a set of three transfer lines for each of the six circuits going up to LSSTCam, and a vacuum insulated heat-exchanger system at the rear of the camera. In addition, the facility's chilled EGW is provided to each of the cabinets for removal of heat from the compressors and electronics.

The cryogenic refrigeration cycle (see Figure 9) begins with the ambient temperature low pressure (~ 30 psia), superheated vapor refrigerant mixture entering the compressor at the suction port (located at position "1" in the figure). It is then compressed to a higher pressure (~ 400 psia) and higher temperature (~ 100 degree C) refrigerant, while remaining in vapor form.




FIGURE 9: Layout of Cryo system.

Leaving the compressor, the refrigerant passes through a hot-oil separator, which returns most of the synthetic polyolester oil to the compressor through a capillary (C1). The refrigerant mix continues through an EGW cooled condenser coil (or after-cooler), bringing its temperature down to ~ 14 degree C. It exits the after-cooler as a vapor-rich mixture which is next separated into its individual liquid and vapor components in the slightly cooler fractionating column (or phase-separator). Now at a temperature of ~ 12 degree C, the high-pressure liquid and vapor phases are transported separately ~ 50 m to LSSTCam heat exchanger system in the Utility Trunk. As the components are at high pressure, the elevation change of ~ 16 m has been shown to have minimal effect. The transfer lines themselves are both rigid and flexible, with the flexible portion at the camera end to allow for rotations and changes in elevation. By mounting the compressor cabinets on the TMA itself, the cabinet and the transfer lines follow the azimuthal motion of telescope, without requiring any additional length to accommodate that 3600 rotation.

The counterflow heat exchangers are vacuum insulated in a separate volume from the cryostat, within LSSTCam Utility Trunk. Each circuit's heat exchanger divides the incoming flow at the warm end into seven smaller diameter ($\sim 1/8$ " OD) copper tubes which fit loosely into a copper outer jacket ($1/2$ " OD) that carries all the return flow. The seven tubes increase the surface area for exchange of heat, and are recombined at the exit of the HeX, at the Cold end.

The incoming liquid phase first undergoes expansion in capillary C3 (0.026" ID X 45' Length). The Cold expanded liquid is injected into the heat exchanger outer jacket at about the one-third point along its length where it pre-cools the incoming high pressure ambient-temperature vapor to approximately -15 degree C. The vapor phase continues through the HeX core and in the second stage is precooled further by the Cold refrigerant returning from the evaporator, which entered the outer jacket at the Cold end. Exiting the HeX at the Cold end, the Cold high-pressure refrigerant is now $\sim 100\%$ liquid. It is next expanded in capillary C4 (0.052" ID X 10' Length) leaving it an admixture of liquid and vapor at approximately -125 degree C just before entering the evaporator within the cryostat. Traveling through the 0.116" ID x 21.5' long tubing of the evaporator, it undergoes a further pressure drop along its path, and boils while absorbing heat from the RTMs, to achieve the required cryogenic temperature around -130 degree C.

Exiting the evaporator and the cryostat, and re-entering the Cold end of the HeX outer jacket as a Cold liquid and vapor mix, the now lower pressure (~ 200 psia) refrigerant exchanges heat with the incoming high pressure vapor in the 7 tube core, before mixing with the returning

Cold refrigerant from the C3 capillary tube. It then re-exits the HeX close to ambient temperature and at low pressure (~30 psia), returning to the compressor as 100% vapor in the large diameter return line.

The use of cascaded cooling in the HeX is a critical feature of the design. The liquid phase expansion at capillary C3, that pre-cools the incoming vapor to ~-15 degree C, allows for the initial cooldown of the massive cryoplate before the returning refrigerant is Cold enough to provide effective lower-temperature pre-cooling at the Cold end of the HeX.

Other important features of the system include the use of a surge tank, a system of replaceable filter-driers to address moisture ingress, in-line particulate filters, and the employment of thermal and geometrical features that are necessary to mitigate oil trapping during transport.

The large volume thermally insulated surge tank (located in the Cryo cabinet on the low pressure side of the circuit plumbing) is used as a temporary reservoir to store refrigerant when the system is cooling down. This prevents high pressures from occurring in the circulating plumbing that could damage the compressor. The surge tank is heated to 40 degree C so refrigerant does not condense in the volume, and it is left fully open to the system during normal operation. A remote-controlled solenoid valve is used to control the isolation of the tank from the suction line.

The moisture in the system must be carefully controlled and monitored both to avoid long term damage to the compressor (moisture reacting with the lubricant to create acids) and also to avoid clogging of the small diameter capillaries as they get cold, due to ice buildup over time.

Throughout the system, all joints are either welded or brazed, while breakable connections utilize VCR fittings with all-metal gaskets. Particulates and flux created during fabrication and installation are carefully pre-flushed, and an in-line 90mm sintered particulate filter is incorporated at the entry point to LSSTCam itself, in each supply line.

After leak checking, the only continuous moisture ingress source should be the sections of flexible transfer line. Accordingly, replaceable filter-drier canisters are installed on the two supply lines and on the suction return line, and an analog (visual) moisture indicator is built into the liquid supply line at the compressor.

Oil management in the system is a critical issue. The synthetic polyolester lubricant [POE-32LT] is used. To reduce mechanical damage to the compressor during startup, oil in the compressor is pre-heated to 40 degree C reducing its viscosity as well as the peak compressor power that is required. As the pour-point of the lubricant is -52 degree C, it is particularly important that it does not propagate into the cryogenic vapor-phase portion of the circuit, in any significant quantity.

The refrigerant mixture's liquid components can however pick up small amounts of residual oil and transport it through the system. While much of the lubricant is removed and returned to the compressor before transport to LSSTCam, a small amount continues to propagate in the system, coating the inner surfaces of the transfer lines. Eventually this goes into an equilibrium state. However, the transfer line system's large inner surface requires the addition of a large volume oil reservoir at the compressor, to prevent the system from operating without sufficient lubricant, thereby damaging the compressor during start up and the initial coating of the lines. In addition to the aforementioned elevation, the system's geometry has been carefully laid out to attempt to eliminate those locations where oil might collect over time. In particular, the three in-line filter-driers are vertically oriented to avoid trapping oil.

Thermodynamics of mixed refrigerants

While still utilizing the concept of a vapor compression refrigeration cycle, the cryogenic system is significantly more complex. It is a two-stage system that uses a novel mixed refrigerant whose ten components each have phase diagrams that are strongly influenced by the presence of each of the other components within the mix.

The complex mixture allows for a greatly expanded range of temperatures and pressures over which refrigeration can be efficiently achieved by two-phase boiling of the refrigerants. In practice, this allows for the use of an efficient two stage auto-cascade refrigeration cycle wherein the high-pressure liquid and vapor components are pre-separated. As indicated, this latter feature is necessary to permit a slow cooldown to the required -130 degree C operating temperature.

The mixture itself (designated M2919C) contains eight conventional refrigerants and two Nobel gases. (See Table 7). The refrigerant composition utilizes two fluorocarbons, five hydro-fluorocarbons and one hydro-chloro-fluorocarbon. The mix is a non-flammable liquid at 20 degree C and atmospheric pressure making its handling and transfer simpler. The moisture content

of refrigerant is typically measured to be 10 to 40ppmv, bottle to bottle.

The final “running” composition of the mixture depends on details of the system’s geometry, and in particular, the volume of the liquid supply line. M2919C was designed for operating the cryoplate with about 5m of liquid supply line tubing (~1/16” ID), while the actual system in the observatory’s two configurations are each about ten times longer and have a total volume (liquid and vapor) of ~ 70 liters. To compensate, the refrigerant charge of the two ambient temperature liquid components of M2919C (R124 and R245fa) are empirically increased after initially charging the system with M2919C, until the desired cooling capacity is achieved, with the appropriate temperature at C3 being the most representative. The mixture itself is then monitored periodically by gas chromatography and the results compared with a reference sample.

A REFPROP analysis of the M2919C refrigerant was performed. The refrigerant mix M2919C is partially condensed at the after-cooler with the vapor portion at 81.755% by weight at the entrance to the fractionating column where the vapor and liquid are separated. The composition in the vapor phase and the liquid phase are calculated using a REFPROP flash calculation and shown in Table 8.

TABLE 7: Composition of M2919C

Component	Chemical Name	Mole Fraction (%) (S=100%)
R14	Tetrafluoromethane	30.77
R124	1-Chloro-1,2,2,2-Tetrafluoroethane	1.54
Argon	Ar	26.92
R218	Octafluoropropane	15.38
R125 R404a	Pentafluoroethane	2.37
R143a R404a	1,1,1-Trifluoroethane	2.80
R134a R404a	1,1,1,2-Tetrafluoroethane	0.22
R23	Trifluoromethane	8.46
Krypton	Kr	3.85
R245fa	1,1,1,3,3-Pentafluoropropane	7.69

The pressure-enthalpy (P-H) diagram of the M2919C refrigerant at the vapor supply is generated by REFPROP 9.1 and presented in Figure 10. At the liquid-vapor transition, or critical point, the M2919C mixture has a temperature of 31.28 C and a pressure of 9.23 MPa.

Unlike a pure substance or an azeotropic mixture where the temperature glide (the difference

TABLE 8: Liquid and vapor phase composition of M2919C (fractionated at 14 °C and 350 psig)

Component	Suction mole fraction (%)	Vapor mole fraction (%)	Liquid mole fraction (%)
R14	30.77	34.42	5.50
R124	1.54	1.04	5.02
Ar	26.92	29.99	5.64
R218	15.38	13.14	30.91
R125	2.37	2.16	3.83
R143a	2.80	2.46	5.15
R134a	0.22	0.18	0.49
R23	8.46	8.89	5.46
Kr	3.85	4.20	1.45
R245fa	7.69	3.52	36.54

between the saturated vapor and saturated liquid temperatures, at a constant pressure) is zero or




FIGURE 10: Pressure-Enthalpy Diagram of the M2919C vapor components, fractionated at 14 degree C and 350 psig.

close to zero, the temperature glide of the M2919C mixture can be > 100 degree C. Accordingly,

the isothermal lines inside the two-phase regime of M2919C are not horizontal lines in the pressure – enthalpy diagram, but angle downward from the bubble point boundary (at the left) to the dew point boundary (at the right). This implies that the pressure drop across the evaporator is expected to be greater than in a typical vapor compression refrigerator, and as anticipated, pressure drops between 20 to 50 psi in the prototype evaporators have been measured.

REFPROP predicts the properties of the M2919C mixture more precisely at room temperatures than at cryogenic temperatures. For example, it fails to generate the isothermal curves inside the two-phase regime when the temperature is -120 degree C or lower, and in particular into the region where the mixture is observed to freeze (~-150 degree C). To illustrate this, dotted lines were added manually in this region, in Figure 10.

The refrigeration capacity can be estimated as the difference between the refrigerant enthalpy at the heat exchanger low-pressure return line, H_{11} , and the sum of the refrigerant enthalpies at the heat exchanger vapor and liquid supply lines, H_5 and H_7 , respectively:

$$Q = \dot{m}H_{11} - (\dot{m}_v H_5 + \dot{m}_l H_7) = \dot{m} (H_{11} - xH_5 - (1-x)H_7). \quad (2)$$

Here, $\dot{m}$ is the refrigerant mass flow rate of the return flow, $\dot{m}_v$ is the vapor refrigerant mass flow rate, $\dot{m}_l$ is the liquid refrigerant mass flow rate, and x is the vapor quality, defined as the ratio of vapor refrigerant mass flow rate to the total refrigerant mass flow rate.

Substituting the refrigerant enthalpy calculated from RefProp at each location to the above equation results in a refrigeration capacity of 29.5 kJ per kg of refrigerant. If the refrigerant flowrate through the compressor is 5 g/sec, we should get a refrigeration capacity of 147.5 W at the evaporator temperature of -130 °C. The actual actual measurement is about 110 W obtained by summing the process heat load on the evaporator and the measured radiative and conductive parasitic losses.

2.2.5 Brief introduction to the temperature sensitivity issue of the Cryo refrigeration system

As shown previously (Figure 9), there are vapor refrigerants and liquid refrigerants coming out at the compressor cabinet and sent to the heat exchanger at the utility trunk via different refrigerant transfer lines (~60 m long and ~16 m elevation change). Among the 60 m long

transfer line, about 50% is rigid stainless steel and 50% of flexible rubber hose with polyamide (nylon) liner. During the transportation, the refrigerant can exchange heat with the environment.

Prototype test at SLAC has demonstrated a successful run over four months long. During the run, the refrigerant transfer lines were 60m long each, with 45m long stainless steel and 15m long flexible rubber hose with polyamide liner. All refrigerant lines are insulated (1/2" thick insulation). The ambient temperature was as low as 6 C and as high as 25 C.

The ambient temperature at the TMA summit has a much larger swing than that at SLAC. The lowest ambient temperature at LSSTCam body was measured at 0.75 C at 05:45:49 UTC on 2025/09/08 while the highest ambient temperature was recorded at 23.16 C at 21:13 UTC on 2025/02/27.

The refrigerant lines at the TMA were not insulated before we observed the ambient temperature sensitivity issue on May 8, 2025. This has created a problem. The refrigerant mixture was fractionated at the phase separator into vapor refrigerant and liquid refrigerant. Therefore, the vapor refrigerant would have a dew point at about the same temperature as the phase separator temperature. The phase separator was cooled by the fixed temperature glycol/water coolant (at a temperature of 9 ± 3 C), which dictates that the dew point of the vapor refrigerant at around 0.2 C warmer than the glycol/water inlet temperature.

Therefore, when the ambient temperature is at or colder than the glycol/water temperature, the vapor refrigerant would condense as it transports from the compressor cabinet to the heat exchanger at the utility trunk. This would cause the performance of the refrigeration system to change (change to bad).

2.2.6 Detail of the cryo-temp sensitivity issue

The first instance of the Cryo ambient temperature sensitivity issue occurred on May 8, 2025. The camera ambient temperature was below 7 C and reached 4.43 C at 8:18:39 UTC. The Cryo trim reduced to zero and the cryoplate temperature warmed up, as shown in Figure 11

All five running circuits have the C3ExitTmp warm up closer to 0 C, indicating that the refrigerant C3 flow was reduced. This is shown in Figure 12

All five running circuits showed an increase in PreC4Tmp, indicating that the C4 flow was reduced. This is shown in Figure 12

All five running circuits showed increase in discharge pressure and reduction in suction pressure, indicating that there was an increase in flow impedance. This is shown in Figure 13

The glycol/water supply temperature was 7 C and the phase separator was at around 7 C at the time being, as shown in Figure 14.

When the Cryo ambient temperature warmed up to 7 C at 15:30 on May 9, 2025, the Cryo plate temperature started to decrease, indicating the refrigeration system resumed working. Cryotrim heat returned to 188 W at 20:32:30 on May 9, as shown in Figure 11. The cryotrim heat reported was higher than that reported on 3:00 UTC on May 8 due to that the CCD camera was off.

This phenomenon was repeated on June 7, 2025. Ambient temperature reduced to 4.43 C at 1:05:30 on June 7, 2025 and Cryo trim heat reduced to zero W, as shown in Figure 15. The refrigeration system started to recover its performance and Cryo trim heat kicked in when ambient temperature warmed up to 6.85 C at 7:50 on June 7, 2025. Cryo trim heat increased to 157 W at 12:33 on June 7, 2025 when ambient temperature was stable at 6.92 C.

2.2.7 Solution to temperature sensitivity issue

The cost of the Cryo trim heat loss when lower ambient temperature is due to the condensation of the refrigerant at the vapor line. Therefore, we need to do the following

- Lowering the phase separator temperature
- Heating the refrigeration lines
- Insulating the refrigerant lines

We have lowered the glycol/water temperature from 7 C to 4 C. This has reduced the phase separator temperature from 7 C to 4 C.

We have installed the self regulating heating cable on the vapor supply line and liquid supply




FIGURE 11: Ambient temperature, Cryo trimheat and cryoplate temperature at around 6 am UTC on May 8, 2025




FIGURE 12: C3ExitTmp and PreC4Tmp of each Cryo circuit at around 6 am on May 8, 2025




FIGURE 13: Discharge and suction pressure of each Cryo circuit at around 6 am on May 8, 2025




FIGURE 14: Glycol/Water Inlet temperature and phase separator temperature of each Cryo circuit at around 6 am on May 8, 2025




FIGURE 15: Cryo system loss cooling capacity due to low ambient temperature on June 6/June 7 in 2025

line on the refrigeration lines from just exiting the compressor cabinet to the 8th floor just before the elevation wrap. This heater has a length of 35 ft. The heating tape is made by RAYCHEM. When the ambient temperature is -40 C, it will generate a heating capacity of 6 W per feet long; When the temperature is 65.5 C, it will reduce its heating power to zero. We only turned on the heating cable on the vapor line currently. However, we want to have the ability to maintain the refrigerant at the liquid line to be above the ambient temperature when we need to.

We also have insulated the refrigerant lines on top of the heating cable using 1/2" thick close cell insulation. The material is K-Flex INSUL-LOCK® DS 6RXLO048058, a closed-cell flexible double-seal pipe insulation with overlap. It has a length of 72 inch, a wall thickness of $\frac{1}{2}$ inch, and an R-factor of 3.3. The material is an elastomeric foam (domestic product), and this insulation helps retain heat within the refrigerant lines.

The Cryo refrigeration behavior when ambient temperature was below 2 C on September 8, 2025 is shown in Figure 16. Cryo trim heat was about 104 W at 18:00 pm September 7, 2025. It was reduced to 78.96 W at 5:45:49 on September 8, 2025 while ambient temperature reduced to 0.75 C. The Cryo trim heat continued to drop down to 74.51 W at 6:13:18 on September 8, 2025 then increased gradually back to 105.4 W at 12:12:30 on September 8, 2025 when the ambient temperature increased to 3.6 C.

This shows that the current solution to Cryo refrigeration system can allow the Cryo refrigeration system provides adequate cooling to LSSTCam when ambient temperature was 3 C lower than the phase separator temperature.

If we want LSSTCam to operate with ambient temperature lower by more than 3 C below the phase separator temperature, we need to improve the refrigeration line insulation further. There are Dyalene lines at the 8th floor that are very close to the refrigerant line. The Dyaline fluid is at 0 C. If there is contact with the refrigerant lines, it will cool down the refrigerant at the refrigerant line, which can condense the refrigerant at the vapor line. We need to insulate the Dyaline line.

2.3 CCS development

- CCS




FIGURE 16: Cryo system behavior when ambient temperature reduced to 0.75 C on September 8, 2025

2.4 Firmware development

Science REB (REB)

Table 9: Science REB Firmware Version History

Version	Release Date	Key Modifications and Notes
0x313B5019	December 2025 (6.4ns: tested briefly on camera; 10ns: tested on ts8)	• Code Refactor Completion: Allows system clock selection on a per-target basis using common functional blocks. → The same VHDL code is used in both cases with only a parameter change. • Expected to definitively fix the ID PROM read instability issue (pending observation). • These libraries are now reusable in GREB/WREB firmware; remaining work is to apply them to devices unique to those boards.
0x313a5015 - 0x313B5018	November 2025 (Not installed on camera)	• A series of releases used on ts8 for the development and testing of the 6.4ns firmware version. • All had various issues, but each was separately versioned before installation on ts8.
0x313a5014	September 2025 (Installed on camera: November 2025?)	• Functional Equivalence: This is the production version that ended up functionally equivalent to 0x31395012. → 0x3139500e with new functions: → – Ability to build firmware with new bias protection thresholds and sequencer overrides. → – Stability of CCD power across REB resets (rms_reset).

Table 9 – Continued

Version	Release Date	Key Modifications and Notes
		• ID PROM Instability: Still suffers from the latent "ID prom instability," a long-standing issue for which the root cause was never fully identified, despite many versions appearing stable.
0x313a5013	August 2025 (Briefly installed on camera)	• Nominally intended to be functionally equivalent to 0x31395012 under the hood. • Technical Stack Update: Migrated to the currently supported SLAC TID-ID firmware build system, SURF libraries, and Vivado 2025.1 compiler. • Code Refactoring: Shared code blocks moved to a dedicated repository (https://github.com/lst-camera-electronics/lst_reb). • Intended as preparation to support both 10ns and 6.4ns clocks. • Major Defect: Introduced a major bug resulting in maximum heater power being applied due to a register map conflict. → This bug was actually introduced in 0x31395011 but went unnoticed.
0x31395012	July 2025 (Not installed on camera)	• Made Bias VUT and HV switch stable across REB resets.
0x31395011	May 2025 (Installed on camera R30/Reb1 only)	• Made CCD bias protection thresholds compile-time configurable. • Added a "sequencer override register" to fix sensor sequencer signals.
0x3139510F	April 2023	• Included fixes discovered during Guider development, only relevant to the multi-sequencer version.

Table 9 – Continued

Version	Release Date	Key Modifications and Notes
	(Not installed on camera)	
0x313950FC - 0x313950FF	(Various Release Dates) (Used on ts8 only)	• A series of 6.4ns test versions built by naïvely modifying 0x3139500e.
0x3139500e	July 2022 (Production until November 2022)	• Baseline Production Release.

Wavefront REB (WREB)

Table 10: Wavefront REB Firmware Version History

Version	Release Date	Key Modifications and Notes
0x1139400B	September 2025	• Updates: Migrated to the current SLAC TID-ID build system, SURF libraries, and Vivado 2025.1. • Code Refactoring: Shared code blocks moved to dedicated repository (https://github.com/lst-camera-electronics/lst_reb). • Intended as preparation to support both 10ns and 6.4ns clocks.
0x1139400A	July 2025	• Ensured stability of Bias VUT and HV bias switch across REB resets.
0x11394009	May 2025	• Added compile-time specification for bias protection thresholds (unused in WREBs).
0x11394008	April 2023	• Baseline Release Version.

Guider REB (GREB)

Table 11: Guider REB Firmware Version History

Version	Release Date	Key Modifications and Notes
0x2139210D	January 2025	• First production release of 2-sequencer firmware for Guider operation.
0x21392x0c	(Not Released)	• Development Dead End.
0x21392x0B	April 2023	• Versions where x=1 are 2-sequencer firmware; x=0 are single sequencer (full frame mode).

2.5 Observation control

- OCS integration kfanning

2.6 Filter Exchange System Commissioning

This section describes the performance of the filter exchange system through commissioning. It details the main hardware issues encountered since first photon and their fixes, the various software upgrades, including the auto-recovery procedures, and exposes the evolution of the system performances from between April and December 2025.

2.6.1 Hardware: problems of youth & fixes

2.6.1.1 Autochanger 1 Online Clamps

2.6.1.2 Carousel clamp friction

2.6.1.3 Carousel collector stuck




FIGURE 17: Histograms of the filter pins Unlock time before (15 August 2025 to 15 September 2025 , left) and after (15 November 2025 to 15 December 2025 ,right), the carousel clamps maintenance in October 2025

2.6.2 Slow Control Upgrade

2.6.2.1 Loader operation

2.6.2.2 FES integration within observatory control

2.6.2.3 FES sleep mode / wake up

2.6.2.4 MCM interaction

- Required configuration for Normal mode
- Filter in place

2.6.2.5 Auto-recovery The implementation of auto-recovery procedures not only improve the reliability of the FES operations (=grant an operation of the FES without interruption) but the work implemented in 2025 significantly improved the repeatability : FES movements execution have smaller execution time jitter. The impact of the last point can be seen in fig 18, for example the long tail visible in the setFilter execution time prior to the FCS October 2025




FIGURE 18: Filter change execution time before (left plot) and after (right plot), the carousel clamps maintenance in October 2025 and the November FCS update. The jitter in the execution time of the filter exchanges has been highly reduced.

update , is now gone : this is mainly due to a change in the fine tuning of the carousel position when it is impacted by filter weights unbalance configuration in carousel.

- Unlock carousel clamp
- carousel rotation fine tuning
- carousel brakes unstuck

2.6.3 Performances

2.6.3.1 Filter change speed and overhead from telescope

2.6.3.2 FES operation tuning

2.6.3.3 Number of filter changes

2.6.3.4 Error rate evolution from April to December




FIGURE 19: Histogram of filter exchange duration between April 15th and July 18th (blue) and after July 18th (red)




FIGURE 20: Daily number of filter exchanges

2.7 Shutter

2.7.1 Shutter Timing issue

Since the first photon observation was made, the shutter system has exhibited issues including blade overrun, leaving the shutter partially opened, and discontinuities in motion profiles (Utsumi et al., CTN-003). This section presents further analysis and results of follow-up tests, and current understanding and remaining open questions.

2.7.1.1 Overview of the shutter operation The camera shutter is controlled by the Beckhoff TwinCAT PLC system. The CCS sends commands to a Beckhoff CX5120 controller, which is an industrial PC running Windows 7 Embedded Standard along with a real-time kernel called TwinCAT 3. This is the gateway and manages the PLC system interconnected by EtherCAT. The PLC system consists of various “terminal” modules that can take different roles, such as controlling digital and analog I/O as well as the timing module, called EL6688. EL6688 provides an external reference to synchronize the clock among terminals with the GPS clock with an accuracy better than $1 \mu\text{s}$.

This system has four different clocks:

1. The CCS clock synchronized by Network Time Protocol (NTP), which is based on UTC;
2. The TwinCAT clock, which is supposed to be synchronized by the host Windows when it is booted;
3. Distributed clock (DC) in the TwinCAT PLC system. When the shutter blade passes the hole sensors placed along with the shutter rail, the PLC is programmed to trigger timing measurements based on DC. This timing information is the foundation of the motion profile as well as the entire timing control of the PLC system;
4. External reference clock from EL6688. The time source is a GPS clock in the observatory and the synchronization is provided by Precision Timing Protocol (PTP), which is expected to provide sub-microsecond timing accuracy through network delay measurements.

We investigate the consistency between the timings retroactively from logging messages, with a focus on the CCS clock and the DC time reported by the motion profile. The UTC timestamps




FIGURE 21: Timing between the CCS time when takeExposure was recorded and the time when the motion profile records the startTime. This does not reflect the actual shutter motion timing.

in the MCM log (“sent takeExposure to shutter1”) were paired with TAI timestamps from motion profiles (“startTime.tai.isot”). Because TAI leads UTC by 37 seconds, the expected difference is approximately: *37 seconds + system/network jitter*. Plotting these deltas reveals three operational periods during 2025 in Figure 21:

April–June 2: The EL6688 was intermittently synchronized (the reported PTP state was unstable, bouncing between LISTENING (4) and SLAVE (9) states). However, evidence suggests that DC time was free-running. The shutter became unstable near the end of this period, including discontinuities in the motion profiles and blade overrun issues. This period ended after power for the shutter system was re-initialized. The underlying cause remains unknown. A reduction in TCPSynRetrans was observed, which may indicate a Linux kernel-level anomaly.

June 3–June 7: The EL6688 synchronized to a MOXA boundary clock, but synchronization quality was poor. Observed scatter reached 50 ms. The MOXA switch is an internal L2 network switch inside the Camera (EDS-G516E-T). As described in Utsumi et al. (CTN-003), the PTP functionality of the MOXA switch does not appear to function correctly with the observatory network. Consultation with the vendor confirmed that this device only supports the software-based PTP and does not support the hardware-based PTP, which can explain the poor quality of the synchronization.

June 8–Dec 2 (present): The EL6688 synchronized to an observatory boundary clock on a leaf switch (10.17.200.33) after three hops from the Grand Master Clock, where the GPS clock is connected. Timing has remained stable.

2.7.1.2 Verification of the Observatory’s PTP Further analysis was conducted using the spare system lsstcam-shutter02 in the clean room. This system is connected to the observatory’s network. This network is specially configured to pass the PTP packets. The network routes through an identical MOXA switch, providing an opportunity to investigate whether the observatory’s PTP system functions as intended and to identify any issues with the MOXA switch. Tests conducted in the clean room revealed successful PTP synchronization using a MOXA operating as an L2 switch by the following command:

```
[yutsumi@lsstcam-shutter02 ~]$ sudo ptp4l -f /etc/ptp4l.conf -i enp0s31f6 -m
[sudo] password for yutsumi:
ptp4l[503690.709]: selected /dev/ptp1 as PTP clock
ptp4l[503690.710]: port 1 (enp0s31f6): INITIALIZING to LISTENING on INIT_COMPLETE
```

```
ptp4l[503690.710]: port 0 (/var/run/ptp4l): INITIALIZING to LISTENING on
    ↪ INIT_COMPLETE
ptp4l[503690.710]: port 0 (/var/run/ptp4lro): INITIALIZING to LISTENING on
    ↪ INIT_COMPLETE
ptp4l[503696.415]: port 1 (enp0s31f6): new foreign master fc59c0.ffff.214803-29
ptp4l[503712.418]: selected best master clock 000efe.ffff.011598
ptp4l[503712.418]: port 1 (enp0s31f6): LISTENING to UNCALIBRATED on RS_SLAVE
ptp4l[503714.237]: port 1 (enp0s31f6): minimum delay request interval 2^5
ptp4l[503714.620]: master offset -311009524 s0 freq +6354 path delay 10727
ptp4l[503715.623]: master offset -311010780 s1 freq +5099 path delay 10727
ptp4l[503716.621]: master offset 474 s2 freq +5573 path delay 10727
ptp4l[503716.621]: port 1 (enp0s31f6): UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED
ptp4l[503717.623]: master offset 1484 s2 freq +6725 path delay 10727
ptp4l[503718.625]: master offset 257 s2 freq +5943 path delay 10727
ptp4l[503719.623]: master offset -26 s2 freq +5737 path delay 10727
ptp4l[503720.628]: master offset -184 s2 freq +5571 path delay 10727
ptp4l[503721.626]: master offset 8 s2 freq +5708 path delay 10727
ptp4l[503722.626]: master offset 443 s2 freq +6145 path delay 10727
ptp4l[503723.631]: master offset 83 s2 freq +5918 path delay 10727
ptp4l[503724.631]: master offset -434 s2 freq +5426 path delay 10727
ptp4l[503725.635]: master offset -149 s2 freq +5581 path delay 10727
```

The offset is in nanoseconds unit and s2 indicates the servo is in the locked state. As indicated in the output, the offset converges quickly. The key in the configuration is that the priorities have to be 255 since the priority of the grandmaster is set to be 255; otherwise lsstcam-shutter takes the grandmaster role in the PTP system which screws up the PTP synchronization. These tests confirm the observatory's PTP functionality.

Another machine in the clean room Dell AIO, aio06, was also used to test the PTP. The hardware clock behaved poorly under frequency adjustment, but acceptable synchronization was achieved using weakened servo parameters:

```
[lsstcam-aio06] % sudo ptp4l -f /etc/ptp4l.conf -i enp0s31f6 -m -H --step_threshold=1
    ↪ --pi_proportional_const=0.01 --pi_integral_const=0.0001
```

2.7.1.3 Verification with the spare shutter system Since the PTP synchronization has been successfully established with linux computers through the MOXA switch, the focus moves onto the actual PTP terminal module EL6688. EL6688 was particularly difficult because the interface only supports a 100 Mbps connection. There was a question whether the link speed conversion interferes with the PTP synchronization. We used several tools to investigate the PTP synchronization comprehensively.

‘readPtpDiag.sh’ is the tool to communicate with Beckhoff controller over the network, which provides a partial view of the variables in the PLC. The following example is taken from Utsumi et al. (CTN-003).

```
[ccs@lsstcam-shutter02 ~]$ /usr/local/bin/readPtpDiag.sh
HCU date and time: 2025-06-09 17:58:35.614918
PTP version: PTPv2 (32)
PTP state: SLAVE (9)
Clock ID (hex): 00 01 05 ff fe 3e 4c 43
Master clock ID (hex): fc 59 c0 ff ff 21 48 03 00 1d
Grandmaster clock ID (hex): 00 0e fe ff fe 01 15 98
Offset from master (ns): -11,662
Mean path delay (ns): 122,992
Steps removed: 4
Sync msg sequence no.: 37388
Time scale: PTP/TAI (1)
Offset from UTC (sec): 37
UTC offset valid? True
Leap61? False
Leap59? False
Epoch no.: 0
External (PTP) time: Wed, 16 Jul 2025 17:40:44.907797 UTC
```

The output has a few distinct states every power cycle. “Offset from Master” is sometimes quite large, as large as 10^9 ns (or 1), and sometimes it is as small as about $10\ \mu$ s or smaller, regardless of the “PTP state”. And sometimes the PTP state never becomes SLAVE(9) but LISTENING(4).

The Windows program TwinCAT XAE shell (TcXaeShell) provides a live and full view of the vari-

ables in the PLC system. Figure 22 shows an example screenshot of TcXaeShell when the vendor-provided example project is running. DcToTcTimeOffset and DcToExtTimeOffset are two of the important variables because we found that they tell which of either TwinCAT or PTP clock is used to synchronize to DC. It appears that only one of them becomes non-zero, and the other becomes zero but randomly. Which clock is used is not deterministic as long as we use TwinCAT version 3.14022.36.

bSynchronized and nTimeDiff are also important parameters. These are unique in the vendor-provided project. These tell us whether EL6688 is synchronized with PTP or not and how much delay is observed. The meaning of synchronization in the vendor implementation is broad. The synchronization is established regardless of whether the DC is synchronized to the Beckhoff Windows or EL6688, since EL6688 gives a reference clock and the actual clock synchronization is done in the Beckhoff Windows system using the built-in functions. nTimeDiff seems to be different from Offset from master.




FIGURE 22: A screenshot of the vendor-provided sample program. DcToTcTimeOffset, DcToExtTimeOffset and bSynchronized as well as nTimeDiff are displayed in this screenshot.

Another diagnostic tool is the generic open-source software, tcpdump.

```
sudo tcpdump -i enp0s31f6 -n -s 0 'ether proto 0x88f7 or udp port 319 or udp port 320' -a
```

The following is example output of tcpdump, focusing on a communication pair between the boundary clock and the client (EL6688) when they exchange "delay request" and "delay response".

```
12:33:01.777635 IP 139.229.175.116.ptp-event > 224.0.1.129.ptp-event: PTPv2, v1
    ↪ compat : no, msg type : delay req msg, length : 44, domain : 0, reserved1 : 0,
    ↪ Flags [none], NS correction : 0, sub NS correction : 0, reserved2 : 0, clock
    ↪ identity : 0x105fffe3e4c43, port id : 1, seq id : 8038, control : 1 (delay req
    ↪ msg), log message interval : 5, originTimeStamp : 1764419582 seconds, 136332373
    ↪ nanoseconds
12:33:01.777892 IP 10.17.200.33.ptp-general > 224.0.1.129.ptp-general: PTPv2, v1
    ↪ compat : no, msg type : delay resp msg, length : 54, domain : 0, reserved1 : 0,
    ↪ Flags [none], NS correction : 0, sub NS correction : 0, reserved2 : 0, clock
    ↪ identity : 0xfc59c0ffff214803, port id : 29, seq id : 8038, control : 3 (peer
    ↪ delay resp msg), log message interval : 5, receiveTimeStamp : 1764419618
    ↪ seconds, 777492282 nanoseconds, port identity : 0x105fffe3e4c43, port id : 1
```

This example confirms that EL6688 communicates with the boundary clock (a leaf switch) directly, regardless of the presence of the MOXA L2 switch. This message exchange is a part of the foundation to measure the delay in the communication, which enables precise timing synchronization as defined in the PTP specification. In this case, the time difference from originTimeStamp to receiveTimeStamp is 36.641159909 sec. Repeated observations of these packets revealed that delays are consistent over multiple days with accuracy better than 1 μ s.

Although the difference is consistent, the difference itself is quite large. This large difference was observed when the large 'Offset from master' was observed (DC is synchronized with the TwinCAT clock). On the other hand, if DC is synchronized with EL6688/PTP, the difference was quite small. This observation suggests that when the DC synchronized with the TwinCAT clock and the DC tries to keep the pace as referenced from EL6688 (PTP) but likely the absolute clock could be wrong if TwinCAT clock is not reasonably synchronized with NTP.

2.7.1.4 Summary

- The observatory PTP system provides a reasonable PTP synchronization.
- MOXA switch does not support the hardware clock, resulting in the poor PTP synchronization.
- Even though the EL6688 supports only up to 100 Mbps network speed, reasonable synchronization can be established.
- TwinCAT's behavior whether DC is synchronized to TwinCAT controller or EL6688 is not consistent and not configurable. Operationally, power cycling is needed until it has a small enough offset from Master.
- **We also note that motion profile timing prior to 2025 June 7 should be considered unreliable; afterward, it appears trustworthy.**
- The jumps in the motion profile can be explained by intermittent adjustment based on the EL6688 to the TwinCAT system appears to play a role here.

Recommendations

- Correct the Puppet-managed chrony.conf on shutter01/shutter02 (LSSTCAM-146).
- Perform the Chris Walter timing test to interdependently check the timing accuracy. (Walter, SITCOMTN-136).
- Consider increasing the priority (lower the number) of the observatory grandmaster/boundary clocks. Any misconfigured computer can mess up the PTP network as ptp4l's default is 128.

Open Questions

- Beckhoff's statement that EL6688 "does not adjust the distributed clock" remains unclear, especially regarding NTP/PTP coexistence.
- Compare MCM command timestamps with motion profile timestamps for deeper latency characterization.
- Determine whether shutter02 can store/export motion profiles for analysis.

2.7.2 Shutter blade trajectory fitting

@andy rasmussen shutter blade trajectory fitting

2.8 Guider star finder

@Aaron Roodman star finder in guider

2.9 Historical telemetry of infrastructure cut events: electricity, glycol, dynalene and air

@Homer Neal telemetries of infrastructure cut events, electricity, glycol, dynalene, and air

3 Camera optimization

3.1 Sequencer changes

This section describes the history of the sequencer changes since Run 7. Run 7 results were obtained with the v30 sequencer. The changes made for v30 are documented in Section 5.2 of LSST Camera Team et al. (SITCOMTN-148). Since Run 7 was concluded, additional optimization targets were defined:

- Eliminating long-range correlations
- Reducing readout noise by switching to 3 s readout timing

Our testing of the resulting optimizations was simplified owing to schedule pressure to begin the Science Verification survey.

3.1.1 Long-range correlations

The measurement of brighter-fatter coefficients reported in LSST Camera Team et al. (SITCOMTN-148) had already shown an anomalously high value of the a_{00} coefficient for one specific e2v

sensor: R23_S02 (see Figure 34 in that note). This measurement is derived from fits over the variance and spatial covariance curves versus the flux. The curves for the channels in that particular sensor were in fact affected by an excess of variance, which affected not only the noise measurement in bias frames, but also introduced an extra flux-dependent component in the variance. This degraded the gain measurements from the PTC (using C_{00}/μ). This issue also manifested as a flux-dependent covariance between amplifiers in the same sensor, and as an excess in the correlation with the next pixel in the serial direction (a_{10}). One way to detect this issue in a systematic manner is to compute the slopes of the covariances between channels versus the flux, and average them over all channels in the sensor (as done in Figure 23). Using this metric, other sensors were found to exhibit a similar issue, but to a smaller degree (same Figure, left plot). Other related metrics (noise correlation between channels, correlations between channels) also pointed to the same sensors.




FIGURE 23: Averages of the slopes of the covariances between channels, for each sensor. Left: during test Run 7. Right: on sky, with the modified ‘RGn20’ readout sequence.

Since the issue affected some sensors uniformly, and exhibited flux dependency, investigations focused on the CCD readout clocks, and in particular on the pixel reset clock (RG). To explain a potential variation from sensor to sensor, it should be noted that while the REB FPGA drives all three sensors with the same readout sequence, the actual clocks are generated on a per-sensor basis. While the jitter of the FPGA timing signal is in the range of 100 ps, the jitter in the chain that translates that signal into a clock can be 1-3 ns. In addition, the RG clock is faster than the serial clocks, due to a lower capacitive load.

For sensors like R23_S02, plots acquired in scan mode at the raft level showed a strong variability (dispersion over consecutive scans) in the timing of the clock transitions around the time that the pixel reset (RG) clock was lowered. In the e2v sequence for pixel readout, this

drop of the RG clock was timed simultaneously with the drop of the S2 clock (which transfers charges to the S3 phase, in preparation for transfer to the readout node). The same scans also showed a much stronger variability in the level of the signal after the pixel reset, up to 1000 electrons instead of the expected 50-100 electrons from kTC noise. Even with the Dual Slope Integration that is performed in the CCD readout chain, such high variability would explain an extra component in the noise.

We then tested a series of modifications to the timing of the readout sequence. In a sensor that did not exhibit the extra covariance with the nominal timing, we were able to induce it by triggering the S2 clock drop 10 ns before the RG drop. Conversely, tests on sensors that exhibited the issue showed that delaying the S2 drop by an increment of 10 or 20 ns after the RG drop removed the issue (Figure 23, on the right). A 20 ns delay was selected as the safer option, and added to the e2v readout sequence (tagged as v30_RGn20, then adopted as v31). A similar modification was tested on ITL sensors, which had no effect on sensors that exhibited correlated noise.

v32 has a stuff related to idle clear, no functional changes in other parts v33 is reordered in a one-part sequencer in timing. Whether inserting a 20ns RG before (as BufferRGp) or after the BufferS+TimeS (as BufferRGn) time slice. No change in the total timing.

3.1.2 3 s sequencer file development

The observing strategy has evolved from two snaps of 15 s exposures to a single 30 s exposure per visit. By eliminating the intermediate readout, the overall timing budget is relaxed: one full readout cycle no longer needs to be performed every 15 s. This reduction frees up readout time that can be reallocated to decrease readout noise. Consequently, the development of a 3 s sequencer (instead of the existing 2 s) was initiated to take advantage of this additional margin.

We examined the time slices in the sequencer files and identified four key parameters for optimization:


- ISO1: Time for ASPIC clamp/reset to start of the first integration to measure the baseline (Ramp Down; RD)
- ISO2: Time between S3 down and start of ASPIC Ramp Up (RU)

- RampTime: The integration time of a pixel read.
- TimeP: Unit element of the Parallel Transfer (parallel timing).
- BufferP: Time period before each parallel transfer starts.

Previous studies (Utsumi et al., 2024; Snyder et al., 2024) examined how variations in ISO1, ISO2 and RampTime affect readout noise; they also demonstrated that altering Parallel Transfer (TimeP) can increase full-well capacity. ISO2 is additionally known to influence crosstalk. For the v30 firmware, BufferP is set to 3360 ns, raising the question of whether the entire parallel clocking should be shifted forward as in the original vendor sequence. In fact this change Because extending ISO1 degraded performance on some e2v sensors, our investigation was confined to the following scenarios:

1. **Longer integration time per pixel (3s-v0 RT)** – RampTime increase= from 380 to 620 ns (factor 1.6)
2. **Extended ISO2 duration (3s-v0 ISO2)** – lengthen the readout interval for each row.
3. **Longer parallel transfer timing (3s-v0 Parallel)** – Allow more time for charge transfer between columns.

Design constraints emphasized minimal deviation from the baseline sequencer, preserving parameters that had been previously tuned.



Sequencer files for each option were produced and peer reviewed. Validation involved bias checks, TS8 hardware testing at SLAC for ITL, and testing at UC Davis for e2v.

Noise measurements indicated that the ISO2 extension (option 2) delivers performance on par with the baseline while achieving the targeted 3 s readout time. Consequently, option 2 was adopted as the new production sequencer (3s-v1).

The total system gain is determined by three parameters: ramp time, ASPIC gain, and ASPIC RC. Since changing RampTime scales the gain roughly proportionally we also need to adjust the ASPIC gain to compensate the gain change from the RampTime in order to achieve a similar level of the output. Retaining a similar output signal was necessary to avoid triggering a so-called bias shift for the high ADC output.

The ASPIC RC parameter changes the selectable resistor in the integrating amplifier stage of

the ASPIC, changing its RC time constant. The necessary compensation can be achieved by changing RC from 14 to 2, which results in a ratio of 1.6.

Related acquisitions are listed below with links to corresponding electro-optical test reports.

- E7662 3s RT sequencers: https://s3df.slac.stanford.edu/data/rubin/lstcam/E7662/w_2025_26/
- E7697 3s Parallel sequencers: https://s3df.slac.stanford.edu/data/rubin/lstcam/E7697/w_2025_26/
- E7714 3s ISO2 sequencers: https://s3df.slac.stanford.edu/data/rubin/lstcam/E7714/w_2025_26/

3.1.3 Natural frequency

Several long-standing features of the LSSTCam bias and flat frames had resisted explanation through the successive sequencer optimizations described above: jumps in the measured gain, a multimodality of the bias frames (in the noise–bias level plane, in the row profile, and in the two-dimensional structure), and a checkerboard pattern in the photon transfer curve (PTC) and overscan covariances. This section describes the origin that was eventually identified for most of them, the firmware and sequencer change that addressed it, and the two testing campaigns that validated the change: one built on bias and flat frames and one on the PTC and linearizer products. The two campaigns look at the same firmware from different directions and are reported together below, phenomenon by phenomenon.

Gain jumps had been known since 2020, when Hyeyun Park found that the measured gain of a given amplifier could move between discrete values from one acquisition to the next, at the sub-percent level (Figure 24). This triggered the RTM-004 optimization runs,² and the effect was partially mitigated at the time by a change of the output drain (OD) voltage.

The bias frames showed a related behaviour: they occupied distinct states in bias level, noise and row profile, which meant that a super-bias built by averaging many frames could not be used reliably. The effect is visible both as separate clusters in the plane of overscan noise versus mean image level and as discrete bands in the overscan noise as a function of time

²<https://confluence.slac.stanford.edu/spaces/LSSTCAM/pages/311547783/RTM-004+optimization+runs>




FIGURE 24: Gain measurements per sensor across the focal plane for two amplifiers (amp08, left; amp09, right). Red boxes mark sensors in which the gain moves between discrete states from one acquisition to the next; the difference between states is at the sub-percent level.

(Figure 25), and it was reported independently in the analyses of Shuang Liang and Thibault Guillemin.




FIGURE 25: Bias multimodality before the natural frequency firmware. Left: overscan noise as a function of time for three sensors of raft R22, showing discrete bands rather than a continuous distribution. Right: overscan noise versus mean image level for R22_S01, segment 10, where the same behaviour appears as separate clusters.

The third feature is the checkerboard, or chessboard, pattern that appears in the noise when differences of pairs of PTC or overscan frames are correlated. In the notation used for the

covariance coefficients, n_{00} is the noise itself while n_{i0} traces CTI or a gain variation, so a regular alternation of the covariance with pixel lag directly pollutes the a_{ij} measurement. The pattern was strong for e2v sensors and essentially absent for ITL.

A presentation by P. Antilogus on 7 May 2024 proposed an explanation and gave the project its name. There are two points in the readout chain at which an uncontrolled phase shift can be introduced between the clocks of different REBs:

- Each REB clock (6.4 ns) is synchronized by the 3.125 GHz common clock. An arbitrary phase shift, a multiple of a unit of 6.4 ns/20, can be introduced at synchronization and persists until a link reset.
- A REB converts its 6.4 ns clock to the 10 ns clock used by the FPGA that runs the sequencer. A random phase shift of up to 10 ns can be introduced at this conversion, and it can differ from one exposure to the next.

If the discrete states of the gain and of the bias arise from these phase shifts, then running the REB at its natural frequency, that is, taking the sequencer timing directly from the 6.4 ns clock instead of the converted 10 ns clock, should remove them.

The checkerboard pattern admits a separate and complementary interpretation, proposed by the French team: the regular pattern can be produced by a pair of clocks that carry a slight time difference between them while their overall timing remains the same. If the ramp time takes an odd number of clock cycles and the number of columns is odd, the alternation between the two clocks maps onto the pixel grid and generates a checkerboard (Figure 26). P. Antilogus produced the even2 sequencer file, in which all the relevant time periods take an even number of cycles. This sequencer, carrying both the even-odd change and the natural frequency timing, was subsequently adopted as the production sequencer under the name 3s-v2 (calibration tag 3s_v2), superseding 3s-v1.




FIGURE 26: Schematic of the even-odd effect. In the ideal case (top) each period of the read-out covers two full clock cycles (2×10 ns). When the number of cycles is odd (bottom), successive periods start on alternating clock phases, so consecutive columns are read with a slightly different timing.

The natural frequency operation is provided by Science REB firmware 0x313B5019, released in December 2025 and listed in Table 9. That release completes the code refactor which allows the system clock to be chosen on a per-target basis using a common set of functional blocks, so that the same VHDL code serves both the 6.4 ns and the 10 ns cases with only a parameter change. The same libraries can now be used in the WREB and GREB firmware, and the remaining work is to apply them to the devices unique to those boards.

Validation was organized by board type. For the Science REBs, the new firmware was tested on TS8 for ITL and at UC Davis for e2v before installation on the camera, with multimodality (100 biases) and small PTC acquisitions; the run records are collected in the RTM-004 optimization runs page. For the WREB, a 100-bias acquisition with LATISS revealed an oddity that paused the deployment (see the LATISS testing discussed below). For the GREB, firmware development is still underway. The configuration adopted for the campaigns reported here is therefore the natural frequency firmware with the 3s-v2 sequencer for the science rafts, and the original firmware for the WREB and GREB. J. Chiang ran the cp/eo pipelines promptly on each acquisition, and the individual analyses below were carried out on those outputs.

3.1.3.1 Testing campaigns. The bias and flat campaign comprised the following acquisitions:

- Biases, 9 December 2025, sequence numbers 13–112.
- Flat stability, 14 December 2025:
 - Gain stability (BLOCK-T661): 11–110, 100 flats targeting 12 ke^- ; 111–210, 100 flats targeting 50 ke^- ; 514–613, 100 flats targeting 12 ke^- ; 614–713, 100 flats targeting 50 ke^- (after sun down).
 - Biases (BLOCK-T659), with an RCE reset between runs: 211–310, 313–412 and 413–512, 100 biases each.
- PTC, white light flats and darks on 27 December 2025 (taken with the mirror cover partially closed), 28 December 2025 (second attempt) and 3 January 2026.
- Science images, taken as a bonus on 3 January 2026.

A second, independent look at the same firmware came from the PTC and linearizer products,

in an analysis by E. Rykoff reported in January 2026. Those data are the acquisitions of 29 December 2025, a full night with 524 flat pairs, processed into the collection LSSTCam/calib/DM-53722 under the tag 3s_v2. The pairs were shuffled in flux order and no photodiode or electrometer reading was available, so that analysis is purely internal to the CCDs; the signal range runs from roughly 50–100 adu up through the PTC turnoff, except on the corner detectors. The comparison set is the equivalent acquisition of 15 July 2025 (BLOCK-T557) with the same signal range and configuration, processed in LSSTCam/calib/DM-53620 under the tag 3s_v1_dp2_v2. Detector 122 is excluded from both.

The linearizer fitter used for those products is an implementation of P. Astier’s linearizer fit with modifications: it includes a global offset term and a linear temporal term, and it does not use modified weights or the REB temperature (TEMPAVG). The fit is split into a “relative” and an “absolute” linearization. The least-squares minimizer was found to get stuck, and applying the temporal fit together with the relative slopes and then re-fitting brought it to a much better state; every one of the 3,024 amplifiers was inspected individually.

3.1.3.2 Noise and gain consistency between configurations. The readout noise measured with the 6.4 ns sequencer (3s_cp_v1_even2_64) and with the 10 ns sequencer (3s_cp_v1) agrees closely, with mean ratios of 1.003 for e2v and 0.991 for ITL (Figure 27). A small rounding in the sequencer timing makes the gains inconsistent between vendors at the 1% level. The outliers are half-dead or high-noise sensors (R10_S01, R43_S20, R41_S21). No geometrically systematic correlation across the focal plane was observed, which addresses the concern that the WREB might inject noise into neighbouring rafts.




FIGURE 27: Readout noise from bias images, comparing the 6.4 ns sequencer (3s_cp_v1_even2_64) with the 10 ns sequencer (3s_cp_v1), for e2v (top) and ITL (bottom). Left: amplifier by amplifier. Right: distribution of the ratio between the two (P. Antilogus).

The PTC products give the same result from the flats, amplifier by amplifier across the whole focal plane (Figure 28). The gains move slightly and in opposite directions for the two vendors,

e2v down a little and ITL up a little, which is the vendor- dependent rounding above seen in the gain itself rather than in the noise. The read noise generally decreases for all amplifiers, and for a subset of e2v amplifiers it decreases substantially — presumably the channels that were carrying the excess variance associated with the checkerboard discussed below.




FIGURE 28: Per-amplifier comparison between $3s_v1_dp2_v2$ and $3s_v2$ from the PTC. Left: gain. Right: read noise in electrons. The dotted line is equality (E. Rykoff).

3.1.3.3 Bias level versus noise. Plotting the overscan noise against the overscan mean for one amplifier per sensor over the whole focal plane (Figure 29) no longer shows multiple distinct populations. The residual systematic drift within each sensor is consistent with gain and temperature changes over the duration of the acquisition rather than with discrete state changes.




FIGURE 29: Overscan mean versus overscan standard deviation for all rafts and sensors (20251214, sequence numbers 211–309), one amplifier per sensor. The distinct populations seen before the firmware change (Figure 25) are absent.

3.1.3.4 Bias row profiles. The row profiles measured before and after the change are compared in Figure 30. With the original firmware the profile of a given amplifier differed between exposures; with the natural frequency firmware the profiles from repeated exposures fall on top of each other.




FIGURE 30: Bias row profiles for three amplifiers (R02_S01 C15, R20_S20 C03, R42_S21 C11) with the original firmware (top) and with the natural frequency firmware (bottom, run E1110, 9 December 2025). The profiles become consistent from exposure to exposure (T. Guillemin).

3.1.3.5 Consistency across resets. Since the phase shift introduced at REB synchronization persists until a link reset, the acquisitions of 14 December 2025 were split into three runs separated by an RCE reset. The mean parallel overscan profiles of the three runs agree amplifier by amplifier (Figure 31); no change was observed between before and after a reset. Together with the noise agreement above, this points at the 6.4 ns to 10 ns conversion rather than at the REB synchronization as the origin of the multimodality.




FIGURE 31: Mean of the parallel overscan for each amplifier of detector 10, for the three runs of 20251214 separated by an RCE reset (sequence numbers 310, 412 and 512). The three curves overlap in every channel.

3.1.3.6 Gain stability. The relative gain change of every amplifier, measured against the focal-plane median, is shown in Figure 32 without and with the natural frequency firmware. Without it the median gain drifts by more than 2×10^{-3} over the run; with it the drift is removed and the residual per-amplifier scatter is at the level of the measurement noise. The amplifiers that remain noisy in the right-hand panel are noisy intrinsically.




FIGURE 32: Difference of the per-amplifier relative gain from the focal-plane median, and the median itself (bottom panel of each plot). Left: without the natural frequency firmware (run E1496). Right: with the natural frequency firmware (2025-12-15).

The same measurement performed on the high-flux flats (Figure 33) is to be compared with the specification of 0.1% RMS over 1 h and 1% RMS over 12 h. The majority of the amplifiers are better than $\mathcal{O}(10^{-4})$. Three classes of outlier remain: non-stable illumination at the edge of the focal plane, sensor-intrinsic behaviour, and a gain-temperature correlation that is not yet understood. The linearity analysis below finds residual patterns of exactly these two kinds — smooth variation over the night on the vignetted edge detectors, and a time and temperature correlation on isolated amplifiers — which suggests the two campaigns are seeing the same residual effects through different measurements.




FIGURE 33: Relative gain change per amplifier for the high-flux flats acquired with the natural frequency firmware (2025-12-15), in the same format as Figure 32.

3.1.3.7 Checkerboard in the overscan. The two-dimensional lagged correlation of the serial overscan, averaged over all amplifiers of R23_S02, is shown in Figure 34 for the original and for the natural frequency configuration. The checkerboard pattern is gone in the latter.




FIGURE 34: Two-dimensional lagged correlation in the serial overscan, all amplifiers of R23_S02. Left: original configuration (20251216, sequence number 366). Right: natural frequency configuration (20251214, sequence number 211).

Figure 35 shows the covariances $\langle \text{cov}_{ij} \rangle$ for all channel pairs of R23_S02, excluding the self-correlations, before and after the change. The average over pairs, shown as an inset, carries a clear checkerboard in the original data and is flat with the natural frequency firmware. P. Astier interpreted the decaying signal at $\Delta y = 0$ as a gain variation, which the natural frequency partially cures: the spread of the relative gain of R23_S02 narrows from roughly $\pm 10^{-3}$ to a considerably smaller range (Figure 36).




FIGURE 35: Covariances for all channel pairs of R23_S02, excluding self-correlations, with the average over pairs and the cross-amplifier C_{00} shown as insets. Left: run 20250715 (original). Right: run 20251214 (natural frequency). The checkerboard in the average of pairs disappears (P. Astier).




FIGURE 36: Relative gain of R23_S02 over the course of a run, on a vertical scale of ± 0.001 . Left: original configuration. Right: with the natural frequency firmware.

3.1.3.8 Checkerboard in the PTC. The PTC products confirm the same result on the flats, through a completely different measurement. Averaged over raft R22, the e2v noise matrix carries an unmistakable checkerboard in 3s_v1_dp2_v2; with 3s_v2 the pattern disappears and the colour scale shrinks by an order of magnitude, from ± 0.08 to ± 0.008 (Figure 37). The A matrix behaves the same way: the alternating structure that was superimposed on the smooth brighter-fatter falloff is absent in the new data (Figure 38), which means the a_{ij} measurement is no longer polluted. ITL, averaged over R02, shows the same improvement in the noise matrix at lower amplitude (Figure 39); the ITL A matrix was already free of the pattern and is essentially unchanged, consistent with the checkerboard having been an e2v effect from the start.




FIGURE 37: Average PTC noise matrix over raft R22 (e2v). Left: 3s_v1_dp2_v2. Right: 3s_v2. Note that the colour scale differs by a factor of ten between the two panels (E. Rykoff).




FIGURE 38: Average A matrix over raft R22 (e2v). Left: 3s_v1_dp2_v2. Right: 3s_v2. The checkerboard superimposed on the brighter-fatter falloff is gone (E. Rykoff).




FIGURE 39: Average PTC noise matrix over raft R02 (ITL). Left: 3s_v1_dp2_v2. Right: 3s_v2 (E. Rykoff).

3.1.3.9 Linearity. The linearity solution improves as well, most visibly on the ITL detectors, where the structure that the old solution left in the residuals is largely gone (Figure 40). With $3s_v1_dp2_v2$ the residual fraction is spread over roughly $\pm 1.5 \times 10^{-4}$ with visible wave-like structure across the full signal range; with $3s_v2$ it collapses to a narrow band that tightens with signal, and the usable range extends past 10^5 adu. Some ITL detectors improve only slightly (for example detector 15), and the e2v detectors, which started in better shape, improve more modestly but consistently (Figure 41).




FIGURE 40: Linearizer residual fraction against input ordinate for every amplifier of ITL detector 2. Left: $3s_v1_dp2_v2$. Right: $3s_v2$. The structure spread over $\pm 1.5 \times 10^{-4}$ in the old solution collapses to a narrow band (E. Rykoff).




FIGURE 41: The same comparison for e2v detector 85, which was already in better shape. Left: $3s_v1_dp2_v2$. Right: $3s_v2$.

A handful of inner amplifiers still show structure in the $3s_v2$ residuals, and the patterns differ from one another. On detector 41 (R22_S12_C14) the residuals appear correlated with the overscan level, which is not typical of the focal plane. On detector 48 (R12_S10_C07) they cor-

relate with time and possibly with the REB temperature, and on detector 152 (R33_S22_C03) with time alone, in both cases for a single amplifier on an otherwise well-behaved detector. These are the linearity counterpart of the gain-temperature outliers noted above.

The amplifiers at the edge of the focal plane show a residual that varies smoothly over the night (Figure 42). The same or similar patterns appear on all partly vignettted and nearly-partly-vignettted detectors, for example detector 178 (R42_S21), while the inner amplifiers and detectors show nothing of the kind. Variable illumination at the edges is the natural explanation, and it accounts for the edge outliers in the flat-field gain stability above.




FIGURE 42: Absolute linearizer residual fraction against elapsed time for detector 188 (R43_S22), amplifier C10, with the 3s_v2 data. The smooth variation over the night is seen on the partly vignettted detectors only (E. Rykoff).

3.1.3.10 PTC turnoff. The PTC turnoffs, expressed in electrons, are broadly consistent between the two configurations, with a set of outliers and a locus of e2v turnoffs that moved up (Figure 43); both are still being followed up.




FIGURE 43: Per-amplifier PTC turnoff in electrons, 3s_v1_dp2_v2 against 3s_v2 (E. Rykoff).

3.1.3.11 Science images. Science images were taken with the natural frequency firmware on 3 January 2026 (for example exposures 655 and 657); two examples are shown in Figure 44. The expected gap in illumination between ITL and e2v sensors, caused by the small gain change, is visible, as is a small crosstalk residue. Both are expected, since no calibration products were regenerated for these frames. No other issue was found.




FIGURE 44: Two science images acquired with the natural frequency firmware on 3 January 2026, around a saturated star. The residual crosstalk and the illumination step between vendors are visible at low level.

3.1.3.12 LATISS testing (WREB). The WREB behaves differently. The noise measurement in the noise-bias plane is noisier than for the science rafts (Figure 45), and print-through was identified as the cause: a section of the serial overscan shows the print-through structure directly (Figure 46). Disabling the ASPIC temperature read made the print-through go away. The pattern in the covariance, however, still exists in a different form, and is present whether or not the science rafts are running the natural frequency firmware (Figure 47).




FIGURE 45: Standard deviation versus mean of the serial overscan for each amplifier of LATISS, for three acquisitions: 20251224 sequence numbers 11–110 (blue), 20251224 sequence numbers 118–217 (red) and 20260112 sequence numbers 27–126 (green).

The PTC products flag a wavefront problem of their own: one amplifier, R40_SW1_C13, produces an A matrix in the new run that is clearly wrong, showing alternating rows that saturate both ends of a colour scale two orders of magnitude larger than expected, where the old run gave the smooth brighter-fatter structure (Figure 48). Whether this is the same underlying




FIGURE 46: A section of the serial overscan for LATISS, 20251226, sequence numbers 207–306, showing the print-through structure.




FIGURE 47: Two-dimensional lagged correlation in the serial overscan for R00_SW0. Left: original configuration. Right: original WREB configuration while the science rafts were running the natural frequency firmware. The pattern persists in both.

WREB behaviour seen in the PTC is not yet established.




FIGURE 48: A matrix for R40_SW1_C13. Left: old run, as expected. Right: new run, with alternating rows and a colour scale two orders of magnitude larger (E. Rykoff).

3.1.3.13 Understanding the wavefront sensors. Scan mode is the natural tool for understanding the underlying issue, since it resolves the signal within the sequencer time slots. The scan-mode data taken on R00_SW0 in run 13108 (Figure 49) show a wild signal after the RST/CL toggling, wiggles in the first ramp, and an increase of the noise for a longer ramp time. These acquisitions have not been repeated since November 2023. Progress requires getting TS8 running and testing there, and then stepping to LATISS and LSSTCam.

3.1.3.14 Summary. Two independent campaigns, one on bias and flat frames and one on the PTC and linearizer products, agree on what the natural frequency firmware and the 3s-v2 sequencer accomplished. The bias multimodality is removed, both in the noise-bias plane and in the row profiles, and the gain variation is now better than $\mathcal{O}(10^{-4})$; the PTC gains and read noise confirm this amplifier by amplifier, with the noise improving markedly on the e2v channels that had carried the excess variance. Since the behaviour is unchanged across REB resets, the conversion of the clocks from 6.4 ns to 10 ns on the REB, rather than the REB synchronization itself, appears to be the source of the inconsistency.

The checkerboard is gone in the overscan lagged correlations, in the cross-amplifier covariances, and in the PTC noise and A matrices alike, so the a_{ij} measurement is no longer polluted. It is likely cured by the even-odd change rather than by the natural frequency, which implies that a small difference remains between the members of a pair of clocks; the sequencer carrying both changes is the one adopted in production as 3s-v2. The linearity solution improves at the same time, most strongly for ITL. No major issue was observed in the science images taken with the new firmware.

What remains is consistent between the two campaigns: residual patterns on the vignettted edge detectors, which both attribute to variable illumination, and a small number of amplifiers whose gain or linearity residual correlates with time and temperature. The PTC turnoff outliers and the locus of e2v turnoffs that moved up are still being followed up. The WREB remains a mystery, with a print-through cured by disabling the ASPIC temperature read but a covariance pattern that persists, and a wavefront amplifier whose A matrix is clearly wrong in the new run. The GREB is next after it. Scan-mode measurements on the focal plane and on LATISS, and the 0/1/2D bias analysis of T. Guillemin, are the natural next steps.




FIGURE 49: Scan-mode data for R00_SW0, run 13108, sampled every 10 ns of the sequencer. Top: mean of the signal for the eight channels. Middle: dispersion of the signal. Bottom: evolution of the signal between scans 800–900 and 200–300. Dotted lines mark the sequencer transitions.

- R30/S12 and R23/Reb10 – I will try to collect all the information from Slack

3.2 Re-implementation of “clear” operation during the idle time

This section describes the implementation of the “idleClear” mechanism used on the LSSTCam to flush charges when the CCDs are otherwise idle.

3.2.1 Introduction

In addition to intentional exposures, the LSSTCam sensors accumulate signal when they are idle, via light leaks, time dependent defects, and dark current. When the sensors are idle, the power consumption changes relative to when they are being read out affect the sensor and electronics temperatures. To keep the sensors ready for exposure and the signal chain at a steady temperature, we require a mechanism which clears the charges semi-continuously at a constant power level commensurate with normal survey observing. This mode is named “idleClearing”.

3.2.2 Background/Context for the sequencer, DAQ and CCS

The core LSSTCam readout is implemented in an FPGA via a firmware engine dubbed “the sequencer” which emits a sequence of values for which each of 14 significant bits drives one of the control signals for the CCD pins, dual slope integration, digitization or data transfer. Each FPGA is integrated into a Raft Electronics Board (REB) with the electronics for up to 3 CCDs and resides with the CCDs inside the LSSTCam cryostat. The sequencer is loaded with a set of “sequencer programs” that define the sequence of output states necessary to accomplish CCD readout. This implementation architecture is limited in that each defined “sequencer program” is triggered, runs to completion and ends in an idle state. The sequencer does not support running a program as an interruptible infinite loop. As a result, the mechanism for idleClearing must be implemented outside of the FPGA and hence at the CCS (Camera Control System) or DAQ (Data Acquisition) level. However, the sequencer is suitable for taking survey image data for which each image is the result of one triggered run of a sequencer main program which reads all or a portion of the CCD. These sequencer programs are parametrized to allow dynamically adjusting readout for overscan and ROIs (region of interest), as one might expect. The sequencer programs are loaded into the FPGA and triggered by or through interfaces provided by the DAQ system (described elsewhere).

The complete LSSTCam readout comprises two distinct sets of sequencers: those that support the 8 guider CCDs and those that support the 189 science and 4 wavefront CCDs. In the DAQ nomenclature, each set is referred to as a “partition” of the complete DAQ. Each partition is operated synchronously by a distributed trigger generated in the DAQ. For science/wavefront sensors, the trigger is initiated by the CCS (via the DAQ interface), once per image at the end of an exposure to read out in ≈ 3 seconds. During science sensor readout, the guider CCDs are intentionally static. During science sensor charge integration (exposure) the guider CCD readout is periodically triggered (nominally at ≈ 9 Hz), using an ROI. The guider readouts are triggered by processes internal to the DAQ using ROI parameters requested by the CCS. When no science sensor exposure or readouts are occurring, all CCDs nominally would be “idle”. This is the issue to be mitigated by idleClearing.

3.2.3 IdleClearing

It is easiest to describe idleClearing in terms of one segment of one CCD. Since all of the segments for all CCDs in a partition are triggered synchronously one segment is sufficient. For this discussion, a single segment has N rows and M columns where N, M include any overscans. The standard values of (N, M) are $(2048, 576)$. To match the goal of constant power, we intend to read out pixels at the same average rate used during survey operation but we do not transfer any data outside of the REB. We defined a period P , and a number of rows n . We set the sequencer parameters for an ROI of nominal width but with only $n < N$ rows. Then we arrange to trigger the sequencer every P seconds. This yields a readout rate of $n \times M/P$ pixels/sec to be compared with $N \times M/(3 + \text{exptime})$ which is ≈ 35000 for the nominal 31 s exptime.

3.2.4 Science and Guider Details

For the science partition we choose $n = 100$, $P = 1.6$ so using $M = 576$ yields 36000 pix/sec. The plots for TotalRebPower during idle periods and observing match up as shown. [insert plot or two here with some explanation]

For the guider partition, we choose $n = 10$, $P = 0.090$ and $M = 576$ to yield 64000. We compare to the rate during nominal guiding where we read a 50×50 ROI at ≈ 9 Hz. Ignoring the flushed pixels this gives 22500 pixels/sec. However, this needs a bit more thought to account for the power used in shifting rows and columns at maximum rates to access the ROI which also

consumes quite a bit of power. [This needs a bit more flesh, but at first glance we should reduce P a bit if we need to match.] [insert plot or two here with some explanation]

With power matched and charges cleared, there is still the issue of latency when the system needs to be used for the intended purpose of taking images or guider stamps. In this case, the loop must be stopped and the system must transition into a state ready to operate. The science and guider cases are slightly different.

Consider the science partition case first, where the loop is implemented as a periodic task in java. The task is scheduled periodically so it just triggers (via the DAQ interface) the sequencer to run. The period is chosen to be longer than one triggered sequencer execution. When it is time to stop idleClearing the scheduled task is canceled and the system waits for confirmation that any in-progress triggered run of the sequencer is completed. Hence an upper bound on the time to stop is the time for one triggered run plus a very small amount of overhead. In the case of $n = 100$ rows this amounts to ≈ 150 milliseconds. This generally will not delay any images during survey imaging.

In the guider partition, the loop is implemented on an embedded linux node where it would not be appropriate to run java. The main loop has two tasks in series: first the sequencer is triggered (reading out 10 rows) and then the system sleeps for 80 ms. Neither task is interruptable but the series can be interrupted at the end of the reading or at the end of the sleep. Since the reading of 10 rows takes ≈ 12 ms, the upper bound on latency to stop the idle clearing in the guider is ≈ 80 ms. As with the science partition, the latency to stop idleClearing in the guider generally does not delay transitioning from idleClearing to survey operations.

Coordination and side-effects: managing the sequencer

Add a discussion of the issue that the sequencer is a single trigger for each partition and when running, has interactions with other aspects of REB/CCD management. So for example we would not want to turn on or off the CCDs with the sequencer running. Similarly if the focal-plane process for idleClearing is periodically triggering the trigger for the science partition and we attempt to power cycle a REB, when the REB reboots it may/will receive a trigger before the sequencer is fully loaded in the REB registers. This is a situation we don't want to happen. More on this topic....

3.3 Electro-Optical Testing

3.3.1 The eo_pipe Software Package

The eo_pipe software package (https://github.com/lsst-camera-dh/eo_pipe) is used to perform a standard set of electro-optical tests on LSSTCam pixel data. The initial set of tests were originally implemented for acceptance testing of the CCDs provided by the ITL and e2V vendors. The criteria and the algorithms used to analyze the CCD acceptance data are described in LCA-128 (<https://ls.st/LCA-128>) and LCA-10103 (<https://ls.st/LCA-10103>), respectively. The acceptance measurements included read noise, non-linearity, serial and parallel CTE, crosstalk, dark current, quantum efficiency, bright and dark pixel defects, traps, pixel response non-uniformity, and point spread function sizes. Over time, as testing progressed from CCD to raft to focal plane level, various other measurements and algorithms were added: photon transfer curve (PTC) generation, tearing detection, divisadero tearing, brighter-fatter analysis, bias shift detection, bias stability, CTI as a function of flux, gain stability, persistence measurement, row-means correlations, raft-level overscan and imaging region correlations, and scan mode analysis.

To ensure consistency with the official calibration products, the cp_pipe pipelines (https://github.com/lsst/cp_pipe) are run first to generate various data products that are inputs to the eo_pipe pipelines, i.e., combined bias, dark, and flat images, catalogs of pixel defects, PTCs, etc.. The eo_pipe tasks also perform other processing of the raw images to fill out its suite of tests. The eo_pipe analyses are implemented as “pipeline tasks” using the Rubin Science Pipelines middleware. This allows them to be run on LSSTCam data that has been ingested into a Rubin data repository and accessed by the Rubin data butler. The Rubin Batch Production Service framework (https://github.com/lsst/ctrl_bps) is used to run the eo_pipe pipelines at scale at the US Data Facility using the slurm batch system.

The eo_pipe results are persisted into the same data repository containing the raw LSSTCam data in collections that are labeled by analysis type, “acquisition run”, and weekly software release. Subsequent specialized follow up analyses can then be run on the eo_pipe data products by accessing those collections. The acquisition runs are identified for each set of LSSTCam configurations, comprising sequencer versions and various other LSSTCam conditions and settings. The data from those acquisition runs are processed together in an “analysis run” using specific versions of the cp_pipe and eo_pipe pipelines. Standard outputs from these analysis runs are available in web-based reports at <https://s3df.slaac.stanford.edu/>

data/rubin/lstcam/

- Correlated noise – Pierre Astier or Pierre Antilogus
- PTC – Pierre Astier

Draft

4 Camera on-sky performance

4.1 CCD Sensitivity to High Flux

Bright sources present a potential stress case for large-format focal plane electronics, motivating dedicated on-sky tests during LSSTCam commissioning. Extremely bright stars can, in principle, induce elevated charge levels that when read out can induce a safety response in the associated REB to power off CCDs as the readout current increases. Monitoring telemetry in response to high-flux signals is therefore a useful diagnostic for characterizing the behavior of electronics in response to bright sources.

The analysis presented here is based on the notebook `brightStarTests.ipynb` in the LSSTCam On-Sky Tests repository and focuses specifically on REB bias currents measured during and around bright-star exposures.

4.1.1 Test Description

The dataset consists of a sequence of on-sky exposures targeting bright stars with approximate apparent magnitudes of 0 (Alpha Centauri), 1 (tet Cen), and 2 (* gam Crv). These exposures were obtained by placing the stars on the central detector of specific REBs in the LSSTCam focal plane.

The following rafts and REBs were included in the study:

- R01/Reb1
- R01/Reb2
- R02/Reb1
- R02/Reb2

These REBs were selected based the results of Run 7, focusing on REBs that had tripped in response to high-flux exposures during the course of testing (LSST Camera Team et al., SITCOMTN-148).

Exposure numbers were used to retrieve precise timestamps from the exposure metadata, and fixed time offsets were applied to define consistent pre- and post-exposure windows.

4.1.2 Data

Telemetry was retrieved from the Engineering and Facility Database using the SAL topic `lsst.sal.MTCamera.re`. From this topic, fields associated with pre-switch bias currents (`hvbias_IbefSwch`) were selected for analysis.

For each exposure, telemetry was queried over a time range extending one minute before and after the nominal exposure start time. This approach allows transient effects to be distinguished from nominal baseline behavior.

4.1.3 Method

For each exposure and REB, time-series plots of the selected bias currents were generated. The analysis focused on identifying:

- Sudden current excursions coincident with bright star illumination.
- Sustained offsets relative to pre-exposure baselines.
- Differences in behavior between rafts or stellar magnitudes.

The time of exposure start and end was used as the reference for identifying correlated behavior. No filtering or smoothing beyond the native telemetry cadence was applied.

4.1.4 Results

Across all examined exposures, stellar magnitudes, and REBs, the bias-current telemetry remains within the acceptable safety limit of 120 mA. No current excursions above 97mA are observed (see Figure 50). No discontinuities, or systematic trends were observed during or after the bright-star exposures. This holds across the entire set of rafts tested, indicating that bright stars does not pose a risk to REBs that have shown previous sensitivity to high flux.




FIGURE 50: The largest bias current excursion observed during the bright star test. The vertical dashed lines indicate the start and end of the exposure. The solid lines indicate the measured current on each REB.

Telemetry inspection shows that any small-scale fluctuations in the currents are comparable to those seen outside of the exposure windows and are consistent with nominal electronic noise rather than illumination-induced effects.

4.1.5 Conclusion

These results indicate that bright star illumination at the tested magnitudes does not induce significant perturbations in REB bias currents that could trigger an automated shutdown of the CCDs. This provides confidence that LSSTCam electronics operate stably even in the presence of bright optical sources up to magnitude 0, and that such sources do not pose a risk to detector safety.

In summary, no evidence is found for anomalous electrical behavior correlated with bright-star exposures. These results verify the robustness of LSSTCam electronics under high-flux observing conditions and validate current operational assumptions for bright-source observations.

4.2 Pinhole repeatability tests




FIGURE 51: Caption

During the commissioning campaign we used the Collimated Beam Projector (CBP) to project a spot through the pinhole filter onto the LSSTCam focal plane. The main focus was on the repeatability of the filter placement by comparing the position of the spot before and after swapping the filter.

The image acquisition was carried out in a systematic sequence that combined three LSST-Cam rotator angles (0° , 18° , and -72°) with two CBP mask sizes (5 mm and 3 mm). For each combination we obtained several exposures before the filter swap and one after, yielding a total of 51 images. The exposures were taken in the order shown in Table 12; the exposure that was marked “IGNORE” was omitted from the analysis.

The left-hand panel of Figure 51 shows a difference image between the first exposure before the filter change and the fourth exposure after the filter change, revealing small but systematic shifts in the pinhole positions. The right-hand panel displays the measured positional offsets relative to the first exposure for images that are taken when the rotator angle was 0° . Scatter within a single condition (e.g., all points from a1) is on the order of a few pixels, indicating that the CBP itself is stable but non-negligible motion could be present.

Systematic offsets are evident when comparing conditions: the average shift between a1 and b1 is ≈ 2 px in X and ≈ 3 –4 px in Y. The four distinct groups (a1, a2, b1, b2) are distributed rea-

seqNum	Rotator angle (deg)	CBP mask (mm)	Before/After swap
31-36	0	5	Before
37-40	18	5	Before
41-44	-72	5	Before
45-50	0	5	After
51-56	18	5	After
57-62	-72	5	After
63-68	0	3	Before
69-74	18	3	Before
75-80	-72	3	Before
81-86	0	3	After
87-92	18	3	After
93-98	-72	3	After

TABLE 12: Sequence of exposures used for the filter-placement test. The table shows, for each exposure, the rotator angle, CBP mask diameter, and whether the exposures were taken before or after the filter change.

sonably consistent but distinguishable distributions. Because the total displacements are well below the 10-pixel (100 μm) repeatability requirement, we conclude that the filter placement is within specification. Nevertheless, the observed systematic trends suggest that both filter positioning and rotator repeatability contribute to the measured offsets, making it difficult to isolate the dominant source without further controlled experiments.

	Persistence	Phosphorescence
Sensor type affected	All e2v sensors	A subset of ITL sensors
Physical driver of effect	Charge trapping at the Si-SiO ₂ interface in the region of saturation, insufficient charge readout due to high parallel clock voltages difference	Fluorescent photo-resist wax trapped under the AR coating
Area of sensor effected	Region of saturated source, trail of persistence in direction of serial register for affected amplifier	Region of photo-resist wax localization
Peak signal level	~ 15 e ⁻	~ 100 e ⁻
Time constant	12.9 s	low: ~ 10s high: ~ 300s
Chromatic dependence	None	Higher signal in response to shorter wavelength exposures
Mitigation status	Mitigated to sub-electron level by decreasing parallel voltage swing in e2v sensors from 9.3V to 8.0V	Masking phosphorescent regions of specific ITL CCDs.

TABLE 13: The general characteristics and differences in the phosphorescence and persistence effects. Persistence time constant from Banovetz et al. (DMTN-276)

4.3 Phosphorescence

The Run 7 persistence optimization process used a short image acquisition sequence with a high flux exposure to saturate the sensors, followed by a series of dark images. The original intent of this script was to quantify persistence in e2v sensors. When the e2v sensors were shown to be nearly free of persistence by reducing their clock swing voltages from 9.3 V down to 8.0 V, a similar memory effect was immediately noticed in a subset of the ITL sensors.

4.3.1 The difference between phosphorescence and persistence

As noted in Sec. 7.4 of LSST Camera Team et al. (SITCOMTN-148), there are many similarities between phosphorescence and persistence. However, the physical origin, characteristics, time constants, signal amplitude, and chromaticity response of each signal is slightly different. We summarize the differences between each effect in table 13.

We describe the on-sky persistence characterization in subsection 4.4. Here, we describe the on-sky characterization of phosphorescence.

	u-band exposures	g-band exposures
Exposure Numbers		2025122400073 - 2025122400397

TABLE 14: The exposures chosen for analyzing phosphorescent response in ITL CCDs.

4.3.2 Previous characterization of phosphorescence

Phosphorescence was first documented in the Run 7 testing campaign (LSST Camera Team et al., SITCOMTN-148). To summarize the phosphorescence analysis from Run 7:

- The phosphorescent signal in ITL sensors can be morphologically classified as either **spot-like** or **diffuse** phosphorescence.
- Removing the HV bias from the sensors can decrease the amplitude of spot-like phosphorescence and increase the amplitude of diffuse phosphorescence
- Independent of morphology, the phosphorescent signal is stronger at shorter wavelengths.
- The phosphorescent signal is triggered at lower illumination levels for shorter wavelengths.
- The time dependence of the phosphorescent signal can have multiple time constants in a single phosphorescent complex.
- The kinetics of a phosphorescent region can be characterized using a multi-time constant decay model.

To characterize phosphorescence in on-sky exposures, we visually inspect series of on-sky exposures taken in g and u band for selected sensors with documented phosphorescence. Exposures are listed in Table 14. We focus here on u and g band, since results from Run 7 testing demonstrated that the amplitude of the phosphorescence signal is larger for short-wavelength exposures.

4.3.3 Phosphorescence in u-band exposures

4.3.4 Phosphorescence in g-band exposures

We visually inspected all of the g-band exposures listed in Table 14, searching for phosphorescent signal on a detector with a large diffuse phosphorescent complex (R43_S11). After visual inspection, we identified the following exposures that have a bright star coincident with the phosphorescent region:

- 2025122400076
- 2025122400078
- 2025122400087
- 2025122400091
- 2025122400130
- 2025122400153
- 2025122400223
- 2025122400237

4.3.5 Conclusion

4.4 Persistence

One of the major successes of the electro-optical testing campaign on the summit (Run 7) was the mitigation of the persistence effect by changing the CCD operation voltage (LSST Camera Team et al., SITCOMTN-148). This change greatly decreased the persistence signal to sub-ADU level. This was confirmed by taking uniform flat illumination images. Here we demonstrate the level of persistence by reducing the on-sky data where much sharper stars play a role.

On the night of May 15, 2025, LSSTCam observed a bright star on R22/S11, which is an e2v sensor, during BLOCK-T417 (CLT-001 Incremental Closed Loop w/ coma bending modes). The first panel of Figure 52 shows the bright star image (Image sequence 470) after applying the standard `isrTask` (correcting bias, dark, flat, linearity, defects, etc.) and `calibrateImageTask` (calibrating WCS, photometry, and removing sky background). The bright star saturated its footprint to the full well and bled charges across entire columns. This star could leave the persistence signal in the subsequent image (471), but there was no obvious persistence remaining (the second panel of Figure 52; the same image process has been applied).

A stronger test is to apply the difference imaging technique. We used image 479 as a template for the difference imaging because the LSSTCam observed mostly the same field and there was no major bright star observed in the preceding images. We further applied `'lsst.afw.math.warpExposure'` with `'lanczos3'` to match the template image to the science image, and `'lsst.ip.diffim.subtractImages.AlardLup'` to solve the PSF in order to match the PSF shapes of two images and convolve the PSF to the best S/N image (here it was the template) and subtract the warped regularized template from the science image.

The result is shown in the right-most panel of Figure 52. No persistence residue is obvious in the region where the bright star was located. There are residual bright stars and dark spots. The bright stars are the stars that are above the saturation, and the dark spots are associated with the saturation bleed in the template image.

The image was obtained in the z-band. The mean sky level for 471 reported by RubinTV was 21296.30 ADU. A study in the u-band, where the sky level is much lower, will be needed.




FIGURE 52: Demonstration of the absence of the persistence effect on an e2v sensor (R22/S11). The strip for channel 14 is cut out from an assembled CCD image. "Bright star" shows how the star appeared in the observation, "Next image" shows the next image after the bright star image, "Template" is a template image, and the Difference image is also shown.

5 Acknowledgements

This material is based upon work supported in part by the National Science Foundation through Cooperative Agreements AST-1258333 and AST-2241526 and Cooperative Support Agreements AST-1202910 and AST-2211468 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSST-DA Institutional Members.

6 References

- [DMTN-276], Banovetz, J., Utsumi, Y., Slater, C., 2024, *Effects of Persistence on E2V Sensors and its Impacts on DC2 Data*, Data Management Technical Note DMTN-276, NSF-DOE Vera C. Rubin Observatory, URL <https://dmtn-276.lsst.io/>
- e2v technologies, 2009, *CCD250 Design Report*, Tech. Rep. LSST-E2V-RP-002, e2v technologies, Issue 2
- Lage, C.S., 2025, E2V output with off-chip (REB) circuit, LTspice schematic
- [SITCOMTN-148], LSST Camera Team, Antilogus, P., Astier, P., et al., 2025, *LSST Camera Electro-Optical Test (Run 7) Results*, Commissioning Technical Note SITCOMTN-148, NSF-DOE Vera C. Rubin Observatory, URL <https://sitcomtn-148.lsst.io/>, doi:10.71929/rubin/2583999
- Marshall, S., 2025a, modify to allow operation with faulty R20/Reb2/S21/RDV voltage reading, SLAC JIRA, LSSTCCSRAFTS-777, Raised 19 April 2025, resolved 7 May 2025
- Marshall, S., 2025b, NTP for implemented solution to running 2 of 3 CCDs, Presentation, LSST-Cam, 19 June 2025
- [CTN-001], Plazas Malagón, A.A., Digel, S.W., Roodman, A., Broughton, A., LSST Camera Team, 2025, *LSSTCam and LSSTComCam Focal Plane Layouts*, Camera Technical

Note CTN-001, NSF-DOE Vera C. Rubin Observatory, URL <https://ctn-001.lsst.io/>, doi:10.71929/rubin/2584019

Snyder, A., 2025, *E2V Operation With VOD Set to Zero*, UCD Report, University of California, Davis,
Last updated 10 June 2025

Snyder, A., Polin, D., Lage, C., Tyson, J.A., 2024, In: Holland, A.D., Minoglou, K. (eds.) X-Ray, Optical, and Infrared Detectors for Astronomy XI, vol. 13103 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1310321, doi:10.1117/12.3020540, ADS Link

Turri, M., 2025, Merge LSSTCCSRAFTS-777 into main, org-lsst-ccs-subsystem-rafts, commit 829a8f2,
7 May 2025; released in org-lsst-ccs-subsystem-rafts 3.0.10

Utsumi, Y., Antilogus, P., Astier, P., et al., 2024, In: Holland, A.D., Minoglou, K. (eds.) X-Ray, Optical, and Infrared Detectors for Astronomy XI, vol. 13103, 131030W, International Society for Optics and Photonics, SPIE, URL <https://doi.org/10.1117/12.3019117>, doi:10.1117/12.3019117

[CTN-003], Utsumi, Y., Schutt, T., Liang, S., et al., 2025, *Configuring the PTP setting on the MOXA network switch*, Camera Technical Note CTN-003, NSF-DOE Vera C. Rubin Observatory, URL <https://ctn-003.lsst.io/>

[SITCOMTN-136], Walter, C., 2025, *Absolute Observatory Timing Calibration using Geo-Synchronous Satellites*, Commissioning Technical Note SITCOMTN-136, NSF-DOE Vera C. Rubin Observatory, URL <https://sitcomtn-136.lsst.io/>

7 Acronyms

Acronym	Description
AST	NSF Division of Astronomical Sciences
AURA	Association of Universities for Research in Astronomy
CCD	Charge-Coupled Device
CCS	Camera Control System
CTN	Camera Technical Note

DE-AC02	Department of Energy contract number prefix
LSST	Legacy Survey of Space and Time
LSST-DA	LSST Discovery Alliance
OCS	Observatory Control System
PTC	Photon Transfer Curve
SLAC	SLAC National Accelerator Laboratory

Draft