



Vera C. Rubin Observatory
Rubin Observatory Operations


The LSST Camera commissioning

LSST Camera Team, Pierre Antilogus, John Banovetz,
Alexandre Boucaud, James Chiang, Seth W. Digel, Alan M. Eisner,
Kevin Fanning, Anthony S. Johnson, Claire Juramy-Gilles,
Travis Lange, Shuang Liang, Sean Patrick MacBride, Stuart Marshall,
Homer Neal, Dmitry Onoprienko, Andrew P. Rasmussen,
Aaron Roodman, Theo Schutt, John Gregg Thayer, Max Turri,
Yongqiang (Brian) Qiu, and Yousuke Utsumi

CTN-007

Latest Revision: 2026-09-12

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
1.1	Issue register and activity timeline for 2025	1
1.1.1	Issue summaries	4
2	Camera operations	36
2.1	Recovery of CCDs	36
2.1.1	R30/S12: dead-looking segment and suspected short	36
2.1.2	R30/S12: running two of the three CCDs on Reb1	41
2.1.3	R20/Reb2/S21: faulty RD voltage read-back	42
2.2	Cryo system	46
2.2.1	Brief introduction to the refrigeration system	46
2.2.2	Architecture and Thermal Design of the Cryostat	49
2.2.3	Architecture of the Cryogenic Systems	52
2.2.4	Cryogenic Mixed-Refrigerant System	57
2.2.5	Brief introduction to the temperature sensitivity issue of the Cryo refrigeration system	63
2.2.6	Detail of the cryo-temp sensitivity issue	64
2.2.7	Solution to temperature sensitivity issue	65
2.3	CCS development	71
2.4	Firmware development	73
2.5	Observation control	77
2.6	Filter Exchange System Commissioning	77
2.6.1	Hardware: problems of youth & fixes	77
2.6.2	Slow Control Upgrade	78
2.6.3	Performances	79
2.7	Shutter	81
2.7.1	Shutter Timing issue	81
2.7.2	Shutter blade trajectory fitting	89
2.8	Guider star finder	89

2.9	Historical telemetry of infrastructure cut events: electricity, glycol, dynalene and air	89
3	Camera optimization	89
3.1	Sequencer changes	89
3.1.1	Long-range correlations	89
3.1.2	3 s sequencer file development	91
3.1.3	Natural frequency	94
3.2	Re-implementation of "clear" operation during the idle time	120
3.2.1	Introduction	120
3.2.2	Background/Context for the sequencer, DAQ and CCS	120
3.2.3	IdleClearing	121
3.2.4	Science and Guider Details	121
3.3	Electro-Optical Testing	123
3.3.1	The eo_pipe Software Package	123
4	Camera on-sky performance	125
4.1	CCD Sensitivity to High Flux	125
4.1.1	Test Description	125
4.1.2	Data	126
4.1.3	Method	126
4.1.4	Results	126
4.1.5	Conclusion	127
4.2	Pinhole repeatability tests	128
4.3	Phosphorescence	130
4.3.1	The difference between phosphorescence and persistence	130
4.3.2	Previous characterization of phosphorescence	131
4.3.3	Phosphorescence in u-band exposures	131
4.3.4	Phosphorescence in g-band exposures	132
4.3.5	Conclusion	132
4.4	Persistence	133

5 Acknowledgements	135
6 References	135
7 Acronyms	136
A Chronology of LSSTCam operations in 2025	137
A.1 January — final clean-room work and preparation for the lift	137
A.2 February — installation on the Top End Assembly	138
A.3 March — the camera reaches the telescope	140
A.4 April — first full focal-plane operation and the R30/Reb1 short	141
A.5 May — cold weather and the first loss of cryo capacity	143
A.6 June — First Look, and completion of the science array	145
A.7 July — operational maturity of idle clear and the FES	147
A.8 August — winter storms, power failures and firmware	148
A.9 September — recovery, then preparation for the shutdown	150
A.10 October — the maintenance shutdown and the loss of Cryo 3	151
A.11 November — control-system instabilities and the Cryo 6 trip	153
A.12 December — refrigeration margin and firmware campaigns	154

The LSST Camera commissioning

1 Introduction

The LSST Camera is the largest digital camera in the world.

The LSST Camera was delivered to the Vera C. Rubin Observatory atop Cerro Pachón in Chile in 2024. Camera re-verification activities were subsequently carried out, as documented in LSST Camera Team et al. (SITCOMTN-148).

This technote describes the commissioning of the LSST Camera on the telescope in 2025, with a focus on camera hardware. Engineering first light through the telescope, marking the first photons detected by the camera on the telescope, was achieved on April 15, 2025. Throughout the year, we encountered and addressed a number of technical challenges and implemented various improvements.

The technote is organized as follows:

Section 1 provides an overview of the events and issues encountered during commissioning. Section 1.1 is constructed by Claude based on all the message exchanges in Rubin Obs Slack channels `#cam-summit-ops` and `#cam-operator`.

Section 2 describes selected topics in greater detail.

Section 3 reports on activities related to system optimization.

Section 4 documents camera performance based on on-sky observations.

Chronology based on the Slack message exchange is also available in Appendix A.

1.1 Issue register and activity timeline for 2025

This section organises the 2025 camera operations record by *problem*: each recurring or distinct issue is given an identifier, summarised once, and placed on a timeline showing when it was actually being worked on. The date-ordered account it is drawn from is reproduced in

Appendix A. The purpose of the organisation is to separate three things that a chronology conflates — issues that were closed, issues that the observatory learned to live with, and issues that were never closed at all.

Identifiers carry a category prefix: CRYO (cryogenic refrigeration), DYN (dynamene (utility-trunk cooling)), GLY (glycol (facility)), PCS (pcs chiller), PWR (site power), AIR (compressed dry air), FP (focal plane), FW (reb firmware), DAQ (data acquisition), CCS (camera control system), SHT (shutter), FES (filter exchange system), OPT (stray light), MEC (mechanical and integration), IT (summit it). Status is assigned on the evidence of the operations channels, using five categories:

R	Resolved	cause found and fixed, no recurrence in the logs
M	Mitigated	recurrence contained by a procedure or design change, mechanism unchanged
W	Workaround	operated around a known defect that was never repaired
O	Open	not closed: still active or recurring at the end of the year, or never resolved and no longer being mentioned
T	Transferred	handed to another group and tracked outside camera operations

2 entries are marked “?”. In those cases the record stops before the outcome, so the classification has to be settled from sources outside these channels; the summary says in each case what is missing.

	Status	Issues
R	Resolved	9
M	Mitigated	8
W	Workaround	6
O	Open	21
T	Transferred	1
	Total	45

TABLE 1: Status of the 45 issues identified in the 2025 camera operations channels. Open covers both ends of a spectrum: issues that were still generating traffic in December, and issues that were never closed and simply stopped being mentioned — the timeline of Figure 1 distinguishes the two, since the latter show a long tail with no activity. Those, together with the 6 workarounds, are the entries most likely to need a decision.




FIGURE 1: Activity timeline of the issues in the register. A solid bar marks a period in which the issue was actively being discussed or worked on in #cam-summit-ops or #cam-operator; the thin grey line joins first to last mention, so gaps in it show how long an issue lay untouched between episodes. Issues that were still open at the end of the year run to December and end in an arrow head. The letter at the right is the status of Table 1 (**R**esolved, **M**itigated, **W**orkaround, **O**pen, **T**ransferred); “?” marks the entries whose outcome is not determinable from the channels. An open issue whose last bar falls well short of December was not being worked on at the end of the year. The identifiers on the left are links to the corresponding entry in Section 1.1.1. The grey band is the three-week maintenance shutdown of October 1–20; the dashed lines mark the installation on the Top End Assembly (February 5), the installation on the telescope (March 6), the First Look event (June 23) and the return to sky after the shutdown (October 20). Absence of a bar means silence in these two channels, not necessarily inactivity.

1.1.1 Issue summaries

Cryogenic refrigeration

CRYO-01 Loss of cryo cooling capacity in cold ambient air. *First to last mention 2025-05-08 to 2025-10-25; status **Mitigated**.* The dominant operational problem of the year. Below a dome air temperature of roughly 7°C the cryo circuits lose capacity, the cryo-plate trim heat falls to zero and the focal plane can no longer be held in thermal control; six distinct events between May and September required powering off the CCDs, the HV and in the worst cases nearly all REBs, with recoveries of seven hours to three days (OBS-931, FRACAS-280). Pointing the telescope to horizon to regain capacity, and adjusting the refrigerant charge, were immediate tweaks. The lasting mitigation only became possible once the mechanism was understood — re-condensation of vapour refrigerant in the line, governed by the relation between the ambient temperature and the phase-separator temperature. With that understanding, insulation and heater tape on the refrigeration lines became the major activity and their implementation was carried through, together with a matching adjustment of the glycol temperature, which is what sets the phase-separator temperature. Treated in detail in Sections 2.2.5 to 2.2.7.




FIGURE 2: **CRYO-01.** Top: the 8 May loss of capacity as it was reconstructed afterwards — cryo temperatures, cryostat and heat-exchanger vacuum and the cryo-plate trim heat, with the moments the dome-closure decision was taken and completed marked (2025-05-27). Bottom: trim heat applied to the cryo plate falling to 0.78 W, the floor at which the plate can no longer be held in control (2025-05-25).

CRYO-02 Cryo 1 heat-exchanger return temperature below the dome dew point. *First to last mention 2025-05-12 to 2025-09-09; status **Open**.* Cryo 1's heat-exchanger return line ran colder than the in-dome dew point, risking frost or an ice sheet on the plumbing, and the circuit had to be stopped nightly for several weeks. Throttling the liquid-line isolation valve (June 1–4) and then removing refrigerant to lower the discharge pressure raised the return temperature to about +3 °C, and in September the circuit was removed from the dome-closure decision criteria. None of that was a fix: Cryo 1 recovered by itself, and the underlying question — whether a low heat-exchanger return temperature at low refrigerant flow is acceptable — has never been settled either way. Experience since, including in high humidity, gives no indication that it is critical, but establishing that would require measuring the tube temperature against the temperature and humidity of the surrounding air. The trim-heat and dome-closure behaviour that accompanies these nights is shown in Figure 2.

CRYO-03 Cryo 6 compressor failure. *First to last mention 2025-04-01 to 2025-06-07; status **Resolved**.* Cryo 6 tripped on a compressor failure on April 1 and stayed out of service. The compressor module was replaced with the Cryo 5 unit from the MRCR between June 2 and June 7 (LCA-14623), and after charging, gas chromatography and additive top-up the circuit returned to service, restoring six of six cryo circuits.

CRYO-04 Capillary clogs (Cryo 1 cycled clear, Cryo 3 lost, Cryo 4 self-cleared). *First to last mention 2025-05-12 to 2025-12-11; status **Open**.* The first of the three is Cryo 1, and it is the one that was never named as a clog at the time. The circuit is believed to have been clogged by the recovery activity of around May 10, after which its heat-exchanger return temperature ran persistently low — the condition registered separately as CRYO-02, since what was managed nightly was the dew-point risk rather than the clog behind it. From May 12 the circuit was power-cycled repeatedly, at times daily, in the hope of cycling the obstruction clear; that continued for some months and the circuit did eventually recover, but no single intervention can be identified as the one that cleared it. During the October recharge Cryo 3 developed a clog — the capillary is the suspected location — that survived repeated power cycles, refrigerant additions and running the telescope horizontal; the circuit was given up on October 20 and the camera has run on five of six circuits since. A second clog, on the Cryo 4 capillary, developed in early December, drove cryo trim heat down to about 1 W, and cleared spontaneously on the evening of December 10 without intervention. The open question is whether contamination was introduced into Cryo 3 during the October maintenance itself; it was on that suspicion that revisions to the evacuation and charging procedure and to filter-dryer replacement (LCA-14625, LCA-17726) were proposed. Cryo 3 was still out of service at year end.




FIGURE 3: **CRYO-04.** The Cryo 4 clog developing through November: the C4 exit temperature (blue) drifts up towards the pre-capillary temperature over five weeks, the signature by which the clog was recognised (2025-12-11, #cam-operator).

CRYO-05 Cryo 6 PLC death after an external setpoint change. *First to last mention 2025-11-18 to 2025-11-19; status **Open**.* On November 19 the Cryo 6 PLC was declared dead at 00:34 UTC with the discharge temperature reading 122 °C and the oil level negative; the PLC recovered by itself 33 minutes later. The overheating of the cabinet that preceded it was attributed to a software update elsewhere on the summit mis-setting the glycol setpoint in response to bad environmental sensor readings (FRACAS-322, OBS-1394, OBS-1395). The PLC failure itself is suspected to be a problem in the PLC, and its root cause is not known. The CCS logic that closes the coolant valve when a PLC is declared dead was flagged as needing revision (LSSTCCSREFRIG-319 and LSSTCCSREFRIG-332).





FIGURE 4: **CRYO-05**. Left: the CCS Cryo 6 panel during the failure, reporting State: PLC_DEAD with every PLC error condition ACTIVE. Right: the discharge temperature and compressor current across the event, the temperature rising while the compressor is off (2025-11-19).

Dynalene (utility-trunk cooling)

DYN-01 Recurrent loss of Dynalene flow to the utility trunk. *First to last mention 2025-01-18 to 2025-12-17; status **Open**.* Dynalene, which cools the utility-trunk electronics, failed repeatedly all year — from control computer (cRIO) reboots and crashes, chiller faults and power events — each time forcing an emergency power-down of CCDs, REBs and often the shutter and filter changer (OBS-1014, FRACAS-321, FRACAS-328). A review of the August 11 event established that the camera tolerates only about 13 minutes without flow before REB power-supply boards approach their 48°C trip limit. The cRIO was replaced on November 4 and a small UPS was added in late November, but outages continued into the second half of December.




FIGURE 5: **DYN-01.** Utility-trunk Dynalene flow collapsing to zero. This is the signature the events have in common: whatever the cause upstream — a cRIO reboot, a chiller fault, a power event — what the camera sees is the flow going away, and from that moment it has about thirteen minutes before the REB power-supply boards approach their 48 °C trip limit (2025-06-03, #cam-operator).

DYN-02 Camera warm-up on Dynalene chiller changeover and filter cleaning. *First to last mention 2025-02-21 to 2025-12-22; status **Open**.* Distinct from an outright loss of flow: routine facility actions — weekly Dynalene filter cleaning, a clogged glycol filter, or switching between chillers — warm the supply enough to push REB power-supply board temperatures towards their limit and trip the HV. A clogged filter drove utility-trunk Dynalene to 22 °C in February, and on December 22 a chiller changeover forced HV and then the CCDs off for about an hour (OBS-1490). No interlock or procedural gate prevents this at year end. The flow signature is the one in Figure 5; here it is the supply temperature rather than the flow that moves.

DYN-04 Condensation and humidity around the Top End Assembly. *First to last mention 2025-06-14 to 2025-12-16; status **Open**.* The camera region of the Top End Assembly runs more humid than ambient, and the compressed dry air supplied there has little authority over it — the June investigation found the supply had been silently off for an unknown period and that the air reaching the camera is heavily mixed with dome air. In November infrared imaging showed hot spots and a suspected dry-air leak near the heat-exchanger return, and in December water droplets were found on M1M3, traced to condensation on uninsulated Dynalene supply pipes at the Top End Assembly (SUMMIT-11795, OBS-1479). Open at year end, with insulation the proposed fix.



FIGURE 6: **DYN-04.** Left: the liquid found below the camera cable wrap on 16 December, a little left of centre looking into the mirror — the drops that were traced back to the Top End Assembly. Right: the source, measured a week later: an uninsulated Dynalene line at the cable wrap reading 1.6 °C on its surface, well below the dew point of the air around it (2025-12-16 and 2025-12-24, #cam-summit-ops).

Glycol (facility)

GLY-01 Glycol flow interruptions to the PCS on power events. *First to last mention 2025-04-28 to 2025-11-28; status **Resolved**.* Every commercial power glitch or generator transfer stops the glycol recirculation pumps for roughly 70–90 seconds, against an observed PCS chiller trip threshold of 90–120 seconds — so each event is a near miss, and several were not (FRACAS-277, FRACAS-278, FRACAS-279). The identified fix, putting the recirculating pumps on UPS, was still pending throughout 2025; in September, observing on generator power had to be restricted in pointing because of it. The issue was closed in early February 2026, after the period covered here.



FIGURE 7: **GLY-01**. One of the near misses, at second resolution. Glycol input flow to the PCS (top) drops from 5.2 gallons per minute to zero at 20:45:30 and is restored about 70 seconds later; the chiller outlet temperature (middle) had already climbed from 12 to 30°C before the interruption and overshoots to 54°C on restart, and the cryo water-out temperatures (bottom) follow it up. The chiller trips if the flow is lost for 90 to 120 seconds (2025-10-07, #cam-operator).

PCS chiller

PCS-01 PCS chiller reliability. First to last mention 2025-03-28 to 2025-08-11; status **Resolved**.

A cluster of PCS chiller problems in spring, all of them traced and fixed: an under-rated 32 A supply breaker against a 50 A service requirement, which tripped the chiller twice in one morning (OBS-826); a stage-1 refrigerant leak that took the chiller out of service for a day and left the camera quiescent overnight; and spurious trips with no real chiller anomaly, traced to an unseated pin in a refrigeration-permit connector on the telescope-side cabling and aggravated by TMA shaking tests (SUMMIT-10642, SUMMIT-10652). The breaker was replaced, the leak repaired and the connector pins and shell renewed.



FIGURE 8: **PCS-01.** The pin face of the refrigeration-permit connector on the telescope-side cabling. An unseated pin here is what produced the spurious chiller trips — no chiller anomaly, a permit signal that came and went as the TMA moved. The pins and the shell were renewed (2025-04-11, #cam-operator).

PCS-02 PCS chiller restart difficulty, and the MRCR unit. *First to last mention 2025-03-28 to 2025-07-02; status **Open**.* Restarting the chiller is a separate problem from the failures above and was never solved: the unit repeatedly trips on low HFE flow because gas bubbles form in the longer line configuration, which procedure LCA-20281 does not cover. In March it took roughly a dozen attempts over two days, and was recovered only by pressurising the coolant tank; a loud bypass-valve noise during those attempts was later explained as the high-pressure bypass cracking at an 80 psi pump differential. The same behaviour blocked the restart of the unit at the MRCR in July. The MRCR system was found to have failed during the year and is tracked as LSSTCAM-119.




FIGURE 9: **PCS-02.** The chiller PLC panel during a restart attempt: stopped, with every error condition clear. Nothing is reported as wrong, and the unit will still trip on low HFE flow when it is asked to start again — which is why the restarts took about a dozen attempts over two days (2025-03-28, #cam-operator).

Site power

PWR-01 Site power: grid glitches, generator failure and UPS capacity. *First to last mention 2025-02-25 to 2025-11-28; status **Open**.* The camera is exposed to the summit's power quality throughout: glitches stick valves closed (FRACAS-282), trip chillers and compressors, and interrupt cooling. The most severe event was August 30, when the generator tripped — its air inlet blocked by ice and snow (FRACAS-299) — and left the camera on UPS alone. The damage was not done by the loss of power to the camera itself but through the cooling chain: the glycol feeding the cryo circuits stopped with the recirculation pumps (GLY-01), the cryo systems were lost, and with them the cryostat pressure began to rise towards the ion-pump shutdown limit, at one point estimated to be about fifteen minutes away. The REBs and REB power supplies were shut down by hand while this was happening; generator power returned first and the pressure recovered. Separately the server-room UPS exhausted its roughly 30-minute battery (FRACAS-300) and the backup diesel generator was found already out of service. UPS batteries were replaced in November; the underlying exposure remains.




FIGURE 10: **PWR-01**. The chain of 30 August, in the order it happened. Top: glycol input flow to the chiller collapsing when the generator trips, and the chiller outlet temperature stepping up as it goes — “No glycol flow to chiller. Tripped.” The cryo circuits go with it. Bottom: the consequence some twenty minutes later — cryostat and heat-exchanger pressure leaving 10^{-7} Torr and rising past 1.1×10^{-6} , with the cryo plate warming from -124 towards -113 °C at about 0.6 °C per minute. The ion-pump shutdown limit was estimated to be fifteen minutes away when generator power came back (2025-08-30).

Compressed dry air

AIR-01 Compressed dry air supply failures. *First to last mention 2025-02-27 to 2025-10-08; status **Mitigated**.* Compressed dry air operates the vacuum gate valves and purges the camera body, so its loss is a direct camera concern — including frosting risk on the L3 lens. Supply was cut for two hours for compressor maintenance in February, both Level-1 compressors degraded in August when their dryers froze at -8 °C, and a compressor was lost to a packing-seal failure in October. Insulating jackets on the dryers addressed the freezing mode; single-compressor operation remains a routine condition.



FIGURE 11: **AIR-01.** The compressed-air dryers at Level 1, whose freezing in August stopped the air: the towers with their exhaust and drain lines, and the insulating jacket being fitted (2025-08-23 and 2025-08-26).

Focal plane

FP-01 R30/Reb1/S12 suspected OD short (subject of this note). *First to last mention 2025-04-07 to 2025-12-10; status **Workaround**.* On April 7, after a routine power cycle and a reorientation of the camera to zenith, R30/Reb1 raised a Bias2 OD2 overcurrent on S12 ($ODI = 0.213\text{ A}$ against an 0.082 A limit); S12 reached bias voltage in two steps and tripped off on both attempts, and the behaviour was read at the time as a real short. A retest on April 25 showed an unexplained digital-current spike of order 1 A , and the REB was left as the only one switched off camera-wide through May. It was returned to service on June 25–26 running two of three CCDs, which required an R30/Reb1 firmware update and relaxed bias-control limits; the CCD current on S12/Seg17 was noted as anomalously low at that turn-on and again on June 29. The REB has carried its own firmware variant through every subsequent firmware campaign. The workaround — running two of the three CCDs on Reb1 — was implemented on June 25–26 and has been in place since; the fault itself is unresolved. Described in detail in Sections 2.1.1 and 2.1.2.

FP-02 R20/Reb2/S21 unphysical RD voltage read-back. *First to last mention 2025-04-02 to 2025-04-24; status **Workaround**.* The April 2 CCD shorts test found R20/Reb2/S21 reading an unphysical $+54\text{ V}$ RD (reset drain) voltage; power cycling the REB and resetting the RCE did not change it, and the working theory was an open ADC input on the board. The sensor blocked power-on of the raft until April 24, when the erroneous RD reading

was configured to be ignored in software and R20 was returned to service. No hardware repair was attempted. Described in detail in Section 2.1.3.

FP-03 Non-linear HV bias behaviour on R21/Reb0 and R12/Reb0. *First to last mention 2025-05-05 to 2025-10-25; status **Mitigated**.* R21/Reb0 repeatedly glitched and tripped its HV bias near 50 V, where its DC-DC converter behaves non-linearly, once producing visible tree rings mid-integration. It was set to run permanently at 46 V, and R12/Reb0 later showed similar sensitivity requiring extra care during ramp-up. Managed by lowered setpoints and careful ramping rather than repaired.




FIGURE 12: **FP-03.** The glitches on R21/Reb0: HV bias voltage, current and DAC setting stepping and spiking as the bias is ramped, the behaviour for which the operating point was lowered (2025-05-05, #cam-operator).

FP-04 R03/Reb1/S11/Seg00 chronic bad amplifier. *First to last mention 2025-04-06 to 2025-05-10; status **Workaround**.* A long-standing bad channel described in the logs as having gone “back and forth from working to not, over the last few years”. Its current stayed pinned at 0.5 mA through repeated power cycles in April and May, and the neighbouring Amp16 also failed to return once. The channel cannot be repaired in place; since the impact of a single amplifier is small, the decision is not to repair it and to continue operating with it as it is.




FIGURE 13: **FP-04.** R03/Reb1 being powered on, 15 April. In the bottom panel every segment current rises to about 1.2 mA except one, which stays flat: S11/Seg00, the channel that has gone back and forth for years and did not come back this time either (2025-04-15, #cam-operator).

FP-05 Corner-raft and wavefront-REB power-on workarounds. *First to last mention 2025-04-29 to 2025-10-21; status **Workaround**.* Corner rafts do not reach the correct dphi level on power-on and need an undocumented DAC-toggle workaround (LSSTCAM-107), and a wavefront REB power-on failure requiring a REB power cycle recurred after the October shutdown (LSSTCAM-37). Both were still being worked around at year end rather than fixed.

FP-06 REB power-supply trips suspected to be electrostatic discharge. *First to last mention 2025-01-21 to 2025-04-29; status **Open**.* Recalled in January: in September 2024 there had been one or two REB power-supply trips that seemed to correlate with people touching or standing near the utility trunk, which is what put the camera-door grounding question (MEC-04) on the agenda in the first place. A concrete case followed on April 29, when R21/Reb0 tripped during carbon-snow cleaning of M2 with no light-leak explanation and was tentatively attributed to an electrostatic discharge. The investigation is tracked in the dedicated channel #cam-issues-esd. The door interface a permanent strap has to use is shown in Figure 35.

REB firmware

FW-01 REB firmware applying maximum heater power. *First to last mention 2025-08-11 to 2025-08-23; status **Resolved**.* The firmware booted during the August 11 recovery made

the CCDs run measurably warmer, which was initially misread as lost cryo capacity during the storm. Power-cycling the focal plane and rolling the firmware back confirmed a real power-draw increase of roughly 30 W, and the cause was a register-map conflict in 0x313a5013 that applied maximum heater power — a defect introduced together with the sequencer-override register in 0x31395011 and unnoticed until then. It was fixed by the firmware upgrade to 0x313a5014. See the version history in Table 10.



FIGURE 14: **FW-01.** Left: CCD temperatures and CCD heater power across the firmware roll-back, the step in heater power being the extra draw the faulty build was applying (2025-08-21). Right: the same measurement annotated during the investigation, with the moment the sensors were turned off marked and the firmware version under test named (2025-08-22).

FW-02 REB serial-number read race leaving REBs INVALID. *First to last mention 2025-10-13 to 2025-12-15; status **Open**.* Whenever REBs are power-cycled or booted from a different firmware slot, some come back reporting a serial number of zero and stay in an invalid validation state, requiring link resets or RCE reboots before they can be used. It appeared with the October candidate-firmware rollout and recurred through every subsequent power cycle to the end of the year; tooling improvements (single-RCE reboot, a re-initialise command) were proposed but not delivered.

```
-- 2025-10-13 19:10 and 21:26 UTC, after the candidate-firmware rollout, #cam-operator

R01/Reb1: CCDsPowerState:OFF HVBiasState:OFF RebDeviceState:ONLINE RebValidationState:INVALID
R01/Reb2: CCDsPowerState:OFF HVBiasState:OFF RebDeviceState:ONLINE RebValidationState:INVALID
R30/Reb1: CCDsPowerState:OFF HVBiasState:OFF RebDeviceState:ONLINE RebValidationState:INVALID
R31/Reb1: CCDsPowerState:OFF HVBiasState:OFF RebDeviceState:ONLINE RebValidationState:INVALID
R32/Reb1: CCDsPowerState:OFF HVBiasState:OFF RebDeviceState:ONLINE RebValidationState:INVALID

-- 2025-11-04 15:31 UTC: the standing recovery, written out for the night crew

Check ID:
rms_read camera 0x800001
Re-read ID:
rms_write camera 0x800000 1
Reset opcode error:
rms_write camera 0x390001 1

-- "Reading 0x800000 shows you the IDs, the firmware bug means that sometimes they come up
-- 0x00000000. If they do, you can cause the REB to re-read the ID by writing 1 to 0x800001."
```

FIGURE 15: **FW-02.** REBs online but refusing validation after a power cycle, and the register recipe that became the standing workaround. The identifier read is a firmware race, not a hardware fault: the board is reachable, it simply reports a serial number of zero until it is told to read it again. No fix was delivered in 2025.

FW-03 Sequencer opcode errors after RCE reset. *First to last mention 2025-10-09 to 2025-12-29; status **Open**.* REBs reset or power-cycled while idle clear was running came up with the sequencer opcode-error bit set, needing a register write to clear. Delaying the start of idle clear until the subsystem is operational was adopted as the mitigation, but the error recurred at the end of December with idle clear confirmed off, pointing at a firmware or boot-sequence cause instead (LSSTCAM-153). Open.

```
-- 2025-12-14 20:17 UTC, #cam-operator, after an RCE restart with idle clear disabled

The RCE restart left three issues. Disabling the idle clear didn't resolve all the issue.

2 Links went down
R22/Reb0: CCDsPowerState:UNKNOWN HVBiasState:UNKNOWN RebDeviceState:OFFLINE RebValidationState:UNKNOWN
R23/Reb2: CCDsPowerState:UNKNOWN HVBiasState:UNKNOWN RebDeviceState:OFFLINE RebValidationState:UNKNOWN

3 REBs have DELTA state in the CCD power
R13/Reb2: CCDsPowerState:DELTA HVBiasState:ON RebDeviceState:ONLINE RebValidationState:VALID
R14/Reb0: CCDsPowerState:DELTA HVBiasState:ON RebDeviceState:ONLINE RebValidationState:VALID
R34/Reb2: CCDsPowerState:DELTA HVBiasState:ON RebDeviceState:ONLINE RebValidationState:VALID

1 opcode error in rms_read
R43/Reb1
```

FIGURE 16: **FW-03.** The observation that reopened the issue. Idle clear — the mitigation adopted earlier in the year — was confirmed off, and an RCE restart still produced a sequencer opcode error, which points the cause at the firmware or the boot sequence rather than at idle clear (LSSTCAM-153).

FW-04 Sequencer optimisation and the 6.4 ns natural-frequency firmware. *First to last*

mention 2025-04-26 to 2025-12-29; status **Resolved**. A year-long programme rather than a fault: new sequencer files to address e2v noise and covariance anomalies (clear improvement on e2v, little change on ITL), the 3 s readout sequencers, and finally the retuning of the sequencer clock to 6.4 ns. The natural-frequency firmware was tested on the summit in December in three campaigns and used for a holiday calibration run, but was not promoted to the default slot; the dense photon-transfer curve that these changes require was repeatedly deferred through the year. The programme was closed in early January 2026, after the period covered here.




FIGURE 17: **FW-04**. Left: read noise per e2v amplifier with the new sequencer against the old one — most amplifiers unchanged, the worst improved (2025-04-28). Right: segment signal with the 3 s corner-raft sequencer against the ITL one, compared with the expected ratio (2025-06-25).

Data acquisition

DAQ-01 Crate, RCE and data-node crashes. First to last mention 2025-03-24 to 2025-11-24; status **Open**. A steady background of DAQ hardware and host faults: bad B-side links traced to DAQ transceivers (March, fixed), memory failures on the data-collection nodes (OBS-859), crashed storage and switch-controlling RCEs (OBS-1208), a data-transport module needing rebooting twice in a month (LSSTCAM-130), and links auto-negotiating down or dropping packets (LSSTCAM-151). Individual instances were cleared by reboots and reseating; the recurrence rate was not. A serial console was subsequently connected to the worst offender to allow it to be debugged, as documented in IHS-9431.

```
-- lsstcam-dc06 kernel log, 2025-04-21 03:12 UTC, found during the 1 May investigation
-- and posted to #cam-operator

[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: Hardware error from APEI Generic Hardware Error Source:
8
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: event severity: corrected
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: fru_text: DIMM_A8
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: section_type: memory error
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: error_status: Storage error in DRAM memory
(0x0000000000040400)
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: physical_address: 0x0000001e772bb000
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: node:0 card:4 module:1 rank:2 bank:9 row:61369
column:512
[Mon Apr 21 03:12:16 2025] {1}[Hardware Error]: error_type: 2, single-bit ECC
[Mon Apr 21 03:12:16 2025] mce: [Hardware Error]: Machine check events logged
```

FIGURE 18: **DAQ-01.** One of the hardware faults behind this entry: a corrected single-bit ECC error on a data-collection node, recorded ten days before anyone went looking for it. Faults of this class were cleared individually — reseating, reboots, transceiver swaps — but the rate at which they arrived was not addressed. A serial console was later connected to the worst offender so that it could be debugged when it next fell over (IHS-9431).

DAQ-02 Dropped guider stamps at high ROI rates. *First to last mention 2025-10-10 to 2025-10-13; status **Mitigated**.* Stress-testing large guider regions of interest with the new DAQ release repeatedly dropped guider stamps on the path from the DAQ to lsstcam-dc10, the host that receives and writes them, and the same episodes coincided with dropped group-communication messages. Extensive network diagnosis of the ATCA switches and the backplane found no definitive cause; the GDS socket buffers were enlarged to 16 MB and a buffer monitor was run on lsstcam-dc10, after which a 20-image run completed with no dropped packets.

```
-- 2025-10-10 18:59--19:15 UTC, guider stress test on lsstcam-dc10, #cam-operator

"I just took one image MC_C-20251010-000004. It seems that there were some dropped stamps."

"N_STAMPS" : 139, "N_STAMPS" : 138, "N_STAMPS" : 139, "N_STAMPS" : 139,
"N_STAMPS" : 138, "N_STAMPS" : 136, "N_STAMPS" : 139, "N_STAMPS" : 139,

("There should have been 139 stamps in that series.")

[2025-10-10T18:59:30.132+0000] WARNING: MC_C-20251010-000004: DAQ stamp count 139 not equal to number
of stamps received 138 (org.lsst.ccs.daq.guider.FitsWriterFactory$FitsWriter fixupStampCount)
[2025-10-10T18:59:30.174+0000] WARNING: MC_C-20251010-000004: DAQ stamp count 139 not equal to number
of stamps received 138
[2025-10-10T18:59:30.174+0000] WARNING: MC_C-20251010-000004: DAQ stamp count 139 not equal to number
of stamps received 136

-- after a buffer monitor was put on the receiving host, 35 minutes later: "That time they all had 139."
```

FIGURE 19: **DAQ-02.** The loss counted at the receiving end. The DAQ reports how many stamps it sent and the image handler reports how many it wrote, and the two do not agree — by one to three stamps per sensor per series. Enlarging the GDS socket buffers to 16 MB and watching them on lsstcam-dc10 removed the loss without ever identifying where it occurred.

Camera control system

CCS-01 Guider state desynchronisation and focal-plane hangs. *First to last mention 2025-05-01 to 2025-11-24; status **Open**.* The guider state held by CCS and by the DAQ drift apart — from dangling initialisation commands, double commanding between CCS and OCS, or new pause timeouts — leaving the guider stuck paused (once for 22 hours, unnoticed, with no guider data written) or the focal-plane subsystem hung in an active command state needing a restart. Individual manifestations were patched repeatedly; the pattern was still recurring on multiple nights in late November (LSSTCAM-148).

```
-- 2025-11-24 09:47 UTC, focal-plane log, posted to #cam-operator the same morning

[2025-11-24T09:47:03.412+0000] INFO: Received command from mcm (user mcm): Command(RawCommand)
    initGuiders with 1 argument ({"common":{"rows":400,"cols":400,"integrationTimeMillis":200}, ...
[2025-11-24T09:47:03.455+0000] WARNING: Accepting initGuiders when guiders already in PAUSED state
    (org.lsst.ccs.subsystem.focalplane.Guiding initGuiders)

... no command response for the next 46 minutes ...

[2025-11-24T10:33:08.699+0000] SEVERE: Raised Alert CMDFAIL (Command failed) at severity ALARM 1 time.
    Last cause : "Execution of command "initGuiders" failed unexpectedly due to: Guider start
    failed, status -1 (rc=-1 GDS - Command Timeout)". Overall AlertState= ALARM
org.lsst.ccs.command.CommandInvocationException: org.lsst.ccs.CommandHelper$CommandFailedException:
    Exception while executing command, fault state entered: Guider start failed, status -1 (rc=-1
    GDS - Command Timeout)

-- 2025-11-02 03:24 UTC, the same disagreement seen from the operator's side

initGuiders failed
Guider entered unexpected state, expected=PAUSED actual=ERROR
```

FIGURE 20: **CCS-01.** The state disagreement in one place. CCS accepts an `initGuiders` for guiders it already believes are paused, gets no answer from the guider, and only declares failure 46 minutes later; three weeks earlier the same disagreement surfaced as a state the operator could not act on. The pattern was still recurring at the end of November (LSSTCAM-148).

CCS-02 Idle clear: from guider clearing noise to default operation. *First to last mention 2025-04-16 to 2025-07-04; status **Resolved**.* Guider idle-clearing was found in April to inject noise into neighbouring sensors during bias readout, making early bias images unusable and producing stripes in wavefront images (LSSTCCSRAFTS-778). The proper fix — a non-guider idle-clear scheme — arrived with CCS 41.0.4 in late June, needed two further locking fixes (skipped invocations piling up, and a sequencer-reload race that showed up as a runaway clock current on R30/Reb2), and was declared operational on June 30. Idle clear has been the default for all on-sky observations since early July, replacing continuous all-day biases.




FIGURE 21: **CCS-02.** A bias exposure with the guider running: the structure across the upper rafts is the artefact that idle clear was introduced to remove (2025-07-01).

CCS-03 ocs-bridge failures to release the mount motion lock. *First to last mention 2025-07-13 to 2025-11-09; status **Open**.* After a filter change the bridge between CCS and the observatory control system intermittently fails to acknowledge the unlock command to the mount and rotator, leaving them motion-locked and the filter command rejected (OBS-1098, OBS-1079, OBS-1119, OBS-1360). A SAL interface version mismatch is suspected. The standing recovery is to take the camera offline and restart the bridge; it cost whole nights in July and recurred in November.

```
-- 2025-04-30 02:40 UTC, ocs-bridge log, posted to #cam-operator

[2025-04-30T02:39:51.323+0000] INFO: MTRotator MotionLockStateEvent MotionLockStateEvent{lockState=2,
    identity=Script:100003}
[2025-04-30T02:39:55.421+0000] INFO: MTRotator MotionLockStateEvent MotionLockStateEvent{lockState=0,
    identity=Script:100003}
[2025-04-30T02:40:22.490+0000] INFO: Received command: SetFilterCommand{name=g_6}
[2025-04-30T02:40:22.491+0000] INFO: Sending MTRotator: LockMotionCommand{}
[2025-04-30T02:40:26.601+0000] INFO: Command complete ack=303
[2025-04-30T02:40:26.601+0000] INFO: Sending MTMount: LockMotionCommand{}
[2025-04-30T02:40:26.603+0000] INFO: MTRotator MotionLockStateEvent MotionLockStateEvent{lockState=3,
    identity=ocs-bridge}
[2025-04-30T02:40:26.627+0000] SEVERE: Command failed: errorCode=1, ack=-302, message=Failed: Currently
    <MotionLockState.UNLOCKING: 2>. Cannot lock motion.
[2025-04-30T02:40:26.627+0000] SEVERE: Failed to get motion lock during setFilter command
[2025-04-30T02:40:26.627+0000] INFO: Reject command: SetFilterCommand{name=g_6} because Command
    rejected: Failed: Currently <MotionLockState.UNLOCKING: 2>. Cannot lock motion.

-- and from the CCS console a quarter of an hour later

ccs>fcs setFilter 6
Error: Error dispatching command: The telescope motion and rotator need to be locked before
performing a filter exchange. Current status of the lock is UNLOCKED
```

FIGURE 22: **CCS-03.** A filter change refused because the mount is stuck between lock states. The rotator reports UNLOCKING when the bridge tries to lock it, the setFilter command is rejected, and the console shows the same lock reported as UNLOCKED. Recovery on this night needed the MTMount CSC to be cycled.

CCS-04 Vacuum-subsystem fault that silently disabled a safety interlock. *First to last mention 2025-02-25 to 2025-03-04; status **Mitigated**.* In February a vacuum-subsystem fault left seven MKS gauges — including the safety-critical cryostat and heat-exchanger (Hx) gauges — offline for over 15 hours and aborted the low-vacuum check, so the refrigeration-not-permitted alarm was never re-raised despite high pressures: protection remained nominally configured while not working. The defect was fixed under LSSTCCSSUB-212 and the corrected build was installed by hand on the summit in early March.

```
-- 2025-02-25 19:31--19:37 UTC, vacuum subsystem log, posted to #cam-operator

[2025-02-25T18:09:58.533+0000] SEVERE: Setting ASM380 offline due to error retrieving initial status
    ASM380: org.lsst.ccs.drivers.common.DriverException: Device not connected
    (org.lsst.ccs.subsystem.vacuum.ASM380Device initialize)

-- and, from the subsystem restart the previous evening:

java.lang.NullPointerException: Name is null
    at java.base/java.lang.Enum.valueOf(Enum.java:271)
    at org.lsst.ccs.subsystem.vacuum.constants.DeviceState.valueOf(DeviceState.java:11)
    at org.lsst.ccs.subsystem.vacuum.VacuumMain.updateCryoTurboPumpVentValve(VacuumMain.java:1118)
    at org.lsst.ccs.subsystem.vacuum.VacuumMain.updateSwitchState(VacuumMain.java:976)
    at org.lsst.ccs.subsystem.vacuum.VacuumMain.updateVacuumState(VacuumMain.java:676)
    at org.lsst.ccs.services.PeriodicTaskComponent$PeriodicTaskRunnableWrapper.run(PeriodicTaskComponent.java:205)
    at org.lsst.ccs.utilities.scheduler.PeriodicTask.run(PeriodicTask.java:186)
[2025-02-24T21:05:33.518+0000] SEVERE: Exception thrown by periodic task vacuum-state
    (org.lsst.ccs.utilities.scheduler.PeriodicTaskExceptionHandler onException)
```

FIGURE 23: **CCS-04.** The February fault as it appears in the log. Devices are set offline at subsystem start-up and the vacuum-state periodic task — the task that evaluates the pressures and raises `RefrigNotPermitted` — dies on a null device state. Seven MKS gauges, including the cryostat and heat-exchanger gauges, stayed offline for more than 15 hours; the interlock was configured throughout and did nothing.

CCS-05 Automatic REB protective shutdown did not fire. *First to last mention 2025-11-20 to 2025-11-20; status **Open**.* During the total Dynalene loss of November 20 the automatic REB shutdown response failed to act, because the REB power-supply board high-temperature limit was set above the value at which the response was expected to intervene, and the camera had to be shut down by hand (FRACAS-321). The high-temperature limit has not been changed since. Although the event presented as a cooling failure, the defect is in the control system's protective logic and its configuration, which is why it is registered here rather than with the Dynalene issues. The board temperatures this response is meant to act on rise as in the Dynalene events of Figure 5.

CCS-06 jGroups multicast storms (“meltdowns”). *First to last mention 2025-02-15 to 2025-11-23; status **Open**.* Retransmit storms on the CCS group-communication layer that break CCS-DAQ communication and, in the worst cases, require nearly every subsystem to be killed and restarted (OBS-950). Several distinct triggers were identified over the year — a console process hitting its memory limit, IGMP snooping enabled on particular switches, guiding at an untested 5 s cadence with high DAQ traffic, a locked console, Kafka broker memory exhaustion — and each was addressed, yet the class of failure recurred in November. No single root cause was established.

```
-- 2025-02-18 17:00 UTC, chiller (refrig) log on lsstcam-pcs02, posted to #cam-operator

Feb 18 17:00:39 lsstcam-pcs02 refrig[2624]: 4561 INFO: [STATUS] NAKACK2: retransmit request from
      CCSAddress s-CCS-Graphical-Console_lsstcam-mcm_FEutSUX uuid=9d3f19e...
Feb 18 17:00:40 lsstcam-pcs02 refrig[2624]: 4562 INFO: [STATUS] NAKACK2: retransmit request from
      CCSAddress s-CCS-Graphical-Console_lsstcam-mcm_FEutSUX uuid=9d3f19e...
Feb 18 17:00:41 lsstcam-pcs02 refrig[2624]: 4563 INFO: [STATUS] NAKACK2: retransmit request from
      CCSAddress s-CCS-Graphical-Console_lsstcam-mcm_FEutSUX uuid=9d3f19e...
Feb 18 17:00:43 lsstcam-pcs02 refrig[2624]: 4564 INFO: [STATUS] NAKACK2: retransmit request from
      CCSAddress s-CCS-Graphical-Console_lsstcam-mcm_FEutSUX uuid=9d3f19e...
Feb 18 17:00:56 lsstcam-pcs02 refrig[2624]: 4567 INFO: [STATUS] NAKACK2: retransmit request from
      CCSAddress s-CCS-Graphical-Console_lsstcam-mcm_FEutSUX uuid=9d3f19e...

-- 2025-04-07 18:39 UTC, mcm log: the other face of the same layer

[2025-04-07T18:39:02.871+0000] WARNING: [STATUS] JGRP000011: s-mcm: dropped message from non-member
      s-image-handling-lsstcam-dc09 (view=[s-camera-mmm|59166] (36) [s-camera-mmm, s-lockmanager,
      s-ocs-bridge, s-mcm, s-cluster-monitor, s-kafkaService, ...]) (org.jgroups.ccs.CCSLog warn)
```

FIGURE 24: **CCS-06.** What a meltdown looks like in the log: an unbroken stream of NAKACK2 retransmit requests, here all provoked by a single graphical console, and — from a different episode — a subsystem dropped from the group view while it is still running and still sending. Neither line names a cause, which is the difficulty with this class of failure.

CCS-07 Kafka broker and rebpower memory exhaustion. *First to last mention 2025-11-01 to 2025-11-24; status **Mitigated**.* The Kafka broker service and the REB-power subsystem exhausted their Java memory, the broker reaching extreme CPU use and the subsystem dropping to tens of megabytes free, which triggered an overnight storm (LSSTCAM-149, LSSTCAM-150). Restarting both services at the start of each night became standing practice — a workaround, not a fix.

```
-- 2025-11-23 04:48--04:50 UTC, rebpower and kafkaService logs, posted to #cam-operator

[2025-11-23T04:48:15.803+0000] WARNING: Raised Alert JVM/memory/available (JVM available memory is
      below 50 MB.) at severity WARNING 1 time. Last cause : "PULSE: Low available memory: 32 MB.
      Max: 1979 MB, Allocated: 1979 MB, Used: 1947 MB.". Overall AlertState= WARNING
[2025-11-23T04:48:18.017+0000] INFO: [STATUS] NAKACK2: retransmit request from CCSAddress
      s-kafkaService uuid=eab348cf... for (1): {96968562}
[2025-11-23T04:48:20.754+0000] INFO: [STATUS] NAKACK2: retransmit request from CCSAddress
      s-kafkaService uuid=eab348cf... for (1): {96969728}
[2025-11-23T04:48:49.643+0000] WARNING: REB PS board (192.168.1.53): Read of seqno a0135940 took
      unexpectedly long time 763ms (org.lsst.ccs.drivers.auxlex.Srp receive)
[2025-11-23T04:48:49.644+0000] SEVERE: Periodic task monitor-update/RebPS/P03 did not finish execution
      by the end of its period of 1000 ms. Still running. ... GC: G1 Young Generation from 254 to 255
      for 1, G1 Old Generation from 255 to 1017 for 762.
[2025-11-23T04:49:59.230+0000] WARNING: [STATUS] JGRP000011: s-kafkaService: dropped message from
      non-member s-rebpower (view=[s-camera-mmm|80414] (36) [...])
```

FIGURE 25: **CCS-07.** Ninety seconds of the November 23 storm, in order: rebpower down to 32 MB of its 1979 MB heap, retransmit requests from the Kafka service, REB power-supply reads taking three quarters of a second, the monitoring task overrunning its 1 s period in old-generation garbage collection, and finally rebpower dropped from the group view. The memory exhaustion precedes the communication failure rather than following it (LSSTCAM-149, LSSTCAM-150).

CCS-08 Silent death of periodic tasks. *First to last mention 2025-03-22 to 2025-10-07; status **Open**?* A separate and more general defect than CCS-04: a periodic task stops executing after reaching its maximum allowed number of exceptions and the subsystem carries on as if nothing had happened. In March a chiller control panel was found frozen, showing the chiller off with the emergency stop asserted while it was in fact running, because the state-update task had died on a null-pointer exception; the resulting failure alert appeared in the log and the alert viewer but did not reach the Slack channels, which is what allowed it to go unnoticed. A similar periodic-task failure recurred in October. Whether the general problem was addressed afterwards is not visible in the operations channels.

```
-- 2025-03-22 18:59 UTC, chiller-ccs log of 2025-03-16, retrieved after the panel was
-- found frozen six days later and posted to #cam-operator

ccs-logs-chiller-ccs-20250316.log.0.gz:
[2025-03-16T22:51:52.603+0000] SEVERE: Raised Alert PeriodicTaskFailure (Alert raised when periodic
    tasks have failures (exceptions) during execution. At ALARM level the task execution is
    terminated.) at severity ALARM 1 time. Last cause : "Task updateCtrlState has experienced the
    maximum allowed number of exceptions: 10 (max:10). It has stopped executing.
    ([updateCtrlState])."

-- the exception itself

[2025-03-16T22:51:45.047+0000] SEVERE: Exception thrown by periodic task updateCtrlState
    (org.lsst.ccs.utilities.scheduler.PeriodicTaskExceptionHandler onException)
java.lang.NullPointerException: Cannot invoke "java.lang.Boolean.booleanValue()" because the return
    value of "org.lsst.ccs.subsystem.refrig.ChillerMaq20Device.isLightOn(org.lsst.ccs.subsystem.refr
    ig.ChillerMaq20Device$LightNumber)" is null
```

FIGURE 26: **CCS-08.** The failure mode in full. A periodic task reaches its allowed number of exceptions, stops executing, and the subsystem goes on serving the last state it managed to write — here a chiller shown as off with the emergency stop asserted while it was running normally. The alert was raised on 16 March and the condition was noticed on 22 March: it reached the log and the alert viewer, but not the operations channels.

Shutter

SHT-01 Shutter blade overrun and stuck states. *First to last mention 2025-03-24 to 2025-10-21; status **Resolved**.* The shutter jammed part-open during a flat calibration in June, with the blade apparently overrunning its limit, and was recovered only by power-cycling both PLCs and the HCU and recalibrating (OBS-1013). It stuck again during a 15 s exposure at a non-zero rotator angle in October, throwing an illegal-duration exception in the controller (LSSTCAM-135), and needed reinitialisation on the first night back from the shutdown. An earlier motion-profile study in March had already shown anomalous Hall-sensor transitions and higher jerk than the 2022 reference set. These stuck and

overflow states are accounted for by the shutter timing analysis and were resolved with it. See Section 2.7.1.




FIGURE 27: **SHT-01**. Left: the shutter motion-profile study of March, fitting model start time, pivot points and jerk for each blade and comparing them with the 2022 reference set (2025-03-24). Right: shutter status and the quadbox 24V supply across one of the stuck-shutter events (2025-04-23).

SHT-02 Shutter PTP time synchronisation. *First to last mention 2025-06-04 to 2025-11-20; status **Mitigated**.* Repeated “timed out waiting for shutter exposure to complete” failures in June were traced not to the mechanism but to PTP clocking: enabling a boundary clock on the shutter’s switch reproduced the fault on the clean-room system, and transparent-clock mode and then the leaf switch as boundary clock removed the offset from master. The analysis and recovery procedure are documented in CTN-003. Synchronisation still falls out of lock after major network events. It did so twice in November: during the recovery from the November 3 jGroups meltdown, where `readPtpDiag.sh` was run repeatedly from 20:53 to 22:42 UTC and the offset from master ranged from tens of milliseconds up to 0.96 s before the shutter was recalibrated in a closed dome, and again on November 20 during the Dynalene recovery, when the shutter was restarted and the synchronisation was reported as bad at 19:52 UTC. Recovery in both cases required PLC power cycles and recalibration. See Section 2.7.1.

```
-- 2025-11-03, recovery from the jGroups meltdown: readPtpDiag.sh on lsstcam-shutter01,
-- run repeatedly between 20:53 and 22:42 UTC. Two of the readings, unedited except for
-- the lines common to both.
```

22:14:25			22:42:06 (after the network settled)
PTP state:	SLAVE (9)		PTP state:
Offset from master (ns):	961,800,987		Offset from master (ns): 43,082
Mean path delay (ns):	9,466		Mean path delay (ns): 26,598
Steps removed:	4		Steps removed:

```
-- and the field that shows what the shutter itself believed the time was:
```

```
External (PTP) time:    Wed, 10 Dec 2025 21:58:45.689932 UTC
(the HCU clock read 2025-11-03 22:14:25 at that moment)
```

FIGURE 28: **SHT-02.** Loss of lock after a network event, in the diagnostic the shutter provides. An offset from master of 0.96 s — against the tens of nanoseconds of a healthy lock — and a PTP time five weeks ahead of the real one; both recover without intervention once the network settles. This is the November recurrence referred to in the entry; the analysis and the recovery procedure are in CTN-003.

Filter exchange system

FES-01 Carousel clamp faults. *First to last mention 2025-05-02 to 2025-12-13; status **Open**.*

The dominant filter-exchange failure mode: a carousel clamp fails to open or lock within the allowed time, faulting the changer and costing observing time. It was attacked from three sides — automated unclamping recovery in July (essentially 100 % success once deployed, error rate below 1 in 400 over 1000 exchanges), mechanical servicing and greasing of the socket clamps during the October shutdown (the grease not having been renewed since 2021), and a warm-up brake sequence in November — yet socket-specific faults were still occurring in December. A filter had to be removed from the camera in July when socket 3 could not be trusted.

FES-02 Carousel socket 3 slip-ring collector. *First to last mention 2025-04-13 to 2025-04-25; status **Resolved**.* Socket 3 lost communication intermittently during carousel rotation, worse at speed and dependent on telescope orientation. The cause was a spring-loaded carbon collector contact that was not sliding freely; it was replaced on April 24–25 and the socket recovered (SUMMIT-10680). The eight untouched collectors, and the effect of the very dry camera atmosphere on them, were noted as a residual risk.

FES-03 Autochanger trucks, homing and the damaged X+ clamp motor. *First to last mention 2025-02-13 to 2025-12-13; status **Open**?* The autochanger truck controllers fault with undervoltage errors after losing homing, a problem that first blocked filter exchange in February and reappeared in December (OBS-1454); a new automated homing procedure

was qualified in hardware in March but the old method was kept for first photon. Separately the X+ online clamp motor, damaged in summer 2023, faulted repeatedly during April exchanges and was never repaired. Alignment test campaigns ran in September; what they concluded is not recorded in these channels.



FIGURE 29: **FES-03**. Left: an autochanger on its stand in the clean room, where the homing procedure was qualified (2025-04-07). Right: the quadbox power panel, whose FES carousel, clamp, loader and brake channels are the supplies behind the truck undervoltage faults (2025-02-13).

Stray light

OPT-01 Stray-light HV bias trips. First to last mention 2025-04-17 to 2025-10-22; status *Workaround*.

Light reaching the focal plane with the shutter closed repeatedly tripped the HV bias: sunlight through an open P-flow door with the mirror covers open, head lamps and light past the filter-loader door during FES work, a laser tracker left running all day, dome floodlights during a public demonstration, dome lights switched on by accident (about 20 channels tripped, all ITL rafts), and specular reflection of the moon within roughly 20 degrees. Each case produced a procedural rule — mirror covers closed when the P-flow is open, a minimum moon-avoidance angle, HV off during dome-shutter and filter operations — and the trips became less frequent, but the mechanism is unchanged.





FIGURE 30: **OPT-01**. Left: a focal-plane image taken while light was reaching the sensors with the shutter closed (2025-04-17). Right: the camera and the open P-flow area on the day light through that route tripped the HV bias (2025-05-06).

OPT-02 “Banana” stray-light feature in darks and flats. *First to last mention 2025-04-19 to 2025-04-20; status **Transferred**.* A banana-shaped feature appeared in dark and flat images, partly rotating with telescope azimuth and partly fixed, with equipment near M1M3 such as a laser tracker suspected as the source; reinstalling the pinhole filter was proposed as a partial mitigation for extended stray light. No follow-up appears in either channel: the investigation was handed over to the stray-light group and is tracked there.





FIGURE 31: **OPT-02**. Left: the banana-shaped feature in the focal plane that gave the issue its name. Right: the dome interior photographed the same evening while looking for its source (2025-04-19, #cam-operator).

Mechanical and integration

MEC-01 Sheared pancake-wrap bolts. *First to last mention 2025-01-28 to 2025-02-04; status **Resolved**.* Several socket-head bolts holding the pancake wrap's central rotating assembly to the adapter plate sheared during the installation campaign and were found loose in the cable wrap. The wrap was removed and rebuilt with new hardware at a reduced torque, and the interface was reworked (pin holes reamed, internal sliding blocks removed) in early February; a subsequent bolt test to 75 ft-lb showed no failure. The joint has given no trouble since and the item is treated as closed.



FIGURE 32: **MEC-01.** Two of the sheared 3/8-16 socket-head bolts recovered from the cable wrap, and the bolt test that followed a few days later: 75.2 ft-lb applied with no sign of failure (2025-01-30, #cam-summit-ops).

MEC-02 Camera cable wrap interlock and telemetry faults. *First to last mention 2025-02-18 to 2025-10-09; status **Open**.* In February the cable-wrap local interlock at Level 3 was found to be entirely isolated from the rotator interlock, so the rotator would not stop on a cable-wrap fault, and the two were joined. Since installation the wrap has instead blocked rotation: loss of its telemetry faulted the rotator in August (OBS-1200) and it stuck in a fault state during October maintenance rotations, recovered by bypassing the interlock at the mount engineering interface. Described in the logs as a known issue.





FIGURE 33: **MEC-02.** Left: the operating note for the interlock, which has to be reset from the mount engineering interface after a fault or power cycle even when the status lights are green (2025-09-23). Right: the camera rotator and cable-wrap region on the Top End Assembly (2025-03-03).

MEC-03 Creaking noise during large camera rotations. *First to last mention 2025-08-04 to 2025-09-07; status **Open**.* A repetitive clicking or creaking picked up by the camera shutter microphone, described as sounding like a Geiger counter. It was first analysed on August 4–5, when it appeared between 23:06:13 and 23:07:18 UTC at 2–6 kHz while the telescope was slewing: the M2 hexapod was found not to be moving but to be holding position with its actuator currents oscillating in a zigzag, which made it the first suspect. On August 12 an alternative explanation — the flapping straps of a temporary safety barrier on the mesh platform, blown by the HVAC fans and heard as the telescope slewed past — was proposed and the straps were taped down. The same sound returned on September 7 during 20-degree rotations every three minutes, which disproved the strap theory: this time the M2 hexapod was stationary while the camera hexapod moved with every rotation, though M2-hexapod motor currents still spiked at each one. The recirculating ball bearings of the hexapod roller-screw actuators, or of the rotator, were then proposed as the source, and whether the sound has grown louder with age was raised and not answered. A ticket was opened (OBS-1158) and nothing further appears in these channels.




FIGURE 34: **MEC-03.** The August analysis. Top: spectrogram of the shutter microphone recording of 2025-08-04, with the click visible as a vertical line across the band at about 18 s into the excerpt. Bottom: M2-hexapod telemetry for the same minute — the hexapod is holding position, not moving, while its six actuator currents oscillate in the zigzag pattern that first made it the suspect (2025-08-05, #cam-operator). The September recurrence, which ruled out the safety-barrier explanation, is not illustrated here: the camera hexapod was the one moving then.

MEC-04 Electrostatic charging of the camera doors. *First to last mention 2025-01-21 to 2025-01-22; status **Workaround**.* The camera doors ride on non-conductive tracks and are not electrically bonded, so they can charge as they slide; the temporary grounding wires used in the clean room had to be removed before the utility-trunk skins went on. The accepted interim measure was an operational step to ground the doors by hand when opening or closing them. A proper workaround using a grounding strap is being implemented. The REB power-supply trips that motivated the concern are registered separately as FP-06.




FIGURE 35: **MEC-04.** Left: one of the temporary grounding wires used in the clean room, which had to come off before the utility-trunk skins were installed (2025-01-21). Right: the utility-trunk door seen from inside, with the chrome-plated lug surface and the door magnet marked — the interface a permanent grounding strap has to use (2025-01-21).

Summit IT

IT-01 Collateral damage from summit OS and Kubernetes patching. *First to last mention 2025-02-06 to 2025-12-16; status **Mitigated**.* Roughly monthly summit patching costs the camera a day and a full reboot cascade across its servers, and has twice done real harm: in August a scripting error included the shutter HCU that was meant to be excluded and an authentication race condition then locked it out for hours — judged serious because the same fault on a hardware-safety host would have been worse — and in December the reboot coincided with a Dynalene outage. The camera team’s response was procedural: take the camera offline first, reboot in a defined order, verify link speeds and images before handing back.

```
-- 2025-08-26 21:25--21:39 UTC, #cam-operator. The shutter HCU had been rebooted by summit
-- patching that was supposed to exclude it. The subsystem came back and kept producing
-- telemetry; nobody could log in to it.
```

```
lsstcam-vw02$ ssh lsstcam-shutter01.cp.lsst.org
Connection closed by UNKNOWN port 65535
```

```
$ nmap -p 22 lsstcam-shutter01.cp.lsst.org
Host is up (0.023s latency).
PORT      STATE SERVICE
22/tcp    open  ssh
```

```
$ ssh lsstcam-shutter01
The authenticity of host 'lsstcam-shutter01' can't be established.
ED25519 key fingerprint is SHA256:enI61SsU+a9SGdEwIPkAF5Lfqh0wTDvi8xLy1PPZEMk.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Server host/lsstcam-shutter01.cp.lsst.org@LSST.CLOUD not found in Kerberos database
```

FIGURE 36: **IT-01**. Patching collateral damage at its most awkward: the host is up, the port is open, the subsystem is running and sending telemetry, and the camera team is locked out of it for hours by an authentication race. The shutter was not used until access was restored. The same fault on a host carrying hardware protection would have been considerably worse, which is why this entry is kept.

1.1.1.1 Scope of these judgements. An issue is called open when the channels record no resolution, whether or not it was still being discussed in December; that is a statement about the record, not proof that nothing was done. Statements of cause are those recorded in the channels at the time. Two issues, GLY-01 and FW-04, were open through the end of 2025 and closed in early 2026, after the period covered here.

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. 37). The read-out section of the CCD250 is organised in a repeating group; in CCS names the pin order is S1 S2 S3 FS OD1 OS1 RD1 c1amp1 RG1 OG1 BS GD FS OD2 OS2 RD2 c1amp2 RG2 OG2, preceded by the parallel clocks P1–P4 (Fig. 38). 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. 40) 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. 39, Table 2). 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

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




FIGURE 37: 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 38: 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.

each trip is documented by only a few samples and the voltage at which the current peaked is loosely constrained.




FIGURE 39: 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.

TABLE 2: 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 $\rightarrow$ 15 V	14.94 V	213.1 mA	~ 140 ms
2	2025-04-07T20:54:32	15 $\rightarrow$ 22.3 V	22.27 V	220.6 mA	~ 100 ms
3	2025-06-03T19:09:08	15 $\rightarrow$ 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 $^{\circ}\text{C}$ excursion of focal-plane/R30/Reb1/Temp9, indicating localised power dissipation at the fault site.




FIGURE 40: 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.

2.1.1.3 Interpretation. A shorted segment reading zero current follows from where the currents are measured (Fig. 40). Each segment current is monitored on the low side of that 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 10) 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 $0.02686 \times 4096 = 110.0$ V, or ± 55 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 ± 5 V, so one code is $10/4095 = 2.442$ mV at the ADC input; RD reaches that input through a divide-by-eleven network, so in units of RD volts one code is $2.442 \times 11 = 26.862$ 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 ≈ 25 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 1100 W. 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 3 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 4.

Finally, Tables 5 and 6 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

TABLE 3: Summit Facility Interface between Telescope and Camera

Requirement	Specification	Notes
Refrigerant transfer line skin temperature	Within $+4/-15$ °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 $+1/-3$ °C of dome's ambient temperature	Foam insulated cabinet walls, and coolant used to control the cabinet interior temperature

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 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

TABLE 4: 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.

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 41 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 independent 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

TABLE 5: 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 6: 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.

from the rest of the cryostat. Figure 42 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 43. 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 44. 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 43).

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 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]




FIGURE 41: Cross section of the cryostat with its principal thermal regions and typical design temperatures during operation (in red).




FIGURE 42: Detailed routing of the refrigeration system at the interface between the cryostat and the heat exchanger vacuum system.




FIGURE 43: Cryoplate (a) Routing of the six refrigerant circuits in the copper cryoplate (b) with stainless steel stiffening frame attached.




FIGURE 44: The physical layout of each Cryo system is shown schematically in Figure 44. 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.

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 7).

TABLE 7: 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 shutdown 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 45) 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 45: 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-

ant 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 8). 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 9.

TABLE 8: 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 46. 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 between the saturated vapor and saturated liquid temperatures, at a constant pressure) is zero or

TABLE 9: 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




FIGURE 46: 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 46.

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 45), there are vapor refrigerants and liquid refrigerants coming out at the compressor cabinet and sent to the heat exchanger at the utility trunk via differ-

ent 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 47

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 48

All five running circuits showed an increase in PreC4Tmp, indicating that the C4 flow was reduced. This is shown in Figure 48

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 49

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 50.

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 47. 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 51. 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 47: Ambient temperature, Cryo trimheat and cryoplate temperature at around 6 am UTC on May 8, 2025




FIGURE 48: C3ExitTmp and PreC4Tmp of each Cryo circuit at around 6 am on May 8, 2025




FIGURE 49: Discharge and suction pressure of each Cryo circuit at around 6 am on May 8, 2025




FIGURE 50: Glycol/Water Inlet temperature and phase separator temperature of each Cryo circuit at around 6 am on May 8, 2025




FIGURE 51: 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 52. 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 52: Cryo system behavior when ambient temperature reduced to 0.75 C on September 8, 2025

2.4 Firmware development

Science REB (REB)

Table 10: 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 10 – 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 10 – 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 11: 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 12: 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 53: 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 54, for example the long tail visible in the setFilter execution time prior to the FCS October 2025




FIGURE 54: 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 55: Histogram of filter exchange duration between April 15th and July 18th (blue) and after July 18th (red)




FIGURE 56: 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 57: 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 57:

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 $10\ \mu\text{s}$ 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 58 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 58: 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 59). 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 59: 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 59, 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 60). 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 60: 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 61), and it was reported independently in the analyses of Shuang Liang and Thibault Guillemin.




FIGURE 61: 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 62). 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 62: 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 10. 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 63). 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 63: 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 64). 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 64: 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).

Overscan Mean vs Std for all Rafts/Segments (20251214, seqnum 211-309)




FIGURE 65: 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 61) are absent.

3.1.3.4 Bias row profiles. The row profiles measured before and after the change are compared in Figure 66. 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 66: 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 67); 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 67: 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 68 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 68: 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 69) 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 69: 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 68.

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 70 for the original and for the natural frequency configuration. The checkerboard pattern is gone in the latter.




FIGURE 70: 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 71 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 72).




FIGURE 71: 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 72: 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 73). 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 74), 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 75); 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 73: 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 74: 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 75: 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 76). 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 77).




FIGURE 76: 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 77: 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 78). 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 78: 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 79); both are still being followed up.




FIGURE 79: 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 80. 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 80: 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 81), and print-through was identified as the cause: a section of the serial overscan shows the print-through structure directly (Figure 82). 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 83).




FIGURE 81: 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 84). Whether this is the same underlying




FIGURE 82: A section of the serial overscan for LATISS, 20251226, sequence numbers 207–306, showing the print-through structure.




FIGURE 83: 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 84: 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 85) 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 vignettied 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 85: 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 86). 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 86: 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 87: 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 13; the exposure that was marked “IGNORE” was omitted from the analysis.

The left-hand panel of Figure 87 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 13: 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 14: 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 14.

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 15: 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 15. 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 15, 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 88 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 88; 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 88. 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 88: 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

A Chronology of LSSTCam operations in 2025

The year 2025 took LSSTCam from the Level 3 clean room to routine on-sky operation on the Simonyi Survey Telescope. The chronology below records the milestones of that transition together with the principal incidents and their resolutions, since much of the operational context relevant to individual focal-plane channels – including the R30/Reb1 events discussed in this note – is only legible against the wider integration and commissioning history. It is compiled from the camera operations Slack channels (#cam-summit-ops and #cam-operator); entries are restricted to what those logs state explicitly. Throughout, “Hx” or “heat exchanger” refers to the heat exchanger inside the camera and its vacuum volume, not to the camera hexapod; where the hexapod itself is meant it is named as such.

A.1 January — final clean-room work and preparation for the lift

Jan 5–6 The camera was powered back up in the Level 3 clean room after the holiday break, opening a year planned around installation, first photon and the start of the survey. Breakers were closed in the sequence main/protection, 24V clean, 24V and 48V dirty, and it was confirmed that the FES PLC power correctly defaults to off.

Jan 7 ComCam was removed from the Top End Assembly (TEA) and relocated onto the camera shipping frame, clearing the top end for LSSTCam.

Jan 6–15 An intermittent utility-trunk leak-detector fault (UtLeakDetcFault) produced three ALARM events and two after-hours callouts with no liquid ever found. A deliberate test with Dynalene drops on the leak cable produced the distinct correct alert (UtCoolantLeak), showing that the recurring fault was a sensor/cable fault; on Jan 15 a load resistor was added to the detector’s 24V-to-5V DC-DC converter, after which the fault behaved correctly.

- Jan 8–14** The TEA was modified to receive the spiral and pancake wraps: holes were drilled and tapped through the first flange of the camera cable wrap (CCW) with chip-capture tooling and the pancake-wrap mounting blocks were installed and torqued.
- Jan 7–22** A sustained refrigeration campaign recovered the six Level 1 MRCR circuits: filter/dryer assemblies were removed for return to SLAC, 2919C refrigerant was recovered from the compressors and surge tanks, and all six cryo control boxes were repaired and reinstalled. Open faults included a Cryo 2 pressure transducer, Cryo 4 PLC communications and a broken Cryo 5 thermocouple.
- Jan 13–16** Dynalene Chiller 02, the unit serving LSSTCam, suffered an internal short circuit and was taken out of service; Dynalene continued to circulate through the camera via Chiller 01.
- Jan 18** Dynalene flow to the camera was lost unexpectedly with facility staff off site. The camera team opened the quadbox switches and stopped the FCS, vacuum and shutter processes as a precaution; flow was restored by powering the Dynalene cRIO back on.
- Jan 10–20** Final filter-exchange-system (FES) and OCS-bridge testing was completed, all filters were confirmed stored in the carousel, and on Jan 20 the camera and pump cart were powered off at the main breakers, removing any dependence on Dynalene.
- Jan 20–24** The pancake wrap was installed on the CCW and mated to the spiral wrap, every utility was disconnected, the utility-trunk skins were installed and the bulkhead limit switches were replaced. Removal of temporary grounding wires exposed an unresolved electrostatic grounding problem with the camera doors, which ride on non-conductive UHMW tracks; grounding the doors by hand was adopted as an interim operational step.
- Jan 27–31** Two hardware problems delayed the lift out of the clean room. Torquing the skin fasteners to a newly specified 34 in-lb repeatedly stripped screw heads, and several 3/8-16 SHCS holding the pancake wrap's central rotating assembly to the adapter plate were found sheared, forcing removal and reinstallation of the wrap with new hardware at a reduced interim torque. A separate bolt test to 75 ft-lb showed no failure, leaving the original failure unexplained.

A.2 February — installation on the Top End Assembly

- Feb 3–4** The pancake-wrap interface was reworked (pin holes reamed, bolts set to half torque, internal sliding blocks removed) and the camera was sealed and moved out of the Level 3

clean room.

Feb 5 LSSTCam was lifted from the saddle stand and installed on the camera rotator of the Top End Assembly — the first time the camera was integrated with part of the telescope. Staged hardware was found to be the wrong grade and a full set of A4-80 stainless bolts was used for the final installation.

Feb 7–11 MTP data fibres at the camera bulkhead were found contaminated with Dynalene residue. The connectors were cleaned and Fluke-tested; many failed certification because of scratches, and the cables were finally accepted as adequate for the DAQ on a best-effort basis (IT-5920, IT-5885).

Feb 6–13 Summit glycol upgrades and computer-room work left the camera unpowered and the control system down for about a week. The first power-on of LSSTCam on the TEA on Feb 13 was aborted minutes after the breakers closed when a back-flange temperature channel ran out of range together with a faint plastic smell; the reading was subsequently identified as a bad sensor and the power-on was repeated successfully.

Feb 13–14 FES carousel verification failed on the first attempt (autochanger truck undervoltage; carousel “Status 3, error register 32”, then a Following Error with no motion) and was recovered the following day by reseating the encoder ribbon cable between controller and motor.

Feb 17–26 All six cryo circuits were pressurised to 150 psig for pressure-hold tests. Leaks on circuits 1, 3 and 5 were tightened out; two Cryo 2 pressure-transducer channels read incorrectly with the fault localised inside the utility-trunk MAQ20 and worked around using spare cables; a Cryo 5 transducer of the wrong pressure range was replaced. All circuits were declared ready for transport by Feb 27.

Feb 18–28 The rotator, camera cable wrap and camera hexapod motion-test campaign ran on Level 3, with the camera rotating for the first time on Feb 18. On the same day the CCW local interlock at Level 3 was found to be entirely isolated from the rotator interlock — so that the rotator would not stop on a CCW fault — and the two were joined that afternoon.

Feb 21–26 REB power-on and fibre-link verification concluded with all 71 camera links declared good; only R20/Reb2 showed link errors during the raft-by-raft cycling, and these did not repeat on retest.

Feb 21–24 A clogged glycol filter at the Dynalene chiller drove utility-trunk Dynalene to 22 °C over a weekend and forced the camera down to minimum power (roughly 400 W); why the alarms did not trigger remained under investigation.

Feb 25–28 A vacuum-subsystem software fault left seven MKS gauges — including the safety-critical cryostat and heat-exchanger (Hx) gauges — offline for over 15 hours and aborted the low-vacuum check, so the REFRIG_NOT_PERMITTED alarm was never re-raised despite high pressures. A fix was prepared under LSSTCCSSUB-212.

A.3 March — the camera reaches the telescope

Mar 1–4 LSSTCam was fully disconnected on Level 3: HFE was purged back to the chiller reservoir, a temporary grounding chain was built, the CCS subsystems and HCUs were stopped and the camera breakers opened.

Mar 3–4 The observatory-wide Cycle 40 update was applied and CCS restarted, but the vacuum subsystem crashed on the new build with a null-pointer exception in the cryo turbo vent-valve update and had to be patched by hand on the summit.

Mar 5–6 The camera was lifted from Level 3 to Level 8 on Mar 5 and installed onto the TMA on Mar 6.

Mar 10–13 Camera utilities were reconnected on the TMA (Dynalene hoses, MTP fibres, main power, compressed dry air) and ComCam was moved back to the white room. The roughing pump cart was swapped after two kinks were found in its hose; a rate-of-rise test confirmed the repair.

Mar 11–13 Image acquisition in emulated mode through OCS was restored after the DAQ was found to be rejecting triggers because the image metadata contained guider sources not present in the partition; the emulation configuration was reshaped to science plus wavefront sensors only. Rapid Analysis went live on the summit on Mar 13, with camera readouts appearing on RubinTV.

Mar 14–18 The first attempt to power the camera on the TMA failed with no network connectivity to any HCU. The two LC fibres of the CCS pair were found swapped relative to the TMA breakout box; connectivity was made by changing which media converter carried the TMA fibre, and power-on was achieved on Mar 18 (IT-6016).

Mar 18–20 The TMA pump cart failed and was replaced by the previously removed unit, whose internal gauge was found miscalibrated for the summit elevation (corrected in the vacuum code). Both the heat-exchanger and cryostat chambers reached roughing vacuum, and repeated USB disconnects of the cart link forced a switch to the serial interface.

Mar 18–24 The FES carousel fault carried over from February was fixed with a 3D-printed retainer on the encoder cable, and five filter exchanges were demonstrated with the camera on the telescope on Mar 24.

Mar 24–26 Systematic REB/DAQ link verification found the whole focal plane good on the OTM-A side and four bad links on the B side, all connected to the same RCE; all four were recovered by swapping DPM transceivers in the DAQ, confirming the fault was on the DAQ side rather than in the camera or the fibre run.

Mar 27 With the PCS chiller cooled from +14.5 to -40°C , all REBs were powered on, the GREB firmware was updated into guider mode, and roughly 350 full focal-plane images with guiding were taken under CCS control.

Mar 28–31 PCS chiller 2 tripped twice when its upstream circuit breaker opened; the installed 32 A breaker was found under-rated against the chiller's 50 A service requirement (OBS-826). Restarting the chiller after the breaker work took roughly a dozen attempts over two days, repeatedly tripping on low HFE flow until the coolant tank was pressurised. Renewed Dynalene instability over Mar 30–31 tripped the REB power supplies on high temperature and led the team to conclude that starting the cryo system was too risky.

Mar 13, Mar 27 Meeting and shift structure was changed for the installation period: a daily Camera Coordination meeting was introduced, daily plans moved to threaded summit posts, and a weekend on-call person and an overnight alert address were named.

A.4 April — first full focal-plane operation and the R30/Reb1 short

Apr 1 The cryostat cooldown and refrigeration restart uncovered failed water solenoid valves on Cryo 3 and Cryo 6 (swapped with MRCR spares) and a Cryo 6 compressor failure; all cryostat and heat-exchanger ion pumps were subsequently turned on successfully.

Apr 2 The CCD shorts test passed for nearly the whole focal plane. The exception was R20/Reb2/S21, whose reset-drain voltage read an unphysical +54 V instead of the expected value; power cycling the REB and resetting the RCE did not change it, and the working theory was an

open ADC input on the board. That sensor was left off while the rest of the focal plane was powered on.

Apr 3–6 A PCS chiller firmware upgrade was followed by a major PCS failure: stage-1 refrigerant leaked out slowly, the control valve would not open and flow could not be sustained. The camera was held overnight in a quiescent state with CCDs and REBs off and the cryo circuits running, and the chiller was restarted successfully on Apr 6 — described by the team as the fastest recovery of its kind.

Apr 7 After a routine power cycle, R30/Reb1 raised a *Bias2 OD2 overcurrent* fault on S12, measuring $ODI = 0.213\text{ A}$ against an 0.082 A limit; S12 reached its bias voltage in two steps and then tripped off on both attempts. The behaviour was assessed in the moment as consistent with a real short. All CCDs were powered off camera-wide and further CCD power-on was suspended pending expert review, in part because the fault appeared after that day's reorientation of the camera to zenith. R30/Reb1 was left powered off.

Apr 7–8 Repeated CCS jGroups multicast storms disrupted trending and control, apparently triggered by starting CCS processes on the spare shutter subsystem whose network port lacked IGMP snooping; the summit Cycle 41 software and OS upgrade fell on the same day.

Apr 8–11 Intermittent PCS chiller 2 trips with no corresponding chiller anomaly were traced to an unseated pin in a refrigeration-permit connector on the telescope-side cabling, aggravated by TMA shaking tests; the pins and shell were replaced (SUMMIT-10642, SUMMIT-10652).

Apr 11 The L1 lens cover and lens cap were removed — a step towards first photon — and a shutter and CCD soak test was authorised with the dome lights administratively controlled.

Apr 13–25 The filter carousel developed an intermittent communication fault on socket 3 that depended on rotation speed and telescope orientation. It was diagnosed as a sticking spring-loaded carbon collector contact rather than a wiring problem, and the collector was replaced on Apr 24–25 (SUMMIT-10680), with the ageing of the other eight collectors noted as a residual concern.

Apr 14–16 The first full focal-plane CCD power-on campaign was run raft by raft, followed by the first HV bias ramp across the REBs, which took about 1.5 hours to reach full voltage. R30/Reb1 again tripped on an OD2 overcurrent (0.087 A) and was left off with R20/Reb2.

Over the same nights the camera took part in TMA soak testing (BLOCK-T447, BLOCK-T449), giving tracking with shutter readouts for the first time, and in the first bright-star tests, whose largest observed CCD current jump was benign.

Apr 16–19 Guider idle-clearing was found to inject noise into neighbouring sensors during bias readout, making early bias images unusable; manual guider sleep/wake was adopted as an interim workaround while a proper non-guider clearing scheme was developed (LSSTCCSRAFTS-778).

Apr 17–29 A series of HV bias trips was traced to stray light reaching the focal plane: head lamps and light past the filter-loader door during FES work, sunlight through an open P-flow door with the mirror covers open, and a laser tracker left running. New procedures were proposed tying HV state to P-flow and mirror-cover status.

Apr 24 R20/Reb2 was returned to service by configuring the erroneous reset-drain reading to be ignored in software rather than repairing the board, and HV was restored to the full camera. The same day a REB-firmware register-size bug was fixed, allowing three previously off corner CCDs to be powered and validated.

Apr 24–25 A renewed attempt to recover R30/Reb1 was abandoned. A shorts test showed an unexplained digital-current spike of order 1 A on the REB, seen in neither repeated retests nor 24 hours of monitoring of any other REB, and it was decided not to attempt powering it on; R30/Reb1 ended the month as the only REB switched off camera-wide.

Apr 26–28 A new sequencer file was tested on sky (BLOCK-T469, BLOCK-T471) to address e2v noise and covariance anomalies; e2v channels improved clearly while ITL channels changed little, and the default sequencer was restored for scheduled active-optics measurements.

A.5 May — cold weather and the first loss of cryo capacity

May 1–7 Memory faults on the data-collection nodes forced image handling off `lsstcam-dc10` onto `lsstcam-dc02`, with a DIMM replacement on another node and a full reboot of the LSSTCam machines on May 7 with the camera taken offline first.

May 4 A filter-repeatability test using the collimated beam projector (BLOCK-T486) measured spot shifts of roughly 5 pixels (about $52\ \mu\text{m}$) before and after a filter swap, within the $100\ \mu\text{m}$ repeatability specification.

- May 6–20** Four in-camera filter exchanges were carried out at horizon pointing, taking the loaded complement through $[g, r, i, z, y]$ to $[g, r, i, z, u]$. Extraction of a single filter took slightly under an hour; swaps required the mirror cover closed against light leaks and the HV switches open.
- May 7** The cryostat pressure reached the 10^{-8} Torr regime for the first time in the reported monitoring.
- May 8–11** The camera lost essentially all cryo cooling capacity overnight (OBS-931, FRACAS-280). The cryo plate began deviating from -127°C in the early hours, all five running circuits appeared to lose capacity, and CCDs and REBs were shut down around 08:10–08:20 UTC as both cryostat and heat-exchanger vacuum climbed; the heat-exchanger gauge reached about 1×10^{-5} Torr against an ion-pump shutdown threshold of 1.1×10^{-5} Torr. Oil was found dripping from a liquid-line filter dryer, three dryers were swapped, and recovery took three days, with all REBs back on by May 10 and HV restored the same evening. Low ambient temperature, with a transition near 7°C , was identified as the likely driver: cryo trim heat went from about 60 W before the event to about 180 W after recovery.
- May 9–11** The focal plane was fully re-powered after the cooling event, with shorts tests on all sensors and REB power-on gated on power-supply board temperature. R30/Reb1 was deliberately skipped, its HV switches left open by direction; R21/Reb0, which had repeatedly glitched near 50 V, was set to run permanently at 46 V.
- May 12–27** A cold-weather operating regime was established. Quantitative “leading indicator” thresholds (Cryo 1 heat-exchanger return temperature below the dome dew point, trim-heat drift, dome ambient below 8°C , rapid warming of circuit exit temperatures) were defined and put on Chronograf dashboards for the observing specialists, later re-worked to use derivatives; the dome-warming procedure was rewritten to use HVAC only. Two further near-crashes (May 16–17 and May 25–27) were managed by closing the dome, running closed-dome calibrations and walking the telescope down to low elevation to keep the Cryo 1 heat-exchanger return above the dew point.
- May 19–23** A refrigerant-additive and insulation campaign recovered a large part of the lost capacity: R124 and R245fa were added to four circuits and 2919C to all five running circuits, after which cryo capacity was assessed as roughly 90–100 W higher, and the Level 8 cryo liquid supply lines were progressively insulated.
- May 21** A power glitch stuck the glycol remote control valve closed, dumping the glycol charge and taking down the PCS and three of four running cryo circuits (FRACAS-282). With the

CCS shell effectively unusable, REBs were powered off with a direct command; the cold plate rose to -16°C before flow was restored via the manual valve.

May 15–16 Two software faults cost observing time: a 0.1 s exposure request left the focal plane stuck integrating because the configured minimum exposure time was below the shutter’s real 1 s minimum (OBS-947), and a broadcast packet storm from three camera hosts faulted the camera (OBS-950).

May 28–30 A new CCS refrigeration release added interlocks and alarms drawn from the month’s failures (blocking coolant flow below a configurable glycol flow, warning on high discharge and low return pressure, fixing reversed Cryo 3 water connections), and approval was given to replace the failed Cryo 6 module with the Cryo 5 unit from the MRCR.

A.6 June — First Look, and completion of the science array

Jun 1–4 Throttling the Cryo 1 liquid-line isolation valve rebalanced liquid and vapour flow and, for the first time, allowed Cryo 1 to run continuously overnight at a workable heat-exchanger return temperature.

Jun 2–3 The shutter jammed part-open in the middle of a flat calibration and could not be recovered by the standard procedure; the blade appeared to overrun its high limit until both PLCs and the HCU were power-cycled and the shutter recalibrated (OBS-1013).

Jun 3–4 Dynalene flow was lost overnight, forcing an emergency power-down of CCDs, REBs and most electronics and a switch to the second Dynalene chiller (OBS-1014). During recovery HV was enabled with the REB back-bias switches still open, producing strong tree rings and large gain shifts in flats; this was corrected the next morning by zeroing HV and closing the switches raft by raft.

Jun 2–7 The Cryo 6 compressor module was replaced with the MRCR unit, and the circuit was evacuated, charged, started and gas-chromatographed, restoring six of six cryo circuits.

Jun 4–11 Recurring “timed out waiting for shutter exposure to complete” failures were traced to PTP clocking rather than to the shutter mechanism. Enabling a boundary clock on the shutter’s network switch reproduced the fault on the clean-room system; transparent-clock mode and then the leaf switch as boundary clock eliminated the offset from master. The investigation is documented in CTN-003.

- Jun 6–7, Jun 16–18, Jun 27–30** Three severe cold-driven cryo degradation events dominated the second half of the month. With dome air down to 3.6 °C, then near 5 °C, then near freezing, cryo-plate trim heat was driven to zero, and each event required powering off the CCDs, the HV and — in the second and third events — nearly all REBs, with recovery taking from seven hours to three days. Moving the telescope from high elevation to horizon was found to produce a step recovery in cooling capacity.
- Jun 7** A multicast storm on the camera network, attributed to enabling IGMP snooping on one switch, forced most CCS services to be stopped and restarted and left the camera in fault.
- Jun 10–11** A screw dropped into the camera body through the FES access during Auto Changer PLC work (SUMMIT-10747) triggered a stop-work and cancelled that night’s observing; the camera was wrapped, scaffolding erected and side panels removed, and the screw was recovered the following morning.
- Jun 12–17** Cold-sensitivity mitigation was attempted with heater blankets and insulation on the refrigeration lines — a thermocouple measurement initially showed the existing insulation had “zero effect” because the lines were bare up- and downstream — and refrigerant was removed from Cryo 1 to raise its heat-exchanger return temperature clear of the dew point. Of four dome air handlers only one was healthy, so dome heaters were ruled out as a mitigation.
- Jun 14–15** The compressed dry air supply to the Top End Assembly was found to have been silently off for an unknown period, and was restored; the TEA air valve was shown to have little authority over humidity in the camera region, and the camera team stated that dry air to the camera was necessary to survive the winter.
- Jun 23** Rubin Observatory’s public First Look event took place. Camera shifts were flexed the previous night around it, and a telescope slew with the mirror covers open and dome floodlights on during the associated demonstration tripped HV on one REB.
- Jun 23–30** CCS 41.0.4 was activated in a staged full-summit restart, delivering idle-clear support and the software needed to turn on the remaining R30 CCDs. Idle clear was tested on sky but initially kept disabled because early measurements suggested it raised the read noise; after two locking bugs were fixed — skipped invocations piling up and a sequencer-reload race — it was declared available for operational use on Jun 30 and nightly operations moved to the 3 s readout sequencer with idle clear enabled.

Jun 25–26 The last two focal-plane sensors, on R30/Reb1, were powered on and biased, completing the science array. The turn-on required an R30/Reb1 firmware update, an NTP fix, and configuration changes that hit a hard-coded firmware parameter range, forcing an emergency change to the bias-control limits and a same-day rebuild of the rafts subsystem. During the subsequent telemetry review the CCD current on S12/Seg17 was noted as anomalously low (0.017), and the channel was again described as the only non-nominal one when R30/Reb1 was checked on Jun 29.

Jun 28–30 During re-power-up after the third cold event, R30/Reb2 showed its clock-high current ramping continuously for about an hour, prompting a near stop-work. The cause was found to be idle clear interacting with the REB because the idle-clear pause had timed out during the shorts test; a power cycle with idle clear correctly paused resolved it, and disabling idle clear during REB, CCD and HV turn-on became standing practice.

A.7 July — operational maturity of idle clear and the FES

Jul 1–4 Idle clear was validated through a series of diagnostic bias and dark acquisitions and adopted as the default configuration for all on-sky observations, replacing continuous all-day bias acquisition.

Jul 1, Jul 15 Filter swaps continued: u out and the pinhole filter in on Jul 1 (followed by an approved pinhole illumination test with HV off), and y out and u in on Jul 15 (SUMMIT-10923).

Jul 4 A power outage briefly interrupted both glycol and Dynalene flow, and glycol chiller 3 showed a rising-temperature anomaly overnight (OBS-1074), prompting discussion of resilience to future outages.

Jul 5–11 A formal shutter and camera-to-OCS-control training programme for observing specialists and commissioning scientists was launched, together with an overhaul of the CCS command-level policy that granted a trained-observer group level-1 access to the filter changer and shutter subsystems.

Jul 6–13 Recurring carousel clamp faults, which had been costing significant observing time, were addressed with an automated unclamping-recovery patch; by Jul 13 recovery was succeeding essentially every time, sharply reducing the need for expert intervention. On Jul 20–21 the FES passed 1000 filter changes since April with the error rate below 1 in 400.

- Jul 11–13** A dangling guider-initialisation command left the guider stuck in a paused state for as long as 22 hours undetected, halting guider file production; the root cause was traced to double-commanding between CCS and OCS.
- Jul 13–28** Repeated ocs-bridge failures to release the mount and rotator motion locks after filter changes cost observing time on several nights (OBS-1098, OBS-1079, OBS-1119), tentatively linked to a SAL/Kafka interface version mismatch and worked around by manual restarts.
- Jul 15–16** Candidate REB FPGA firmware was loaded into the inactive EEPROM slot on all 71 xREBs in subsets — wavefront and guider REBs, R30/Reb1, and the remainder — in preparation for a later activation during a planned power cycle.
- Jul 23–29** A serious clamp fault developed at carousel socket 3, with repeated position-sensor errors and failed locking recovery suspected to be an autochanger alignment problem. The socket was taken out of service and the *g*-band filter was removed from the camera on Jul 29 for bench inspection.
- Jul 16–31** Cold-weather mitigation continued with heater tape and insulation on the cryo liquid lines, refrigerant top-ups on Cryo 1, 2 and 6, and the adoption of quantitative dome open/close temperature criteria derived from analysis of the four earlier cryo crashes.

A.8 August — winter storms, power failures and firmware

- Aug 1–3** A winter storm closed the mountain road with outside temperatures down to -5°C ; most of the crew descended before the closure and monitored remotely while Cryo 1 and Cryo 2 lost some capacity.
- Aug 5–12** Filter operations were disrupted twice: the $u \rightarrow y$ swap was rescheduled after a clamp timeout during a new interface-cleaning step and completed on Aug 7 (SUMMIT-11029), and a defective batch of filter pins was found to trip FCS limit errors and was replaced before the *g* filter was reinstalled (SUMMIT-11092, LCA-13232).
- Aug 4–12** A repetitive clicking picked up by the shutter microphone, described as sounding like a Geiger counter, was analysed for the first time. It ran from 23:06:13 to 23:07:18 UTC on Aug 4 at 2–6 kHz while the telescope was slewing; M2-hexapod telemetry showed the hexapod holding position with its actuator currents oscillating, which made it the first suspect. On Aug 12 the flapping straps of a temporary safety barrier on the mesh platform, blown by the HVAC fans, were proposed instead and taped down.

Aug 11–16 Power glitches during summit UPS and generator inspection tripped the PCS chiller twice on a drive fault (FRACAS-292) and forced an emergency camera shutdown. The outage was used to boot the xREB EEPROMs into the candidate firmware slot, which preserves HV bias across REB resets. Recovery was then complicated by days of jGroups retransmit storms that broke CCS-DAQ communication whenever an image was taken with the guider running; by Aug 15 the trigger was traced to guiding at an untested 5 s exposure cadence combined with high DAQ node traffic, and reverting to the normal cadence stopped the storms.

Aug 15 The cryo system was deliberately driven towards its crash threshold, with the dome air handler stopped and the glycol chiller lowered, to validate a lower glycol setpoint. The system held — camera air temperature reached about 5.7°C with trim heat above 90 W — and the test was curtailed only because of a condensation risk.

Aug 18 A cold night pushed ambient temperature more than two degrees below the phase-separator setpoint, causing refrigerant condensation in the vapour lines of all six cryo circuits and warming the cryostat plate; the glycol chiller was lowered further and a portable 60 kW heater was deployed as mitigation.

Aug 21–23 CCDs had been running measurably warmer since the Aug 11 firmware change. Power-cycling the whole focal plane and rolling the firmware back confirmed a genuine power-draw regression of roughly 30 W, tentatively attributed to a new sequencer-override register; the apparent loss of cryo capacity during the storm had in part been this self-heating. R30/Reb1 was left running its own firmware version, required for its bias-threshold adjustment.

Aug 23–27 A new automated HV-bias ramp-up and regulation subsystem ran through a full night on sky for the first time and was then enabled camera-wide, with the procedure trained to a new operator. An anomalously high HV current on R03/Reb1 was managed by temporarily lowering its setpoint (LSSTCAM-124, LSSTCAM-125) and resolved by Aug 26.

Aug 23–30 Both Level-1 air compressors degraded in the cold — the dryers froze at -8°C , blocking airflow — and were fitted with insulating jackets. Loss of the dry-air purge was a direct camera concern, both for the vacuum gate valves and for frosting on the L3 lens.

Aug 26–27 Large-scale summit OS and Kubernetes patching (SUMMIT-11198) forced a reboot cascade across the camera servers. A scripting error included the shutter HCU, which had been meant to be excluded, and an authentication race condition then cost access

to it for several hours; the incident was judged serious because the same fault on a hardware-safety host would have been far worse.

Aug 30–31 The summit generator tripped, cutting glycol flow and leaving the camera on UPS alone. REBs and their power supplies were deliberately powered off while cryostat pressure rose towards the ion-pump shutdown limit; recovery of the refrigeration, shutter, FES and DAQ took the rest of the day, and the full camera power-on was completed on Aug 31 after overnight snow.

A.9 September — recovery, then preparation for the shutdown

Sep 1–5 The camera recovered from the generator failure, whose root cause was ice and snow blocking the generator air inlet (FRACAS-299); the summit server-room UPS had separately exhausted its roughly 30-minute battery and dropped camera network communications (FRACAS-300). Recurring generator voltage glitches caused repeated minute-scale glycol interruptions — close to the observed 90 to 120 s PCS trip threshold — and observing on generator power was restricted in pointing. Commercial power was confirmed restored on Sep 5.

Sep 3–9 Cryo 1 again showed cold-ambient degradation on several nights as outside temperatures approached the phase-separator setpoint. It was decided that the circuit need not be switched off, and Cryo 1 was removed from the dome-closure decision criteria.

Sep 6 An HV trip occurred with the moon only about 20 degrees away, and after recovery the guider was left in an inconsistent state between CCS and the DAQ; the case prompted discussion of a proper fault path for HV bias trips at the control-system level.

Sep 7 The same sound returned over the shutter microphone during 20-degree camera rotations every three minutes, which disposed of the safety-barrier explanation of August: this time the M2 hexapod was stationary while the camera hexapod moved with every rotation, though M2-hexapod motor currents still spiked at each one. The recirculating ball bearings of the hexapod roller-screw actuators, or of the rotator, were proposed as the source and a ticket was opened (OBS-1158).

Sep 14–18 A cluster of DAQ and control-system faults occurred: a data-transport module on one crate needed rebooting for the second time in a month (LSSTCAM-130), leaving the DAQ data store inconsistent and image acquisition failing; a storage-node switch failed; and a duplicated guider-initialisation left the system in an error state.

- Sep 17–25** The refurbished vacuum pump cart was commissioned on the camera elevation platform, and both the cryostat and heat-exchanger turbo pumps were run for their semi-annual maintenance cycle ahead of the October cryo work. Entanglement of the cart hoses during platform motion was flagged as a safety hazard.
- Sep 20–21** Cryo-line heater-tape circuits repeatedly blew fuses after termination work — five to six failures across several attempts, including with slow-blow replacements — and the cause was not resolved within the month.
- Sep 21–22** Dome lights were switched on briefly by accident, tripping roughly 20 REB HV bias channels, all on ITL rafts; recovery followed the documented HV turn-on procedure without incident.
- Sep 23–25** The r and g filters were extracted ahead of the October carousel maintenance (SUMMIT-11343), leaving i , z and the pinhole filter in the camera, and setting a record of 2 h 10 min for a two-filter operation.
- Sep 9–30** The camera state for the coming shutdown was defined and executed — REBs on for thermal monitoring, CCDs off, HV off — while the CCS database underwent a raw-data partitioning migration and CCS Cycle 41 was rolled out to all summit nodes.

A.10 October — the maintenance shutdown and the loss of Cryo 3

- Oct 1** Starting a new turbo-test-station subsystem broadcast a “refrigeration not permitted” message that the refrigeration subsystem accepted as real, shutting down all six cryo cabinets. Three circuits then refused to restart on low discharge temperature and needed a heat gun on the lines before they would run.
- Oct 1–20** A roughly three-week planned camera downtime ran in parallel with facility work: the FES team serviced and reinstalled the carousel clamps on sockets 1–4 and fixed an autochanger overcurrent problem, cryo filter dryers were replaced on three circuits, CCS was upgraded to cycle 42, the MCM, DAQ-management and diagnostic-cluster machines were replaced, and the HCU operating systems were patched.
- Oct 3, Oct 9** The camera cable wrap interlock faulted and blocked rotator motion during FES maintenance rotations, recovered by bypassing the interlock at the mount engineering interface and re-aligning the wrap to the rotator; the same problem recurred on Oct 9 and cost the FES team over an hour.

- Oct 3** A major fibre-optic failure on the AURA base to Pachón link, ongoing since Sep 30, left base-to-summit communications fluctuating between the primary fibre and a microwave backup, with no estimated repair time; camera systems were not reported as directly affected.
- Oct 14–15** The +X deployable platform jammed about 42 cm short of full extension after a hose-drape test, the extension encoder having mis-reported position because of motor slip under excess load. Both platforms were taken out of service and returned the next day under new hard load limits with no personnel aboard during motion.
- Oct 15–20** Cryo 3 developed a capillary or expansion-valve clog during recharge that could not be cleared by power cycling, refrigerant addition or running the telescope horizontal. The circuit was given up on Oct 20 and the camera continued on five of six circuits, with the cryo trim setpoint raised to recover about 10 W of margin; procedure revisions on isolation-valve handling were planned in response (LCA-14625, LCA-17726).
- Oct 13–16** Candidate REB, WREB and GREB firmware was loaded to the spare EEPROM slot and booted across the focal plane, including the special R30/Reb1 variant. The rollout repeatedly exposed a REB serial-number read race that left REBs in an invalid validation state, and a sequencer opcode-error bug when REBs were power-cycled while idle clear was active, mitigated by delaying the start of idle clear until the subsystem reached its operational state.
- Oct 20–21** The camera returned to on-sky operation. The first night surfaced a stuck shutter requiring reinitialisation (OBS-1013), incorrect guider bias voltages in the new guider firmware, which was rolled back, a recurrence of a long-standing wavefront-REB power-on issue (LSSTCAM-37), and corrupted FITS headers introduced by a recent software change — the last fixed and verified end to end the same day.
- Oct 22–23** The final science-filter complement was loaded via the FES after re-qualification at horizon and zenith, and a new rule was deployed requiring the FES to be in a loaded state before it can be handed to observatory control. With cryo capacity reduced by the loss of Cryo 3, the CCD and raft-support-assembly setpoints were raised by about one degree for thermal margin.
- Oct 28–30** Two Dynalene cooling-loss events forced protective shutdowns. In the first, REB board temperatures approached the HV-trip limit and REBs were powered off; on power-up they reported a zero serial number and stayed invalid, needing repeated RCE resets

and clearing of a DAQ opcode-error register. The second, caused by a crash of the Dynalene control computer, required a full camera shutdown and a two-to-three hour recovery. A review of the earlier Aug 11 event had established that the camera tolerates only about 13 minutes without Dynalene flow before board temperatures approach the trip limit.

A.11 November — control-system instabilities and the Cryo 6 trip

Nov 1–3 A subsystem network fault pegged the REB-power process at full CPU and filled the management host’s filesystem, and on Nov 3 a major jGroups meltdown during a dark sequence forced nearly every subsystem to be killed and restarted. In the aftermath the shutter’s PTP synchronisation fell badly out of lock and required repeated PLC power cycles and a full recalibration in a closed dome (LSSTCAM-144; recovery procedure in CTN-003).

Nov 4 A planned swap of the Dynalene control computer required a roughly three-hour full camera power cycle, with REBs powered off in groups to protect the power-supply boards while flow was cut. Several REBs returned in an invalid state and a corner REB tripped because its heater had not been switched off first.

Nov 5 A power glitch with no UPS backing tripped the PCS chiller and the Cryo 6 compressor and briefly dropped Dynalene flow; systems recovered within about ten minutes, and a second smaller event later the same day again disturbed the chillers.

Nov 6 A dynamic guider region-of-interest test crashed the focal-plane subsystem: an out-of-range ROI produced a parameter exception and a core dump rather than a rejected command, and recurred on a second attempt (LSSTCCSRAFTS-821).

Nov 10–21 Weekly planning surfaced thermal problems in the Top End Assembly — infrared imaging showed hot spots near the shutter electronics box and cavity humidity above ambient, attributed to a dry-air leak near the heat-exchanger return — and the FES carousel developed an intermittent following-error and overcurrent fault on the first move of the night, addressed incrementally with a warm-up brake sequence and automated recovery.

Nov 17–18 A scheduled summit OS and Kubernetes upgrade took the control system offline for a morning, after which the camera team rebooted essentially every LSSTCam server and verified the DAQ link speeds before returning the camera to observatory control.

Nov 18–19 Hours after the Top End Assembly air nozzle was replaced, Cryo 6 was lost: the circuit's PLC was reported dead, the discharge temperature spiked to 122°C and the oil level read negative. The PLC recovered by itself after about half an hour. The root cause was traced to a software update elsewhere on the summit mis-setting the chiller setpoint in response to bad environmental temperature readings, and the control logic that closes the coolant valve on a dead PLC was flagged for revision (FRACAS-322, OBS-1394, OBS-1395).

Nov 20 A generator-to-commercial power switchover caused a total loss of Dynalene, and the automatic REB shutdown response failed because the board-temperature limit had been left set too high. The team shut the camera down manually and carried out a multi-hour recovery including RCE reboots, invalid REB states and a shutter PTP resynchronisation (FRACAS-321).

Nov 21–24 The focal-plane subsystem repeatedly hung in an active command state after guider initialisation on several nights, once timing out after 45 minutes and each time requiring a subsystem restart (LSSTCAM-148); a parallel memory-exhaustion event in the message broker triggered an overnight jGroups storm (LSSTCAM-149, LSSTCAM-150). On Nov 24 the DAQ link on one data-collection node was found to be dropping packets, with register-read retries up by one to two orders of magnitude since Nov 20, and the link was bounced and the switch power-cycled (LSSTCAM-151).

Nov 24–25 A newly trained loader operator completed a filter swap using the simplified loader procedure without camera-team intervention for the first time, following weeks of FES software fixes.

Nov 24 The “natural frequency” project, retuning the REB sequencer timing to 6.4 ns, neared completion and was revalidated on the test stand, with dedicated dark-dome engineering time requested ahead of a deployment decision.

A.12 December — refrigeration margin and firmware campaigns

Dec 3–11 With Cryo 3 out of service and failing to restart, a capillary clog developed on Cryo 4 and drove cryo-plate trim heat down to about 1 W whenever the telescope moved to high elevation. The raft-support and cryo-plate setpoints were raised in steps as mitigation; the Cryo 4 clog cleared spontaneously on the evening of Dec 10, trim heat recovered to tens of watts, and the setpoints were restored. An average cryostat temperature of −120°C was documented as the hard safety boundary against sublimation.

- Dec 8–9, Dec 13** Two FES faults required expert intervention: repeated PLC faults on an on-line clamp overnight on Dec 7–8, and autochanger truck controllers going to fault with undervoltage errors after the carousel was found out of standby, reset remotely (OBS-1454).
- Dec 9, Dec 14–15** The 6.4 ns natural-frequency REB firmware was tested on the summit in two phases, booted from the spare EEPROM slot. Both attempts left several REBs hung or in an invalid state after RCE resets, and the planned bias and flat sequences (BLOCK-T659, BLOCK-T661) were only partly completed before the production firmware was restored.
- Dec 15–16** Water droplets and streaks were found on M1M3, and a stone under the telescope near the camera. Inspection with a cherry-picker and an infrared camera traced the water to condensation on uninsulated Dynalene supply pipes at the Top End Assembly, where the dew point exceeded the heat-exchanger return temperature (SUMMIT-11795, OBS-1479).
- Dec 16–17** Two Dynalene chiller outages, the first coinciding with a summit-wide OS patching reboot, caused REBs and CCDs to be powered off as an emergency response — which worked as designed — and put the camera into fault; a few REBs needed a focal-plane restart to clear data corruption (FRACAS-328).
- Dec 22** A changeover between Dynalene chillers warmed the supply to the camera and drove REB power-supply board temperatures towards their safety limit, forcing HV bias and then the CCDs off for about an hour until cooling was restored by switching back (OBS-1490).
- Dec 21–25** A natural-frequency firmware test on the auxiliary telescope’s wavefront electronics failed, producing anomalous bias levels at both clock rates that no reset sequence would fix. The cause was an inverted-polarity reset signal continuously holding the AS-PIC in reset; the fix was applied and retested successfully on Dec 24, although a new stripe-pattern noise appeared and required further study (SUMMIT-11822).
- Dec 27–29** The science REBs were switched again to the 6.4 ns firmware for an extended holiday calibration campaign, completing a partial photon-transfer curve on the night of Dec 27–28 before reverting to production firmware; opcode errors recurred on three REBs after RCE resets even with idle clear confirmed off, pointing at a firmware or boot-sequence cause (LSSTCAM-153).

Dec 12, Dec 18 Operations structure changed at year end: the camera weekly planning meeting moved from Monday to Thursday, and from Dec 28 the summit support scientists moved to a single daily shift, with active support from 14:00 to 22:00 local time and on-call support thereafter until sunrise.

A.12.0.1 Sources and method. The entries above were extracted from the complete 2025 message history of the Rubin Observatory Slack channels #cam-summit-ops and #cam-operator, including all thread replies (roughly 2400 and 11 700 top-level messages respectively, with about 5400 thread replies in #cam-summit-ops). Where the logs give a ticket identifier it is reproduced verbatim so that the underlying record can be retrieved; where the logs are ambiguous about a date the entry gives a range. Statements of cause reflect the assessment recorded in the channel at the time and are not independent confirmations.