A LOOK BACK ON 2012 LHC AVAILABILITY

B. Todd*, A. Apollonio, A. Macpherson, J. Uythoven, D. Wollmann, CERN, Geneva, Switzerland.

Abstract

This paper presents the conclusions of studies into the availability of LHC during 2012.

Initial and root causes of beam aborts above 450GeV have been identified and downtime associated with each event has been determined.

A comparison of operation's indications from e-logbook entries versus equipment group experience has been made for those systems having the largest indicated downtime.

Topics for further study are then introduced, including a weighting of systems by complexity, and a comparison of the observed versus predicted failure rates of the principle elements of the Machine Protection System (**MPS**).

Conclusions are drawn and proposals are made, including new approaches to information capture in the post-LS1 era to improve the quality of availability calculations.

POST-MORTEM - DUMP CAUSE

Every beam abort leads to the creation of a post-mortem event and corresponding post-mortem database entry. This contains raw information pertaining to the dump event. It is completed and a root cause is identified, by experts following investigations. These determine whether it is safe to continue LHC operations and make the post-mortem database one of the most reliable sources of information concerning LHC operation.

Post-Mortem Dump Cause Evolution 2010-2012

Considering only beam aborts that took place above injection energy, between March and November, then classifying dump cause into five categories (external, beam, equipment, operations or experiment) leads to the following distribution of beam aborts for 2010 [1]:

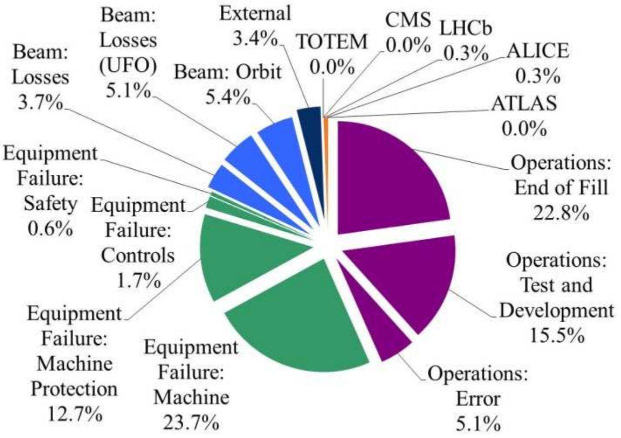

Figure 1: Distribution of Beam Aborts in 2010 (total 355)

The same analysis for 2012 [3]:

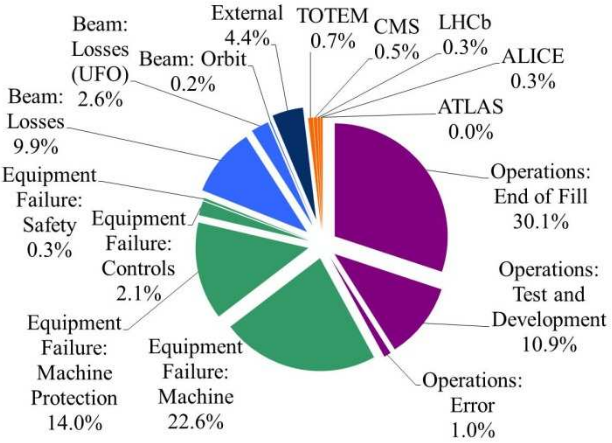

Figure 2: Distribution of Beam Aborts in 2012 (total 585)

Significant changes from 2010 to 2012 are:

- An increase in the number of fills from 355 to 585.
- An increase in the ratio of fills reaching the normal "end of fill" from 22.8% to 30.1%.
- An increase in the ratio of beam aborts due to "beam losses" from 3.7% to 9.9%.
- A slight increase in the ratio of beam aborts due to failures of the machine protection system from 12.7% to 14.0%.

Details of Dump Causes

Considering 2012 operation, excluding "Operations: End of Fill" and "Operations: Test and Development", leads to the following table of dump causes with their occurrences [4]:

Table 1: Root Cause versus Occurrence for 2012

Dump Cause	#	Dump Cause	#
Beam: Losses	58	BPM	8
Quench Protection	56	Operations: Error	6
Power Converter	35	SIS	4
Electrical Supply	26	LBDS	4
RF + Damper	23	TOTEM	4
Feedback	19	CMS	3
BLM	18	BCM	2
Vacuum	17	Water	2
Beam: Losses (UFO)	15	Access System	2
Cryogenics	14	LHCb	2
Collimation	12	ALICE	2
Controls	12	Beam: Orbit	1

DOWNTIME

The occurrence rate alone is not sufficient to determine the complete impact of each cause on the availability of

---

*benjamin.todd@cern.ch

the LHC. Downtime is an interpretation considering impact-on-physics. It is formed from two parts: lost-physics and fault-time.

Lost-Physics

Lost-physics (t_{lp}) indicates the length of stable-beams time lost due to a fault occurring. The calculation for this metric is to considers

1. The length of time the machine has been in stable-beams at the time the fault occurs (t_{stable}).
2. The average duration of stable-beams for physics fills in 2012, **nine hours**.
3. A pessimistic turnaround time (the time taken to get from beam abort of one fill to stable-beams of the next) of the LHC of **three hours**.

Every occurrence of a dump cause is assigned up to **three hours** of lost-physics if t_{stable} is less than **nine hours**. For example, if a water fault occurs after seven hours of stable beams, it is assigned two hours lost-physics. Table 2 indicates lost-physics due to each root cause [5]:

Table 2: Lost-Physics by Root Cause for 2012

Root Cause	**t_{lp} [h]**	**Root Cause**	**t_{lp} [h]**
Beam: Losses	147	Collimation	27
Quench Protection	126	BPM	17
Power Converter	80	Operations: Error	12
RF + Damper	53	LBDS	12
Vacuum	51	SIS	9
Electrical Supply	47	TOTEM	6
Feedback	45	BCM	6
BLM	43	Water	6
Beam: Losses (UFO)	42	LHCb	6
Cryogenics	33	ALICE	6
Controls	33	CMS	4

This gives a total of **812 hours** lost-physics for the LHC in 2012, due to failures.

Fault-Time

Once a beam abort has occurred, corrective action may be needed to restore operational conditions to the LHC. These actions are recorded by the operations team as faults in the e-logbook. Faults are not exclusive to the system which first caused the beam abort; they can run in parallel and can have interdependencies. Moreover, a fault can occur independently of LHC operation, not causing a beam abort, but delaying the next fill.

The time a system is unable to carry out its function, halting LHC operation (t_f) is shown in Table 3 [6].

The total fault time for the LHC in 2012 is **1524 hours**.

Downtime = Lost-Physics + Fault-Time

For each of the causes identified, t_{lp} and t_f can be combined to give downtime (t_d). Table 4 is a breakdown of downtime associated with each cause for operation in 2012, having a grand total of **2336 hours** downtime, equivalent to around 98 days.

Table 3: Fault-Time by Fault for 2012

Fault	**t_f [h]**	**Fault**	**t_f [h]**
Cryogenics	358	PSB	37
SPS	155	BLM	37
RF + Damper	119	Cooling & Ventilation	31
Quench Protection	112	Controls	26
Power Converter	106	ATLAS	17
Injection	86	ALICE	17
PS	82	Feedback	6
Vacuum	75	LHCb	4
Electrical Supply	70	CMS	4
Collimation	48	BPM	4
LBDS	44	BCT	2
BSRT	41	BIS	2
Access System	39	SIS	1

Table 4: Downtime by Cause for 2012

Cause	**t_f [h]**	**Cause**	**t_f [h]**
Cryogenics	391	BSRT	41
Quench Protection	238	Access System	39
Power Converter	186	PSB	37
RF + Damper	172	Cooling &Ventilation	31
SPS	155	ALICE	23
Beam: Losses	147	BPM	21
Vacuum	126	ATLAS	17
Electrical Supply	117	Operations: Error	13
Injection	86	SIS	10
PS	82	LHCb	10
BLM	80	CMS	8
Collimation	75	TOTEM	6
Controls	59	BCM	6
LBDS	56	Water	6
Feedback	51	BCT	2
Beam: Losses (UFO)	42	BIS	2

In this case, impact-on-physics is expressed as the number of hours of physics lost due to each cause. A more relevant metric would be the impact on integrated luminosity due to each cause.

SYSTEM BY SYSTEM

There are several shortcomings in the methods outlined for calculating downtime. The principle source of error is the use of the e-logbook for indication of fault occurrence and duration. E-logbook entries are not systematically completed, are not retrospectively corrected, do not account for multiplicity of faults and do not indicate dependency between faults.

Equipment-level fault information concerning t_{lp}, t_f and t_d for the three largest contributors to downtime was considered against operations' viewpoint. Faults were split into three categories. External – faults being outside of the control of the system in question. Internal – faults due to the system in question. Radiation Induced – faults due to the system in question with radiation induced effects as a root cause.

Power Converters

The post-mortem database indicated **35 beam aborts** due to power converters, with operations indicating **59 faults** totalling **106 hours**. Table 5 shows equipment group records for these events, with the cause, total and average fault times given [7].

Table 5: Power Converter Operational Experience

Cause	#	Total [h]	Average [h]
External	2	2.5	1.3
Internal	38	64.8	1.7
Radiation Induced	12	25.2	2.1
combined	**52**	**92.4**	**1.8**

Equipment group experience, excluding external events, gives **50** faults for a downtime of **89.9 hours**.

Remote reset capabilities are increasingly being employed to cure faults without needing access to the machine. In addition, some faults take longer to repair than predicted due to the time needed to get piquet teams in place. Power converter problems are often linked, or in the shadow of other faults, further effecting these figures.

Quench Protection

The post-mortem database indicated **56 beam aborts** due to quench protection, with operations indicating **57 faults** totalling **112 hours**. Table 6 shows equipment group records for these events [8]:

Table 6: Quench Protection Operational Experience

Cause	#	Total [h]	Average [h]
CMW / WorldFIP	3	2.7	0.9
DFB	6	17.3	2.9
EE (600A)	6	18.9	3.2
EE (13kA)	1	4.7	4.7
QPS Acquisition Failure	10	27.4	2.7
QPS Acquisition Radiation Induced	7	11.8	1.7
QPS Detector failure	7	12.0	1.7
QPS Detector Radiation Induced	15	14.2	0.9
combined	**55**	**109**	**2.3**

Equipment group experience, excluding external CMW, WorldFIP and DFB events, gives **46** faults for a downtime of **89 hours**.

Quench protection functions during 2012 were 100% successful, Table 6 is not exhaustive however, as some faults took place in parallel.

Cryogenics

The post-mortem database indicated **14 beam aborts** due to quench protection, with operations indicating **37 faults** totalling **358 hours**. Table 7 shows equipment group records for these events. This table includes all known faults for the cryogenic system, not limited to those indicated by operations [9].

Table 7: Cryogenic Operational Experience

Cause	#	Total [h]
Supply (CV / EL / IT)	17	19
User	28	25
Internal	46	233
Radiation Induced	4	57
combined	**95**	**334**

Equipment group experience, excluding external events, gives **95** faults for a downtime of **334 hours**.

These figures represent 14 days downtime over 263 days operation, giving around 95% availability. Cryogenic downtime has halved between 2010 and 2012, however, in the post-LS1 era, recovery from a “User” fault (e.g. Quench) is expected to take 10-12 hours.

FURTHER CONSIDERATIONS

Relative Complexity

A direct comparison of failure rates leads to poor results, as system complexity has a significant influence on reliability. A large, complex system made of many components will be expected to fail more often than a smaller system with fewer components.

An example of this diversity is shown in Table 8, which details the number of channels leading to an interlock for several LHC systems [10].

Table 8: Number of Interlocking Channels by System

System	Interlocking Channels
RF	800
BIS	2000
Cryogenics	3500
Quench Detection	14000
BLM (surveillance)	18000
BLM (protection)	48000

Predicted and Observed Rates

The MPS ensures that the LHC operates with an acceptable risk, failure rates of the key LHC systems were predicted in 2005 and were used to determine the residual risk related to the operation of the LHC.

Note that it is not feasible to study unsafe events to track the safety of the MPS, as they are predicted to occur too infrequently. Instead, the MPS reliability must be well understood and closely monitored, and the safety of the MPS inferred.

Table 9 outlines the observed number of failures of key elements of the MPS in 2010, 2011 and 2012 [11].

Table 9: Downtime by Cause for 2012

System	prediction 2005	observation 2010	observation 2011	observation 2012
LBDS	6.8 ± 3.6	9	11	4
BIS	0.5 ± 0.5	2	1	0
BLM	17.0 ± 4.0	0	4	15
PIC	1.5 ± 1.2	2	5	0
QPS	15.8 ± 3.9	24	48	56
SIS	not studied	4	2	4

Figures shown in orange exceed predictions. These observations have not been adjusted for radiation induced effects, which were outside of the scope of the original study. Considering this, these figures show a good correlation between the predicted and observed rates. From this it can be inferred that the safety of the MPS is close to the original predictions, if it is assumed that the ratio of safe failures to unsafe failures holds.

CONCLUSIONS

On the LHC Availability

When referring to availability in the LHC context the ultimate meaning is impact-on-physics. This paper has attempted to quantify availability with impact-on-physics by determining a so-called downtime, related to lost-physics and fault time. In this case, a system referred to as high-availability should be understood as having a low impact-on-physics.

As the operations team and equipment experts have increased their understanding of the LHC and its sub-systems, so the availability of the LHC has improved. Proof of this is the evolution in the ratio of fills reaching "end of fill" between 2010 and 2012. Equally, the machine has been seen to reaching the limits of beam operation more often, shown by the increase in ratio of beam losses as a dump cause. It can be concluded that equipment is operating well enough to allow operators to spend more time exploring the physical limits of the machine.

LHC equipment has been shown to have a somewhat stable influence on availability between 2010 and 2012. This gives an indication that systems appear to be in their normal operating life, having passed the wear-in stage, before reaching the end-of-life stages. This stability is despite the increase in radiation fields which lead to higher failure rates of exposed equipment. The effort groups have made to compensate for these increased fields have successfully mitigated the impact of radiation on global machine availability between 2010 and 2012.

Equipment group records do not match operator information in the LHC e-logbook. With significant effort it has been possible to consolidate these differing opinions to a certain extent. This was made more difficult by both the variation in tracking techniques between equipment groups, and the e-logbook used for principle fault tracking. The e-logbook is not completed in a systematic manner, provides little scope for relating faults and is not retrospectively corrected. It must be noted that the primary function of the e-logbook is not fault tracking.

Comments on Future Work

If information on availability is to be used to drive investment by the organisation, it is vital that an adequate fault tracking tool be developed and implemented for the LHC restart after LS1. This tool needs to be sufficiently detailed to capture information in an unobtrusive manner yet must fit with the research and development style of running a machine such as the LHC.

To implement such a tool the fundamental definitions used must first be consolidated, as several are open to subjective interpretation.

LS1 presents an opportunity to implement new methods to improve availability studies once LHC restarts. Three recommendations (**R1-3**) emerge from this work, with two suggestions (**S1-2**):

R1. A new LHC fault tracking tool and fault database is needed.

R2. Defined and agreed reference metrics are needed to consolidate views on definitions used in availability calculations.

R3. Reliability tracking of the critical elements of the MPS is needed to ensure that LHC machine protection integrity is acceptable.

S1. A means to convert the hours of physics lost to the impact on integrated luminosity should be investigated.

S2. A metric for weighting reliability by complexity should be introduced, giving a so-called per-unit reliability.

ACKNOWLEDGMENTS

This work is a combination of information from many different sources, which took a considerable effort to consolidate, in particular the authors wish to express their gratitude for the input from the Availability Working Group (**AWG**), S. Claudet, K. Dahlerup-Petersen, R. Denz, E. Duret, V. Montabonnet, I. Romera and Y. Thurel in the completion of this work.

REFERENCES

[1] PM database. Extracted from 23rd March to 6th December 2010, only for fills >450.1 GeV, ignoring entries marked as "no input change".

[2] PM database. Extracted from 17th February to 13th December 2011, only for fills >450.1 GeV, ignoring entries marked as "no input change".

[3] PM database. Extracted from 1st March to 6th December 2012, only for fills >450.1 GeV, ignoring entries marked as "no input change".

[4] Sort [3] by MPS Dump Cause, Discarding EOF, MD and MPS Test.

[5] Calculate stable beams per fill from TIMBER, assign lost physics by dump cause from [4]

[6] e-logbook. Extracted from 1st March – 6th December 2012. Supress duplicate entries, MISCELLANEOUS is ignored, except for the BSRT related events. No correction for entries rolling-over between shifts.

[7] Courtesy V. Montabonnet, Y. Thurel.

[8] Courtesy K. Dahlerup-Petersen, R. Denz, I. Romera.

[9] Courtesy S. Claudet, E. Duret.

[10] Courtesy O. Brunner (RF), S. Claudet (Cryogenics), R. Denz (Quench Detection), C. Zamantzas (both BLM).

[11] Table 1: R. Filippini et al., "Reliability Assessment of the LHC Machine Protection System" PAC 2005, Knoxville, Tennessee May 2005, TPAP011, http://accelconf.web.cern.ch/AccelConf/p05/PAPERS/TPAP011.PDF