

The ATLAS Alarm Helper

Florian Haslbeck^{1,2,*}, *Carlos Solans Sánchez*¹, *Ignacio Asensi Tortajada*¹, *Daniela Bortoletto*², *André Rummler*¹, and *Gustavo A. Uribe*³

¹CERN, Geneva, Switzerland

²Department of Physics, University of Oxford, United Kingdom

³Universidad Antonio Nariño, Colombia.

Abstract. The Detector Safety System is the last line of defence to protect the ATLAS detector against abnormal and potentially even unforeseen situations. It is designed to return the detector to a safe state based on predefined actions triggered by alarms which are triggered on their part by specific sets of conditions. Every alarm whether it results in an action taken or not is followed up by the operations team that assesses the criticality, takes countermeasures and identifies the point of failure. From experience abnormal situations can result either from faults or from side effects of planned interventions which were either not properly identified despite the mandatory planning and review or where a mistake during execution occurred. In many cases there are multiple interventions ongoing simultaneously in order to profit from shutdown periods. The rapid analysis of alarms while the incident is ongoing is often complicated due to the complexity of the ATLAS detector and its infrastructure and the large number of responsible groups and experts.

A new Alarm Helper tool was designed to assist the operation team, particularly the operator in the control room responsible for infrastructure and safety (SLIMOS – Shift Leader in Matters of Safety), by providing real-time information about ongoing interventions and the possible related causes of failure. The new tool will combine historical events, documentation, and limited knowledge about ongoing interventions. It extends the Expert System which visualizes and simulates infrastructure inter-dependencies and allows to trace faults or alarms to a list of potential points of failure. The new tool also proposes which experts should be contacted in the particular circumstances.

1 Introduction

The operation of the ATLAS Experiment [1] requires handling of sudden and sometimes even unforeseen failures including those which could potentially compromise the safety of the detector. In order to avoid delays due to human reaction time under circumstances when time is critical an automatized and autonomous system called Detector Safety System (DSS) [2] is used which is connected to all parts of ATLAS. DSS is based on Programmable Logic Controllers (PLC) and it returns the detector into a safe state by taking predefined actions. For example, DSS can interlock the power to a rack upon receiving a fault signal from a cooling system that provides cooling to this rack. At the same time it alerts the operators both

*e-mail: florian.haslbeck@cern.ch

on-site and off-site by issuing alarms. Each alarm indicates that an abnormal or even critical condition was detected and must subsequently be followed up immediately with the goal to achieve again nominal operation conditions. The ATLAS operators must infer the cause of the alarm and identify the point of failure to assess the criticality of the situation. Actual faults must be distinguished from harmless alarms that were triggered by interventions, such as the restart of equipment as those are signaled as alarms but do not threaten the detector. Part of the operation procedure is that planned interventions are presented and discussed in the weekly operation meeting in order to evaluate which alarms could be triggered during the upcoming interventions and hence to inform the whole operations team. Despite the best effort of all stakeholders sometimes interventions and their impact on other systems, especially those mitigated by DSS, are not always fully understood. In such cases the cause of the alarm and potential failure can only be identified afterwards. The analysis of all alarms is presented by DSS and other experts during the weekly operation meeting or in case of severe incidents special meetings are called. A quick follow up is complicated by the large number of different groups responsible for the various parts of the detector which are involved in the planning, execution and review of interventions; the complexity of the ATLAS detector infrastructure as well as the loss of institutional knowledge due to turn over. The current procedure has the disadvantage of limited available knowledge depending on individual experience and the reporting only one week later which also results in a number of false negative alarms occurring due to interventions.

The operators in the control room use various information sources and communication channels to analyse an alarm and to assess its criticality including the use of the Expert System and direct contact with members of the ATLAS operations team and experts which are off-site. The Expert System [3–5] targets to describe as many domains of the ATLAS infrastructure as possible, namely detector control and safety systems, gas, water, cooling, ventilation, cryogenics, and electricity distribution by an object orientated abstraction layer that has been rule based validated thoroughly and is constantly updated and extended. The database currently describes more than 13000 elements with more than 89000 relationships. The graph-based inference engine allows propagating a state of one or multiple objects across the simulated infrastructure. For example, the Expert System enables simulation of the impact on the detector and its infrastructure when several racks are switched off or when a specific DSS alarm is triggered. The inference engine is furthermore used to trace back faults or DSS alarms by reversing the relationships and providing a list of possible points of failure given a list of faulty objects [6] within several minutes. DSS inputs, alarms and actions can be inhibited to mitigate the impact of interventions, which can be imported and applied to the Expert System description to allow a realistic simulation. Naturally, the simulation is blind to ongoing interventions for which no inhibit was requested. Furthermore, the lessons learnt from previous interventions are not considered by the inference engine.

In this proceeding a new alarm helper tool is presented that could aid the operators in assessing the criticality of an event based on the combination of alarms, their time of occurrence, their history, and finished and ongoing interventions. If possible, the alarm helper also lists the most likely points of failure using pre-calculated Expert System simulations. Furthermore, the tool proposes which experts should be contacted and provides links to related documentation. The tool's knowledge base can be updated by dedicated users on the spot, and will be otherwise updated weekly. The tool must be fast, reliable and easy to understand as it is supposed to reduce the need to call several experts, particularly for alarms that occur outside usual working hours.

2 Event analysis

The alarm helper knowledge base is derived from all DSS alarms that occurred between 2018 and 2023 and which were reviewed during the ATLAS operation meetings by the Technical Coordination team. The alarms are linked to their respective causes and grouped into events by parsing the weekly DSS reports with python. An event is defined as all DSS alarms and actions that were triggered by the same root cause. For example, if failure *X* triggered alarm *A1*, which triggered on its part alarms *A2* and *A3* five and ten minutes later, respectively, then all three alarms *A1*, *A2* and *A3* are grouped as belonging to failure *X* even though *A2* and *A3* are not *directly* triggered by *X* but as a consequence of alarm *A1*. As another example, if a failure *Y* occurs multiple times within a very short period and each time triggers alarm *A4*, then the alarm occurrences are grouped together into one DSS event. This can be the case for example for an intervention which requires restarting a sensor multiple times. Two separate events are defined if the cause might be the same, but it could not be confirmed conclusively or the time between the occurrences is too long.

Table 1. DSS alarm criticality description and distribution.

Criticality	Description	Follow-up	# alarms
Level 3	Immediate threat to detector or personnel.	Immediate intervention	145
Level 2	Medium impact to detector.	Immediate intervention.	595
Level 1	Contained impact to other systems.	Weekly review.	65
Level 0	Purely informative.	Weekly review	45
			850

DSS alarms are manually categorised by criticality into four levels from 0 to 3, indicated in Table 1 based on their impact to personnel and the detector. Level 3 represents the highest criticality and indicates an immediate threat to the detector or personnel. They require immediate reaction by operators. Some alarms are shared with alarms that are under the responsibility of the Occupational Health and Safety and Environmental Protection Unit. Level 2 alarms have a medium impact on the detector and trigger DSS actions that return the machine to a safe state that are followed up by experts immediately. Level 1 alarms indicate a minor fault and trigger less severe actions and have thus a limited impact on the whole detector. They are followed up timely in any case. Lastly, Level 0 alarms merely inform the operators about the change of certain detector states and do not need to be followed up immediately. Level 0 alarms do not trigger any actions. They are nevertheless followed-up and reviewed weekly. In total, ATLAS DSS features 850 alarms (May 2023), out of which the majority is classified as Level 2, followed by a number of Level 3 alarms. The criticality of a DSS event is defined by the highest criticality of the associated alarms.

DSS alarms are not solely triggered by faults but may also result from interventions during which equipment is switched off while their DSS signals are not masked. It is vital that the operators correctly identify the reason for an alarm and assess the criticality of the situation and not purely rely on the alarm criticality. The statistical analysis of DSS events with python of the last five years showed that the vast majority of events resulted from interventions, out of which almost half of the events contained at least one alarm with a high impact (Level 3) to the detector and about a third with a medium impact (Level 2), illustrated in Figure 1. This is highly indicative of the complexity of the infrastructure and the (unforeseen) possible




Figure 1. Criticality distribution of the foreseeable and unexpected events from January 2018 to May 2023. The categorisation is based on the DSS reports delivered in the weekly Operational Management Meeting. In total, 3071 events are analysed.

consequences caused by these events on the rest of the experiment. Moreover, false alarms not only distract the operators from “real” alarms but can also damage the detector, particular subsystems that were unaware of the intervention. Unexpected events represent about a quarter of the total number of events. Approximately half of the unexpected events are classified as Level 2 and about a quarter as Level 3. The new Alarm helper tool is designed to reduce particularly the number of alarms related to interventions and to speed up the clarification of unexpected alarms.

Figure 2 shows the top five triggered alarms classified as interventions or errors. Three out of the five top most triggered alarms during an intervention and four during an unforeseen event are related to cooling. The other two most frequent alarms during an intervention are related to a stop of the Uninterruptible Power Supply (UPS) or gas distribution of the ATLAS Transition Radiation Tracker (TRT) [7] sub-system. Furthermore, the severity and number of occurrences of interventions is typically higher than of unexpected events, demonstrating the need to address the triggering of alarms by interventions. Interestingly, the evacuation alarm for the US15 cavern is at the top of the unexpected alarms. The alarm has no impact on the detector and has no actions, thus the seven events with zero criticality indicate the sole triggering of the alarm. Given that the other evacuation alarms for the USA15 service gallery and UX15 experimental caverns are not triggered even nearly as often demonstrates that the US15 evacuation alarm is often triggered by non-ATLAS infrastructure from the LHC side.




Figure 2. Top five alarms between January 2018 and May 2023 caused by interventions (left) and unexpectedly (right).

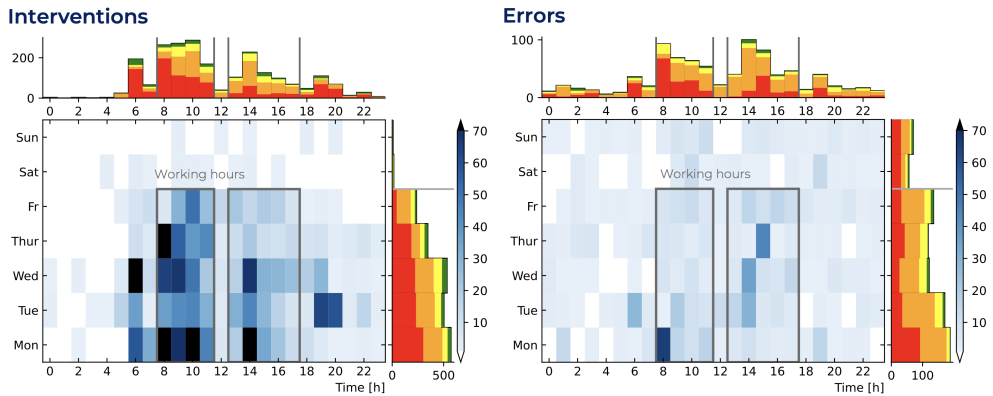


Figure 3. Frequency distribution as a function of time of day and day of the week for expected and unexpected events.

Figure 3 shows the time distribution of DSS events as a function of the day and the time. The criticality distribution per day or hour is also shown on the minor axes. As expected, interventions are concentrated during extended working hours (8:00 to 17:00) with one hour break, while the unexpected events are more evenly spread over the day, with a significant peak at the start of working hours (8:00 and 14:00). These are the events we want to address as soon as possible since they could be mistaken for an ongoing intervention. Very high-impact interventions, such as the annual testing of all DSS alarms are executed before working hours at 06:00 for safety reasons.

3 DSS alarm helper

The DSS alarm helper tool is designed to assist the operation team, particularly the Shift Leader in Matters of Safety (SLIMOS) in analysing incoming DSS alarms. The final decision on the appropriate reaction is left to the operator. If the expert cannot confirm a relation with an ongoing intervention, the presented information assists in identifying the alarm trigger. Each analysed event is used to improve the helper tool further. The tool is accessible to users via a web interface, for which a prototype version is shown in Figure 4. The tool will have access to the Expert System, historical DSS events, DSS documentation, and limited information about ongoing interventions. The Expert System graph-based database will be used to identify causality between actions, alarms and inhibits in real-time, as well as to




Figure 4. Prototype of the web interface of the ATLAS alarm helper




Figure 5. Mapping of DSS alarm combinations to unique possible points of failure and their criticality distribution which is taken from the highest criticality of the alarms.

describe the alarms. The historical DSS events will be extracted from the DSS logbook. The Expert System Object Kernel Support (OKS) [8] database schema is extended with two new classes: Log and Event. For each entry in the DSS logbook, a new object of type *Log* is created and linked to the corresponding Expert System alarm which is used to simulate the impact on the detector. The second class *Event* is links *Logs* with a common cause and allows classification of the event in Intervention and Error. This allows to reverse search all occurrences and events of a given alarm and to make informed decisions in the alarm follow-up.

The interface provides a short description of ongoing alarms and links to related documentation. Furthermore, it displays the criticality of the individual alarms and a list of historical events that caused the same (combination of) alarms, including their times. The recent events are ranked by their similarity to the situation.

Moreover, the tool lists which expert might have more information in order to make an informed decision of the actions to take. Moreover, for some alarms the point of failure can be immediately retrieved as several hundreds of DSS alarm combinations can be directly mapped to a significantly shorter list of possible points of failure using the Expert System inference engine (see Figure 5). In this context, a point of failure is defined as any object that is described in the Expert System. As of May 2023, 412 alarm combinations can each only be caused by one specific point of failure. 172 alarm combinations can be solely caused by two distinct points of failure. Notably, the most easily resolvable alarm combinations have a high or medium impact on the detector. The alarm helper further also checks for subsets of the mapped combinations if alarms are inhibited.

4 Impact on operations management




Figure 6. Proposed work-flow using the alarm helper to forecast the impact of interventions.

We propose to extend the new ATLAS alarm helper capabilities by further automatising the gathering of knowledge about planned and ongoing interventions. Based on announced interventions, project management tools like Jira or other forms of electronic communication, we propose to develop a tool that can extract the interventions that are planned for the next week. This list could be augmented by the alarm helper tool, where knowledge can also be used to predict the impact of the interventions on the detector. In addition to existing predictive tools like the Expert System, the new tool profits from "real events" and becomes smarter with every intervention. The predicted impact on the experiment, particularly its subsystems could then be communicated effectively to all relevant stakeholders. Furthermore, the forecasts can be used in a post-intervention analysis to efficiently generate documentation that is subsequently used to update the tools. If successfully implemented and tested, the weekly forecast could be used by the alarm helper to inform the operator whether a DSS alarm is expected to be linked to a specific intervention, and significantly ease the criticality assessment and follow-up as sketched in Figure 6. The new tools could increase the transparency of interventions to the whole experiment even non-experts. In addition, to the weekly postmortem analysis in the operation meeting, the weekly forecast could be discussed and adjusted based on the experiments. The new tool aims at improving the understanding of the experiment and prevent unforeseen negative impacts.

5 Conclusions

A novel tool to assist the ATLAS operators in assessing the criticality of DSS alarms and their cause has been presented. The tool extends the knowledge base gathered by the Expert System by comparing the current state of the detector to past events. The implementation is on the way and ready to be used by the end of the year shutdown (YETS) 2023. The alarm helper will be used to improve the planning of interventions and forecast and prevent possible alarms that can decrease the uptime of the detector. Furthermore, it will increase the transparency of alarms and their causes to the experiment and simplify documentation of interventions and follow-up while constantly updating the expert knowledge database.

References

- [1] ATLAS Collaboration, *Journal of Instrumentation* **3**, S08003 (2008)
- [2] S. Luders et al., *Proceedings of ICALEPCS2003* p. TH214 (2003)
- [3] I. Asensi Tortajada et al., *EPJ Web Conf.* **214** (2019)
- [4] I. Asensi Tortajada et al., *Proceedings of ICALEPCS 2019* pp. 106–109 (2020)
- [5] I. Asensi Tortajada et al., *EPJ Web of Conferences* **251** (2021)
- [6] G.A. Uribe et al., *Journal of Physics: Conference Series* **2438**, 012045 (2023)
- [7] The ATLAS TRT collaboration, *Journal of Instrumentation* **3**, P02014 (2008)
- [8] R. Jones et al., *IEEE Transactions on Nuclear Science* **45**, 1958 (1998)