



# **External Review of the CERN Beam Interlock System - Report**

October 10, 2009  
FINAL

<p>Originator: Critical Systems Labs, Inc. #618-475 Howe Street Vancouver, B.C. Canada V6C 2B3</p>	<p>Prepared For: CERN European Organization for Nuclear Research CH-1211 Genève 23 - Switzerland</p>
--	--

This page intentionally left blank.

## Revision History

<b>Revision</b>	<b>Sections</b>	<b>Document</b>	<b>Approval Date</b>
-	Document creation with draft versions of chapters 1 to 5.	DRAFT	2 October 2009
A		FINAL	10 October 2009

This page intentionally left blank.

## Executive Summary

The LHC Beam Interlock System (BIS) is an integral part of CERN's approach to the protection of the LHC. The safety function of the BIS is to expeditiously and reliably propagate a beam dump request from a user system to the LHC beam dump system. A failure of the BIS to perform this function could contribute directly to a sequence of events that results in significant damage to the LHC or possibly even the total loss of the LHC. The decision to allow beam to be injected into the LHC must be made with a careful consideration of the extent to which potential safety risks associated with the BIS have been adequately mitigated.

CERN engaged Critical Systems Labs, Inc. (CSL) to review the design of the BIS in September 2009. In addition to a 5-day site visit to CERN, CSL has studied a substantial volume of technical documents relating to the design of the BIS. This document presents the findings of this review.

CSL has found nothing in the design of the BIS that suggests a possible weakness with the potential to jeopardize the safety function of the BIS. Subject to the limitations of this review with respect to scope and resources, CSL concludes that there is sufficient reason to be confident that the BIS will perform its intended safety function. The design of the BIS is a product of very impressive engineering skill combined with very substantial knowledge about machine protection. CSL is particularly impressed by the depth of thought exhibited by the BIS developers in their consideration of potential failure modes that might jeopardize the safety function of the BIS and by the thoroughness of the measures taken to mitigate these potential failure modes. Notwithstanding this conclusion, CSL has made several recommendations in this report that should be considered and addressed if accepted by the BIS team and CERN before the LHC resumes operation.

In particular, attention should be paid to any interface the BIS system has with other systems. CSL concluded that while the design of the core of each system component is extremely sound, a somewhat lesser level of rigor or focus appears to have been used to address any interfaces between the BIS and other systems at the equipment level or at a procedural level.

Additionally, CSL suggests using a more formal and systematic approach to address the sources of safety risk in the system. Such an approach would provide evidence that the safety analysis is complete and comprehensive. It would also facilitate any future safety assessment in the situation of a system upgrade or a design modification.

This page intentionally left blank.

## TABLE OF CONTENTS

EXECUTIVE SUMMARY .....	5
1. INTRODUCTION .....	9
2. SCOPE .....	11
3. REFERENCES .....	11
4. OBSERVATIONS AND RECOMMENDATIONS .....	12
4.1 Introduction .....	12
4.2 Documentation .....	13
4.3 BIS Design .....	13
4.3.1 User Permit Processing .....	14
4.3.2 CPLD and VHDL code .....	14
4.4 Interfaces .....	17
4.4.1 Interfaces to User Systems .....	17
4.4.2 Interface to the LHC Beam Dump System .....	19
4.5 FMECA .....	19
4.6 Verification .....	20
4.7 System configuration and pre/during/post operation .....	21
5. DIFFERENCES WITH INDUSTRY .....	24
5.1 Common Practices in Industry .....	24
5.2 Practices at CERN .....	27
6. THOROUGHNESS OF THE SAFETY APPROACH .....	27
7. ASSESSMENT FRAMEWORK FOR FUTURE PROJECTS .....	31
7.1 Process Objectives .....	32
7.1.1 System concept and lifecycle is defined .....	32
7.1.2 Sub-system design and development life cycle is defined .....	33
7.1.3 Safety Analysis .....	35
7.1.4 System safety is assessed .....	35
7.1.5 Sub-system safety is assessed .....	36
7.2 Product Objectives .....	36
7.2.1 Evidence of correctness is available .....	36
7.2.2 Evidence of robustness is available .....	37
7.2.3 Evidence that product non-conformance / problems are tracked .....	38
7.3 Operational Objectives .....	38
7.3.1 System initialization is controlled .....	38
7.3.2 Static data is protected .....	39
7.3.3 Guidelines for periodic maintenance operations are defined .....	39
7.3.4 Safety assessment process for system change is defined .....	39
8. SUMMARY .....	39
APPENDIX A .....	41

This page intentionally left blank.



## 1. Introduction

The Large Hadron Collider (LHC) at CERN (The European Organisation for Nuclear Research) is one of the world's largest and most complicated machines. The LHC has been designed and implemented to create collisions of sub-atomic particles with extremely high energies and intensities for the purpose of observing the inner workings of the quantum world.

Inside the LHC accelerator two beams of particles travel in opposite directions, they are accelerated from injection energy of 450 GeV to collision energy of 7 TeV, being held in a circular orbit by magnetic fields. A 7 TeV circulating beam within the 27km circumference of the LHC requires a dipole magnetic field of 8.33 Tesla. Superconducting dipole magnets generate this field, operating with a forward current of almost 12kA at temperatures just above absolute zero (1.9 K or -271°C).

The energy stored in each of the two circulating beams in the LHC reaches a maximum of 360 MJ at collision energy and design intensity ( $3.2 \times 10^{14}$  protons per beam). Beam impact into the material of the accelerator produces a cascade of particles due to nuclear and electromagnetic interactions, which deposit energy into the accelerator equipment. For *fast beam losses*, over a millisecond to second timescale, losing as little as  $10^{-8}$  of the beam into one of the superconducting magnets will lead to a quench, where the magnet becomes normal conducting and has to be switched off before it destroys itself. A fast beam loss of  $10^{-4}$  of the beam into any part of the machine will cause damage, such as rupturing the machine vacuum, which in the best case results in costly repairs and weeks of downtime, in a worse case the destruction of one or more dipole magnets would mean many weeks of repairs to return the machine to operation

In order to prevent damage to the accelerator due to beam losses, a Machine Protection System (MPS) has been implemented, which detects emerging dangerous situations, and ultimately extracts the circulating beam onto a purpose built graphite target, safely depositing its energy. Several subsystems make up the MPS, at the heart of which lies a Beam Interlock System (BIS).

In the LHC the BIS connects USER SYSTEMS to the LHC Beam Dumping System (LBDS). USER SYSTEMS can make one, or many connections to the BIS, giving USER\_PERMIT signals. The BIS then derives a BEAM\_PERMIT signal from these, and relays it to the LBDS. When BEAM\_PERMIT transitions from TRUE to FALSE, the LBDS initiates a controlled extraction of beam from the LHC, and simultaneously further injection of beam is inhibited.

As there are two circulating beams in the LHC, there are two LBDS, and two BIS. Some User Systems distinguish between LHC Beam-1 and LHC Beam-2, others do not, acting on Both-Beams of the LHC.

Safety requirement: The path from USER SYSTEM to LBDS through the Beam Interlock System should have a very high safety, it is required that the probability of unsafe failure of this link should be  $< 10^{-7}$  per hour.

Availability requirement: The Beam Interlock System should not adversely affect the availability of LHC,  $< 1\%$  of LHC missions should be aborted due to failures of the BIS

As the CERN prepares to resume operation of the LHC, the Machine Protection and Electrical Integrity Group engaged Critical Systems Labs, Inc. (CSL) to review the design of the BIS. CSL is an engineering consultancy based in Vancouver Canada that specializes in the specification, analysis and verification of safety-critical system. To the task of reviewing the design of the BIS, CSL brings a wide range of experiences from a variety of industry sectors including aerospace, automotive, rail signaling, medical device technology and defense.

The purposes of this review have been to:

- identify possible weaknesses in the mission-critical BIS before LHC reaches high intensity beam operation
- assess the adequacy of the external and internal mitigations for critical component failure in the BIS
- provide a general comparison of the BIS with approaches in industrial systems.
- suggest potential improvements of the BIS
- review and comment on the pre/during/post operational software sequences that verify the integrity of the BIS
- provide CERN with a model for future assessments of mission-critical systems

Two senior CSL engineers, L. Fabre and J. Joyce, visited CERN for five days during the week of September 7, 2009. This site visit was facilitated by B. Todd (Beam Interlock System Hardware Engineer) and B. Puccio (Machine Interlock Section Leader). Prior this site visit, CERN provided CSL with a very substantial volume of engineering documents that were studied by CSL in preparation for the site visit. Following the site visit, CSL has corresponded with B. Todd and B. Puccio using electronic mail to clarify certain technical details related to the above purposes of this review.

This report is organized as follows:

- Section 2 describes the scope of this review.
- Section 3 lists all the documents that are referenced within this report.

- Section 4 describes observations and recommendations made by CSL.
- Section 5 describes the most salient differences between the methodology used to develop the BIS and common methodologies used to develop safety-critical systems in other industries.
- Section 6 focuses on the review of the safety analysis approach used by the BIS team.
- Section 7 provides an assessment model for mission-critical systems.

## 2. Scope

The scope of this independent review is limited to the LHC BIS including the pre/during/post operational software sequences that verify the integrity of the BIS.

The scope of this review is limited to a consideration of:

1. Potential sources of safety risk within the BIS, where a request from a USER SYSTEM is not relayed to the LBDS, or is delayed by more than acceptable amount of time, resulting in a ‘missed dump’ and potentially machine damage.
2. Potential sources of unavailability, where failure of the BIS leads to removal of the circulating beam, without a request from a USER\_SYSTEM, resulting in a ‘false dump’.

While motivated by CERN’s interest in the protection of the LHC from damage, this review is not a safety analysis of the LHC. In particular, the scope of this review does not include the task of identifying additional hazards. (For the purpose of this review, the only hazard of interest is a ‘missed dump’, where the BIS fails to expeditiously propagate a request for a controlled extraction of a beam from the USER SYSTEM to the LBDS.

A quantitative assessment of residual risk is also outside the scope of this review.

## 3. References

1. B. Puccio et al ., The Beam Interlock System for the LHC, February 2005, EDMS 567256.
2. B. Todd, Beam Interlock System Standard CIBM Matrix specification, December 2007, EDMS 884688.
3. B. Puccio, Slide Presentation, CIBU Connection Review.12Dec08.ppt.

4. R. Denz et al., The Beam Interlock System – Report on the audit held on September 18th-25<sup>th</sup>, 2006, October 2006.
5. B. Todd, Failure Modes, Effects and Criticality Analysis of the Beam Interlock System, July 2009, EDMS 762129.
6. B. Todd, LHC and Injection Beam Interlock Systems, August 2008, EDMS 952098.
7. M. Kwiatkowski et al., Automated Testing of the User System Interfaces to the LHC Beam Interlock System, TE Technical Note, 0.3.
8. B. Todd, Configuration Verification, Fault Diagnosis and Monitoring of the Beam Interlock Systems, July 2007, EDMS 855069
9. RTCA, Inc. Design Assurance Guidance For Airborne Electronic Hardware. RTCA/DO-254, April 2000.
10. RTCA, Inc. Software Considerations in Airborne Systems and Equipment Certification. RTCA/DO-178B, December 1992.
11. B. Todd, CIBM Technical Design Report, June 2007, EDMS 761333.

## 4. Observations and Recommendations

### 4.1 Introduction

During a visit to CERN, CSL observed that the knowledge of the BIS system design was presently available in the CERN design team. Any questions CSL had about the intent of some of the BIS features could be answered in a timely and concise manner. This made a very positive impression to CSL.

CSL observed in their interactions with the BIS team and through the many Technical and Design notes that thorough problem analysis has been performed for the design of the various system components. A depth of thought is evident in all aspects of this work and helps to justify the important design decisions.

CSL realized in the early days of the review that, contrary to their initial understanding, the purpose of the BIS system is greater than just machine protection contrary: it appears that the BIS system also helps to guarantee operational readiness while at the same time it is paramount to the protection of the LHC and to the LHC experiments.

The other sections in this chapter list all the suggestions and **recommendations** that were deemed important enough by CSL to be stated in this report. The implementation of these **recommendations** would further increase the level of confidence that the BIS performs its intended function in its operational environment. In addition the implementation of both suggestions and **recommendations** would further increase the maintainability of the system without compromising safety aspects.

CSL advises implementation of the recommendations included in this report before the LHC resumes its operations in late 2009.

## **4.2 Documentation**

CSL received a large amount of BIS related documentation. CSL found these documents to be clear, detailed and very complete. However CSL noticed that some of these documents have not been updated to reflect the current situation. For example, the number of BICs in the currently configuration of the BIS is greater than what is described in [1].

**S 1: Some documents should be updated so that a snapshot of the most up-to-date BIS documents can be easily retrieved at any time. This would identify what documents are current and what documents are obsolete.**

CSL also observed that the functional behavior of the BIS system is not centrally described but that the information is distributed among several documents. While all the current system knowledge is presently available in the BIS team and does not present an issue at this point, a functional specification would facilitate the maintainability of the system as well as any future design changes. Such a document would also be an excellent tool to exchange information with external systems.

**S 2: A functional technical specification of the BIS should be developed. This document should focus on the functionality of the BIS and should include limited design information. This document should also include interface specifications for any external systems that connect to the BIS.**

## **4.3 BIS Design**

Many aspects of the BIS system have applied proven concepts, protocols and use a simple design<sup>1</sup>. The number and the kind of interactions between components and across interfaces is limited and reduces unforeseen system behaviors. This is certainly a very favorable factor for the safety assessment of the overall BIS design.

As the BIS design uses proven concepts and components, there is reason to expect that the number of “unknown unknowns” in the BIS system is very limited.

The BIS system uses mostly independent channels in the evaluation of the conditions to trigger the beam dump system. This full redundancy starts from the gathering of the user inputs all the way to the provision of the BEAM\_PERMIT information by two different loops to the LDBS. Also the voting system to combine this redundant information is not part of the BIS system but is implemented in the LBDS. The implementation of any voting system typically requires a very high-level of rigor. The

---

<sup>1</sup> “One of the most important aspects of safe design is simplicity”, N. G. Leveson, Safeware, p.405.

fact that this voting system is not part of the BIS simplifies, to some extent, the BIS design.

### 4.3.1 User Permit Processing

If a user permit is considered maskable then the value of this user permit does not affect the determination of the loop beam permit when the `Safe_Beam_Flag` is equal to `TRUE`. As a result, this maskable / non-maskable categorization is essential since it can lead to the omission of some user permit values.

In Appendix B of [1], user permits are organized in the maskable / non-maskable categories. However, CSL was not able to identify a document that defines the rationale to decide whether a user input should be categorized as maskable or non-maskable.

**R 1: The rationale to make a user permit maskable / non-maskable should be documented. If no systematic rationale exists then the justification to make any specific user permit maskable should be documented.**

CSL understands that the Machine Protection Panel, a group that meets on a monthly basis, is responsible for the above decisions. However this responsibility was not described in any document.

### 4.3.2 CPLD and VHDL code

Since the CPLD implementing the Matrix code is one of the most critical components of the BIS, CSL dedicated particular attention to its design and development during the review.

The physical CPLD (Xilinx XC95288XL) used to implement the Matrix is highly reliable, but has a very limited capacity for functionality.

Two different implementations of the VHDL exist: Matrix A and Matrix B. This diversity extends the concept of redundancy with two independent detection channels to this critical component.

CSL noted that the last modification of either Matrix code was made at least two years ago. Therefore the VHDL code used is very stable and the latest version of the code has been under test for many months. In the experience of CSL, this degree of stability is very unusual for a complex system. It is more typical for a system of this complexity to have undergone changes up to a time very close to its deployment, or at least, the period of stability is much more likely to be measured in weeks or months rather than years. The stability of the BIS combined with extensive testing during this period of stability is a strong positive factor in an argument that the behavior of the BIS is well understood for the purpose of assessing its dependability.

The version number of the VHDL code is hard-coded within the file. This theoretically does not prevent someone from modifying the code and preserving the same version number. However this VHDL code hardly ever changes and therefore the risk associated to this sequence of events is small. Currently the constraint of space on the actual CPLD does not allow the BIS team to store the version number in another way that would prevent the above mentioned risk.

It appears that most parts of the VHDL code for Matrix A are similar in structure to VHDL code for Matrix B and these parts could not be implemented any differently. The only true difference between Matrix A and Matrix B is the implementation of the glitch filter (one for Matrix A and the other for Matrix B).

CSL performed an extensive review of the VHDL code and informally provided feedback to the BIS team. CSL has identified some limitations that were already known to the BIS group. In general while several improvements could be made in terms of readability and clarity of the VHDL code, CSL has not identified changes that could enhance the dependability of the processing of these two matrices.

CSL has noted that certain minor anomalies of the VHDL implementation of Matrix A. Some of these anomalies are stylistic in nature, e.g., certain signals in the VHDL have names such as “clk1m00” that imply that these signals are clock signals when they are really just signals that enable an action. Other details may be consider deviations from best practices because they might introduce some unnecessary uncertainty about particular aspects of the behaviour of the synthesized design. For example, the clk1m00 signal is not initialized upon reset which means that it is in an unknown state immediately after reset. Because of these anomalies, extra care should be taken if and when changes are made to the VHDL implementation of Matrix A. However, CSL is not aware of any reason why these anomalies could have an adverse affect on the propagation of a beam dump request.

The glitch filtering performed by the matrix is intended to avoid excessive false trips. It has no benefit to safety. A glitch is defined as a period when the user permit is FALSE for X nanoseconds or less. The diverse implementation of this glitch filtering (i.e., one implementation for Matrix A and a different implementation for Matrix B) is intended to avoid a common failure mode for glitch filtering by both Matrix A and Matrix B.

The glitch filter in Matrix A and Matrix B uses a threshold value of 1.6 micro seconds as described on page 11 of [2]. However the justification and origin of this value is not documented. This assumption about this essential value being the right value for filtering needs to be reviewed.

**R 2: The origin of the value of 1.6µs used in the glitch filter should be documented and reviewed.**

The VHDL implementation of the glitch filtering for Matrix A is more complex than the implementation of glitch filtering in Matrix B. The implementation of glitch filtering in Matrix A is based on an algorithm that, in our opinion, is more complex than more standard implementations of a glitch filter. Due to some particular timing-

related details, it is possible that some “non-glitches” might be erroneously filtered, where the term “non-glitch” refers to a situation where the input to the glitch filter is continuously FALSE for a duration that exceeds 1.6 $\mu$ s. However, the worst case known to CSL is when the duration of the non-glitch is just slightly less than 2 $\mu$ s. This finding has been discussed with the BIS team. CSL understands that this aspect of the behaviour of the Matrix A glitch filter is known to the BIS developers and that it can be tolerated without concern that the safety function of the BIS might be jeopardized.

In general, CSL believes that non-critical concerns about the VHDL implementation of the Matrix A glitch filter are outweighed by the potential benefits of using a diversely redundant implementation of glitch filtering in the implementation of the BIC.

The glitch filter, at least for Matrix B, does not filter out "up-glitches": this would be a scenario where the input is continuously FALSE for more than 1.6 $\mu$ s except for very short (50 nanosecond) up-glitches when the input is TRUE. If an up-glitch occurs at least once every 1.6 microseconds, the output will remain TRUE. Technically, this is the correct behavior. But what assures us that the input will not contain these up-glitches? This finding has been discussed with the BIS team. CSL understands that this aspect of the behaviour of the Matrix B glitch filter is known to the BIS developers and that it can be tolerated without concern that the safety function of the BIS might be jeopardized.

CSL notes that a single bit is used to connect the output of the AND function in the matrix to the switch that gates the re-transmission of the beam permit loop. A single event upset that affects this single bit could prevent the propagation of a beam dump request from any one of the user systems. CSL suggests using at least two bits to propagate the result of the AND function in the matrix to the switch. This change would make the most essential output of the matrix less susceptible to single event upsets<sup>2</sup>. If this change is to be implemented, the designer should then ensure that the Xilinx synthesizer does not optimize this data structure away.

**S 3: CSL suggests using at least two bits to propagate the result of the AND function in the matrix to the switch.**

CSL noted that another potential concern is a clock freeze. However the output of the matrix CODE block called ‘Beam\_Permit\_Out’ witnesses the proper behaviour of the clock. This output is transmitted to the FPGA on the CIBM board that monitors several inputs and feedbacks to the DIAMON BIS monitoring system.

CSL observed that the Xilinx synthesis requires some options to be manually configured at synthesis time. Two options are essential for this particular CPLD.

---

<sup>2</sup> “Safety-critical systems shall not employ a logic 1 or 0 to denote the safe and armed (potentially hazardous) states. The armed and safe states shall be represented by at least a four bit unique pattern” Michael L. Brown in Software systems safety design guidelines and recommendations. Technical report NSWCTR 89-33, March 1989.



These configuration parameters cannot be formally saved. As an alternative the BIS team inserted these settings as comments in the VHDL code in an attempt to keep some record of this information. For critical systems, the build or synthesis process should ideally be fully deterministic.

CSL observed that no formal tool qualification process of the Xilinx ISE was performed by CERN. CSL also observed that the version of the tool used by the BIS team is an earlier version than the version officially used in CERN. This decision was made when the BIS team realized that using later versions of the tool would create difficulties, if not errors, in the synthesized code.

**S 4: The rationale and benefits to use a Xilinx ISE tool version different from the official CERN recommended version should be documented. Any future update to this tool should be at least documented and tracked. Finally tool assessment and potentially qualification for the Xilinx SE should be considered. See section 5 about tool qualification.**

CSL understands that the BIS group was originally told that radiation would not be a factor, but more recently they learned that the radiation levels could be a problem for the electronics in the BIC. The BIS team is currently working to address this issue.

Radiation on the Xilinx CPLD could cause a single event upset which might adversely affect the persistent states of the matrix. However, radiation could also alter the flash memory that contains the functionality synthesized from the VHDL. This flash memory is what controls the CPLD. Altering this would obviously change the functionality of the matrix. CSL questions what protection exists for this problem.

During their radiation tests, the BIS team has never seen a single event upset in the flash memory in the power-on phase, when the contents of the flash memory are loaded into the logic of the device. Therefore a single event upset in this flash memory can only appear when the device has been power cycled.

The BIS team suspects that Xilinx power-on process includes some level of checking before loading and executing a program on the CPLD. In this context the BIS team does not currently check that the contents of the flash are free from errors and makes the assumption that the CPLD would not allow itself to start operations if there had been a single event upset in the configuration.

**S 5: The BIS team should check with Xilinx that an error-check is performed before loading and executing the CPLD program. Alternatively the BIS team should consider a CRC check on the synthesized code.**

## **4.4 Interfaces**

### **4.4.1 Interfaces to User Systems**

CSL noted that the purpose of the user permit is either:

- failure detection (e.g., Beam Loss Monitor), or
- operational readiness (e.g., Radio Frequency).

A general principle for the design of safety protection function, such as the BIS, is to isolate the safety protection function as much as practically possible from the rest of the system. This principle aims to minimize the possibility that a problem in the rest of the system might jeopardize the dependability of the safety protection system. At a very early stage in this review, CSL initially assume that most of the inputs to the BIS would be from sources dedicated to machine protection, e.g., beam loss monitors. In this regard, CSL was surprised to discover the extent to which the BIS accepts other inputs and especially the extent to which the relevance of some user inputs to the goal of protecting the LHC appears to be uncertain, or at least, not rigorous established and documented. Indeed, it seems clear that some sources of inputs exist for the purpose of operational readiness rather than failure detection.

CSL has some concern that part of the response to the September 19, 2008 incident might have been decisions (outside the control or visibility of the BIS team) to add more and more sensor sources to the user systems that generate inputs to the BIS without the necessity of these additional inputs being well understood by the machine protection group. Experience with the operation of the LHC will soon reveal the impact of these additions to the availability of the LHC, as some inputs might generate more “false trips” than what is considered acceptable. CSL has no evidence of a problem or concern about unavailability. But in the event that too many false trips are experienced once the LHC becomes operational, CSL is concerned that there will be a “counter-response” to eliminate some sensors sources without a clear understanding of their safety relevance. This counter-response, especially if it is made with haste in the absence of a documented record of the safety relevance of each input to the user systems, might result in the disconnection of an input which in fact is vital to machine protection.

**R 3: Every user condition that contributes to a user permit input should be justified, in particular, the inputs that come from the experiments and other sources which are outside the BIS. In particular, the safety relevance of each such condition should be documented.**

The importance of the maskable / non-maskable grouping was previously highlighted in this report. CSL initially understood that Experiment User Input could not be masked and later realized that this statement was always entirely true.

In principle each input can fall into one of three categories:

- 1) Detector electronics – Essentially Beam Loss Monitors inside the experiment. These are unmaskable.
- 2) Magnet – A signal from the magnet control system. These are maskable.

- 3) Moveable Devices – Signals from the hardware which can move close to the beam (normally within tens of microns of the beam itself). These are unmaskable.

**S 6: The partitioning of Experiment User Inputs between maskable and non-maskable should be documented and justified.**

Following the UJ33 incident, there is a procedure that an experiment (or any other group that provides user permits) must follow before their user input can be connected to the system. One purpose is to show that the input will not harm the BIS.

The recommendations made following the UJ33 incident, slide 32 of [3] are being systematically followed up. CSL heard that the Machine Protection Panel has agreed with the strong position that ‘no redundancy = no connection to the BIS’. This apparently resulted in some hardware modifications on the user side.

CSL understands that the BIS team asks users to supply the schematics of the interface and verify the conformance of these schematics. In addition, the BIS team performs the CIBU test (measure V & I, check redundancy) for this particular interface.

**R 4: Continue to follow the recommendations made following the UJ33 incident and ensure that these recommendations are incorporated into life cycle processes for maintenance of the LHC.**

#### **4.4.2 Interface to the LHC Beam Dump System**

The Beam Dump System receives a copy of the Beam\_permit\_loop\_out signal for each loop. This system has implemented its own detector of the beam permit frequency to decide whether or not it should trigger a dump. CSL understands that the implementation of this detector has been done independently from equivalent detect functions in the BIS system and that the BIS team provided the LDBS team all relevant and necessary information.

CSL was also told that the LDBS team is performing an audit of their system.

**R 5: CSL recommends that a member of the BIS team participates in the review of the optical beam permit detector developed by the LDBS team. In particular this person should identify whether any assumptions were made by the LDBS team for the development of this function.**

This recommendation echoes the recommendation made by the internal audit held Sep. 18-25, 2006 [4] that stated that “exceptional care has to be put on its [i.e., the detector function] integration and functional testing”.

#### **4.5 FMECA**

The BIS team performed a very comprehensive set of Failure Modes Effects and Criticality Analysis (FMECA) for all the possible boards in a BIC and for the different elements of the user interface.

All these results were consolidated in document [5].

CSL observed that the BIS team was able and took the initiative to compare predicted reliability figures with actual figures. Results were gathered over a period of more than 1000 days of operation. The results of this comparison were then analyzed when actual results were worse than predicted results. This effort led to a new PCB design for the CIBM board.

CSL underlines that this reliability analysis is quite important as the BIS is mostly a protection system and the safe operation of the LHC is directly connected to the reliability of its protection, i.e., the BIS.

CSL recognizes that this analysis has been performed to a high standard.

## **4.6 Verification**

CSL's overall impression is that a very comprehensive multi-level approach has been taken to verify the BIS. It matches or exceeds what would be done for very high integrity systems in industry, with a few qualifications as noted below.

The levels of testing and verification include:

- 1) Component testing - This level of verification is typically a very exploratory effort that attempts to verify specific characteristics of individual components. One example is an effort by the development team to verify the attenuation of the frequency in the fibre optic cable as it passes from BIC to BIC. Another example is EMC testing of digital electronic components used to implement the BICs.
- 2) VHDL verification (which involves three distinct kinds of verification) as described below
- 3) Board level testing (done for every board)
- 4) BIC testing (performed in a rack), done for every BIC
- 5) System testing (while in test mode)

VHDL verification involved:

- 1) Model Behaviour. This is performed using the ModelSIMSE tool from Mentor Graphics. It uses the VHDL code directly (without any processing from Xilinx). Files were tested both individually and as a set. Coverage options have been selected by the BIS team to ensure a highly thorough level of

coverage i.e. Statement, Branch, Expression and Condition. The ModelSIM SE tool provided the coverage results. The coverage results that appear in [2] page 22 are not directly under configuration management in the form of test results, as it is usually performed in other industries. These coverage results are extremely important and give a high level of confidence that no unintended functionality exists in this code.

- 2) The Post-fit simulation from Xilinx generates synthesized design information that has been reviewed by the BIS team from a correctness point of view.
- 3) Board testing covers testing of the Hardware i.e., CPLD with synthesized code.

However it was noted that all these various verifications and levels of testing are not described in a verification & validation plan or a master test plan.

CSL did not receive nor review any test procedures from the BIS group. CSL saw a limited number of associated test results. For instance CSL believes that the BIS team verified that the use of the CPLD by CERN conforms to all the usage-rules that the manufacturer has specified such as a maximum clock rate or signal rise time. However this validation evidence may not be documented.

Overall CSL notes that documentation evidence has been developed to a great extent for the design of the system, but to a lesser extent for the verification phase of the system. CSL recommends ensuring that proper test procedures are available with archived test results.

In particular, test procedures and results for the glitch filter that were collected using a signal generator should be fully documented and archived.

**S 7: Test procedures should be available for all testing levels and should be detailed enough so that existing test results can be reproduced in the future. These test procedures would also help to ensure the integrity of the BIS system is not compromised over its lifetime.**

#### ***4.7 System configuration and pre/during/post operation***

All the elements relevant to the configuration of the BIS are stored in the BIS configuration database (Oracle database). This data in the BIS configuration database is used by the pre-operational testing sequence.

Currently there is no verification of the BIS configuration database, with the exception of access control, and versioning. However the Controls group supplies a log of changes made between two versions that could potentially be reviewed.

For instance, someone making changes to the database could accidentally assign a user input that should never be masked to a connector that is maskable. If a

subsequent error is made at the physical connection level, the BIS may not generate a beam dump signal when it should.

**R 6: A verification process for changes to the BIS configuration database should be defined. This verification process could be a review of the changes log between two versions.**

The integrity of the BIS configuration database does not seem to be checked. In CSL's experience, industries with critical systems that involve a database with essential configuration information always have a way to ensure the integrity of the information in the database before using it. In the context of the BIS, CSL does not know if there is any reason why this configuration database could become corrupted and believes there is only one copy of the database i.e., the same replica of the database is used for edit and for the pre-operational sequence.

**R 7: A means to check the integrity of the database before the pre-operational sequence is recommended.**

CSL noticed that the BIS team commonly uses a user input table maintained in an excel spreadsheet. This user input table describes the list of users and their inputs to the BIS. This table is loosely connected to the BIS configuration database. CSL remarks that this spreadsheet that gives a useful and convenient summary of BIC occupancy may be used in meetings to make decisions about changes to the system while the information may not be faithfully identical to the BIS configuration database. The process could be improved by automatically generating this spreadsheet from the contents of the BIS configuration database. In this context, CSL underlines the importance of recommendation R 6.

The pre-operational test is a Java program that is invoked by the "sequencer", which is a program that automates the process of starting up the machine. The BIS group is developing the pre-operational test program, but they will not have control over this software once it is handed over to the controls group. Since the command to arm the BIS is embedded in the pre-operational test program, it is not possible to simply skip the pre-operational test (except for the temporary "arm button" alternative mentioned below). However, it may be possible for someone, without the knowledge of the BIS group, to modify the pre-operational test program such that some (or even) all of the checks are skipped. The possibility that such a modification might be performed due to malicious intent is a security concern outside the scope of this review. However, it is easy to imagine the possibility of such a modification being made for non-malicious reasons (e.g., to expedite the process of getting beam into the LHC) without a full understanding of the safety implications of making this modification.

**R 8: A procedure should exist to ensure that the BIS portion of the pre-operational program run by the Control group is identical to the program handed-over by the BIS group to the Control group.**

CSL observed that the correct maskable or non-maskable configuration of user inputs is enforced via the pre-operational check sequence that will ensure that maskable user inputs are connected to the right CIBP ports.

CSL observed that there is a temporary capability that allows a system operator to arm the BIS directly without pre-operational testing. It is crucial to make sure that this temporary capability is removed before the system becomes operational, and the BIS modified so that it would be very difficult for anyone to restore this capability at some time in the future.

**R 9: The short-term “re-arm” (without checks) button provided to the system operator is a source of risk that should be removed before the LHC resumes operation.**

During operation, history logs from the BICs are retrieved and stored in a central place. The amount of information is considerable and it is difficult to locate parts of interest unless one knows where to look. CSL suggests developing analysis tools to post-process history logs and particularly to verify a list of pre-defined properties. These properties would be conditions, states or sequences that should hold TRUE at any time based on system design knowledge. One such property could be that if a user input changes to FALSE both Matrix A and Matrix B should interrupt the propagation of the signal on the respective loops to which they are connected. Another property could be checking that the BIC processing of a user input (between detection of user input change to cutting the loop frequency signal) does not exceed 10us.

**S 8: CSL suggests post-processing the operational history logs to ensure some system properties remain TRUE.**

The documents that describe the pre-operational / operational and post-operational phase include [6], [7] and [8]. CSL understands that these documents are currently updated as the nature and the number of the checks performed in each phase is fully determined.

**S 9: The documents describing pre-operational, operational and post-operational checks should be finalized and any overlap between these documents should be removed.**

CSL understands that the BIS team initially intended the pre-operational sequence to:

1. test completely the BIS, in using an IST [6]
2. test completely the User Systems in using the User Test [7]

However the BIS team observed that the length of time that Step 2 takes might not justify a test before each mission.

It should not be left undetermined what user inputs are proof-tested and how often they are proof-tested. CSL recommends finalizing the list of user inputs that are required to be tested before each mission and identifying the frequency of testing for other inputs. In addition the rationale used for the frequency should be documented as it may be useful for the evaluation of other user inputs in the future.

**R 10: The test frequency of each user input should be specified.**

## **5. Differences with Industry**

This section reports on the differences that CSL auditors have observed between the development of the BIS system and safety-critical systems in other industries (aerospace, defence, railway, medical devices).

### **5.1 Common Practices in Industry**

This section describes practices commonly adopted by industries developing safety-critical systems for which CSL has found no direct equivalent within the BIS project.

In general this section does not include recommendations from CSL. Instead, the observations provided here about the differences between industry and the approach taken by CERN to develop the BIS are intended to be neutral statements. CERN and the BIS team can assess whether some of these practices would be of some benefit in the CERN environment.

In particular, CSL very much appreciates that CERN is a research institution with a research-oriented culture that is very different from the typical culture of an industrial organization. It would be unreasonable, and also undesirable, to expect the development of the BIS to exactly follow the practices of industry used for the development of complex safety-related systems. Moreover, CSL has noticed how the research-oriented culture of CERN has benefited the development of the BIS. This is especially evident in the richness of various technical documents that have been produced in the course of developing the BIS. It is not common to see the same deep of thought in the technical documentation typically produced by industry for the development of complex systems. Nonetheless, it is useful for CERN to be aware of differences between the practices used by CERN for the development of the BIS and the common practice of industry for system similar to the BIS in regard to complexity and criticality.

#### **1) Engineering Planning documents**

The design and development lifecycle of a critical system is usually planned and described in a hardware or software planning document. This plan describes the various activities involved and the inter-relationships between these activities. For instance the software development project may include the following activities: requirement development, design development, coding, integration, testing. The plan documents the entrance and exit criteria for all these activities. The development plan also addresses how change is managed during the project lifecycle. Additionally, the development plan describes roles and responsibilities that are relevant to the project. Roles and responsibilities may apply to an individual, to a group or to a panel.



The benefits of the engineering plan include:

- ensuring repeatability,
- ensuring a common understanding of the way forward by various project stakeholders.

## **2) Reviewing process**

The reviewing process of engineering artifacts produced for a safety-critical system is formally described: the description includes the kind of review (peer, milestone, client, external, etc.), the frequency or timing of the reviews and how comments will be tracked and addressed. Often these reviews are part of the exit criteria that allow moving from one engineering activity to the next.

The reviewing process is usually documented in the engineering planning documents.

CSL has observed that different types of reviews have been performed for the BIS system but that the review process was not formally documented. In addition, a system wide internal audit involving a group of reviewers was performed in 2006 and resulted in a significant number of recommendations that were documented and subsequently addressed.

## **3) Use of Style Guide / Standards**

The use of standards or an internal style guide for requirements, design and coding is a common practice for safety-critical systems. The standards or style guides ensure consistency between work performed by different individuals. This is particularly true for large-scale projects.

The standards / style guides contribute to the readability of engineering artifacts and allow for faster reviews. Standards and style guides can also contain “Do / Don’t Do” practices that are relevant to a specific technology, environment, or programming language.

When a reviewing process exists, the reviewing goals ensure that standards/ style guides have been applied correctly.

## **4) Traceability**

Another cornerstone of the development of safety-critical systems is evidence of traceability between the various engineering artifacts. Traceability means some recorded evidence of a mapping between different engineering artifacts.

The intent of traceability is to ensure that the meaning of each level of specification or description has not been corrupted or lost in the next lower level of specification or description. As well, the intent of traceability is to minimize the

possibility that unintended functionality has been introduced at some level of specification or description.

The existence of traceability also enables "drilling down" from requirements to implementation and follows a specific thread of processing from initiation to completion.

If a system is modified, the existence of traceability allows for assessing the impact of a change in a systematic manner and allows for an easy verification that a change is still compatible with higher level system description.

Traceability can exist between most of the engineering artifacts. Important traceability evidence includes the mapping between:

- System requirements and software/hardware requirements
- Software / hardware requirements and artifacts of architectural design decision
- Software requirements and source code
- Hardware requirements and HDL code
- Requirements and verification results.

CSL did not see traceability evidence in the BIS system.

## 5) Tool Qualification

Production of software and hardware is becoming increasingly more sophisticated. In this context tools are heavily used to develop and verify software and hardware. However errors could also exist in tools and tools could produce erroneous results. Therefore industries involved in critical systems often require that tools be assessed and qualified for their intended usage in a given system.

RTCA DO-254 / EUROCAE ED-80 [9] defines the purpose of tool qualification as follows:

"The purpose of the tool assessment and qualification is to ensure that the tool is capable of performing the particular design or verification activity to an acceptable level of confidence for which the tool will be used".

RTCA DO-254 [9], section 11.4 and RTCA DO-178<sup>3</sup> [10] describe the tool assessment process and also what the qualification of a tool entails.

---

<sup>3</sup> The draft version C of DO-178 dedicates an entire supplement to the topic of Tool Qualification.

In the context of the BIS system, the tool qualification question would apply to the Xilinx ISE synthesizer and the MentorGraphics ModelSIMSE. The use of these tool should be assessed to determine if a qualification is required.

DO-254 provides a flow chart to help with tool assessment considerations. In particular, one of the essential steps is to determine whether the tool output is independently assessed. If so, then no further assessment is necessary. If not, then the organization should first look for relevant tool history. If this is insufficient, DO-254 requires a proper tool qualification.

CSL notes that a Xilinx white paper dedicated to “Meeting DO-254 and ED-80 Guidelines When Using Xilinx FPGAs” states that their ISE suite is robust and adequate for usage in a DO-254 context. The Xilinx white paper gives some basic level of confidence that the Xilinx ISE synthesizer is appropriate for its usage at CERN.

## **5.2 Practices at CERN**

This section describes practices that are used at CERN and that are not common in industry. CSL considers these practices as positive factors in the development of dependable systems.

### **1) Design Analysis documents**

CSL observed that many analysis documents and technical notes have been produced. The use of any technology has been analyzed to show that it meets the intended system requirements. These analysis documents detail the adequate configuration and parameterization of a given component or technology. In some respects, CSL has found more explanation for the many choices made during the system design than in some other safety-critical industries.

### **2) Design and Development performed “in-house”**

CSL observed that most of the design, development and verification is performed by the CERN team i.e., not many of these engineering tasks are outsourced to third-party companies. In general most industries tend to keep the ownership of system design. However there is a noticeable trend towards outsourcing some of these activities, particularly the development and verification activities. Outsourcing requires additional layers of communication. It is often more complex to understand how the system works when design and development of system components are distributed among several different contracting entities. In addition it is often more complex to make any design changes to the system.

## **6. Thoroughness of the Safety Approach**

In general, the safety of a complex system such as the LHC BIS is achieved by a process that seeks to identify sources of safety risk and then eliminate or mitigate these sources of safety risk. From this perspective, CSL has assessed the safety approach taken by CERN for the development of the BIS in terms of the following two questions:

1. To what extent has CERN done everything reasonable to identify potential sources of safety risk for the BIS?
2. To what extent has CERN addressed known sources of safety risk adequately?

CSL understands that the BIS has a single safety goal, which is to minimize the possibility that the BIS continues to provide a TRUE beam permit to the beam dump system after one or more FALSE beam permits have been received by the BIS from user systems. This is a part of a more general safety goal to protect the LHC against loss or damage.

It is obvious to CSL that every significant decision made about the design of the BIS has been made with a thorough consideration of the above mentioned safety goal. This extends from large-scale decisions such as the use of two fully redundant channels to very specific details such as the choice of frequencies used for the beam permit signal transmitted on each loop. It is also obvious that these decisions have benefited from experience and knowledge within CERN and more generally, the worldwide accelerator community. However, CSL has some concerns about the extent to which the safety of the BIS has been documented in a manner that will allow the safety of the BIS to be maintained over its entire lifecycle.

### **Identification and Understanding of Potential Sources of Safety Risk**

The project documentation for the BIS does not explicitly enumerate a set of identified hazards or potential causes of these hazards. There is no evidence that a systematic process was used to identify hazards for the BIS. Instead, it seems that the development of the BIS has relied on the specialized knowledge of key individuals at CERN associated with machine protection for an implicit understanding of the hazards and their potential causes. CSL also notes that CERN has drawn from the specialized knowledge of the worldwide accelerator community in its design of the BIS.

It is easy to deduce from the project documentation that one possible consequence of a failure of the BIS is the possibility that LHC beam dumping system might not be activated when required, e.g., upon detection of a beam loss condition. However, the project documentation does not indicate whether a failure of the BIS could cause harm in other ways beside the failure to dump the beam. For example, CSL was unable to determine from the project documentation the potential effect of a BIS failure on the injection point where particles are transferred from the SPS ring to the LHC. In the course of the site visit to CERN, CSL learned that the possibility that particles might be injected into LHC during a beam dump had been studied by several PhD students at CERN. However, the relationship between a failure of the BIS and

the injection of particles from the SPS is not explained in the BIS documentation. Are the design decisions taken to minimize the possibility that a failure of the BIS might be the reason that a beam dump is not initiated sufficient to also address safety concerns about the injection of particles when conditions warrant a beam dump? This is just one example of the uncertainties that arise when there is no explicit enumeration of known hazards or the potential causes of these hazards. Such uncertainties would be avoided by a systematic process to identify and define hazards, with the results recorded in the project documentation. Ideally, hazards would have been identified and defined at an early stage in the development of the BIS. Even at the current stage in the lifecycle of the BIS, the identification and definition of a set of hazards for the BIS will be useful for the purpose of assessing the safety impact to potential changes to the system over the lifetime of the LHC.

**S 10: By means of a systematic process, a set of hazards for the BIS should be identified and defined.**

The BIS project documentation describes many of the design decisions taken to eliminate or mitigate potential causes of hazards. However, the project documentation does not systematically identify the safety-related purpose of these design decisions. In general, the linkage between design decisions and hazard causes is not systematically recorded. In fact, it is mostly left for the reader of the project documentation to discover the relevance of individual design decisions to safety and to other properties such as availability.

For example, [11] offers the following statement about glitch filtering:

Glitches in the input signal are also to be removed; this is carried out by the ‘filter’ circuits. It is very important that the specification of this circuit be clear, as it performs an extremely delicate function, potentially delaying a beam dump request.

The reader can see from this short description that glitch filtering is relevant to safety (i.e., “potentially delaying a beam dump request”). However, this description does not provide linkage to a potential cause of the hazard. Without some knowledge about algorithms for glitch filtering, it is quite possible that the reader would wrongly infer from this brief description that the failure of glitch filtering to eliminate some glitches could cause a beam dump request to be delayed. In fact, the opposite is true. The safety relevance of glitch filtering is the possibility that glitch filtering might eliminate a “non-glitch”, i.e., an input which is not a glitch. Of course, with enough time and energy a sufficiently motivated reader will discover the causal link between a potential flaw in glitch filtering and the hazardous situation that a beam dump request is delayed. But this should have been documented explicitly.

Another example is the design to make each BIC invert the beam permit loop output, i.e., when the local permit is TRUE and the input from the beam permit loop is TRUE, then the output re-transmitted on the beam permit loop output is FALSE. It is not possible to determine from the project documentation the relevance of this design decision to a potential hazard cause. As before, with enough time and energy a sufficiently motivated reader will discover how this detail is relevant to safety.

Understanding the causal linkages between such details, along with dozens of other such details, should not be a process of discovery. Instead, such linkages should be explicitly documented.

For reasons similar to the need to identify and define hazards, explicit documentation of the linkages between design decision and hazard causes will be useful for purpose of assessing the safety impact to potential changes to the system over the lifetime of the LHC. Furthermore, the documentation of these linkages is the foundation of how a “safety case” could be developed for the BIS. Finally, the documentation of such linkages would very likely be useful to CERN and the rest of the worldwide accelerator community as a means of systematically transferring safety design knowledge to other future projects.

Appendix A of this report provides an illustration of one possible approach to addressing the following suggestion.

**S 11: The relationship between every safety-related design decision to one or more potential causes of a hazard should be documented.**

In addition to documenting design decisions that reduce safety risk, the project documentation should identify design decisions that increase exposure to safety risk. In addition to safety, the designers of complex systems such as the BIS typically need to satisfy other goals related to properties such as availability. For most complex systems such as the BIS, it is quite normal for some design decisions to increase exposure to safety risk. For example, the decision to include glitch filtering in the functionality of the matrix is motivated by the desire for availability rather than safety. This decision increases exposure to safety risk, but the BIS project documentation does not explain why glitch filtering is warranted or why it is acceptable from a safety perspective. In the case of glitch filtering, the argument that the glitch filtering, as it is implemented in each of Matrix A and Matrix B, is acceptable from a safety perspective is not trivial and depends on details that do not appear in the project documentation. For example, it depends on certain assumptions about the quality of the beam permits received by the BIS from the user systems. Once again, the motivated reader could discover why glitch filtering is warranted and possibly why it is acceptable. However, this should not be a matter of discovery. Instead, the project documentation should identify design decisions that potentially increase exposure to safety risk and provide an argument for why the increased exposure to safety risk is both warranted and acceptable.

**S 12: Design decisions that might increase exposure to safety risk should be documented along with an explanation for why the increased exposure is both warranted and acceptable.**

In summary, it is reasonable to expect CERN to have used a more systematic approach to identify potential sources of safety risk for the BIS, or in other words, it is reasonable to expect CERN to have done more to document hazards, their potential causes and the relationship between design decisions and the potential causes of hazards. It might be argued that CERN, along with the rest of the worldwide accelerator community, has had more than a half century of experience with

accelerator technology and, as a result, the hazards are so well known that a formal process to identify and understand hazards and their causes is unnecessary. However, such an argument may not take sufficient account of the fact that the energies involved in the operation of the LHC are significantly greater than the past experience of CERN and the rest of the accelerator community. While this past experience is extremely important, it should not be used in this instance as a reason to avoid a more formal process to identify and analyze hazards associated with the BIS.

### **Adequacy of Measures Taken to Eliminate or Reduce Known Sources of Safety Risk**

Although CSL has not performed a safety analysis of the BIS, CSL has examined the design and implementation of the BIS for the purpose of assessing the extent to which known sources of safety risk have been eliminated or reduced to an acceptable level. It must be emphasized that CSL has not performed a safety analysis of the BIS and that this report does not offer a conclusion about the safety of the BIS. Nevertheless, CSL has sought to understand as much as practically possible about the design of the BIS with respect to the safety goal of minimizing the possibility that the BIS continues to provide a TRUE beam permit to the beam dump system after one or more FALSE beam permits have been received by the BIS from user systems. Additionally, CSL has sought to understand as much as practically possible about the processes used to develop the BIS (e.g., verification and validation), processes used to maintain the BIS (e.g., configuration data) and relevant aspects of the use of the BIS by operators (e.g., pre-operational testing).

During the site visit by CSL to CERN as well as during subsequent follow-up interactions with CERN personnel, CSL has probed deeply into many details of the BIS design and the above mentioned processes. In particular, CSL asked numerous “what-if” style questions to assess how thoroughly CERN has anticipated possible failures as well as human errors. Without exception, key CERN personnel responsible for the development of the BIS were able to provide a very clear and satisfactory response to our questions. Relevant to the general experience of CSL with a variety of clients across different sectors of industry (e.g., aerospace, defense, rail, medical technology, advanced automotive control), it is very clear that the developers of the BIS have been remarkably thorough in their effort to anticipate possibly failures and human errors. Notwithstanding the previously mentioned concerns about the absence of a formal process for the identification and analysis of hazards, CSL is genuinely impressed by the thoroughness of the measures taken by CERN to eliminate or control potential failures and errors that could impact the safety of the BIS.

Section 4 of this report documents several concerns and observations about the adequacy of existing mitigations for known sources of safety risk. Several of the recommendations presented earlier in this report are directly relevant to the adequacy of these mitigations.

## **7. Assessment framework for future projects**

This section provides a sketch of a general framework that could be used by CERN to assess programmable electronic systems that implement critical functions related to machine protection. The framework is expressed in terms of objectives that leave considerable flexibility for the user of this framework to decide how the objectives may be best achieved. There are three groups of objectives, namely, process objectives, product objectives and operational objectives.

## **7.1 Process Objectives**

The following process objectives aim to determine if evidence is available that the system has been developed in a manner appropriate for a system of this level of criticality.

### **7.1.1 System concept and lifecycle is defined**

#### **7.1.1.1 System requirements are defined**

The intent of this objective is to specify the functional and non-functional requirements for the system in a manner that satisfies the specific needs of the project and gives a unique understanding to all project stakeholders.

Assessment criteria should include requirement properties such as completeness, consistency, unambiguousness and verifiability. The origin of safety related problems is often traced back to problems in requirement clarity, completeness and consistency.

#### **7.1.1.2 External interfaces are defined**

The intent of this objective is to ensure that all the external interfaces of the system are identified and described. The intent is to achieve a common understanding of an interface to the developers of the system as well as to the developers of the external systems.

Assessment criteria should include data format, data protocol and data timing.

#### **7.1.1.3 System design is documented**

The intent of this objective is to ensure that all the system requirements are allocated to hardware and software components. Some requirements may be allocated to several sub-systems.

Previously developed sub-systems should be identified. Identification of new technology should be identified. The selection of sub-systems may require analysis and / or justification.



#### **7.1.1.4 System integration is performed**

The intent of this objective is to ensure that all the various system components are integrated in a manner that satisfies the project needs. This objective aims to ensure the correct behaviour of internal interfaces.

Assessment criteria should include considerations for the normal behaviours of each internal interface as well as abnormal situations (loss of connection, intermittent failures).

#### **7.1.1.5 System verification and validation is performed**

The intent of this objective is to ensure that verification and validation activities are performed at the system level. These activities should include functional testing as well as performance testing and robustness testing.

#### **7.1.1.6 Integration with external interfaces is performed**

The intent of this objective is to ensure that a systematic integration of each external interface is performed.

Assessment criteria should include considerations for the normal behaviours of each external interface as well as abnormal situations (loss of connection, intermittent failures).

#### **7.1.1.7 Life cycle is considered**

The intent of this objective is to ensure that engineering change can be handled by the project during its lifecycle as well as during operations.

Assessment criteria should include the tracking of changes as well as the determination of the impact of a change. Examples of the nature of changes that may be considered could be a requirement change or the change of a selected hardware component.

### **7.1.2 Sub-system design and development life cycle is defined**

#### **7.1.2.1 Sub-system requirements are defined**

The intent of this objective is to specify the functional and non-functional requirements for a specific sub-system that satisfies the specific needs at the system level.

Assessment criteria for hardware sub-systems should include performance considerations as well as architectural considerations such as built-In test, environment, power, physical characteristics and interfaces.

Assessment criteria for software sub-systems should include off-nominal behaviours, performance considerations as well as architectural considerations such as fault-tolerance or downgraded modes.

### **7.1.2.2 Architecture and detailed design are documented**

The intent of this objective is to ensure that architecture and design activities are performed based on the sub-system requirements.

Assessment criteria for a hardware sub-system should include a justification of the technology / component and a reliability analysis.

Assessment criteria for a software sub-system should include aspects of fault-tolerance, consideration of Worst-Case Execution Time and robustness as applicable to the software sub-system.

### **7.1.2.3 Implementation is performed**

The intent of this objective is to ensure that a sub-system is produced based on the detailed design data.

Assessment criteria should include a repeatable implementation process that is documented and easily retrievable.

Assessment criteria for a programmable sub-system should include how vulnerabilities of the method used to program the sub-system (e.g., use of a programming language such as Java or a hardware description language such as VHDL) are addressed, e.g., a styleguide that constrains how the method may be used.

### **7.1.2.4 Verification is performed**

The intent of this objective is to ensure that for each phase of design and development of a sub-system there is a corresponding verification activity.

Verification may be performed via simulation, testing, analysis or reviews.

Assessment criteria for a hardware sub-system developed using HDL may include some test coverage criteria for HDL simulation tailored to the criticality of the project.

Assessment criteria for a software sub-system should include some test coverage criteria tailored to the level of verification and the criticality of the project.

### **7.1.2.5 Configuration is managed**

The intent of this objective is to ensure that an appropriate level of configuration management is used for any file or artifacts produced during the sub-system design and development lifecycle.

Beyond assurance of proper identification and version control of each artifact, the assessment criteria should also include the creation of baselines and the existence of a problem reporting system. Assessment criteria should be extended to cover the control of software installation and load on the target platform.

### **7.1.2.6 Tools are assessed and qualified**

A tool may be used to eliminate, reduce or automate some of the activities of the sub-system design and development lifecycle. These activities can be design / development activities or verification activities.

The purpose of the tool assessment and qualification is to ensure that the tool provides a level of confidence compatible with its intended usage in the lifecycle of the sub-system component.

A tool assessment should be performed and should identify if the tool output is independently assessed. If not independently assessed, a proper qualification may be required whether achieved through internal means or by the way of an audit of the tool vendor.

### **7.1.3 Safety Analysis**

#### **7.1.4 System safety is assessed**

The intent of this objective is to ensure that the safety aspects have been considered from a system perspective. In particular, hazards need to be precisely identified so that the scope of the subsequent analysis is well defined.

Assessment criteria include:

- Hazards are identified and defined
- Hazards are analyzed using appropriate technique and knowledge/expertise.
  - This analysis should result in
    - requirements that are flowed down to the sub-system levels
    - identification of all the components of the system architecture that contain functionality that might affect any control action or data item that is relevant to the hazard.

- Safety validation is performed.
- Conclusions are developed with argumentation, limitations.
- System life cycle safety consideration exist.

### **7.1.5 Sub-system safety is assessed**

The intent of this objective is to ensure that all the sub-system behaviours that can contribute to a hazard at the system level are identified and assessed.

For a hardware sub-system, this objective requires all hardware components that can contribute to a hazard be identified. As a result, for all the safety-critical hardware components, the design assurance process needs to address potential anomalous behaviors and potential design errors of these hardware functions.

For a software sub-system, this objective requires the identification of all the safety-critical functions. In addition this objective requires that an appropriate level of rigor is used for the development of such functions. For this purpose, the assessment criteria will ensure that:

1. All of the software functions that can contribute to a hazard are identified.
2. A development assurance level is associated to each safety-critical software function.
3. The software development lifecycle tasks for each software development assurance level are defined.

## **7.2 Product Objectives**

The following product objectives aim to determine and demonstrate certain qualities of the system / sub-system which contribute to the safety and reliability objectives.

These objectives are equally applicable to the system and its sub-systems.

### **7.2.1 Evidence of correctness is available**

The intent of this objective is to ensure that review, analysis, simulation or test evidence is available that demonstrate that the system is fit for its functional purpose.

Assessment criteria should ensure that the correctness evidence can be systematically reproduced.

Assessment criteria should ensure that correctness evidence can be traced to the corresponding requirements. In particular assessment criteria should ensure that correctness evidence can be traced to the requirements describing safety mitigation functions such as interlock, monitoring function, alarm, etc.

Assessment criteria should ensure that the coverage of the evidence is appropriate for the level of system / sub-system criticality. This is to ensure that the decision logic of the implemented system has been verified to a known degree of thoroughness. As well, this is to minimize the possibility that unintended functionality has been introduced at some level of specification or description.

### **7.2.2 Evidence of robustness is available**

The intent of this objective is to ensure that review, analysis, simulation or test evidence is available that demonstrate that the system / sub-system is robust and responds in a predictable way to abnormal inputs and conditions.

The assessment of this objective shall include, but not be limited to, the following criteria:

1. For each requirement, there are verification results for the response of the system to invalid inputs, including testing for “off-by-one” defects using boundary value analysis.
2. For each real-time requirement, there are verification results for the real-time response of the system under conditions that exceed the specified or expected maximum performance load of the system.
3. For each requirement related to the capacity of the system, there are verification results for the response of the system under conditions that exceed the specified or expected maximum capacity of the system.
4. There are verification results for the behaviour of the system for an interval of continuous operation that exceeds the specified or expected maximum interval of continuous operation.
5. There are verification results for the start-up of the system in a state when initialization data is corrupted, stale or incomplete.
6. If the system includes a capability to switch from a primary to “hot” standby processor, there are verification results for the behaviour of the system after a “switch-over” in a state when the state of the new primary is corrupted, stale or incomplete.
7. For each of the external interfaces of the system, there are verification results for the response of the system when the interface has failed, including intermittent failures.

8. For each of the external interfaces of the system, there are verification results for the response of the system when a previously disconnected interface is re-connected.
9. If the system has degraded modes of operations, there are verification results for the response of the system for every possible transition between different modes of operation, including restoration of the system to its normal operating model.
10. If the system contains multiple clocks (e.g., a distributed system where each target machine has its own clock), there are verification results for the behaviour of the system when the clocks are not fully synchronized.

### **7.2.3 Evidence that product non-conformance / problems are tracked**

The intent of this objective is to minimize the possibility that a known defect or other problem is forgotten or left unresolved. If an identified defect is not reported and/or left unresolved because “it is unlikely to happen in the real environment” there is a chance that a similar condition may arise as the result of a set of circumstances unforeseen during development (e.g., in a slightly different environment) causing an unsafe situation.

Assessment criteria shall ensure that problems are recorded and tracked by a problem/defect tracking system. In addition assessment criteria shall ensure that a defined process exists for the resolution of (suspected) defects and problems that specifies acceptable resolution conditions.

## **7.3 Operational Objectives**

The operational objectives aim to determine if the system will be used in its intended manner.

### **7.3.1 System initialization is controlled**

The intent of this objective is to ensure that:

- All of the prerequisite conditions for normal operations are verified before start-up
- all of the sub-systems and system initialization steps are performed in a predictable sequence.

The assessment criteria should ensure that prerequisite conditions for normal operations are identified. The assessment criteria should ensure that the start-up

process of safety-critical sub-systems has been properly reviewed and has a level of predictability commensurate with the level of criticality of the system.

### **7.3.2 Static data is protected**

The intent of this objective is to ensure that static data used by the system during its execution is treated with the same degree of rigor as if the same behaviour was implemented in a software or hardware function. In this regard, the term “static data”, sometimes called “adaptation data” in some organizations, refers to data that exists prior to system startup and does not change during execution of the system.

This objective recognizes the fact that off-line data and site configuration data determines the functionality of a system as much as software or hardware, and therefore, should be subjected to the same controls.

The assessment criteria should ensure that any element of static data that has the potential to affect the behavior of the system is identified. In addition assessment criteria should ensure that each element of static data has a corresponding verification process and that each element of static data follows the organization configuration management rules.

### **7.3.3 Guidelines for periodic maintenance operations are defined**

The intent of this objective is to ensure that any periodic maintenance operations at the hardware level do not compromise the integrity of the system.

The assessment criteria should ensure that maintenance procedures or guidelines exist, are communicated to the maintenance personnel and that evidence of maintenance operations is available.

### **7.3.4 Safety assessment process for system change is defined**

The intent of this objective is to ensure that changes to the system between phases of operation or during operation are controlled and assessed from a safety perspective. Changes could include a new interface, the replacement of equipment with newly designed equipment, and/or the use of a new version for a given hardware device.

The assessment criteria should ensure a safety analysis process exists to assess and verify any design change to a system in its operation phase.

## **8. Summary**

Within the limitations of this review with respect to scope and resources, CSL has completed a thorough review of the BIS design. CSL concludes that there is sufficient reason to be confident that the BIS will perform its intended safety function.

The design of the BIS is a product of very impressive engineering skill combined with very substantial knowledge about machine protection.

This report offers a total of ten recommendations that should be addressed before the LHC resumes operation in late 2009. Additionally, the report offers twelve suggestions for improvement that could be addressed with less urgency.



## Appendix A

This appendix presents an illustrative example of one possibility of how Suggestion 11 could be addressed by CERN. In particular, it illustrates an analysis of the relationship between potential failure modes of the CERN LHC Beam Interface System (BIS) and design decisions that mitigate safety risk. The contents of this appendix are strictly illustrative. This information contained in this appendix is incomplete and unverified. No conclusion about the safety of the BIS should be derived from the contents of this appendix without verification of its accuracy and completeness.

One aim of the analysis documented by the table shown below is to identify possible failure modes of the BIS that could adversely affect the safety function of the BIS, i.e., to expeditiously propagate a request by a user system to dump the beam to the beam dump system. The other aim of this analysis is to identify design decisions or possibly other external factors that mitigate the failure mode by either elimination of the failure mode, reducing the likelihood that an instance of this failure mode could occur or reducing the likelihood that an instance of this failure mode could prevent, delay or otherwise adversely affect the propagation of beam dump request. An additional step to this analysis, which is not shown in the table, is to assess the adequacy of the mitigations for each failure mode.

Ideally, this analysis would have been developed iteratively in parallel with the development of the BIS. CSL is quite sure that the developers of the BIS have followed the “thought process” that underlies this style of analysis throughout the development of the BIS. Suggestion 11 of this report is simply motivated by the fact that there is no documentary record of the output of this thought process. Even at this stage in the lifetime of the BIS (when it has been deployed), it is important and useful to document the relationship between failure modes and mitigations, especially to aid in the assessment of the safety impact of any changes to the BIS that might be proposed at some point in the remainder of its lifecycle.

As shown by the table below, the analysis is developed systematically by tracing “backwards” throughout the pathway between the point where user systems connect to the BIS and the point where the BIS provides an output to the beam dump system. The direction of this analysis, i.e., “backwards”, is not significant. The analysis could also have been developed by tracing in a “forward” direction through the same path. One way or another, the analysis should be developed systematically.

The most important content of the table shown below is contained in the columns labeled “failure mode” and “mitigations”. The column “potential causes” is supplementary information that only serves to help the reader imagine how an instance of the failure mode might occur. For the purposes of this analysis, it is unnecessary to exhaustively identify all of the potential causes of each failure mode. (However, another analysis to more thoroughly identify potential causes may also be warranted as part of an overall safety process.)

In reference to the state of a beam permit loop, the word “TRUE” is used to describe when a signal with a frequency within the valid range of frequencies is present on the beam permit loop, e.g., “the beam loop is TRUE”. Otherwise, the word “FALSE” is used to describe when a valid frequency is not present on the loop. Similarly, the words “TRUE” and “FALSE” are used to describe the state of an input from a user system where “FALSE” means that the user system is requesting a beam dump; otherwise, the state of the input from the user system is TRUE.

For the purposes of this analysis, it is assumed that a beam dump will be initiated for Beam X whenever one or both of the beam permit loops for Beam X is FALSE.

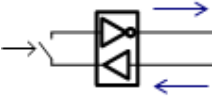
The fully redundant channel from user system to the beam dump is an implicit mitigation for every failure mode identified in the following table. For the sake of conciseness, it is not explicitly shown in this list of mitigations for each failure mode.



# Critical Systems Labs

Strategic Insight for Safety

ID	BIS component	Failure	Some Potential Causes	Mitigations
1.	Beam permit loops	The local permit of at least one BIC is false for the beam. However, a valid frequency is generated (by some unknown means) at some downstream point on the loop between this BIC and the beam dump.	None known	<ol style="list-style-type: none"> <li>1. There are no known sources of a valid frequency for this beam other than the frequency generator at start of the loop.</li> <li>2. The only known failure mode (i.e., signal is not propagated at a valid frequency) will result in a beam dump, which is a safe reaction.</li> <li>3. The frequencies used for each loop are unique, i.e., the frequency of each loop is not a resonant frequency of any other known frequency in the LHC (except for the corresponding loop of the BIS for the other beam).</li> <li>4. The frequency used by each loop is different and not overlapping with the other loop for this beam.</li> </ol>
2.	Same as above	Similar to above	The optic fibre cables are connected to the wrong connection points such that the cable meant to carry the output of Matrix A for Beam 1 is mistakenly connected	<ol style="list-style-type: none"> <li>1. While there is no direct prevention of this possibility (such as opposite gender connections), it seems impossible that the BIS could be successfully armed, especially if the arming sequence for each beam is not performed simultaneously.</li> </ol>

			to the connection point for the output of Matrix A for Beam 2 (rather than Beam 1). Or some other analogous misconnection.	
3.	Optical Receiver	Electrical output does not consistently track optical input, e.g., stuck at 1, stuck at 0, output varies	Manufacturing defect	<ol style="list-style-type: none"> <li>1. Proven technology (?)</li> <li>2. If output is not tracking input consistently, there is no known way in which the output could produce a signal within the valid frequency range.</li> </ol>
4.	Matrix CPLD	<p>The switch allows the input to pass to the output when the local beam permit is FALSE.</p> 	Damage to CPLD due to radiation	<ol style="list-style-type: none"> <li>1. Proven technology.</li> <li>2. Pre-operational testing should only pass if the switch in both Matrix A and Matrix B of each BIC is observed to open such that the re-transmission of the beam permit is interrupted. (?)</li> <li>3. FPGA Monitor monitors beam permit output (?) and compares with this with the local beam permit (?).</li> </ol>
5.	Same as above	Failure in the connection from the matrix to the switch, i.e., output of matrix is FALSE, but input to EO switch is TRUE		<ol style="list-style-type: none"> <li>1. The connection is designed to fail safe, i.e., it is connected to PWR through a resistor so that it is "pulled up" if there is failure in the connection. As this is active low logic, this requests a beam dump as a safe reaction.</li> </ol>
6.	Matrix CPLD	CPLD does not respond to changes to inputs from user systems	Clock fails, i.e., CPLD halts computation	<ol style="list-style-type: none"> <li>1. Clock is monitored by the Monitor FPGA which should detect when the clock fails. Then DIAMON would raise an immediate alarm. In turn, the SOFTWARE_PERMIT to stop the propagation of the frequency, as this is not registered. This signal has a combinational path to the switch inside the CPLD.</li> </ol>

7.	Same above	as	Reset condition is TRUE during operation	None known	<ol style="list-style-type: none"> <li>1. Output is forced to fail safe value, i.e., beam dump is requested, during reset.</li> <li>2. Several independent conditions must all be TRUE before a reset will occur (?).</li> </ol>
8.	Same above	as	Glitch filtering causes beam dump request from user system to be filtered	Potential design error	<ol style="list-style-type: none"> <li>1. Static and dynamic verification during development</li> <li>2. Anomalies in monitoring data detected during post-operational analysis (?)</li> <li>3. Interface requirements for user systems</li> </ol>
9.	Same above	as	Glitch filtering causes beam dump request to be delayed	Potential design error	<ol style="list-style-type: none"> <li>1. Static and dynamic verification during development</li> <li>2. Anomalies in monitoring data (?)</li> <li>3. Interface requirements for user systems</li> </ol>
10.	Same above	as	Glitch filtering causes beam dump request to be truncated, i.e., FALSE then TRUE again before FALSE is detected by beam dump	Potential design error	<ol style="list-style-type: none"> <li>1. Whenever a BIC outputs FALSE, the BIC maintains the FALSE until the BIC is re-armed regardless of changes to other inputs.</li> <li>2. Static and dynamic verification during development</li> <li>3. Anomalies in monitoring data detected during post-operational analysis (?)</li> <li>4. Interface requirements for beam dump</li> </ol>
11.	Same above	as	Inputs or intermediate results of processing corrupted while inside CPLD	Single bit upset	<ol style="list-style-type: none"> <li>1. Most of the data in the CPLD is transitory; for most data, the maximum duration of its existence in the BIS after possible corruption is a relatively small number of microseconds</li> </ol>
12.	Same above	as	Aside from glitch filtering, other CPLD processing causes beam dump request to be delayed		<ol style="list-style-type: none"> <li>1. Static and dynamic verification during development</li> <li>2. Anomalies in monitoring data detected during post-operational analysis (?)</li> </ol>
13.	Same above	as	The function performed by the CPLD is different than the function specified by the	A single bit upset occurs while the BIS is	<ol style="list-style-type: none"> <li>1. Depending on the nature and severity of the corruption, anomalous behaviour of the CPLD might</li> </ol>

		currently authorized VHDL	armed and operational causing the function performed by the CPDL to deviated from its VHDL source	be detected by the Monitor FPGA such that the SOFTWARE_PERMIT is denied.
14.	Same as above	Same as above	Human error, such as loading the wrong version of the "code" onto the CPLD	<ol style="list-style-type: none"> <li>1. Each version of the "code" includes a version ID that should be unique to this version. This version ID is compared during pre-operational testing to the version specified in the configuration database.</li> <li>2. The task of loading the "code" into the CPLD is labour-intensive, i.e., must be done individually for each BIC at the location of each BIC (not over a network). The labour-intensive nature of this task will lower exposure to risk due to a single human action (as would be the case if a new version of the code could be downloaded at once by a single human action over a network).</li> </ol>
15.	Same as above	Same as above	Human error, such as a "quick fix" to the VHDL without thorough analysis, verification and validation.	<ol style="list-style-type: none"> <li>1. Each version of the "code" includes a version ID that should be unique to this version. This version ID is compared during pre-operational testing to the version specified in the configuration database.</li> <li>2. The task of loading the "code" into the CPLD is labour-intensive, i.e., must be done individually for each BIC at the location of each BIC (not over a network). The labour-intensive nature of this task will lower exposure to "quick fixes" or attempts by authorized personnel to modify the code.</li> </ol>
16.	Safe beam	Maskable input is ignored when	Safe beam flag is TRUE	1. <i>Additional analysis of the safe beam flag pathway</i>

	flag	beam is operating at energy level capable of causing harm	when it should be FALSE	<i>should be performed tracing from the use of the safe beam flag by the CPLD back to the point at which the safe beam flag is provided to the BIS. This could be a separate analysis or an elaboration of this analysis.</i>
17.	BIC User Interface	Wrong input is masked	Non-maskable input is connected to an input connection point for a maskable input	<ol style="list-style-type: none"> <li>1. Masking only affects output of the BIS when the beam is not operating at an energy level that could cause damage, assuming that the safe beam flag is valid.</li> <li>2. Physical assess to connection points is limited, which reduces exposure to this risk, e.g., less likely that someone will accidentally disconnect and then incorrectly re-connect an input.</li> <li>3. Pre-operational testing includes checking that non-maskable inputs are not connected to a connection point for a maskable input.</li> <li>4. Additional analysis of the pre-operational testing should be performed to search for potential failure modes that could affect the integrity of this checking.</li> <li>5. Additional analysis of the procedures for maintaining and protecting the configuration database should be performed to search for potential failure modes that could affect its accuracy.</li> </ol>
18.	Same as above	Wrong input disabled	Erroneous jumper connection	<ol style="list-style-type: none"> <li>1. Physical access to CIMB is very limited which reduces opportunities for accidental or unauthorized changes to the configuration of</li> </ol>

				jumper cables.
19.	Same as above	Same as above	Input that should not be disabled is connected to a disabled connection point	<ol style="list-style-type: none"> <li>1. Physical assess to connection points is limited, which reduces exposure to this risk, e.g., less likely that someone will accidently disconnect and then incorrectly re-connect an input.</li> <li>2. Pre-operational testing includes checking that inputs are connected to the correct correction point, as per the configuration database.</li> <li>3. Pre-operational testing includes checking that no connection point is erroneously disabled, as per the configuration database.</li> <li>4. Additional analysis of the pre-operational testing should be performed to search for potential failure modes that could affect the integrity of this checking.</li> <li>5. Additional analysis of the procedures for maintaining and protecting the configuration database should be performed to search for potential failure modes that could affect its accuracy.</li> </ol>
20.	Same as above	Input from user system is not propagated from connection point to the CIBM	Loss of electrical connectivity	<ol style="list-style-type: none"> <li>1. Fail safe design of the circuit (e.g., for active low logic, circuit is connected to PWR through a resistance) causing a beam dump request as a safe reaction.</li> <li>2. The circuitry is designed at an electrical level so that there is a "narrow window" of voltage values for which USER_PERMIT = TRUE, i.e.,</li> <li>3. There is some fault detection such that detection of</li> </ol>



				fault will cause USER_PERMIT to be FALSE, which is fail safe.
21.	Same as above	Same as above	Electrical damage or interference caused by user system (e.g., such as UJ33 incident)	1. Review process in which the BIS group inspects the user system to check for possible sources of harm to the BIC.
22.	Same as above	Same as above	Electrical short that forces the connection to TRUE	1. ??
23.	Same as above	Same as above	Aside from loss of electrical connectivity or short circuit, some other interference with the propagation the input	1. Simple design implemented using highly reliable technology with well known performance characteristics.
24.	Same as above	Input from user system for this beam is connected to the BIS for the other beam		<ol style="list-style-type: none"> <li>1. Connection points for one beam are the opposite gender of the connection points for the other beam, making this physically impossible using existing cables.</li> <li>2. Physical assess to connection points is limited, which reduces exposure to this risk, e.g., less likely that someone will accidently disconnect and then incorrectly re-connect an input.</li> <li>3. Pre-operational testing includes checking that inputs are connected to the correct correction point, as per the configuration database.</li> </ol>
25.	Power Supply	One power supply fails		1. There is a redundant power supply that seamlessly continues to provide power to the BIC

26.	Frequency Generator	Frequency Generator continues to re-generate the signal even though the frequency is not present on the input (while the system is operational)	External condition that forces this input to be TRUE when the system is armed remains TRUE, thus masking FALSE received from any of the BICs. (In particular, consider the case when the beam dump is upstream of BIC requesting beam dump.	1. Is there something in the BIS design that would detect this? Otherwise, the only mitigation seems to be the other loop.
27.	Entire BIS for one loop of one beam	BIS becomes operational in a faulty or otherwise unready state	Pre-operational test not fully performed	1. As currently implemented, the command to arm the system is built-in the pre-operational test procedure. However, this procedure could be modified without the approval or knowledge of the BIS team.
28.	Same as above	LHC becomes operational when BIS is not operational	BIS provides operators with a false indication of being armed	1. If the BIS is not operational, there is no beam permit and so it should be impossible to even start injecting beam into the LHC
29.	Same as above	Same as above	Operator use (temporary) capability to directly arm the system rather than using the pre-operational test	1. A potential future mitigation (which should be implemented before the LHC resumes operation with beam) is to remove this capability in such a way that it would be very difficult for someone to re-introduce this capability in the future. However, this mitigation does not currently exist.
30.	Entire BIS for both loops of	Some other general failure of the BIS that is not covered by above rows and that totally	unknown	1. The BIS for one beam provide some limited redundancy for the BIS of the other beam. For the normal operation of the LHC (i.e., when doing

	one beam	impairs the ability the ability of the BIS to propagate a beam request on either loop.		physics), the BIS for both beams will be armed. Some of the conditions that should trigger a beam dump for one beam will also likely trigger a beam dump for the other beam, e.g., beam loss detection. However, in general, the interlocking provided by the BIS for one beam should be considered separately from the interlocking provided by the BIS for the other beam.
--	----------	--	--	--