

BALANCING STABILITY AND AMBITION: LHC AVAILABILITY IN RUN 3

L. Felsberger*, T. Argyropoulos, A. Asko, H. Dostmann, J. Heron, M. Hostettler, B. Mikulec, C. Roderick, J. Uythoven, D. Wollmann
European Organization for Nuclear Research, Geneva, Switzerland

Abstract

Run 3 represents the final operational phase of the LHC before the transition to the High-Luminosity (HL) LHC era. It covers the years 2022 to 2026 and ends with Long Shutdown 3, starting in July 2026. This period is characterized by the successful restart of beam operation following the maintenance and upgrades during Long Shutdown 2 (2019-2022), the establishment of a stable operation scheme enabling record-breaking integrated luminosities, and the exploration of performance limits in preparation for HL-LHC. This has all been achieved with infrastructure and equipment soon approaching two decades since first commissioning.

Run 3 availability was closely monitored using CERN's Accelerator Fault Tracking tool, which records fault source, duration, and cross-system impacts. The analysis presented here extracts lessons for HL-LHC from these fault statistics. Machine unavailability is dominated by long-duration faults arising from latent weaknesses, often introduced by recent upgrades and exposed when performance is pushed to its limits. While individual systems require targeted mitigation, there is little evidence of accelerator-wide aging. Radiation-induced faults show a strong impact, demanding specific attention for HL operation. These insights support future strategies for efficient machine exploitation.

INTRODUCTION

Run 3 (2022–2026) acts as a critical bridge from early LHC operation to the high-luminosity HL-LHC era, characterized by unprecedented performance. Pushing the existing machine to its limits has increased stress on equipment, making operational monitoring paramount. Using Accelerator Fault Tracking (AFT) data [1] from 2022 to 2025, this paper analyzes the interplay between performance gains and accelerator availability. Primary downtime contributors and their root causes are identified, the impact of performance demands on availability is assessed, and key lessons for future HL-LHC operation are extracted.

Accelerator Fault Tracking

Fault tracking quantifies short-to-medium-term issues to prioritize mitigation via standardized metrics. The Operations (OP) team records faults semi-automatically in the AFT tool, followed by independent reviews by equipment experts and a Reliability and Availability Working Group.

This framework relies on three core definitions. A *fault* is any event blocking nominal operation as defined in the accelerator schedule. *Root-cause statistics* attribute downtime to the initiating system; for example, a cooling failure

triggered by a power outage is assigned to the power distribution system. *Availability* is the fraction of fault-free time within the scheduled operational window. Given the slow-cycling nature of the LHC, only a portion of its up-time is spent in “Stable Beams” to produce physics data. Fault tracking spans from annual beam commissioning to the Year-End Technical Stops (YETS). While such scheduled stops are excluded from fault tracking, this analysis incorporates significant events outside formal tracking periods.

RUN 3 AVAILABILITY

Figure 1 displays LHC unavailability (2022–2025), excluding injectors, categorized by fault duration. Availability remains stable across the years when excluding two events with major impact (>10 days) in 2022 and 2023. The following sections review the primary drivers of these statistics.




Figure 1: Unavailability 2022-2025 by duration of fault.

Long Shutdown 2 Recovery

Long Shutdown 2 (LS2), lasting over 3 years, ended when beam commissioning commenced in April 2022, significantly later than initially planned. During an LS, superconducting magnets undergo a thermal cycle and need a **training campaign** to reach nominal performance via controlled training quenches. During the 2021 campaigns, a magnet coil in Sector 78 and a cold diode in Sector 23 short-circuited. These incidents necessitated additional thermal cycles and re-training, resulting in a three-month delay of the first beams. Consequently, the target beam energy was reduced from 7 TeV to the previous operating beam energy of 6.8 TeV to mitigate the risk of further failures [2].

An injection beam test in October 2021 then revealed an aperture limitation within a vacuum **plug-in module** (PIM). This was traced to a buckled RF finger, flexible connectors ensuring electrical continuity across vacuum interconnections. The defect had bypassed routine inspections, which were subsequently refined [3]. Resolving the issue required a thermal cycle of Sector 23, which concluded in April 2022.

“Black swans” On the Rise? - 2022 and 2023

As magnet training neared completion in March 2022, a loss of cryogenic conditions triggered the premature opening

* lukas.felsberger@cern.ch




Figure 2: Unavailability (left) and failure rate per uptime (right) statistics by contributing system captured during all beam-operation periods from 2022 to 2025. The ten most contributing systems are sorted by average downtime contribution; 2023 Magnet circuits and 2022 Radio Frequency unavailability cropped at 5%. LBDS: LHC Beam Dump System.

of a **burst disk** in an radio frequency (RF) cryostat. The resulting warm-up, cool-down, and subsequent conditioning of RF equipment further delayed the first beam by several days [4]. A similar failure occurred in August 2022 during the physics run, causing a four-week delay. Both incidents involved burst disks installed during LS2. Although these units were replaced during YETS 22/23, a third event occurred in April 2023 due to a production non-conformity. Improved operational procedures limited downtime to under two days. While residual risk remains due to narrow pressure margins, no further burst disk failures occurred.

In July 2023, the most significant Run 3 downtime occurred when an insulation vacuum leak developed in the **inner triplet assembly** left of Insertion Point (IP) 8. A dormant mechanical non-conformity in the magnet interconnection was identified as root-cause, which was revealed by a quench induced by a power grid glitch. To preserve the physics run, an in-situ repair was performed, though the event still resulted in 50 days of unavailability. The identified non-conformity will be resolved in LS3 and no additional events occurred during the remainder of Run 3.

Run 3 Unavailability - Statistical Analysis

Following the analysis of the longest events, Fig. 2 illustrates the unavailability and failure rate statistics by system for all beam-operation periods from 2022 to 2025. The injector complex is marked in the top plots with 5% unavailability on average. The ten most contributing LHC systems are ranked by their average unavailability across Run 3, with magnet circuits and RF systems leading due to the major events in 2023 and 2022, respectively. **Magnet circuits** show decreasing failure rates, attributed to the stabilization of flat-top training quenches following the LS2 training campaign. These are highlighted in the left panel of Fig. 3 together with the unavailability caused by beam losses due to Unidentified Falling Object (UFOs), which also decreased as operation

progressed. This behavior is consistent with trends observed after previous long shutdowns [5, 6].

RF systems experienced increasing failure rates while maintaining stable unavailability, excluding the burst disk events. **Cryogenics** unavailability and failure rates peaked in 2024 and 2025, driven by limited operational margins due to increased heat load and Radiation-to-Electronics (R2E) issues due to higher luminosity. Quench Protection Systems (QPS) reduced unavailability in 2025 through targeted R2E mitigation. **Power converter** metrics remained stable throughout the run. **Experiments** primarily impacted operation via a delayed LHCb upgrade in 2022. The **electrical network** infrastructure saw elevated downtime in 2025, partially due to leaking water-cooled cables and aging 18 kV equipment from the LEP era. External network perturbations, counted separately, peaked at 0.8% unavailability in 2023 and 2024. The 2025 peak of **cooling and ventilation** unavailability stems from a busbar water leak.

While latent weaknesses have been identified as causes for downtime, the following sections study how hardware aging and performance optimization impact system availability.




Figure 3: Left: Flat-top training quenches and UFOs caused more downtime at the beginning of runs. Right: Heat load, RF finger and R2E events led to increased downtime as performance ramped-up, but were mitigated in 2025.

Signs of Aging? Despite much of the equipment approaching two decades since first commissioning, stable aggregate failure rates (top right in Fig. 2) do not indicate accelerator-

wide aging effects. However, system-level fluctuations reflect proactive management of emerging hardware issues.

IP1/5 triplet magnet radiation lifetime concerns were addressed by deploying reverse polarity optics during the 2023/24 and 2024/25 YETS. This strategy redistributed the integrated dose, ensuring the magnets remain operational until their replacement during LS3.

Several cases continue to require targeted mitigation. Six vacuum leaks occurred in **Edge-Welded Bellows (EWB)** of collimation systems during Run 3. Investigations revealed reduced fatigue lifetime under lateral movement due to inadequate specifications [7]. Despite mitigation via blocking of collimators in parking positions and reduced strokes, several collimators installed since LS1 have reached the expected end of life for their bellows and require replacement [8].

Recurring quenches and Higher-Order Mode coupler overheating in **RF module M1B2** necessitated a voltage reduction. Following its scheduled replacement in LS3, this module will be refurbished and is foreseen as the accelerating system’s sole spare for the HL era [9].

These examples illustrate the management of assets by equipment teams. It kept the impact of aging on Run 3 availability limited and remains crucial for HL-LHC success.




Figure 4: Evolution of CMS integrated luminosity per week during proton Run 2 and Run 3 (> 1200 bunches) and fraction of scheduled time spent in Stable Beams.

Balancing Stability and Ambition

Run 3 serves as a bridge towards the high-luminosity configuration. A balance was struck between reliable luminosity production and probing performance boundaries in view of HL-LHC, causing some downtime (Fig. 3, right).

Heat Load pushed cryogenics equipment beyond operational margins in 2024, leading to several short stops for system recovery. Optimized filling schemes and proactive management limited the total unavailability impact to 0.6%.

Beam-Induced Heating damage to RF fingers caused notable downtime, peaking at 1.9% in 2023 due to a damaged module in sector 4L1. The four-day repair and subsequent performance loss were traced to production non-conformities; while within original baseline specifications, small irregularities enhanced impedance and power dissipation at higher proton bunch intensities. Mitigation included replacing critical components with improved deformable RF bridge designs and streamlining operational procedures,

such as operating with longer bunch lengths. Consequently, similar events in 2024–2025, such as in sector 6L2, had significantly reduced impact on availability.

Radiation to Electronics (R2E) impact peaked at 4% unavailability in 2024, particularly affecting QPS. Increased luminosity in LHCb further induced events around IP8 affecting the cryogenics systems. Operational experience confirms that machine parameter modifications—for both proton and ion runs—strongly influence R2E performance [10]. While hardware teams deployed timely mitigations, coordinated strategies between machine exploitation and hardware upgrades remain essential for HL-LHC operation.

Figure 4 illustrates the weekly integrated luminosity for CMS and the Stable Beam fraction from 2015 to 2025. Excluding 2022 and 2023—years heavily impacted by RF burst disk and inner triplet vacuum leak events—weekly integrated luminosity increased from approximately 3 fb⁻¹ in Run 2 to over 6 fb⁻¹ in 2025. Throughout this period, Stable Beam fractions remained constant or improved, demonstrating that availability can be maintained even as performance exceeds LHC design parameters [11]. High-intensity tests scheduled for 2026 aim to identify any last limitations before LS3.

Table 1: Summary of Notable Events Impacting Run 3 Availability

Event(s)	Root Cause	Run 3 Impact	Year(s)
Inner Triplet Vacuum Leak	Production NC & Rand. HW failure	> Month	23
RF Burst Disks	Production NC & Rand. HW failure	> Weeks	22, 23
R2E effects	Beam induced	> Week	22–26
LHCb upgrades	Project delay	> Days	22
Edge Welded Bellows	Spec. NC & Wear-out	>Days/Perform.	22–26
Vac. RF Fingers	Production NC & Beam induced	>Days/Perform.	23, 25
Heat Load	Beam induced	>Day/Perform.	24
SC Mag. Training	Rand. HW failure	> Months*(TC)	21
PIM RF Finger	NC & Procedural	> Months*(TC)	21

* events outside fault tracking periods
TC: thermal cycle
NC: non-conformity

SUMMARY AND CONCLUSIONS

LHC fault data systematically collected throughout Run 3 indicates that a small number of notable events, summarized in Table 1, account for a significant fraction of total unavailability. Root-cause analysis suggests these incidents reveal latent weaknesses, relate to LS2 activities, or stem from the performance push of Run 3. Evidence of accelerator-wide aging remains limited, indicating successful handling of emerging challenges by equipment teams. Effective management of heat load, beam-induced heating, and R2E effects has enabled substantial performance gains—nearly doubling weekly integrated luminosity compared to Run 2—while maintaining high availability, establishing a robust operational baseline for the HL-LHC era.

These findings underline the utility of systematic fault tracking. By moving beyond anecdotal observations, it enables objective prioritization based on quantitative data.

REFERENCES

- [1] C. Roderick *et al.*, “Accelerator fault tracking at CERN”, in *Proc. ICALEPCS'17*, Barcelona, Spain, pp. 397–400, Oct. 2017. doi:10.18429/JACoW-ICALEPCS2017-TUPHA013
- [2] A. Apollonio *et al.*, “Summary of the Post-Long Shutdown 2 LHC Hardware Commissioning Campaign”, in *Proc. IPAC'22*, Bangkok, Thailand, pp. 335–338, Jul. 2022. doi:10.18429/JACoW-IPAC2022-MOPOPT040
- [3] G. Bregliozzi, Radiography and Tomography Results in 21L3, Presentation at the LHC Machine Committee (LMC #428), CERN, 2021, <https://indico.cern.ch/event/1096812/#2-radiography-and-tomography-r>,
- [4] A. Butterworth, Impact on LHC RF and Recovery Planning, Presentation at the LHC Machine Committee (LMC #436), CERN, 2022, <https://indico.cern.ch/event/1142994/#3-impact-on-lhc-rf-and-recover>,
- [5] D. Mirarchi *et al.*, “Special Losses during LHC Run 2”, in *Proceedings of the 9th LHC Operations Evian Workshop*, Evian-les-Bains, France, pp. 213–220, Jan. 2019. <https://cds.cern.ch/record/2750296>
- [6] M. Lamont, “LHC Performance in Run 2 and Beyond”, in *Proceedings of the 27th International Symposium on Lepton Photon Interactions at High Energies*, pp. 1–8, Aug. 2016. doi:10.22323/1.245.0001
- [7] J. Guardia Valenzuela, Final Report of Edge-Welded Bellows Task Force, Presentation at the LHC Machine Committee (LMC #500), CERN, 2025, <https://indico.cern.ch/event/1487244/#4-final-report-of-edge-welded>,
- [8] D. Mirarchi, Report from the Task Force for mitigation of collimator end-of-life in the HL era, Presentation at the LHC Machine Committee (LMC #526), CERN, 2026, <https://indico.cern.ch/event/1679425/#3-report-from-the-task-force-f>,
- [9] K. Turaj, Update on RF Module M1B2, Presentation at the LHC Machine Committee (LMC #521), CERN, 2026, <https://indico.cern.ch/event/1651225/#3-update-on-rf-module-m1b2>,
- [10] G. Lerner, Report from the Monitoring and Calculation Working Group on the 2025 LHC Run and 2026 Prospects, Presentation at the LHC Machine Committee (LMC #521), CERN, 2026, <https://indico.cern.ch/event/1651225/#4-report-from-the-monitoring-a>,
- [11] R. Tomas *et al.*, “Towards a High Luminosity LHC with even higher performance”, in *Proc. IPAC'25*, Taipei, Taiwan, pp. 278–281, Nov. 2025. doi:10.18429/JACoW-IPAC2025-MOPM008