THE FIRST FULLY AUTOMATED ANALYSIS OF THE HARDWARE COMMISSIONING TESTS FOR SUPERCONDUCTING CIRCUITS AT THE LHC

A. Chmielińska*, A. Barba, M. Cerqueira Bastos, S. Mazenoux, A. Mnich, A. Murariu,
M. Medeiros Pereira, A. Verweij, CERN, Geneva, Switzerland

Abstract

The superconducting magnet circuits of the Large Hadron Collider (LHC) at CERN undergo a commissioning campaign at regular intervals, referred to as Hardware Commissioning (HWC), to validate magnet powering and protection functions. During each campaign, more than 1,500 circuits must be commissioned, requiring the execution and analysis of several thousand powering tests. Historically, this process relied on legacy or manual tools, resulting in high workload and sometimes inconsistent results.

To streamline the validation process, a Python-based automated analysis framework has been developed within the Signal Monitoring (SigMon) project over the past two years. The 2026 HWC campaign marked the first time this automated analysis was applied to all powering tests across the superconducting magnet circuits of the LHC. This paper describes the SigMon analysis framework, reports on the 2026 HWC campaign, evaluates the impact of automation, and presents key circuit performance findings.

INTRODUCTION

The LHC superconducting magnet circuits operate under cryogenic conditions and store significant energy, which can reach 1.3 GJ in the case of the Main Dipole circuits. Their reliable operation — including power converters, superconducting magnets and all connected components (current leads, busbars, diodes, instrumentation), as well as associated protection systems such as quench heaters, Quench Detection System (QDS), Energy Extraction (EE) and related interlocks — is critical. Undetected faults may lead to equipment damage and accelerator downtime.

The LHC comprises approximately 8,000 superconducting magnets along its circumference of 26.7 km, divided into eight sectors and independently powered sub-sectors [1]. The superconducting circuits are grouped into eight families: low-current circuits (60 A, 80–120 A, and 600 A) and high-current circuits — Individually Powered Dipoles (IPDs), Individually Powered Quadrupoles (IPQs), Inner Triplets (ITs), Main Quadrupoles and Main Dipoles — with nominal currents for the latter ranging from 4 kA to 11.6 kA. Before each LHC restart, after the annual Year-End Technical Stop (YETS) or after a Long Shutdown (LS), all circuits must be qualified for safe operation. This is performed during the HWC campaign. The number of circuits commissioned during the 2026 HWC is shown in Table 1.

Table 1: Superconducting Circuits Commissioned during the 2026 HWC Campaign and the Maximum Test Current

Family Name	**Number of Circuits**	**Max. Test Current (A)**
Main Dipole	8	11 600
Main Quadrupole	16	11 300
IT	8	6850
IPQ	76	2900
IPD	16	6050
600 A	389	590
80-120 A	279	105
60 A	749	55
Total	**1541**	—

During powering tests, the circuit current is increased in steps to its nominal value in strict sequential order — a test may only commence once all preceding tests have been validated. Figure 1 shows the post-LS test program, consisting of Power Converter Configuration (PCC) and Powering Interlock Controller (PIC) tests at minimum operating current, followed by intermediate current (PLI1–3) and nominal current tests (PNO). Suffixes denote test variations, e.g. a – current cycling, b – EE discharge triggered by the QPS, d – power converter failure, e – Slow Power Abort (SPA) via PIC. Following a YETS, only a subset of these tests is required. The total number of tests ranges from approximately 5,000 (YETS) to 9,000 (LS). Tests are scheduled and tracked via the Accelerator Test Tracking framework (AccTesting) [2] and executed by the HWC Sequencer [3]. Test analysis must be signed off by experts from different domains, namely MP3 (LHC Magnet Circuits, Powering and Performance Panel), MI (Machine Interlocks), and PO (Power Converters), with responsibilities shown in Fig. 1.

Historically, powering test analysis relied on a fragmented ecosystem of tools [4]: MP3 used legacy automated tools (LabVIEW and eDSL [5]) or manual analyses executed via the Service for Web-based Analysis (SWAN), while other teams maintained their independent tool-chains. Validation required signatures across two platforms — AccTesting and the Post Mortem Event Analyser (PMEA) [6]. This fragmentation led to inconsistent acceptance criteria and reduced traceability of the validation process. Legacy tools were difficult to maintain and evolve, while manual analyses and multi-party sign-offs imposed a significant workload on experts and prolonged campaign durations. These limitations motivated the automation effort described in this paper.

---

* agnieszka.chmielinska@cern.ch

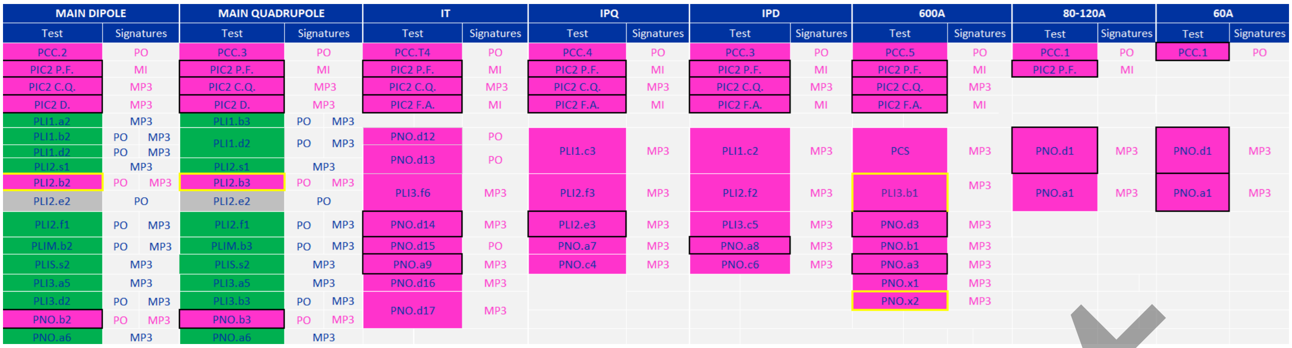

MAIN DIPOLE		MAIN QUADRUPOLE		IT		IPQ		IPD		600A		80-120A		60A	
Test	Signatures	Test	Signatures	Test	Signatures	Test	Signatures	Test	Signatures	Test	Signatures	Test	Signatures	Test	Signatures
PCC.2	PO	PCC.3	PO	PCC.T4	PO	PCC.4	PO	PCC.3	PO	PCC.5	PO	PCC.1	PO	PCC.1	PO
PIC2 P.F.	MI	PIC2 P.F.	MI	PIC2 P.F.	MI	PIC2 P.F.	MI	PIC2 P.F.	MI	PIC2 P.F.	MI	PIC2 P.F.	MI		
PIC2 C.Q.	MP3	PIC2 C.Q.	MP3	PIC2 C.Q.	MP3	PIC2 C.Q.	MP3	PIC2 C.Q.	MP3	PIC2 C.Q.	MP3				
PIC2 D.	MP3	PIC2 D.	MP3	PIC2 F.A.	MI	PIC2 F.A.	MI	PIC2 F.A.	MI	PIC2 F.A.	MI				
PLI1.a2	MP3	PLI1.b3	PO MP3												
PLI1.b2	PO MP3	PLI1.d2	PO MP3	PNO.d12	PO	PLI1.c3	MP3	PLI1.c2	MP3	PCS	MP3	PNO.d1	MP3	PNO.d1	MP3
PLI1.d2	PO MP3			PNO.d13	PO										
PLI2.s1	MP3	PLI2.s1	MP3												
PLI2.b2	PO MP3	PLI2.b3	PO MP3	PLI3.f6	MP3	PLI2.f3	MP3	PLI2.f2	MP3	PLI3.b1	MP3	PNO.a1	MP3	PNO.a1	MP3
PLI2.e2	PO	PLI2.e2	PO												
PLI2.f1	PO MP3	PLI2.f1	PO MP3	PNO.d14	MP3	PLI2.e3	MP3	PLI3.c5	MP3	PNO.d3	MP3				
PLIM.b2	PO MP3	PLIM.b3	PO MP3	PNO.d15	PO	PNO.a7	MP3	PNO.a8	MP3	PNO.b1	MP3				
PLIS.s2	MP3	PLIS.s2	MP3	PNO.a9	MP3	PNO.c4	MP3	PNO.c6	MP3	PNO.a3	MP3				
PLI3.a5	MP3	PLI3.a5	MP3	PNO.d16	MP3					PNO.x1	MP3				
PLI3.d2	PO MP3	PLI3.b3	PO MP3	PNO.d17	MP3					PNO.x2	MP3				
PNO.b2	PO MP3	PNO.b3	PO MP3												
PNO.a6	MP3	PNO.a6	MP3												

Figure 1: Overview of powering tests per circuit family and corresponding signatures as of February 2026. Black borders indicate mandatory tests performed after YETS, while the remaining tests are additionally required after an LS. Yellow borders denote on-request tests performed during the 2026 HWC. Magenta color indicates tests that are automated within SigMon; green indicates tests for which a manual analysis notebook exists for the MP3; and grey indicates tests analyzed exclusively by power converter experts. For PIC2 tests, the acronyms P.F., C.Q., D., F.A. denote Powering Failure, Circuit Quench via QPS, Discharge via PIC, Fast Abort req. via PIC, respectively.

SIGMON AUTOMATED ANALYSIS

The SigMon framework consolidates the analysis of powering tests into a unified Python environment [7]. Each analysis is structured in three steps: *query*, *analysis*, and *verdict*. In the *query* step, relevant signals are retrieved from the Next CERN Accelerator Logging Service (NXCALS) [8] or Post Mortem (PM) system [9]. In the *analysis* step, checks are applied to data in accordance with the official test procedures, including timestamp checks (data availability), analog checks (current levels, plateau duration, and comparison with the previously accepted reference waveform), and digital checks (transitions of boolean signals). The *verdict* step combines all individual check results into a single overall outcome: pass or fail.

Over the past two years, 60 analysis classes were developed, as shown in Fig. 1. Rigorous code quality is ensured through unit tests that validate individual check functions and integration tests that verify end-to-end signal retrieval on real LHC data. The correctness of each analysis was validated against results from previous HWC campaigns using a dedicated *validation notebook*, enabling systematic comparison with expert decisions and legacy tool outputs.

Each analysis has an associated Jupyter notebook that displays analysis results and logs alongside interactive plots, following the same order as defined in the analysis class and official procedures. Since the last campaign, experts no longer need to manually run notebooks in SWAN to investigate failed tests. Automatically executed analyses now generate interactive HTML reports using the Plotly library.

During HWC, automated analysis execution and report generation are handled by a dedicated SigMon server. SigMon has been fully integrated with AccTesting, supporting storage of analysis results, log messages, and direct links to interactive reports — providing a single environment for analysis, expert review and sign-off.

THE 2026 HWC CAMPAIGN

The 2026 campaign took place from 4 to 20 February. All mandatory powering tests (indicated by black borders in Fig. 1) were performed, except the Bipolar Cyle (PNO.a3) test on 600 A circuits, as the Bipolar Powering Failure (PNO.d3) test was sufficient for their validation. On-request tests (highlighted in yellow in Fig. 1) were also conducted, namely energy extraction on Main Dipoles (PLI2.b2) and Main Quadrupoles (PLI2.b3), as well as energy extraction (PLI3.b1) and combined decapole/octupole powering (PNO.x2) on selected 600 A circuits. Two so-called snapshot tests, outside the standard test scope — PLI2.r2 at 2 kA and PNO.r2 at 11 kA, following a fast power abort at 10 A/s on the Main Dipoles — were conducted to support simulation studies. Finally, the PGC.2 test — combined powering of all circuits to the current level needed for operation of the LHC at 6.8 TeV beam energy, followed by a fast power abort — was performed in each arc.

In total, 5,095 standard and on-request tests were planned, along with 8 PNO.r2, 8 PLI2.r2 and 8 PGC.2 tests. All standard and on-request tests were analysed in SigMon, achieving 100 % automation coverage for the first time in the LHC's history. The key milestone was the extension to high-current circuits, building on low-current automation introduced in the 2025 HWC campaign. The distribution of planned and executed tests (including re-executions) is shown in Fig. 2.

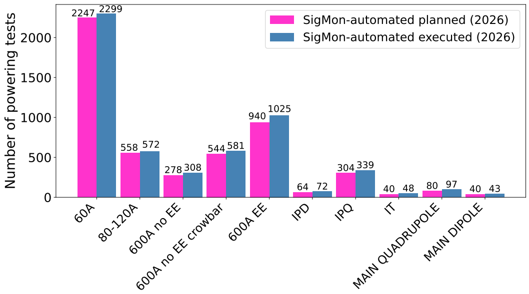

Figure 2: Planned and executed powering tests, including re-executions, analyzed in SigMon during the 2026 HWC.

ANALYSIS PERFORMANCE AND IMPACT

The total number of powering tests (including re-executions) reached 5,409, with 5,384 analyzed in SigMon (see blue bars in Fig. 2) and 25 analyzed manually (PLI2.r2, PNO.r2 and PGC.2). There were 289 re-executions of SigMon-automated tests. In 106 cases (37 %), the sequence succeeded, the automated analysis returned a fail verdict and the expert signed the test as failed, confirming a genuine issue during the test. In 96 cases (33 %), both the sequence and analysis failed, indicating execution-level problems. The remaining 87 cases (30 %) were repeated for other reasons.

False Positives (FPs) — cases where no execution issue is present, the automated analysis returns a fail verdict, but the test is later validated as successful by experts — are a key indicator of analysis quality. Table 2 shows the number of FPs together with the number of successful cases for tests that were automatically analyzed with eDSL and LabVIEW in 2024 and with SigMon in 2026. Thanks to the optimization of the acceptance criteria and transition to SigMon, the FP rate went down from 10 % to 1.5 %, substantially reducing unnecessary expert intervention.

Table 2: Number of false positive results among tests analyzed in eDSL and LabVIEW (2024) and in SigMon (2026), compared to the number of successful tests. The PIC2 C.Q. test on IPQ circuits was already analyzed in SigMon in 2024 and is not included in the table.

	2024		**2026**		
Test Name	**Success**	**FP**	**Success**	**FP**	**FP Diff**
PIC2 P.F.	813	240	797	0	-240
PIC2 F.A.	524	95	491	0	-95
PIC2 C.Q.	413	74	525	28	-46
PLI3.b1	188	14	187	10	-4
PNO.d1	1051	25	1029	22	-3
PCC.1	763	2	749	2	0
PNO.a1	748	0	750	8	8
Total	**4500**	**450**	**4528**	**70**	**-380**

Another major improvement in 2026 was the consolidation of signatures into a single signature per test for most SigMon-automated tests. This was achieved by merging cross-domain checks into a single analysis, resulting in fully coherent acceptance criteria and a significant reduction in manual review time. Only four tests retained two signatures to avoid last-minute updates to AccTesting; these were consolidated after the campaign, completing the transition to one signature across all SigMon-automated tests.

In total, 5,392 SigMon-automated signatures were granted in 2026. As shown in Table 3, between 2024 and 2026, the number of signatures manually entered by equipment experts decreased substantially, demonstrating a significant reduction in workload. Analysis time is now no longer a limiting factor during HWC, the duration of which is instead governed by cryogenics availability, test duration, and sequence execution quality.

Table 3: Number of Manual Signatures in 2024 and 2026

Signatures	**2024**	**2026**	**Reduction [%]**
MP3	1159	259	77.7
PO	568	94	83.5
MI	993	28	97.2

CIRCUIT PERFORMANCE

The automated analysis correctly identified various anomalies, all of which were subsequently investigated and either resolved or formally acknowledged. Representative examples of detected issues include:

- Current instability induced by electromagnetic interference, leading to converter trips in several PNO.d3 tests on 600 A circuits.
- Incomplete negative current cycles, identified in approximately 20 PCC.1 tests on 60 A circuits, leading to the replacement of several power converters.
- Timing mismatches between QPS signals and the PIC in several PIC2 tests on 600 A circuits.
- Signals marginally exceeding prescribed limits, e.g. elevated earth leakage currents.
- Inconsistencies with reference voltage waveforms, including a case where a previously approved reference was found to be incorrect.
- Wrong QDS settings for the voltage ramp on current leads in PNO.b2/b3 and PLI2.b2/b3 tests.
- Inconsistencies with reference current decay, including one case where circuit operational parameters were modified to prevent degradation.

CONCLUSIONS

The 2026 hardware commissioning campaign marked a milestone in the automation of powering test analysis, achieving 100% coverage for all planned tests, excluding special tests and PGC.2 tests. Acceptance criteria were reviewed, optimized, and fully validated. The SigMon analysis was integrated with AccTesting, enabling a unified environment for analysis, validation, and expert sign-off. Operational experience confirmed the high analysis quality, significantly reducing the number of false positives and enhancing coherency through consolidation of signatures. Efficiency gains were substantial, due to a significant reduction in expert analysis workload across MP3 and collaborating teams, as well as a reduced overall campaign duration. All circuits were successfully validated for LHC operation in 2026. The developed analysis framework remains valid for the vast majority of circuits to be commissioned after the major upgrade of the LHC taking place from 2026-2030. Future developments will focus on completing automation for the remaining tests and modernizing additional analysis modules.

REFERENCES

[1] O. S. Brüning *et al.*, “LHC Design Report”, CERN, Rep. CERN-2004-003-V-1, 2004. `https://cds.cern.ch/record/782076`

[2] D. Anderson *et al.*, “Testing Quality and Metrics for the LHC Magnet Powering System Throughout Past and Future Commissioning”, in *Proc. IPAC’14*, Dresden, Germany, Jun. 2014, pp. 3995–3997. doi:10.18429/JACoW-IPAC2014-THPRI095

[3] V. Baggiolini, R. Alemany-Fernandez, R. Gorbonosov, D. Khasbulatov, and M. Lamont, “A Sequencer for the LHC Era”, in *Proc. ICALEPCS’09*, Kobe, Japan, Oct. 2009, paper THC003, pp. 670–672. https://cds.cern.ch/record/1215886/files/CERN-ATS-2009-114.pdf

[4] A. Chmielińska, “Hardware Commissioning 2025 of the superconducting circuits with focus on automated test analysis and circuit performance”, presented at the 220th TE-TM Meeting, CERN, Geneva, Swizterland, Sep. 2025, unpublished. https://indico.cern.ch/event/1580034/#44-hardware-commissioning-2025

[5] M. Audrain *et al.*, “Using a Java Embedded Domain-Specific Language for LHC Test Analysis”, in *Proc. ICALEPCS’13*, San Francisco, CA, USA, Oct. 2013, paper THPPC079, pp. 1254–1257. https://proceedings.jacow.org/ICALEPCS2013/papers/thppc079.pdf

[6] H. Reymond *et al.*, “The Post-Mortem Analysis Software Used for the Electrical Circuit Commissioning of the LHC”, in *Proc. EPAC’08*, Genoa, Italy, Jun. 2008, paper THPC151, pp. 3345–3347. https://proceedings.jacow.org/e08/papers/thpc151.pdf

[7] A. Chmielińska, A. Barba, S. Mazenoux, A. Mnich, A. Murariu, and M. Pereira, LHC SigMon Analysis, 2025. https://gitlab.cern.ch/LHCData/lhc-sm-analysis

[8] J. Wozniak and C. Roderick, “NXCALS – Architecture and Challenges of the Next CERN Accelerator Logging Service”, in *Proc. ICALEPCS’19*, Brooklyn, New York City, NY, USA, Oct. 2019, pp. 1465–1469. doi:10.18429/JACoW-ICALEPCS2019-WEPHA163

[9] JF. Barth *et al.*, “A Modernized Architecture for the Post Mortem System at CERN”, in *Proc. IPAC’22*, Bangkok, Thailand, Jun. 2022, pp. 1557–1560. doi:10.18429/JACoW-IPAC2022-TUPOMS055