HOW TO ACHIEVE SATISFACTORY PERFORMANCE OF THE ACCESS SYSTEM: STABILITY, EFFICIENCY, OPERATION, FLUIDITY

T. Hakulinen (GS/ASE)

With contributions from the LHC Access Team (GS/ASE)

Abstract

Heavy utilisation of the access system during LHC hardware commissioning has uncovered shortcomings in the performance of the system. While generally available to operation as required, response of the system has sometimes been below expectations of its users and operators. Reasons for these problems are both technical and procedural. Issues have been found in both hardware and software as well as in the system's response to external factors, such as network problems. Real-world usage of the access system has also not always reflected its original design. Possibilities for improvement exist: It should be possible to automate some of the administrative checks by the operators in the restricted mode. The access cycle may be streamlined, and new lighter-weight access modes for specific situations may be investigated. User interface improvements are possible to facilitate management of multiple access points at busy times. The exact actions to be taken will need to be evaluated between the access and operational teams. The most important issues affecting the performance of the system should be addressed first to ensure the best possible service to the users during the next shutdown.

LHC ACCESS AND SAFETY SYSTEM

General Description of LASS and LACS

The access and safety system of the LHC consists of two complementary systems: The LHC Access Safety System (LASS) [1] and LHC Access Control System (LACS) [2]. The LASS controls a number of Elements Important for Safety (EIS). The EIS are subdivided into access-related elements (EIS-a: doors, key distributors, patrol boxes, etc.), and machine/beam-related elements (EIS-m/f [f=faisceau]: electron-stoppers, access safety blocks, beam dumps, etc.). By interlocking the EIS, LASS enforces safety conditions on the LHC access zones to allow or deny access or beam operation. Computing in LASS is carried out by Programmable Logic Controllers (PLC), which are connected by a private optical network for maximum independence. Redundancy is provided by a separate cabled loop.

The function of the LACS is to manage the physical access barriers and to provide the access control of the LHC, subject to permission from LASS. The LACS consists of a number of access point devices, Personal Access Devices (PAD), Material Access Devices (MAD), card readers, intercoms, video subsystem, etc. All the access point devices are supervised by LACS servers and the overall status of the LACS is constantly kept up to date in the access database (centrally managed Oracle instance). All communication between the access point devices and the operator posts takes place via the database. The LACS database connects to the external HR/Adams Oracle database, which manages the access authorizations of all CERN users. A rough schematic of the LHC access and safety system is given in Figure 1.

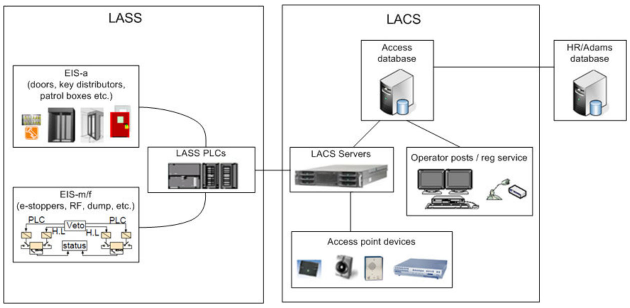

Figure 1: Schematic of the LHC access and safety system

LHC Access Modes

The LASS manages also the different access modes of the LHC: General, Restricted, Patrol, Closed, Veto, and Test.

- *General mode* is a mode, where access is granted automatically without explicit operator action, provided that the person has a pre-approved access authorization to the zone in question. The access procedure is simply:
  1. User badges.
  2. User enters the PAD, which carries out checks to verify that the user is alone in the PAD and not carrying material. If the check succeeds, the first PAD door closes behind the user.
  3. User looks into the iris scanner for the biometric verification.
  4. If the scanned iris matches the person to whom the badge is assigned, the second PAD door opens and the user enters zone.
- *Restricted mode* and *Patrol mode* are both operator controlled access modes with key, where an additional verification is required from an access operator in the control centre (either CCC or experiment control room). The key has an attached safety token, whose function is to interlock the beam system when people are inside. To control activities in the LHC, an approved *Avis d'Intervention* (ADI) in EDH is normally required. However, the ultimate responsibility for granting access lies with the Engineer in Charge at the control centre, who may authorize or deny an access at his/her discretion. The

Patrol mode differs from the standard Restricted mode in it being used to patrol the machine, i.e., to verify that there are no people inside the zone and to arm the patrol boxes thereby allowing the zone to be switched into a safe state for beam operation. The normal access procedure in restricted mode is:

1. User calls operator via intercom and gives the ADI number.
2. Operator checks the ADI in EDH and if it checks out, instructs the user to badge.
3. User badges making his/her name appear on the operator interface and allowing the operator to give a key or reject the access if necessary.
4. Operator gives a key and the user takes it.
5. User unlocks the PAD with the key.
6. User enters the PAD, which carries out checks to verify that the user is alone in the PAD and not carrying material. If the check succeeds, the first PAD door closes behind the user.
7. User looks into the iris scanner for the biometric verification.
8. If the scanned iris matches the person to whom the badge is assigned, the second PAD door opens and the user enters zone.

In case of a group with the same ADI, steps 3 to 8 are repeated until all the users have passed.

- *Closed mode* and *Veto* are modes where no access is possible. Technically Veto is not an access mode but a system state, set automatically by LASS (EIS interlock) or manually by operator (access off), which denies all access to the zone. These modes are used when in beam operation or for any reason that makes access to a zone unsafe (high-power hardware tests, radiation delay, etc.).
- *Test mode* exists for purposes of specific groups at specific access points (e.g., RF at UX451). It allows taking an access key without authorization by the access operator. The Test mode is rarely used in the operational LHC system.

Goals of the LHC Access Control System

The purpose of the LHC access control system (LACS) is to ensure safe access of the personnel to the controlled areas of the machine and the experiments. The general design goals of the LHC access system can be enumerated as follows:

- Be reliable, meaning that the system should not cause users to be exposed to danger. It should also not cause the beam to stop due to spurious alarms.
- Offer good performance to both users and operators.
- Offer flexibility to change and reconfigure things when necessary.
- Allow traceability via logging of events and the operational history.
- Automate as many things as possible and reasonable.
- Offer best possible interface to manually carry out things that cannot (or should not) be automated.

EXPERIENCE FROM LACS IN 2009

Some Statistics

To understand the scale of the usage of the LHC access control system it is instructive to look at some key statistical figures from the last 6 months. This was a very busy time period with lots of accesses to the LHC. The observation period spans from August 1, 2009 to January 23, 2010:

- There were 181893 valid accesses total over all LHC access points in all access modes. This means on the average 1033 accesses per day to the controlled areas of the LHC. Most of these accesses were in General mode, mainly at experiments and non-interlocked areas, where operator supervision was not required.
- Out of the above total figure, there were 33676 keys taken in restricted mode, which translates to 191 keys given by operators per day, on the average. Most of these restricted mode accesses were managed by CCC operators with a smaller portion managed by experiment control rooms. Figure 2 shows the distribution of keys taken per day over the observation period. The week structure is clearly visible with weekends showing little activity as well as the beam test period during the month leading to Christmas break. It is notable, however, that even during holidays, there were on the average tens of accesses per day in the LHC with activity picking up strongly during the first weeks of January.

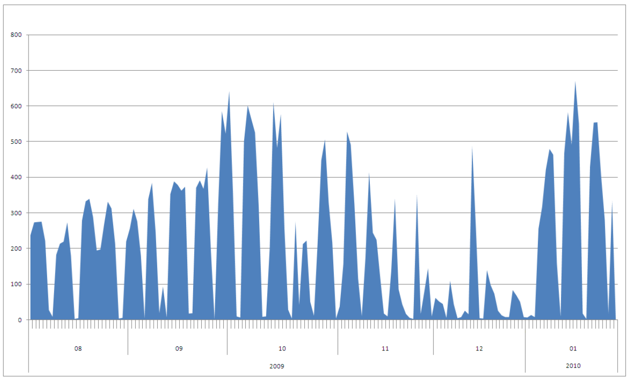

Figure 2: Keys taken / day (Aug 1, 2009 to Jan 23, 2010)

- The busiest day for the operators fell on January 14, 2010, when there were 670 keys given in restricted mode. The access distribution by access point, presented in Figure 3, gives an idea of the hot spots of operation, in this case points 6 and 4 followed by the Atlas experiment at point 1. Days with a comparable number of accesses can be spotted in Figure 1 in the first half of October 2009.

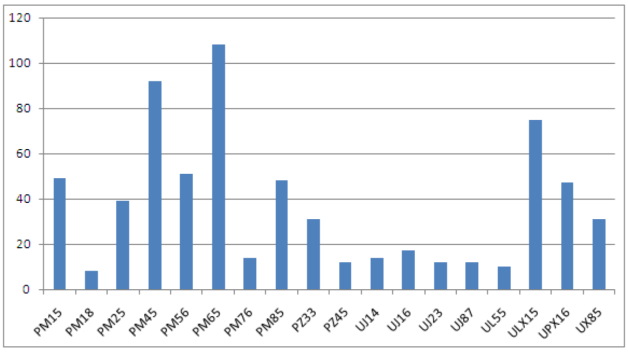

Figure 3: Accesses by access point (Jan 14, 2010)

- LACS logs don't give an easy indication of user waiting times from the moment he/she first tries to contact the operators to the moment that the key is given. This is due to the fact that intercom calls are not logged, and the first indication of an access request by the user is when he/she badges, which according to the protocol, happens only after the operator has already acknowledged the user and checked the validity of the ADI. A subjective estimate based on operator experience can be made, however:
  - The best case occurs when there is no rush, the user's ADI checks out without problems, and the access system is working nominally. In this case, the user can expect to get the key in less than 1 minute.
  - Given a normal operator load, meaning more than one call coming in, or the operator having to babysit one access point while answering a call from another, or the user's ADI requiring extra verification, delays from 1 to 5 minutes can be expected.
  - The worst case can occur when several factors coincide: technical problems in the access system during a big rush of large groups of people calling from several access points simultaneously. In a few cases waiting times of over 30 minutes have been seen, sometimes leading to some users abandoning their access attempt altogether. Intervention by the access team has also been required a few times to repair a malfunctioning piece of equipment (e.g., key distributor boxes with jammed shutters).

A Typical Busy Day

To demonstrate typical operator experience during a relatively busy day during the observation period, a synthesis of two separate shifts is presented below. In reality, the two shifts described occurred on two separate days, but they should nonetheless be quite representative. It is also to be noted, that at peak times a second operator managed some of the accesses, which are not taken into account in the present analysis.

- There are two single-operator shifts: 1st from 7:30 to 12:30 and 2nd from 12:30 to 17:30.
- Two periods of peak activity can be seen: morning (8:45-10:30) and after lunch (13:15-15:00).
- During a peak period one would normally have 3-5 calls from access points in the queue all the time.
- Following types of events occurred during the shifts:
  - Morning: 99 intercom calls, 170 user accesses.
  - Afternoon: 3 patrols, 97 intercom calls, 210 user accesses.
  - There were 2 persons per call on the average (size of the group), while the biggest group of the day consisted of 16 persons. While the ADI only needs to be checked once per group, each key must still be given separately, which obliges the operator to follow each access.
  - One system problem requiring operator intervention was seen, where a user could not exit a zone requiring an access maintenance intervention.
  - One hardware problem was seen requiring maintenance intervention.

All calls to the operator over the course of the day are shown in Figure 4 with the number of persons per call on the y-axis. From experience it can be deduced that an experienced operator can expect to manage 1 call per minute on the average during a time of normal operator load.

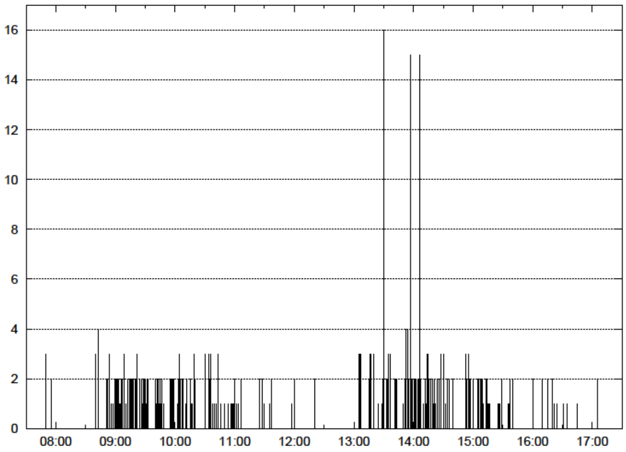

Figure 4: Calls to the operator during the two shifts with the number of accesses given per call on the y-axis

ISSUES AFFECTING ACCESS PERFORMANCE AND POSSIBLE TECHNICAL SOLUTIONS

Over time, a number of issues affecting the performance of the access system have been identified. These can roughly be divided into straightforward technical malfunctions in the system or its environment, shortcomings of the original system design with respect to today's realities, and administrative overhead adding complexity to the access procedure. In the following sections, the proposed solutions to the presented problems are shown in italics.

Technical Malfunctions

Hardware

The access system has had its share of hardware problems. A type of unreliable position contacts was identified as the source of incoherent signals (ambivalence detected by LASS) from the PAD doors. Several key distributor boxes have had to be changed and/or repaired due to shutters getting blocked or key slots getting damaged. Relays are always the weak spot of any hardware cabled loop and several have had to be changed.

Hardware problems have been dealt with as they have appeared. In the cases where a clear fault in design or implementation is identified, an alternative design needs to be proposed. The ongoing rigorous preventive maintenance program of the access system has already addressed some of these issues, for example, a campaign was launched to change the PAD position contacts in 2009. Another case is the campaign to upgrade the key distributor boxes to solve the recurring problem of damaged key slots.

To solve the issues with the video from access points, a redesigned video architecture with new video recorder hardware and software is likely to be necessary. This will imply a major redesign of the video subsystem.

The access system includes basic monitoring of the key hardware components. However, on some equipment improved monitoring will be added.

Software

Numerous problems due to software have also been seen. It is notable that these problems have occurred mainly in the parts of the access software written specifically for CERN, while the parts indigenous to the vendor's original software solution have given relatively little trouble. A bug in the commercial access software was identified as a source of persistent operator interface problems in the autumn. The biometry subsystem suffered from considerable problems during the first half of 2009, which mainly affected the biometry enrolment process. The video subsystem has shown signs of instability, for which a temporary correction exists (restarting the video recorder in question).

Software problems are mainly to be solved by correctives from the vendor. Sometimes this will require (and has required) considerable debugging and analysis by the CERN access team. Improvements in the software managing the PAD cycle are a good example.

Workarounds to some immediate problems have been implemented by the access team while waiting for a definitive fix (e.g., automatic monitoring of the biometry database to spot inconsistencies to be manually corrected by the access team).

The biometry subsystem can be made simpler and more fault tolerant by including the user's iris imprint on the badge. This should considerably simplify and speed up the biometry subsystem.

Improving monitoring of the various software components of the access system is also underway.

External factors

The correct functioning of the access system is also strongly dependent on a number of external systems. The most important of these prerequisite systems is the CERN TCP/IP network infrastructure. LACS devices are connected directly to CERN Technical Network, which houses most of CERN's control systems. Considerable network-related problems were seen during the last trimester of 2009. The difficulty in resolving these issues stemmed from the fact that according to standard network monitoring metrics, the network appeared to be functioning correctly. A closer collaborative analysis between the access team and the CERN network team revealed a hardware problem related to a firmware update and a configuration issue in one of the routers managing connections from access servers to equipment.

Other external factors that may occasionally have an effect on the access system are the CERN central Oracle service as well as the HR database service, even though these services have generally proven very resilient due to their fault-tolerant design. It has also happened that simple human interventions, intentional or not, have caused unavailability of parts of the system (e.g., a bent MAD door that wouldn't close properly).

External factors are by definition not directly under the control of the CERN access team. In this case, the only viable approach is a close collaboration with the respective services (example: the analysis with IT of the network router problems mentioned earlier).

Again, in order to be able to react promptly, some monitoring of the prerequisite systems by the access team is necessary.

Shortcomings of the System Design

The access system was designed in the first half of the 2000's based on the technology of the day. It is also apparent today that the usage of the system differs to a certain extent from what was forecast and on what the original design assumptions were based.

LACS internal communications

The communication protocol of the LACS from access devices on sites to the operator posts follows a fundamentally four-tier model: Access devices are connected to local processing units (UTL), which talk to access servers, which update the status of the system in the database, with which operator-posts interact. Correct operation of the system is fundamentally dependent on the presence of the database, without which restricted or patrol modes cannot be operated, changing of operating mode is not possible, and event logging and monitoring of the system is disabled. In this case, events are kept in local buffers of the devices and servers, from where they will be uploaded into the database once it comes back online, provided that the break is not overly long (of the order of days). Performance bottlenecks have also been identified related to this architecture, mainly in the communication between the access servers and UTLs.

The LACS communication protocol is a fundamental system feature, which cannot be modified at will. Therefore, only limited improvement is possible. It is possible to optimize the server processes somewhat by reorganizing the allocation of the UTLs between the servers. It is also possible to add server machines or to simply upgrade them to increase processing speed. One of the most important things is to make sure that network and database are always in good shape.

LACS operator interface

One part of the system, where the age of the design of the system shows, is the LACS operator interface, which has suffered from scaling limitations and want of responsiveness. The user interface is relatively complex with all the access points individually modelled with details from both LASS and LACS. This slows down the update speed of graphical items on screen. Adding even more elements (for a more intuitive interface) has also run into internal limitations of the maximum number of graphical elements in the program.

Improvement of the LACS operator interface is a long time operator request. In principle the system should allow a considerably faster processing of incoming access requests. The standard interface can be further streamlined, but this is a limited approach.

The most logical approach would be to go towards access software, which is as close to the vendor standard release as possible. This would mean taking out as many of the CERN specific features as possible, which would also make it much easier to follow vendor's standard software releases. A special-purpose high-performance interface for access-operation only without generic overhead facilitating management of multiple access points could be developed. The vendor software would still be used for all the other functions that it does well. This would be a development project involving CERN for the interface development and the vendor for the parts that communicate with the rest of the access system.

Key distribution

Key distribution is currently part of the access cycle, which makes it a bottleneck at access points in restricted mode. As presented earlier, the operator has to follow each access of a group of people to assure a smooth passage of everyone.

One improvement with a great potential for speeding up accesses is separating the key distribution phase from access entry cycle. In this case operator gives out all the keys of a group and lets the users pass through access point at their own pace.

Administrative Issues

ADI mechanism

One of the major differences from the original assumptions on how the access system would be operated is the use of the EDH-based ADI mechanism to filter user accesses. While the access system operates on zone-based access models to which individual users have access authorization or not, given all the standard prerequisites, the ADI acts as an additional administrative filter, which is not integrated into the access system. As the ADI was never designed with this kind of use in mind, it has proved to be somewhat inflexible in practice. This is mainly due to the EDH approval mechanism, which locks an approved ADI document making no modifications possible. As a need for last minute modifications in case of urgent changes will frequently occur, automatic check of the ADI by the access software could not be made reliable.

To deal with the ADI issue, a wider approach is necessary than what can be accomplished by modifying the access system only. First it must be known what the future "ADI" mechanism will look like. This is primarily operational business, but input from access team will be necessary. The most likely candidate is the proposed AET mechanism [3]. Once the new mechanism is in place, this information can be better integrated into the access interface for restricted mode: When user badges, the system can check and show (all) his/her valid AETs for the access point. For this to be useful, the new AET mechanism must be enforced. A mechanism to modify AETs rapidly will also be necessary.

Another approach is to add new [partial] access modes to the system. Two examples:

- *General mode with AET: This would be an extension of the current General mode but with the ability of the access software to check automatically if the user has a valid AET for this zone. A downside is that automatic operation would not allow the system to treat exceptions.*
- *General mode with operator confirmation (supervised without key): This would be a mode similar to the Restricted mode, but without a key.*

Any of the proposed modes would only require modifications in LACS. LASS would not be modified in any case.

Scheduling conflicts

Another issue, which has risen occasionally, involves scheduling conflicts between different activities, such as hardware tests and access maintenance, where accepted ADIs may have to be rejected on the spot without an easy method of informing affected users in advance. While strictly speaking not an issue with the access system *per se*, properly automated management of the administrative authorizations would alleviate this problem as well.

As mentioned above, scheduling conflicts are mostly out of scope of the access system. However, improvement even in this regard would be possible with the new AET mechanism.

PRIORITIES AND TIMETABLES

Best estimates at this time of the cost, complexity, and required lead times of the various proposed modifications are presented in Table 1. However, while the items are in principle shown in the rough order of their priority, this ordering is not definitive due to uncertainty of some of the items.

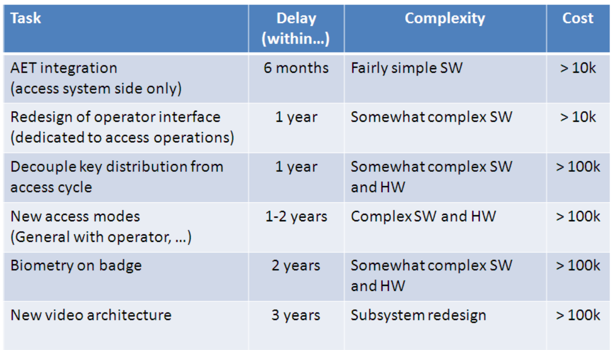

Task	Delay (within…)	Complexity	Cost
AET integration (access system side only)	6 months	Fairly simple SW	> 10k
Redesign of operator interface (dedicated to access operations)	1 year	Somewhat complex SW	> 10k
Decouple key distribution from access cycle	1 year	Somewhat complex SW and HW	> 100k
New access modes (General with operator, …)	1-2 years	Complex SW and HW	> 100k
Biometry on badge	2 years	Somewhat complex SW and HW	> 100k
New video architecture	3 years	Subsystem redesign	> 100k

Table 1: Cost / complexity estimates by task

CONCLUSIONS

Heavy utilization of the LHC access system has uncovered shortcomings, which have been analyzed by the access team. Some of the issues have already been corrected either by modifications to the system by the vendor or through CERN-specific workarounds. Some issues mainly related to the overall performance of the system still remain.

To achieve a better performance from the point of view of users and operators, both technical and administrative issues will need to be addressed. Several technical improvements are possible, but depending of the scope of the modifications required, considerable lead times and cost may be involved. Work is underway to identify and implement the most effective modifications to permit them to be available for the next machine shutdown.

Lessons learned are also being applied in the design of the future access and safety system upgrades (PS, SPS), of which the specification of the PS Personnel Safety System upgrade is already well advanced.

ACKNOWLEDGEMENTS

Contributions to the presentation in Chamonix as well as to this paper came from the whole LHC access team (GS/ASE). Thanks are also due to the LHC operation team (BE/OP) for all their input.

REFERENCES

[1] LASS; https://edms.cern.ch/document/707874

[2] LACS; https://edms.cern.ch/document/693084

[3] Julie Coupard’s presentation, Session 5