LHC OPERATION: THE HUMAN RISK FACTOR

A.L. Macpherson, CERN, Geneva, Switzerland

Abstract

Issues associated with the human risk factor for the machine protection and operation of the LHC are discussed, with examples taken from the 2010 run. Emphasis is placed on risk factors that are present in the current modus operandi, and areas of improvement, both procedural and otherwise, are addressed. In addition, the The potential sources of human risk factors that lie outside the standard operations envelope and protective procedures are also considered.

INTRODUCTION

This paper takes a look at the human factors in LHC operation and discusses the human risk factors both for LHC operation and for machine protection. Given that at time of writing, the very successful 2010 run has only recently finished, the focus of this paper is on universal human risks factors and observations from the 2010 run rather than attempting to provide a list of operations errors from the first full a very year of running.

Human risk factors in LHC operation can take a variety of forms and can cause a wide range of issues ranging from weaknesses in the machine protection system to loss in operational efficiency through to risk oriented behaviour or operational mistakes. For machine protection the key issue with the human factor is whether the shift crew can damage the damage the machine. Clearly, for LHC operation, it is essential for the shift crew to exercise full control over the LHC and its systems, so by default, the possibility exists for the shift crew to drive the machine to a working point outside the machine protection envelope. However, the machine protection system, the operational procedures, the expertise of the shift crew, and the attention to the human factor greatly mitigate this risk.

HUMAN RISK ASSESSMENT CULTURE

When dealing with the human risk factors for LHC operation, the goal is not minimise risk by a post-problem reaction or pathological culture, but rather, by instilling a clear proactive risk assessment culture that respects, anticipates and responds to risks. This notion of a developing human risk assessment culture is one that is adopted in disciplines such as the nuclear and aeronautics industries, and can be defined in five broad categories[1]:

- GENERATIVE: Respects, anticipates and responds to risks. A just, learning, flexible, adaptive, prepared & informed culture. Strives for resilience.
- PROACTIVE: Aware that 'latent pathogens' and 'error traps' lurk in system. Seeks to eliminate them beforehand. Listens to 'sharp enders'.
- CALCULATIVE: Systems to manage safety, often in response to external pressures. Data harvested rather than used. 'By the book'.
- REACTIVE: Safety given attention after an event. Concern about adverse publicity. Establishes an incident reporting system.
- PATHOLOGICAL: Blame, denial and the blinkered pursuit of excellence (Vulnerable System Syndrome). Financial targets prevail: cheaper/faster.

The task for an effective human risk assessment culture is to evolve toward a Generative culture the promotes resilience, where resilience is defined as the ability of a system to adjust its functioning to sustain operations during expected conditions and in the face of escalating demands, disturbances, and unforeseen circumstances [1].

To assess the human risk factors associated with LHC Operation in 2010, a preliminary survey of the post mortem data and logbook statistics can be made. From approximately 500 global post mortem events that occurred over the last 4 months of running, 204 were with beams above injection energy. From these 204 events only 8 were classified as operational errors; a 4% rate of operational errors that led to beam dumps. These beam dumps were typically provoked either by hidden interlocks which were not cleared prior to the setup beam flag energy threshold being reached during the ramp, or the incorrect configuration of a setting during the commissioning with beam.

In addition to this, the logbook reveals a number of instances where operational irregularities also resulted in beam dumps. Examples from the logbook include accidentally switching off with the Equip State application and playing the wrong squeeze function in the squeeze.

What is clear is that in the 2010 run there were more operational errors than were documented, and a significant fraction could be associated with human risk factors. Unfortunately a significant number of the operational errors went untagged, thereby making it difficult to get a representative assessment of the human risk factor. However, it is reassuring that to date, the operational errors incurred have been caught by the machine protection system and dumped the beam immediately. This reduces the risk of damage to the machine but does not completely remove the risk of damage due to the risk from operational errors coupled with an asynchronous beam dump or an equipment failure. It is insufficient to rely solely on the hardwired machine protection system, and it is clear from the 2010 run that improvements can be made in the culture of human risk assessment.

NORMAL OPERATION

Normal LHC operation is defined in terms of a nominal operational procedure, which is mapped to a nominal LHC operational sequence. However, this is not a one-to-one mapping, as not all steps in the nominal procedure can be encapsulated in the nominal sequence. This then opened up several possibilities for human risk factors.

- Not all tasks are integrated into the sequencer, so there was the risk of tasks not being done. e.g. running through the collision beam process without switching off the tune and orbit feedbacks.
- Missing or skipping required steps in the nominal sequence.
- Playing an out of date sequence.
- Resorting to special procedures or workarounds that only have a limited duration validity or are not well documented.

In an attempt to curb these types of errors, an LHC State Machine (based on machine protection guidelines and the Beam Mode states) has been developed and is to be deployed for the 2011 run.This state machine will work in conjunction with the LHC sequencer and will help enforce that there is an adherence to the nominal procedure and that tasks are not performed out of order. Also, as part of the state machine, there is an incorporated checklist view that allows an overview of the performed task within a given state. This state machine should aid aid the shift crew in ensuring that all the required task have been performed before a state transition is performed.

Yet there is still the risk of that the wrong commands are sent or that a trim is too large and moves the working point outside the machine protection envelope. Such situations are difficult to catch automatically, as it is primarily an issue of operator competency. As seen from the 2010 run, the level of operator competency is extremely high, but that for whatever reason such errors have occasionally crept in. At present, the way to programmatically combat these errors is to implement settings checks and validation on operator initiated write commands. This can at best be only partially successful, as it is difficult to define a machine protection envelope that covers all the operational phases of the machine, without becoming so restrictive that the operations flexibility of the shift crew is compromised. The process of defining a machine protection envelope will continue in 2011.

NON-STANDARD OPERATION

The possibility for human risks in operation is naturally increased when there is need to move away from standard operational procedures. In particular, two specific cases were identified: the use of special interim procedures for the resolution of short term problems, and the use of low level applications at the operations level.

In 2010 the first case was highlighted with the case of bent RF fingers causing an obstacle in the beam pipe in at the end of the beam 1 injection. In order to avoid this obstacle steering was performed in the transfer line and the obstacle was successfully bypassed, but the steering induced significant injections oscillations. However over time the obstacle drifted and the steering had to be adapted, which resulted in unacceptable injection oscillations. For this case, there was no a clear definition of an operational envelope, and as the initial steering was set up at the limit of tolerable injection oscillations, the was no margin for fluctuations or for diagnostic probing of the problem by the shift crew.

As an example of the latter case, the use of the Equip State application is mentioned. Equip State is a low level application that allows the operator to directly set properties on the hardware, and the is not machine protection check on the settings being sent. This, coupled with the fact that some of the naming conventions for beam processes and setting are not always obvious or adhered to, means that there is a real risk of sending the wrong settings. It is only the vigilance of the operators that prevents such errors (e.g. when changing collimator settings during loss maps etc). In the 2011 run, when there is beam in the machine, the access of low-level applications such as Equip State is to be restricted or if possible, prohibited.

APPLICATIONS AND CONTROLS

Human risk factors in LHC operation are not solely linked to the LHC Operations team, but are also related to the LHC applications, controls interfaces and experts.

For the applications and controls interfaces, there is an obligation to present operational information at the top level in a clear and understandable way. In the 2010 run there were occasions where the information from an application was not clear yet was needed in order for the shift crew to react to a beam related problem. The loss of the tune feedback during the squeeze due to large coupling is a good example, as the tune feedback application gave significantly different coupling values depending on the tune fitter filter selected, leaving the shift crew unsure of the actual value of the coupling. This is an example of an extremely powerful application that sometimes failed to clearly deliver the information needed by the shift crew.

In addition to the presentation of monitoring data, there is also need for clarity in design and layout of setting controls in applications. Having a clear and responsive control interface is needed both for routine operation and for situations where immediate response is needed. For the risk from control interfaces in 2010, the proximity of the ON/OFF buttons in the Kicker application (normal operation) and the slow response and poor state selection of the tune feedback fixed display are examples where the interface can be improved.

From an operations point of view, it is clear that applications should provide an operations view, but they should also allow for an expert view. However, in order to reduce the risk of operator error the two view should remain separated, and where possible, both views should be documented. It is also crucial that after a commissioning or machine development session, equipment experts re-establish the operations view and do

not leave unvalidated settings or configurations in the operator applications running in the CCC. When this happen in 2010, it only helped to complicate the diagnosis of problems.

Included in this issue of settings and configurations is the updating of front-end firmware, which at present is not controlled by the standard RBAC security checks[2]. Standardisation of firmware version tracking is not foreseen for 2011, and so the minimisation of risk from this source relies on clear communication between equipment teams and the operations team, and well prepared scheduling of updates.

As part of the issue of information transparency for the operations, one key issue is the presentation of alarm information through the LASER and DIAMON applications [3]. For the 2010 run the operations team did not have a clear picture of the alarms information and alarm flow from the LASER system simply because it was swamped with alarms. This made the monitoring of problems via LASER untenable, and as such greatly reduced the ability of the shift crew to respond to warnings an alarms flagged within the LASER system. For 2011 it is imperative that the alarm definitions be cleaned up and here the responsibility lies primarily with the equipment teams, but also with operations.

Similarly, the DIAMON application which is used to diagnose and monitor and front end servers, the 2010 run showed that the configuration of alarms within DIAMON is not yet optimised, and in addition that the operations monitoring view was not restricted to just the operational front-ends (i.e. it also included non-operational front-ends, which often showed alarms, and so made the monitoring of real alarms from operational front-ends difficult). Again the clean up of the DIAMON configuration lies primarily with the equipment teams.

COMMUNICATION

One of the primary areas for improvement that has been identified from the 2010 run is the area of communication and coordination. The lack of clear communication and coordination across the operations teams can and has resulted in a direct increase in the risks of human error and the potential for jeopardising the machine protection envelope. Lack of clear communication can create inconsistencies at the program level that can be consistent at the level of individual tasks, but may result in an overall working point that is outside the machine protection envelope.

In terms of communication, it is essential that a clear line of communication and chain of command be maintain between the machine coordinators, Engineers in Charge, and LHC operators, so that the programme is clear and the operational steps co-ordinated and well defined. As was seen in 2010 this line of communication needs to extend not only to the LHC but also to its injectors, the technical infrastructure, and the cryogenics shift crews, to avoid misunderstandings that unnecessarily stress the machine protection system.

As part of the communication issue there needs to be an improved passage of information and summary of decisions taken during the 8:30am meetings to the shift crews. Once the program is clear, it is also necessary that people in the LHC island respect the defined roles of the LHC operators and Engineers in Charge and permit them to carry out their functions, as it is the shift crew that is responsible for the safe an efficient running of the LHC during the shift.

OTHER FACTORS

In addition to all the above mentioned sources of human risk factors, there are other factors that can potentially affect machine protection, and these are the environmental factors. Environmental factors cover a wide range of topics ranging from:

- Working conditions in the CCC
- Operator fatigue
- Unbalanced work loads across the equipment and piquet teams
- Unnecessary pressure for fast turnaround times and rapid re-establishment of stable beams.
- Simple typing mistakes due to too many keyboards in the LHC Island.

For these environmental factors the responsibility to minimise there effect lies solely with the operations team, and as seen from the 2010 run, the influence of such environmental risk factors is being progressively reduced.

REDUCTION OF HUMAN RISK

The first step in reducing human risk factors is to realise that we are moving from a beam commissioning period into one of routine operation, and as such there is a need to tag instances of operational errors, in order to gather statistics and the analyse the manor and degree of the human risk factors. Implicit in this is the commitment from the operations team to tag any operation situations that involve error or risk, and also the support the management team in addressing operational errors so that a real human risk assessment culture can evolve.

As we move to routine operation, the robustness of the machine protection system is to provide the first line of defence against human error, such that deviations from normal operational procedure will initiate a beam dump. Beam conditions should then only be re-established once the reasons for the deviation are understood. In this way a more comprehensive machine protection envelope will be developed.

To aid in the reduction of operational errors, the operations team needs to build on the experience from the 2010 run, refine the machine protection envelope, and increase the degree of self assessment and evaluation of the operational procedure. This coupled with a balanced shift load, and clear lines of communication will help in reducing the operational errors as well as further help moving the LHC risk assessment culture from a Calculative level toward a fully Proactive and Generative human risk assessment culture.

REFERENCES

[1] "Human Risk Factors Culture" J. Reason; Workshop on Management of Human Factors Risk in Safety-Critical Industries, 11/05/2006. http://www.raes-hfg.com/reports/...hfrisk/11may06-hfrisk_sample.htm

[2] Role-Based Access Control for the Accelerator Control System at CERN. P. Charrue et al. Proceedings of ICALEPCS07, 15-19/10.2007. http://neutrons.ornl.gov/conf/icalepcs07/

[3] http://ts-project-tim.web.cern.ch/ts-project-tim/doku.php?id=documentation:daq:laser

The DIAMON Project - Monitoring an Diagnostics for the CERN Controls Infrastructure. M. Buttner et al. Proceedings of ICALEPCS07, 15-19/10.2007. http://neutrons.ornl.gov/conf/icalepcs07/