July 1st 2008
Revision 1.1

The LHC Beam Loss Monitoring System

Report on the Audit held from June 10th to July 1st 2008

Auditors[1]: Miguel Anjo (IT/DM), Joachim Bächler (PH/DT), Philippe Farthouat (PH/ESE), Stefan Haas (PH/ESE), Stefan Lüders (IT/CO), Javier Serrano (AB/CO)

Distribution: Bernd Dehning (AB/BI), Roland Garoby (AB/BI), Eva Barbara Holzer (AB/BI), Steve Myers (AB), Hermann Schmickler (AB/CO), Rüdiger Schmidt (AB/CO), Jörg Wenninger (AB/OP)

1 Executive Summary

The LHC Beam Loss Monitor System (BLM) has been audited by a team of experts external to the BLM team. Generally, the auditors found that the design and implementation of the BLM electronics and the threshold management is sound, complete, straight-forward, and, in particular, conform to the requirement of high inherent level of safety, reliability and availability. However, quite a number of substantial recommendations have been made.

From the auditors' point-of-view, the initial determination of the threshold values is very critical. Despite the large number of sophisticated simulations, dedicated measurement runs and a subsequent and iterative adaptation of the threshold values during the early running of the LHC must be conducted. Neither the auditors nor this report endorse blindly relying on the BLM system at this stage.

Furthermore, although the auditors agree on the data driven approach, which manages all threshold values centrally in an Oracle database, concern has been the management of the threshold values themselves. The current procedures are found incomplete (partially since still under development), and adequate protections against tempering or human errors are missing.

It is also recommended to perform additional radiation tests in order to determine the total ionizing dose (TID), the non ionizing energy loss (NIEL) and the flux of particles able to produce single event effects (SEE).

Finally, the auditors would encourage the BLM team to reanalyze the FPGA code, once it has been finalized.

[1] With contributions from Benjamin Todd (AB/CO)

2 Contents

3 Scope

This audit is supposed to verify the design and implementation of the LHC Beam Loss Monitoring System (BLM). It should cover the fundamental design of the BLM ionization chambers, the front-end electronics in the LHC tunnel, the threshold comparator and combiner electronics as well as the means to determine, store and manage the threshold matrices in the corresponding BLM data base. In particular, it should include shower simulation and threshold determination, PCB schematics and layouts, FPGA programming, interfaces to other systems, mainly the Beam Interlock System (BIS), and the procedures for testing, “As-Good-As-New”-recovery and commissioning as well as proper documentation.

Particular focus should be put on the reliable detection of beam losses and the safe transmission to the BIS, without causing too many false beam dumps. The audit should reveal also single points of failures and failure modes leading to blind faults (i.e. non-detection of beam losses or failure of beam dump transmission). The audit will be restricted to single failures but will consider double-failures where they have been evident.

However, this audit does neither cover the in-depth verification of the Failure Mode, Effects and Criticality Analysis (FMECA) [1], the decision on where to place the ionization chambers around the LHC aperture nor the cross-correlation of signals from multiple ionization chambers and the subsequent consequences for dumping the beam. Finally, high-level control aspects and system software running on the PowerPC are covered only as far as needed for reviewing the management of the threshold databases and the “As-Good-As-New”-approach.

4 General Comments

The auditors are convinced that the fundamental implementation of the BLM (as defined by the scope of Chapter 3) is sound and properly executed. The system as such makes a mature and solid impression. The requirements for the very reliable detection of beam losses, but also for a high availability of the system itself have been adequately defined. The current implementation of the beam loss monitors, the front-end and back-end electronics, and the FPGA code fulfils to a very large part these requirements. A consequent application of "As-Good-As-New"-test during fills and the implementation of a redundant signal chain, where necessary, allows to obtain the required Safety Integrated Level 3 (SIL 3). In addition, the auditors are confident that the implementation is fail-safe, since no failure mode has been identified that does not lead to a beam abort.

In particular, the auditors agree on the deployed "data-driven"-approach, which handles all threshold values of the beam loss monitors offline in a central data base. The values are downloaded on demand to the BLM electronics, and — vice versa — regularly uploaded for integrity checks.

However, major concern has been the initial determination of the threshold values. The auditors strongly encourage a very pessimistic assumption of these values, and a subsequent and iterative adaptation during the early running of the LHC.

Since the final installation has not yet been fully commissioned, and the final threshold values have not yet fully defined, neither the auditors nor this report endorse blindly relying on the BLM system at this stage.

A second concern has been the management of the threshold values themselves. The current procedures are found incomplete (partially since still under development), and adequate protections against tempering or human errors are missing. This holds also for the management and configuration of the "maskable" and "disabled" flags of the monitors.

Furthermore, the auditors suggest conduction of further radiation tests, in particular on the power supplies used in the straight section.

On the other hand, the auditors are very positive on the current implementation of the front-end and back-end electronics as well as the FPGA programming, where only a few minor recommendations have been made.

As a final side remark, the review itself has been well prepared by the BLM team. A summary of the BLM team on their actions should be given in 6 months time.

5 Recommendations by the Auditors

Having reviewed carefully the basic design principles and functionalities, and having gone thoroughly through beam loss simulations, threshold determination and management, PCB schematics and FPGA code, and through available documentation, several areas for improvements have been identified.

This chapter lists all recommendations the auditors consider important enough to be mentioned. Quite some more comments have been directly made during the audit and in dedicated discussions with the BLM team. Other issues have been directly communicated to the corresponding experts.

Major points and issues are marked in bold.

5.1 Determination and Management of Thresholds

About 4000 beam loss monitors will finally be deployed. The signal of each monitor is compared to a 32×12 matrix of thresholds, since the impact of the beam loss depends on the LHC beam energy and the energy deposition per time interval. The full set of matrices will be stored in an Oracle database, while subsets will be stored in the corresponding trigger cards.

5.1.1 Simulations of Loss Signals and Response Functions

In order to determine the proper threshold values, very detailed simulation studies have been performed and compared to measurements. Several studies were done with the simulation packages GEANT4, FLUKA and Garfield. The performances of the ionization counters and Secondary Emission Monitors (SEM), and the associated electronics were studied for special cases. These simulation and measurements show the relation between the damage potential of the losses for LHC components and the measured signals in the monitors. This allows giving an estimation of the threshold settings.

1. However, it has not been clear to which extend all possible beam loss scenarios were fully computed. Therefore, the auditors would like to encourage the BLM team to summarize the results of the impressive simulation studies and measurements done so far. It might be beneficial to describe the different measurements and results using the same physical units, which are relevant for the BLMs under the specific conditions of LHC operation.

   It should be pointed out in which domains improvements are ongoing, like the activity by the BLM & FLUKA teams on the implementation and computation of the LHC setup in the collimation regions.

   Also, the limits of the computations should be discussed. This relates to the difficulty to get all information of correct geometry and material compositions of some LHC components.

2. The signal development as function of different loss scenarios was studied. In particular, the response functions of the monitors and the electronics response have been calculated. Saturation effects should be studied in more detail, and the limits of the current monitor design should be summarized.

3. As the monitors are exposed to high radiation, the activation/de-activation as function of LHC operation should be computed in order to determine the baseline shifts of the BLM signal.

4. A more detailed description of the materials and geometry of the monitors and a summary of the total material budget should be added.

5. Quite a series of sophisticated simulations on shower propagation due to beam losses have been made. However, the systematic errors on the results are quite large (up to 50%).

   Therefore, and due to the aforementioned points, the auditors share doubts that the monitors would guarantee a safe and efficient operation of the LHC without a re-adjustment of the thresholds.

   The initial threshold settings have to be sufficiently conservative in order not to damage the LHC magnets. During the initial runs of the LHC, they must then be iteratively adjusted.

6. The first operation of LHC will give the BLM team the possibility to compare the behavior of the monitors with the predictions. As the monitors have to cover an enormous range of radiation intensity, specific tests during LHC operation have to be performed.

   Dedicated test procedure should be proposed by the BLM team. Sufficient time should be assigned to make those tests.

   For example, tests with provoked beam losses should be conducted in order to verify the proper detection of those beam losses.

5.1.2 Management and Storage of Threshold Values

The individual 32×12 thresholds for the monitors are derived from a Master Table, which holds all basic threshold values derived from simulation, and from an individual factor ("Applied Factor" with a value range of]0…1[) defined globally for all 32×12 thresholds of a particular monitor. The result of this calculation is stored in the "Applied Table". In addition, the Master Table holds information on which monitors are "maskable" (e.g. can be ignored during Safe Beam runs), and which are disabled.

Both tables and the Applied Factors are stored in the Oracle LSA DB. The Master Table can be altered by the BLM experts, e.g. adjusting values due to the new experience while running the LHC. The Applied Factor, however, can be altered by dedicated Operators through the "Trim" application.

7. **Clear documentation must be produced for the procedures on how the initial values of the Master Table are defined, how the values in the Master Table can be altered, and how these changes are propagated to the individual tables.**

   This documentation should be approved by and made known to all parties involved (i.e. AB/BI and AB/CO/DM). Exercises with all parties involved should be conducted to verify these procedures.

8. Changes to the Master Table are cached in the so-called "Stage Tables" of LSA.

   An application should be deployed that provides means to minimize the introduction erroneous values to this table, e.g. through human errors.

   Technically, this might be done through sophisticate value and cross-correlation checks.

9. The aforementioned procedures should cover the conditions how and when the Stage Tables can be committed and merged into the final tables. Time-outs might be useful to avoid a too long latency between change of values and their commitment.
10. **An application should be deployed to safely handle the "maskable" and "disable" flags in the Master Table.**
11. Such application(s) should issue alerts if the safety for a particular region (i.e. a quadrupole) is undermined by too few "unmaskable" or too many "disabled" monitors. They should also allow for regular audits providing means to easily assess which monitors are "masked" or "disabled".
12. These applications should also allow for authentication and should register who changed which values when and why.
13. (This point has been removed after clarification with the experts.)
14. Human errors by operators/experts can be a cause of unavailability of the system. Currently the database is well configured to keep data changes up to 24h in the past. It is not protected, however, against the drop of database objects, which would require a recovery of the full database to a point in time. As the database is backed up only to the tape servers, this brings recovery time to several hours (depending on the volume of the database) and furthermore a dependency on the tape system (which is supported only during working hours). It is suggested to implement on-disk backups and to have a written procedure on how a restore of the data can be requested and how it will be done. This addresses both ACCCON and ACCMEAS databases. Due to very high volume of data and less criticality, it does not address the Logging database.
15. Currently, two additional databases are used for configuration management. The "MTF" database holds hardware descriptors and their history, while the "Layout DB" holds system descriptors (no history). This set-up adds complexity in terms of (common) data synchronization, update, and management.

    Means should be investigated for merging and combining those two databases.
16. Either a Service Level Agreement (SLA) or a Memorandum of Understanding (MoU) stating the responsibilities of the IT department in case of database failure is recommended. This document should include the procedure to contact the Service Manager (phone/email) and whom can contact (CCC operator, BLM expert, other), time to response and expected time to recover the database. The SLA/MoU should list elements which the database is dependent on and which can influence the recovery time (tape servers, network switches); how intervention agreements are settled; etc.
17. A similar SLA/MoU should be set-up with the AB/CO/DM section covering the support and maintenance of the BLM specific databases.

5.2 Monitors & Electronics

The two types of beam loss monitors, ionization chambers and Secondary Emission Monitors (SEM) are read-out via a dedicated read-out chain:

- The Current to Frequency Converter (CFC) electronics beneath the LHC arc quadrupole magnets or in the LHC service galleries is frequency-modulating the initial signals and transmitting them to

- The Threshold Comparator (TC) electronics, which determines if a monitor signal is above threshold;
- The signal combiner finally combines all signals and triggers the beam dump via the Beam Interlock System (BIS).

5.2.1 Ionisation Chambers

18. The HV stability as function of ionization currents should be studied. Voltage variations could be induced by a voltage drop over the RC. Furthermore the provoked pulse induction could lead to an error of the current measurement.

19. It has not been clear how the calibration of the SEMs will be done in the mixed radiation field of the LHC. Therefore, the BLM team is asked to produce a clear documentation including a discussion on the consequences.

5.2.2 PCBs and Choice of Components

20. Currently, the BLM holds about 5% of spares for the major PCBs.

    It is suggested to increase this stock to at least 10% including spares of the major components like GOL and FPGAs. Particularly critical or difficult-to-find components are all the ASICs from the PH/MIC group used on the CFC electronics (i.e. ADC, CRT), and the optical components (i.e. the GOH mezzanine card including the GOL, and the laser and PIN receiver on the TC).

21. An accelerated thermal aging test of one system might be conducted as well, in order to check that the computed lifetime is not completely wrong.

22. The lack of possibility for a full remote reset or power cycle of the front-end electronics is unfortunate, since e.g. having the FPGA entered into a locked state will be difficult to mitigate.

    Mitigations should be developed.

23. The Failure Mode, Effects and Criticality Analysis (FMECA) [1] indicated that among the most unreliable sub-systems in the BLM system are the low and high voltage power supplies in the straight sections with an MTBF of 525 000 hours (i.e. probability of failure of 1.9 10^{-6} h^{-1}).

    Since this has a high weighting on the overall system availability (the probability of a failure for one power supplies in an octant is about 10^{-3} $mission^{-1}$), and the power supplies in the straight sections are not redundant, it is recommended to perform an accelerated testing of a few power supply units in order to verify these values.

24. In addition, the sensitivity analysis of [1] has pointed out a large dependency of the number of false beam dumps on the failure rate of the arc power supplies[2]. Thus, it is also recommended to perform an accelerated testing of a few power supply units in order to determine these values.

25. Depending on the conclusions from points 23 and 24, the BLM team should plan for a contingency for redundant power supply units in the arcs and in the straight sections.

[2] On should add that this sensitivity analysis is very brief and, thus, difficult to digest. A more detailed discussion would have been beneficial.

26. However, one should also consider the effects of redundant power supplies. It must be ensured that the partial load is well-below 50% and that a break-down of one power supply does not affect the other.

27. Furthermore, an overview of all the electrical circuits (HV, R-C circuits, CFC, digital electronics) would be important.

28. In particular, the CFC input stage uses a low-pass filter which limits the incoming signal to about 1 MHz maximum. Also other R-C circuits introduce time constants, which might spoil the original signal, when one R-C component fails. Thus, an analysis should be conducted in order to estimate this impact on the signal quality.

5.2.3 FPGA Programming

29. Generally, a commonly agreed body of knowledge for safe digital design exists [2], which includes concepts such as systematic synchronization of all asynchronous inputs before using them anywhere, making sure unreachable states in state machines are properly handled, etc. It has been found that these techniques have not been systematically applied in all designs. For example, the lack of a proper synchronizer at the output of the frame receivers in the TC electronics might create a race condition.

    Although the auditors have not identified any possibility of malfunction due to these design choices[3], they recommend agreeing on a set of design rules, and upgrading the design in the future so that systematic reviews become easier to perform.

30. The final state machines (FSM) in the CFC FPGA use one-hot state encoding with many undefined states. This is particularly worrying since none of these FSMs are protected in any way against SEU, and, in addition, no way of remotely resetting is foreseen (see also comment 22).

31. In order to ensure complete testing of future changes in FPGA designs, a PASS/FAIL set of regression tests should be designed, preferably using VHDL assertions or equivalent. Each time changes are introduced in the design, the tests should be run on the new design, and a PASS result should be all that is needed to grant release. This avoids the error-prone human activity of looking at waveforms in a simulator, and is inherently more easily extended.

32. In addition, there a more complete review of the FPGA designs should be conducted once these have been finalized.

33. Proper documentation of the FPGA code inside a central software repository like CVS or EDMS is recommended.

5.3 Environmental Aspects

Since the CFC electronics and their power supplies are located beneath the LHC arc quadrupole magnets or in the LHC service galleries, they are susceptible for signal cross-talk, varying magnetic fields, and radiation.

[3] In the aforementioned example, the race condition shall never be problematic due to the sequencing of actions inside the design.

5.3.1 Electromagnetic Compatibility (EMC)

34. "Walkie-Talkie"-type or RF susceptibility test following IEC-61000 should be conducted, e.g. in order to confirm the non-susceptibility of the floating ground of the first ten meters of signal lines between the monitors and the CFC.

35. In particular, the sensitivity of the front-end to noise induced by the power supplies should be determined by injecting common mode and differential mode current on the power leads.

36. Also, the impact of the remaining fringe fields on the BLM electronics should be determined.

5.3.2 Radiation

37. An impressive series of radiation tests on batches of the front-end electronics have been performed using proton beams at PSI and CERN. However, the exposure was short, and with a high-intensity beam.

    It is recommended to redefine the requirements in terms of total ionizing dose (TID), non ionizing energy loss (NIEL) and flux of particles able to produce single event effects (SEE).

    Independent tests for these effects would give more information on the failure modes. TID tests must be done with a gamma source, NIEL with neutrons, and SEE with a beam of hadrons of more than 20MeV. Details can be found in [3].

38. A single CFC electronic board is subject to only relatively low particle fluxes (10^{10} to 10^{11} particles per year). However, due to the large number of boards, the overall system might be susceptible to Single Event Effects (SEE).

39. In particular, it is recommend performing a more in depth analysis of the effects of potential SEUs on the behavior of the CFC and it's FPGA as well as the corresponding recovery scenarios (see also point 29).

40. Furthermore, SEEs in the power supplies of the arc and in the straight sections can lead to their complete failure.

    The failure rate should be determined and the power supplies should be verified to sustain the radiation in the LHC tunnel.

5.4 Commissioning, Testing, and Documentation

The high availability and reliability of the BLM systems highly depends on the proper commissioning and on regular testing ("As-Good-As-New"-tests) [1].

41. Although all ionization chambers and their read-out chain will be tested during commissioning with radioactive sources, the auditors also recommend using flexible input testers with variable current (10pA to 1mA). Such a tester will allow better verifying the full read-out chain including the several integration windows. The subsequent risk due to disconnection the monitor's signal cable is thought to be acceptable.

42. **Large parts of the BLM system are already installed and in place. It is recommended to take benefit of this and start as soon as possible full scale test including the full BLM read-out chain.**

43. In addition, it is encouraged to expand those tests as soon as possible including the BIS. Running the combined systems (e.g. in point 6 or 8) for several weeks, and monitoring the beam permit signal should provide sufficient experience on the behavior of the overall system. Such tests should also include the repeated triggering of beam dump signals.

44. Also, the reception and usage of energy values use for the different threshold levels can already be tested today by using simulated values provided by the BIS team.

45. For bug-tracking, further development, and future upgrades, the auditors recommend setting up a "vertical slice"-test bench, which covers both types of monitors (ionization chambers and SEMs), the full chain of read-out electronics as well as a test database. This will avoid using the final system as a test set-up, and the consequential dangers of accidental mis-configuration.

    Eventually, such a vertical slice might be installed in the SPS in order to benefit from its non-critical environment.

46. Finally, an impressing number of documents, test results, schematics, photos, etc. is stored on the web (http://cern.ch/blm). However, for the long-term, it is recommended to use a more formalized archiving scheme like EDMS.

6 References

[1] G. Guaglio, "Reliability of the Beam Loss Monitors System for the Large Hadron Collider at CERN", Université Clermont Ferrand II – Blaise Pascal, 2005

[2] M. Berg, "Methodologies for Reliable Design Implementation", NASA, 2004, http://klabs.org/mapld04/tutorials/vhdl/presentations/methodologies.ppt; Synplicity Inc. "Designing Safe VHDL State Machines", 1999, http://www.synplicity.com/university/pdfs/designing_safe_vhdl.pdf

[3] ATLAS Collaboration, "ATLAS Radiation Hard Electronics Web Page", http://atlas.web.cern.ch/Atlas/GROUPS/FRONTEND/radhard.htm, accessed 6/2008