NATS System Failure 12 December 2014 – Final Report
Independent Enquiry

Final Report dated 13 May 2015

Authors

Robert Walmsley
Timothy Anderson
Clay Brendish
John McDermid
Martin Rolfe
Joseph Sultana
Mark Swan
Michael Toms

Secretariat

Peter Whysall, NATS

NATS System Failure on 12 December 2014 – Final Report dated 13 May 2015

CONTENTS

ACKNOWLEDGEMENT

In the course of the Enquiry, the Panel held discussions with airlines, airports and other organisations, most notably NATS and the CAA. All were willing providers of information, advice and their time. For that and for the secretariat support provided by Pete Whysall, the Panel is most grateful.

NATS System Failure 12 December 2014 – Final Report

EXECUTIVE SUMMARY

Introduction

ES1. Following a failure of some United Kingdom air traffic control (ATC) services on 12 December 2014 (the Incident), the Civil Aviation Authority (CAA) and NATS (formerly National Air Traffic Services) established an independent enquiry into the cause of the failure, the recovery and other relevant factors. After the appointment of the Panel members, the Enquiry formally started on 13 January 2015.

ES2. The Incident started with the failure at 1444 UTC of a computer system used to provide information to Air Traffic Controllers managing the traffic flying at high level over England and Wales. This traffic includes aircraft arriving and departing from London airports as well as aircraft transiting UK airspace. The Controllers put agreed procedures into action so as to limit traffic entering their area of responsibility and adopted manual methods for decision-making to ensure aircraft continue to maintain safe separation.

ES3. At 1455 all departures were stopped from London Airports and at 1500 all departures were stopped from European airports that were planned to route through affected UK airspace. The computer system was restored to the Controllers at 1549, but without its normal level of redundancy. By 1900, the Engineering staff believed they understood the cause of failure and full redundancy of the computer systems was restored at 2010. Traffic restrictions were gradually lifted from 1555 as confidence increased, and the final restriction was lifted at 2030. The disruption caused by the restrictions affected some airlines, airports and passengers into the following day.

The Impact of the Incident

ES4. The Incident occurred at 1444 on a Friday afternoon in the run up to Christmas. By 1500 there was information available on news broadcasts and social media suggesting that there was a UK air traffic control issue and this evolved into the story that UK airspace was closed. At Gatwick, the Controller managing take-offs had received a telephone call at 1448 from NATS at Swanwick to "Stop all departures" and relayed this information to the 3 aircraft queuing to line up for take-off. At about 1500 she was called by the pilot of the leading aircraft along the lines of: "My passengers are telling me that they're hearing on Sky News that there's an air traffic problem. Can you tell me something?" The Incident had quickly become a cause célèbre with the media.

ES5. The primary NATS duty in delivering ATC services is to maintain safety. The safety criteria are embodied in defined minimum separation distances and heights between aircraft. The Enquiry found no evidence or suggestion that these safety criteria had been breached and, given the abrupt and broad nature of the failure, has no hesitation in commending this achievement. Beyond this, NATS' performance management regime is focused on minimising the aggregate flight delays attributable to its ATC services. The total aggregate delay attributed to NATS on 12 December was calculated by NATS as "about 15,000 minutes, broadly equivalent to a normal bad weather day event"; this figure was later refined to 14,863 minutes applied across 353 flights. In accordance with the agreed arrangements for calculating such delays, there is no attribution from flights cancelled as a result of the disruption, of which NATS estimate there to have been approximately 150 with up to a further 20 being diverted away from UK airspace.

ES6. Under normal circumstances, all aircraft and airports are treated equally by NATS so that neither aircraft carrying large numbers of passengers nor congested airports are given priority. No arrangements are made by NATS to collect or estimate the impact on passengers as to numbers or delays, this being seen as the business of the airports and the airlines. The enquiry asked for an estimate of the number of passengers affected by direct delay (i.e. not including cancelled flights or missed connections); NATS provided a figure of 65,000 passengers affected by the direct delays, cancellations and diverted flights described in the previous paragraph. It is recognised by both NATS and the Panel that this figure will be a significant underestimate of the total number of passengers affected due to factors such as the delayed arrival of one flight often leading to the delayed departure of another. NATS estimates that, if such factors are taken into account, a maximum of 1900 flights and 230,000 passengers were affected during the afternoon and evening of 12 December. Additionally several airlines reported some level of cancellations and flight disruption running into 13 December with approximately 60 aircraft and 6000 passengers affected.

The Systems Failure

ES7. The systems at the NATS Swanwick operations centre entered service in 2002 but were in development during the previous decade. Failure occurred on 12 December 2014 because of a latent software fault that was present from the 1990s; this is referred to as the proximate cause of the failure. The fault lay in the software's performance of a check on the maximum permitted number of Controller and Supervisor roles (known as Atomic Functions). These Atomic Function identifiers are used to index (access) the table of data held in the System Flight Server (SFS) and distribute the correct and relevant data to individual roles. The check should have been whether or not the limit of 193 (the total of civil and military roles) had been reached; instead the check was performed against a civil limit of 151.

ES8. The total number of Atomic Functions in use at the time of the Incident was 153, a figure that was reached for the first time because a change was put into operation on 11 December to include further military Controller roles. This change was not sufficient, in itself, to cause the failure; there was a specific trigger that led to the failure.

ES9. The workstations in the Operations room are normally left "Signed On", even when unattended, so that they are readily available for use. However, if a Controller presses the "Select Sectors" button (to begin controlling aircraft) on a workstation that is not "Signed On", that workstation enters Watching Mode. When the SFS receives a command to enter Watching Mode it generates a table to hold a copy of some system data (it is a copy, as it is just "Watching" not in control). Due to the inclusion of the extra military Controller roles on 11 December and the way in which the airspace was being managed at 1444 on 12 December, an unintentional request to enter Watching Mode led to the generation of a table of Atomic Functions (representing all the then current Controller and Supervisor roles) with 153 entries. This failed the check on the permitted maximum size of the table (as the limit of 151 was used, not 193). This led to an internal error, known as an "exception", being raised within the execution of the software.

ES10. Discrepancies between the state of the workstations and the SFS are potentially unsafe as Controllers may be presented with the wrong data. Thus SFS was designed and programmed to shut down the primary SFS in response to this "exception". However, it is important to preserve availability of the SFS, so control is transferred to a secondary. The secondary reprocesses the commands from the workstation as would be entirely appropriate if a hardware fault had precipitated failure of the primary SFS. In this case, as the fault was in the SFS code, and it was triggered by the command from the workstation, the same exception was raised in the software in the secondary SFS, and that too shut down. This "double failure" of the SFS occurred at 1444 GMT, was reported on each Controller screen and led to the imposition of limitations on aircraft movement (regulations) as described below.

The Systems Recovery

ES11. The standard practice in NATS is that engineering recovery is coordinated through a group of designated engineers, known as the Engineering Technical Incident Cell (ETIC) and drawn from those available in the Systems Control Centre adjacent to the Operations Room. While some recovery actions are automated, ETIC manually control all key recovery actions, e.g. the restoration of data, to ensure that decisions are made with due and careful deliberation;

this is important, as the wrong decisions could have further downgraded the ATC performance.

ES12. Identifying a software fault in such a large system (the total application exceeds 2 million lines of code), within only a few hours, is a surprising and impressive achievement. This was made possible because system logs contain details of the interactions at the workstations, for example selecting Watching Mode, messages sent between the workstation and SFS, records of significant events, e.g. system failures, and software exceptions. Some of these functions run continuously, and produce error reports automatically – for example the responsible engineers and managers received email notification of the SFS failures. There are also queries and other functions that can be applied to the logs by the engineers on duty, enabling particular circumstances to be analysed. These enabled the events that triggered the failure – sending a command to enter Watching Mode and the replay of the commands to the secondary SFS – to be quickly identified.

ES13. Part of the recovery philosophy is for the software automatically to restart SFS following a failure. The SFS servers (A and B) accordingly restarted within a few minutes of the failure. ETIC asked for a further manual restart of the SFS to increase their confidence in the state of the hardware. A successful system restart then requires restoration of the flight data to achieve full SFS capability and this was duly done from the unaffected National Airspace System (NAS). ETIC decided that only one SFS server should be updated with flight data, as this gave more options should an error be made in the process of returning the system to service.

ES14. Two critical factors enabled this rapid fault detection and system restoration. First, the Lockheed Martin engineers (who had played a major role in the development of the code) in the UK and USA had secure real-time access to data logs, and thus contributed fully to the diagnosis of the Incident. Second, the NATS team at Swanwick, as exemplified by the ETIC, operates a collaborative culture, and their working is not hindered by organisational or commercial boundaries.

Systems Management

ES15. Design of software of any significant complexity is difficult, and it is unrealistic to expect that software faults will not be introduced in development. The design process conforms to appropriate standards to ensure that the level of faults constitutes an acceptable risk. There are many standards for safety related software development and the CAA promulgates regulatory requirements that constrain the standards and approach used by NATS. Although these post-date the development of the faulty software, NATS has subsequently demonstrated the adequacy of its processes with these requirements and continues to do so for each build.

ES16. In summary, NATS' processes are thorough and professional and are believed to have been ahead of their time when first developed and they meet the requirements of the CAA. There is a strong and effective process for controlling software releases, requiring signatures from key stakeholders. The resultant integrity appears better than would be expected for software of this importance. However, there will be major changes in NATS systems over the next few years with the deployment of new pan-European systems (known collectively as SESAR). The current processes will need updating to address new factors that will arise in this international collaboration, including: achieving a stable requirement agreed with other European Air Navigation Service Providers (ANSPs); a new focus on NATS capabilities in a system and software programme management role; and significant changes to suppliers.

ES17. The Watching Mode is not needed operationally. It was initially used for training and familiarisation (in system "work up"), and was retained for presentation purposes, but cannot be used for controlling aircraft. This leads to the question of why the Watching functionality was implemented in such a way that there was the potential for complete loss of the SFS. However, removing the capability is not straightforward, and there has to be an evaluation of the costs and risks of making such a change, as it is likely that there will be unintended consequences that could impact other functionality.

ES18. The decision to make a procedural change to limit the use of "Watching" is defensible, given its potential complexity and cost of removing it. Finally, NATS have now fully resolved the problem by correction of the code implementing the check of the maximum number of permitted Atomic Functions, rather than making a more major change to disable "Watching". This is the cost-effective and timely solution to the problem; it was possible to make this correction in days, rather than the months of design and testing work required to disable "Watching" altogether.

Operational Response

ES19. On being alerted to a system failure by a warning at their workstations at 1444, the London Area Control Supervisory team were directed by the Operations Supervisor (OS) to follow fallback Checklist 4 for 'SFS Unavailable'. ATC tower supervisors at airports within the London Terminal Manoeuvring Area (TMA) were instructed to stop all departures immediately and a tannoy summoned all operational staff present in Swanwick Centre to the London Area Control (LAC) and London Terminal Control (LTC) Operations rooms. At 1450, anticipating a shutdown of the NAS Flight Data Processing system as part of the technical recovery actions, the LAC OS advised TC of the possible need to revert to full manual control. Such a procedure is not referenced in the relevant fallback checklists, but it appears that the perception of a NAS failure and/or of the recovery involving shutting down the NAS may have led the OS and Airspace Capacity Manager (ACM) to initially implement more conservative restrictions than those required.

ES20. In addition to the immediate suspension of departures at London TMA airports, the ACM took action to reduce traffic levels. A zero flow rate restriction (ZRR) was applied to the London Flight Information Region (FIR) from surface to unlimited altitude. The ZRR was initially prescribed for a 4-hour period, with the intention of ensuring that it immediately stopped all departures from across the wider European region, including UK airports. However, Heathrow, Gatwick and Manchester departures were dealt with individually via 3 standard contingency restrictions applied in parallel, each defined as taking effect from +45 minutes from implementation. Notwithstanding, the initial suspension of departures communicated verbally direct to tower supervisors was still in place at this stage, so the period of grace built into the formal airport-specific regulations did not take effect on the day.

ES21. The failure warnings presented to controllers were correct. Nonetheless, uncertainty over which system was affected and the type of failure that had occurred continued until shortly after ETIC was convened, when it was concluded that the NAS was in fact serviceable. Only after ETIC had confirmed this to the OS and ACM, around 40 minutes after the initial warning indications at the workstations, was consideration given to managing LAC traffic tactically via contingency routes in LTC and Prestwick Centre's (PC) respective airspace areas of responsibility, a lesson identified after the December 2013 system failure.

Operational Recovery

ES22. Full LAC functionality was restored one hour after the initial failure. At 1555, regulated departures were authorised by Swanwick from Heathrow, Gatwick and Manchester airports, initially at an extended Minimum Departure Interval (MDI) of one every 5 minutes (both northbound and southbound for Heathrow). During an Air Traffic Incident Coordination and Communication Cell (ATICCC) customer conference call at 1605, Eurocontrol's Network Management Operations Centre (NMOC) was informed that the ZRR and London airports' departure suspensions could be cancelled.

ES23. As many of the aircraft already airborne when the failure happened had continued to arrive at Heathrow, parking and stand availability became a critical issue and, following a request from Heathrow, the arrival rate was reduced to 20 per hour by NMOC at 1626. By 1730 all departure restrictions were cancelled for Heathrow, Gatwick and Manchester, but it was not until 1935 that the remaining Heathrow arrival restriction was cancelled.

Operational Control and Communication

ES24. During the first internal ATICCC call at 1545, Silver command (NATS crisis management operational command level) confirmed that the secondary SFS was recovered and that Ops would commence a graduated lifting of restrictions at 1605. However, at this stage, the root cause of the double SFS failure had not been identified. Subsequently, at 1619, Silver team were apprised of the possibility that exceeding the permissible number of Atomic Functions was related to the system failure. At the Engineering Bronze (NATS crisis management tactical command level) teleconference at 1630, the risk of recurrence of the failure was assessed as High, as the root cause was not fully understood and the engineering design team could not yet be specific with recommendations. At 1715, after further analysis by the engineering design team, the risk of recurrence was reassessed as Low, as the role that Watching Mode had played in the failure had been identified. Nevertheless, the root cause had still not been identified at this stage.

ES25. At 1936, Silver Chair was briefed on the maximum number of available Atomic Functions and the latent defect that reduced this to 151 if a terminal is in Watching Mode. At the time of this brief, it was reported that 142 Atomic Functions were active on the system. During an Engineering teleconference at 2045, the number of Atomic Functions active was revised up to 153 and the fact that the system was therefore still exposed to the risk of SFS server shut-down if Watching Mode was entered resulted in Silver Chair direction to immediately remove 7 of the currently operating Atomic Functions.

Effect

ES26. Five specific aspects of Swanwick Centre's actions appear to have had consequences for the scope and severity of the Incident's impact, and the ease of the recovery. These were:

(1) Stopping departures at Heathrow, Gatwick and Manchester airports;

(2) Perceiving a need to conduct a NAS recovery from SFS data;

(3) Instituting a comprehensive ZRR for all London airspace;

(4) Initially applying all contingency regulations for 4 hours;

(5) The NATS-led 'generic' recovery.

ES27. **Suspension of Departures**. The initial verbal suspension of departures was not rescinded when the 4 formal contingency regulations were applied and, consequently, the window in the latter designed to accommodate continued departures at Heathrow, Gatwick and Manchester was unavailable. Moreover, the initial confusion over the status of the NAS appears to have distracted Swanwick supervisors from substantive consideration of implementing contingency routing procedures quickly. Up to 1 hour 15 minutes of potential departures from Heathrow, Gatwick and Manchester were therefore lost, accelerating congestion significantly and making the recovery more challenging than it could have been.

ES28. **Status of London Airspace**. "EGT1ACC", the standard contingency ZRR requested by the Swanwick ACM with effect from 1500, is defined as applying to all Swanwick airspace from surface to unlimited. It was not until 1535 that the ACM confirmed formally to NMOC that the Centre was able and willing to continue accepting arriving traffic. It may therefore be argued that, until then, it was not an unreasonable inference by NMOC that this, coupled with an immediate and enduring suspension of all London departures, and a zero traffic rate applicable to all London FIR airspace, effectively amounted to closure of that airspace. During the initial 45 minutes of the Incident, it is likely that this ambiguity reinforced perceptions that London airspace was closed and resulted in up to 20 aircraft being diverted pre-emptively to alternative airports and around 150 flights being cancelled.

ES29. **Duration of the Regulations Applied**. NMOC's initial assessment of the airspace being closed and the fact that the contingency regulations had been applied for 4 hours duration triggered a number of procedural responses by NMOC, including immediately suspending the Flight Plans of all affected flights. This action should have been accompanied by NMOC releasing an associated ATFM Information Message (AIM) to inform operators but, unhelpfully, this was omitted by the NMOC supervisor on duty, causing further confusion during the recovery.

ES30. **Designing and Controlling the Recovery**. ATICCC's four customer calls during the Incident were predominantly structured around 'push' communications informing customers of actions taken, or planned, by NATS. There does not appear to have been a formal process to receive, triage and prioritize customer information and requests. It may therefore be concluded that the NATS-led recovery was largely generic in nature and focussed on re-establishing LAC operations in the round. In addition, the initial suspension of all affected Flight Plans by the NMOC in response to the 4-hour ZRR prompted mass cancellations of Calculated Take-Off Times (CTOTs) in the recently introduced Eurocontrol Airport Collaborative Decision Making (A-CDM) system and, to some airports and operators at least, gave the appearance of A-CDM not being able to 'keep up' with the crisis. Both Heathrow and Gatwick Airport Ops Cells dispensed unilaterally with the A-CDM system and resorted to managing departures locally within the MDI rates set by Swanwick and extensions to slot times agreed by the NMOC. This was less efficient than it may have been and had the effect of removing key data and communication pathways to and from the NMOC.

Previous Lessons

ES31. A previous NATS' investigation into a serious communications system failure that occurred on 7 December 2013 identified a number of lessons and prompted associated recommendations by NATS and the CAA most of which were reported as closed off and in place ahead of this most recent incident. However, amongst these recommendations were three of particular note in the context of the 12 December 2014 failure. The first was to review with stakeholders the industry's ability to respond to service failures and identify required changes to NATS' crisis management capabilities, resilience of systems, procedures and service continuity plans. The second, made by the CAA, encouraged NATS to make best use of all means by which a crisis can be handled from an operational standpoint, including exploring the more effective use of and interactions with the Eurocontrol Network Manager (NM). Despite being assessed by NATS as complete before 12 December, it is evident that neither of these recommendations had been addressed fully. Finally, a review of the wider industry crisis response and resilience arrangements was recommended. Invitations to participate in the crisis response exercise were extended by NATS to major stakeholders in May 2013 and the event was anticipated to take place in February / March 2015, although that date has now been postponed until after this Enquiry reports.

Safety

ES32. There were no safety events recorded within LAC and LTC during the period of fallback operations or during the recovery phase. Notwithstanding, post-incident technical analysis revealed that Watching Mode had been selected accidentally by LAC controlling staff multiple times a week on average in the months leading up to the 12 December Incident, despite the existence of a Temporary Engineering Instruction stating that Watching Mode should not be selected. NATS' Safety Management System includes a facility to lodge a 'Safety Incident' Mandatory Occurrence Report for significant safety occurrences and the reporting of lower level safety events is also encouraged. However, there is no distinct Error Management System (EMS) of the sort employed in other high-hazard industries. Such a system can offer considerable benefits to risk identification, risk management and safety assurance. It is also possible that, had it existed, such a system would have highlighted the relative propensity for staff to miss-select Watching Mode and that this may have prompted earlier action to mitigate the hazard more effectively.

Systems Requirement, Management and Delivery

ES33. The New En Route Centre (NERC) Operational System has been in service since 2002. It has been upgraded over its lifetime in terms of hardware, to support operational changes and to implement a range of system enhancements and problem resolutions. The technology, with its origins in the 1990s is naturally dated and, although it was reported to be "leading edge" in its time, requires more "hands on" involvement than modern systems to address amendments. The change lifecycle for the NERC System follows established best practice processes and procedures for a system of this complexity although some of the change processes are manpower intensive. Overall the Panel accepted that the current process for low level changes, implemented at the operational level, prioritising safety risks and impacts, is both appropriate and well executed.

ES34. The requirements for NATS future systems and the international context for their delivery are significant changes. In the past, NATS had considerable flexibility in how their requirements were met. Since the passing of Single European Sky legislation (SES) in 2004 and its adoption into national rulemaking, the primary regulation of ATM has been from Europe through the European Commission (EC) that has introduced the performance regime the concept of Functional Airspace Blocks and SESAR, the technological dimension. Through SESAR, the European ATM industry has developed a European ATM Master Plan that defines capabilities and concepts for deployment across Europe to deliver both interoperability and challenging performance standards.

ES35. Despite this, neither the SES regulations nor the SESAR deployment plans specify details on resilience requirements or how resilience should be measured. It is considered necessary to set out contingency, resilience and business continuity performance requirements in a clear and unambiguous way that will help to manage and align stakeholder expectations. These will need to be agreed by NATS and CAA in consultation with other stakeholders and ideally aligned within Europe to avoid driving different requirements and costs across the network.

ES36. NATS' "Deploying SESAR" programme is intended to deliver many of these capabilities and to transform the overall operation of its business through phased transitions over the next 5 years. If it is to be successful NATS must find effective ways of working with its suppliers, partners and regulators that recognise the challenges of a large scale multi-year collaborative project. In particular:

(1) The Deploying SESAR programme will be very different in nature from the NERC programme and NATS must ensure that they are fully aware of these differences as they put in place the plans, expertise and governance necessary to deliver a successful programme.

(2) A key risk in long term, large scale collaborative programmes is poor control of the requirements. A well articulated concept of operations, agreed with all stakeholders, is the first step towards achieving a stable requirement and must be complemented by rigorous control of changes.

(3) The obligation on NATS is to deliver the required capabilities safely into service, complying with the technical requirements of SESAR and meeting the SES performance standards. Effective programme assurance arrangements, including periodic formal scrutiny by NATS staff outside the programme, are key to good risk management.

ES37. "Deploying SESAR" is a large programme with already tight timescales and challenging collaboration aspects. Acceleration in a search for earlier benefits would be likely to lead to shortcuts being taken in the early requirements and specification phases, additional costs and further delay. The Panel accepts that it would not be sensible for NATS to attempt to accelerate this programme beyond the currently defined plan as this is already ambitious.

The CAA NATS Relationship

ES38. Oversight of safety, operations, investment and charges to airlines by NATS is the responsibility of the UK CAA. The CAA is a well-established regulator led by a board with a wide range of relevant experience and supported by a highly regarded, experienced and expert management. It has clearly defined responsibilities and organisational structures. It derives its authority from a number of sources. The framework for the licencing and regulation of NATS is provided by the Transport Act 2000, but its safety and performance are also the subject of a number of provisions of European law. The CAA's primary duty is to maintain high standards of safety but it also has a number of explicit secondary duties. The CAA's approach to the regulation of NATS generally follows established best practice for UK regulators modified by European requirements and the particular features of the NATS business. The regulation of safety takes place primarily through the oversight of the NATS Safety Management System, people, systems, operational procedures and safety cases for changes. Oversight of other aspects of the NATS licence is more 'light touch' and is based around the principle of encouraging NATS to agree the key elements of its plans, including performance targets, capital and operating expenditure with its airline customers.

ES39. The panel has found no suggestion that any failure of the CAA's oversight contributed to the events of 12 December or posed any threat to safety. However there are aspects of the CAA's oversight arrangements which could usefully be brought further in to line with regulatory best practice to minimise the risk of further incidents in the future and ensure that recovery takes place with the minimum inconvenience to passengers (whilst maintaining safety).

ES40. These measures include greater engagement by the CAA in the NATS investment programme and steps to bring the interests of airline passengers more directly into the regulatory equation. The CAA should require NATS to submit and maintain an operational resilience plan, as is required for major airports. To achieve these objectives the CAA should be given enforcement powers, including power to levy fines for breaches of the NERL licence, comparable with those of other regulators (although an incident of this scale would not have been of a sufficient magnitude to result in a fine unless it had formed part of a sustained pattern of performance failure).

ES41. The performance related element of the remuneration of the top management of NATS would be expected to reflect both the annual and the longer term corporate objectives of NATS, including the delivery of high and consistent levels of service. The CAA should monitor the performance related component of NATS remuneration policy and if necessary be prepared to influence the board of NATS on this matter.

Recommendations

Recommendations are numbered sequentially through the document from R1 to R31 and are not in any priority order. Key recommendations representing a subset of those presented within the body of the Report are provided below in the order they appear together with their reference number.

R1 NATS should retain, for their deployment of SESAR:

- The system architecture approach, including hardware redundancy and the fault management capabilities to provide resilience in the presence of hardware and software failures and operator errors associated with configuring the system;
- The automatic logging of system behaviour, including Controller commands, software failure conditions ("exceptions") and hardware failures;
- The provision of real-time access to these logs and other system data by software development and support staff;
- The ETIC and its important role in managing technical failures and their recovery;

- The collaborative culture between NATS and its suppliers, integrating new suppliers into established information sharing mechanisms. (Para 2.8.2, 2.8.8)

R4 NATS should ensure that contracts or other suitable arrangements provide a complete, continuing evidence base for the current operational software. This should be demonstrated throughout the remaining life of the system by audits of the software development records and NATS should ensure that identified discrepancies or omissions are resolved. The reviews should be timed to support the five-year planning cycle, and instigated as necessary to respond to a perceived risk, e.g. an accumulation of change. (Para 2.8.2)

R8 NATS should consider the costs and benefits of adopting any of the identified modern testing and software assurance methods on a targeted basis for the current LAC software environment, including determining whether or not the return on investment, e.g. from using modern static analysis of software code, is likely to be worthwhile, and the benefits of de-risking new tools on the current software before using these tools within SESAR. (Para 2.8.11)

R11 NATS should consider introducing a formal Error Management System (EMS) to capture anomalous occurrences that fall below the safety event threshold, but which may indicate where changes in systems, procedures or training would benefit the management of risk. (Para 3.6.2)

R13 NATS should review their hierarchy of fallback procedure checklists for completeness, coherence and consistency, so that they support controllers via an intelligent checklist architecture that leads intuitively through conditions-based options, including making clear where the controller has discretion to adjust and refine responses as circumstances dictate or allow. (Paras 3.7.2 & 3.7.3)

R15 NATS, in conjunction with the Eurocontrol NM and key customers, should review their ‘standard’ contingency routing, flow rate and departure regulations to ensure they are suitably responsive, precise, effective and sensitive to their impact on the wider aviation system. (Para 3.7.4)

R21 NATS and the CAA should agree on how to provide assurance that the evolving capability meets the functional and non-functional requirements of SESAR while complying with the performance regime of the Single European Sky regulations. (Para 4.7.2)

R25 NATS should include, within their phased approach to SESAR deployment, scrutiny and control of the concept of operations, clear requirements and exit criteria for each phase defined in advance with strong governance of initial approval, management of change and phase completion. (Para 4.7.8)

R27 The CAA should ensure that they have sufficient internal expertise to enable them to complement, select and manage external consultants in analysing and assuring the NATS capital programme, and overseeing its evolution through the annual Service and Investment Plan (SIP) (Para 5.12.3)

R29 The CAA and NATS should develop systems to estimate, monitor and publish the scale and direct impact to passengers of serious events causing air traffic control disruption (Para 5.12.7)

Chapter 1. Introduction

1.1 Context

1.1.1 Following a failure of some United Kingdom Air Traffic Control (ATC) services on 12 December 2014 (the Incident), the Civil Aviation Authority (CAA) and NATS[1] (formerly National Air Traffic Services) announced the establishment of an independent enquiry into the cause of the failure, the recovery and other relevant factors. After the appointment of the Panel members, the Enquiry formally started on 13 January 2015.

1.1.2 Terms of Reference (ToRs) for the Enquiry, which include a list of Panel members, are at Annex A and were published on the CAA website on 16 January 2015. The ToRs called for an Interim Report by 31 January 2015 and stated that this should be focused on the NATS internal investigation[2] of the 12 December Incident. The Interim Report was submitted on 28 January 2015[3]. This Final Report, which was due no later than 14 May 2015, addresses both the subject matter of the Interim Report and the remaining and generally wider issues specified in the ToRs, including the Panel's views on the root causes lying behind the Incident.

1.1.3 The Incident started with the failure at 1444 UTC (this and all subsequent times are reported in the 24 hour format at UTC) of a computer system used to provide data to Air Traffic Controllers to assist their decision-making when managing the traffic flying at high level over England and Wales. This traffic includes aircraft that have departed or are planned to arrive at major London airports (Heathrow, Gatwick, Stansted, Luton and City) as well as aircraft transiting UK airspace. The Controllers put their pre-agreed operating procedures into action for the particular computer system failure; these included adopting manual methods for decision-making to ensure aircraft continue to maintain safe separation and restricting air traffic entering their area of responsibility.

1.1.4 At 1455 all departures were stopped from London Airports and at 1500 all departures were stopped from European airports that were planned to route through affected UK airspace. The engineering experts were able to determine the nature of the failure and agree a safe recovery procedure so that the computer system was restored to the Controllers at 1549, but without its normal level of redundancy (back-up). By 1900, the Engineering staff believed they understood the cause of failure and full redundancy of the computer systems was restored at 2010. Traffic restrictions were gradually lifted from 1555 as confidence increased, and the final restriction was lifted at 2030. The disruption caused by the restrictions affected airlines, airports and passengers into the following day.

1.2 Enquiry Process

1.2.1 The Enquiry followed traditional lines so that some or all Panel members undertook visits to relevant facilities for briefings and fact-finding. These visits were supplemented by discussions with organisations and individuals, and with meetings of Panel Members. A full list is provided at Annex H but prominent events were:

(1) Visits to the NATS operational centres at Swanwick (where the relevant Air Traffic Controllers and the operational Systems Engineers work), and to the nearby NATS Corporate and Technical Centre (CTC) which contains representative elements of current computer systems, development laboratories and administration.

(2) Visits to Heathrow, Gatwick and Luton airports.

(3) Discussions with British Airways, EasyJet, Flybe, Ryanair and the International Air Transport Association (IATA).

[1] A high level overview of the NATS and CAA organisations can be found in Annexes E and F respectively.

[2] SP301 Major Incident Investigation: Preliminary Report Version 2.0 January 2015.

[3] NATS System Failure 12 December 2014 – Interim Report; Version 3.0 January 2015.

(4) Discussions with CAA and NATS officials.

(5) Panel Meeting with the Director Aviation, Department of Transport.

1.2.2 The ToRs for the Enquiry included an invitation to any person or organisation impacted by the events being addressed by the enquiry to send their comments or suggestions to the Enquiry Secretariat, although no such independent submissions were received.

1.3 Background

1.3.1 The framework for carrying out UK air traffic management is provided in the Transport Act 2000. The Act designates the CAA as responsible to the Secretary of State for Transport for, inter alia, providing:

(1) The CAA must exercise its functions so as to maintain a high standard of safety in the provision of air traffic services; and that duty is to have priority over the application of subsections (2) to (5).

(2) The CAA must exercise its functions in the manner it thinks best calculated—

(a) to further the interests of operators and owners of aircraft, owners and managers of aerodromes, persons travelling in aircraft and persons with rights in property carried in them;

(b) to promote efficiency and economy on the part of licence holders;

(c) to secure that licence holders will not find it unduly difficult to finance activities authorised by their licences;

(d) to take account of any international obligations of the United Kingdom notified to the CAA by the Secretary of State (whatever the time or purpose of the notification);

(e) to take account of any guidance on environmental objectives given to the CAA by the Secretary of State after the coming into force of this section.

(3) The only interests to be considered under subsection (2)(a) are interests regarding the range, availability, continuity, cost and quality of air traffic services.

(4) The reference in subsection (2)(a) to furthering interests includes a reference to furthering them (where the CAA thinks it appropriate) by promoting competition in the provision of air traffic services.

(5) If in a particular case there is a conflict in the application of the provisions of subsections (2) to (4), in relation to that case the CAA must apply them in the manner it thinks is reasonable having regard to them as a whole.

(6) The CAA must exercise its functions under this Chapter so as to impose on licence holders the minimum restrictions which are consistent with the exercise of those functions.

1.3.2 NATS provides UK air traffic management in two adjoining regions, The Scottish Flight Information Region (FIR) and the London FIR. The London FIR is divided into:

(1) London Area Control (LAC), which handles civil aircraft over England and Wales in flight at high level.

(2) London Terminal Control (LTC)[4] which is a smaller area, including the five main London airports, and covers aircraft generally flying below 21,500 feet, with the precise height demarcation with LAC depending on the location.

[4] Note also that lower level airspace in the north of England is controlled from the Prestwick Centre.

1.3.3 These areas are shown diagrammatically at Annex D. Aircraft passing through UK airspace (principally between Europe and North America) transit LAC en route; aircraft destined for the London Airports transfer from LAC to LTC as they descend and vice-versa for departing aircraft.

1.4 The LAC Operation

1.4.1 The Incident on 12 December abruptly affected ATC throughout London Area Control at 1444. Air traffic services for both LAC and LTC are operated by NATS and, together with military aircraft services for the UK, are provided from separate control rooms within the same building at Swanwick, near Southampton. LAC is divided into a maximum of 32 sectors that can be combined ("band-boxed") at times of light traffic or separated or sub-divided ("split") when the traffic is heavier. The number of staff varies through the day, week and season but broadly depends on the number of aircraft expected to be flying in or through the London FIR. There are five "watches" of Controllers to manage the Operations Rooms on a continuous basis.

1.4.2 Each Controller can operate for up to 90 minutes without a break and controllers are rostered throughout the day to meet this requirement. When staff are not required at a workstation because of lighter traffic conditions they are encouraged to leave the Operations Room, partly so as not to distract those engaged in operational duties. At the time of the Incident there were 26 Controllers in the LAC operations room with some further 42 (LAC) Controllers on duty elsewhere at the Swanwick site. At the time of the Incident there were also 6 operational engineering staff in Systems Control (which oversees the status of the technical systems supporting the Swanwick site) adjacent to the Operations Room. 244 aircraft were expected to be under control of LAC during the hour following the Incident.

1.4.3 Controllers normally work in pairs: a Tactical Controller who communicates with the aircraft under control and a Planning Controller who manages the flow of traffic into and out of their area of responsibility through liaison with adjoining NATS or other national ATC areas. An Air Traffic Services Assistant provides support to the controllers when required. The primary safety objective of these arrangements is to ensure a height separation of at least 1000' between aircraft or, where aircraft are within this limit, to maintain a lateral separation of at least 5 miles.

1.4.4 Each pair of Controllers is assigned to a particular sector or combination of (band-boxed) sectors. They are supervised in groups of 5-8 sectors by Local Area Supervisors. An Airspace Capacity Manager is focused on the overall flow of traffic in the LAC and supports the Local Area Supervisors in managing the band-boxing or splitting of sectors. The Operations Room as a whole comes under the charge of the Operations Supervisor. Both the Operations Supervisor and the Airspace Capacity Manager have designated Assistants.

1.4.5 NATS operates a network of radar stations that provide the position and height of all aircraft flying in the LAC. A data fusion system determines the best estimated position when an aircraft is detected by more than one radar so that the aircraft appears only once on the workstation screen; a label adjacent to the aircraft icon gives its height and can give the heading and other related information.

1.4.6 The Controller can call up all other necessary data associated with a particular aircraft, derived from its flight plan information. The flight information derives from a flight data processing system, also operated by NATS and known as NAS (or National Airspace System), and this is routed to a System Flight Server (SFS) that delivers the right information to each workstation.

1.5 The Impact of the Incident

1.5.1 The Incident occurred at 1444 on a Friday afternoon in the run up to Christmas. By 1500 there was information available on news broadcasts and social media suggesting that there was a UK air traffic control issue and this evolved into the story that UK airspace was closed. At Gatwick, the Controller managing take-offs had received a telephone call at about 1448 from NATS at Swanwick to "Stop all departures" and relayed this information to the 3 aircraft queuing to line up for take-off. At about 1500 she was called by the pilot of the leading aircraft along the lines of: "My passengers are telling me that they're hearing on Sky News that there's an air traffic problem. Can you tell me something?" The Incident had quickly become a cause célèbre with the media.

1.5.2 The primary NATS duty in delivering ATC services is to maintain safety. The safety criteria are embodied in defined minimum separation distances and heights between aircraft. Beyond this, the NATS performance management regime is focused on minimising the aggregate flight delays attributable to its ATC services. The total aggregate delay attributed to NATS on 12 December was calculated by NATS as "about 15,000 minutes, broadly equivalent to a normal bad weather day event"[5]; this figure was later refined to 14,863 minutes applied across 353 flights[6]. In accordance with the agreed arrangements for calculating such delays, there is no attribution from flights cancelled as a result of the disruption, of which NATS estimate there to have been approximately 150 with up to a further 20 being diverted away from UK airspace.

1.5.3 The European Network Manager has recorded a total delay of 18,433 minutes caused by the NATS Systems failure of which 3,983 minutes was recorded by neighbouring Air Traffic Control Centres (mainly Paris and Brest). The figures calculated by NATS and Europe are therefore broadly consistent. According to the European Network Manager, the aggregate delay reported by airlines was, however, more than double these ATC figures. He suggests that an explanation for this divergence may derive from the European Network Manager calculating individual flight delays from the datum of a new estimated (and delayed) departure time, whereas the airlines would continue to use the original departure time[7].

1.5.4 The NATS licence granted by the Secretary of State for Transport states that "the Licensee shall not unduly prefer or discriminate against any person or class of person in respect of the operation of the Licensee's systems, after taking into account the need to maintain the most expeditious flow of air traffic as a whole without unreasonably delaying or diverting individual aircraft". Thus in general all aircraft are treated equally by NATS and, in particular, aircraft carrying larger numbers of passengers are not given priority. No arrangements are made by NATS to collect or estimate the impact on passengers as to numbers or delays, this being seen as the business of the airports and the airlines. The Enquiry asked for an estimate of the number of passengers affected by direct delay (i.e. not including cancelled flights or missed connections); NATS provided a figure of 65000 based on an average number of passengers per delayed flight. It is recognised by both NATS and the Panel that this figure will be a significant underestimate of the total number of passengers affected.

1.6 The Panel and Report Structure

The Panel included four independent members with, between them, deep expertise in software design; the management of aviation and other emergencies; the procurement and delivery of large computer based systems; and the statutory regulation of UK entities. The Panel also included a senior official from each of NATS, CAA and Eurocontrol (the Brussels based international organisation for pan- European air traffic management). This arrangement was designed to minimise the possibility that the Panel failed to understand either the facts or the implications of specialist information. An independent Chairman completed the Panel and

[5] NATS Preliminary Report into 12 December 2014 v2.0 paragraph 5.7
[6] P. Whysall email of 28 April 2015
[7] J. Sultana email dated 24 March 2015 enclosure "NATS ENQUIRY - ECTL Network Manager Input"

helped to ensure that any conflicts of interest between the official members and their normal employment were scrupulously managed.

1.6.1 The report structure reflects the specialist capabilities of the Panel members so that it first addresses the detail of the underlying computer systems failure and restoration before going on to describe the operational response to the Incident and the longer term recovery. The following chapter addresses the evolving requirement for NATS systems together with their management and delivery. The final chapter covers the relationship between the CAA and NATS and the effectiveness of the oversight arrangements.

1.6.2 Each chapter contains recommendations and the Panel agreed the key items that should be included in the Executive Summary.

Chapter 2. Systems Failure and Recovery

2.1 Introduction

2.1.1 This chapter considers the software factors in the Incident on 12 December 2014. It introduces sufficient context to understand the system failure and the technical recovery from the failure. Having considered the events themselves the chapter goes on to consider the software design and the development processes to shed light on why the critical fault was not detected and removed prior to the Incident. It sets out conclusions and makes recommendations primarily for the development and management of the next generation of systems being introduced by NATS, which derives from the European initiative known as SESAR (for Single European Sky ATM Research).

2.1.2 The chapter is presented in non-technical terms, so far as is possible, and is supported by a more detailed analysis of the software and design issues in Annex G.

2.2 Technical Context

2.2.1 Figure 2.1 shows the hardware systems that support the Air Traffic Controllers in LAC. Information on civil flight plans comes from the National Airspace System (NAS) and is routed by the System Flight Server (SFS) to the appropriate Controller workstation. The SFS also augments flight plan data with dynamic information, including clearance and coordination data from the Controller. For safe operation of the LAC it is important that the information on the workstations remains consistent with that in the SFS.

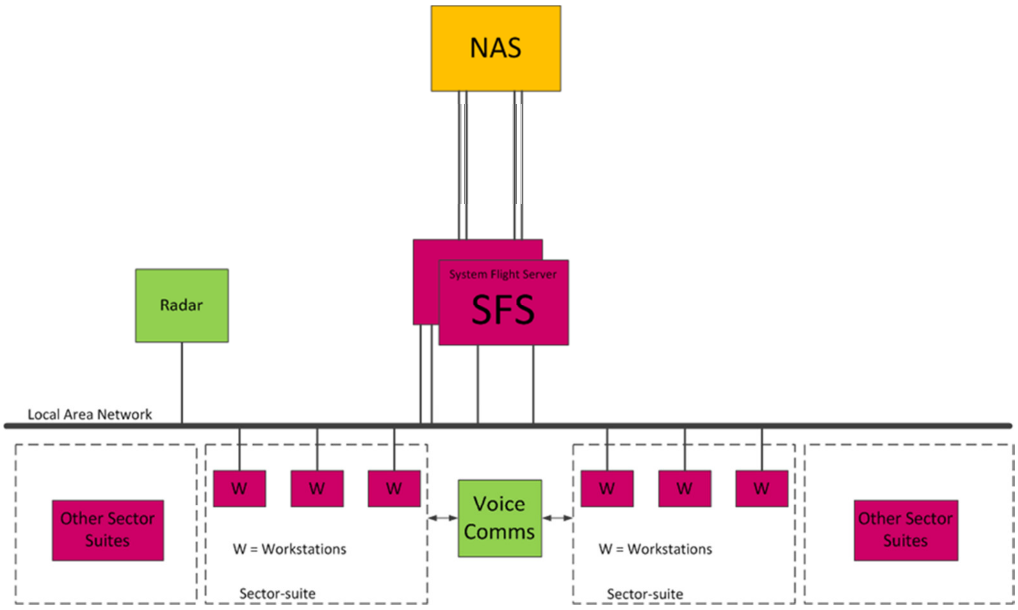

Figure 2.1: Major Systems supporting LAC (Simplified)

2.2.2 The SFS can route information to the correct workstation because the activities undertaken by each Controller are labelled with a unique identifier known as an Atomic Function[8]. There is an Atomic Function for each role for each sector, and for supervisor roles. The design of the airspace (civil plus military), and of the systems, is such that a maximum of 193 Atomic Functions will be required and can be supported by the system. The software design means that there can be a maximum of 151 Atomic Functions for the operation of the civil sectors (civil Controllers plus supervisors).

2.2.3 There are two SFS; one is primary and the other is secondary. The primary is normally in control; if there is a failure of the primary then the secondary takes over. The system will operate with a single SFS until the primary can be restored.

[8] Controllers can carry out the same role for more than one sector at once (this is what is meant by band-boxing) and the Atomic Functions are used to represent these responsibilities; the details of how Atomic Functions are defined and allocated through band-boxing and splitting are not relevant to understanding the Incident so are not discussed in detail.

2.3 Cause of the Systems Failure

2.3.1 The system failure occurred because of a latent fault in the SFS software, present since the 1990s. This fault was certainly present in the software in 1998 (prior to the system going operational), and may have been originally introduced in 1994; this is referred to as the proximate cause of the failure.

2.3.2 The fault lay in the performance of a check on the maximum permitted number of Atomic Functions. The check should have been whether or not the limit of 193 (the total) had been reached; instead the check was performed against the civil limit of 151.

2.3.3 The fault had been present since the time the SFS system became operational in 2002. It was triggered because a set of circumstances arose that had not occurred previously in the system operation (they can not have arisen previously, as the failure would inevitably have occurred due to the nature of the fault and the wider system design).

2.3.4 The total number of Atomic Functions in use at the time of the Incident was 153, a figure that was reached for the first time because of a November 2014 system change to allow the inclusion of further military Controller roles within the system and which was put into operation on 11 December, i.e. the day before the Incident. This change was not sufficient, in itself, to cause the failure; there was a specific trigger that led to the failure. To discuss the trigger requires further context to be given on the way that Controllers interact with the system.

2.3.5 When a Controller signs on to a workstation in its initial powered state, it changes from “Base Mode” to “Prepare Mode” but the workstation cannot be used to control air traffic. The Controller then selects his or her designated sector thereby notifying the SFS of the aircraft data required by the workstation (interaction between the workstation and SFS software determines the Atomic Functions for that sector and the associated roles); the workstation moves into “Elected Mode” and displays a copy of the data being used at that time to control the selected sector. If the Controller then selects “Open Sectors”, the workstation goes into “Controlling Mode” and becomes fully operational while the workstation previously controlling that sector moves into the “Elected Mode”; this transfer of responsibility is managed by the Local Area Supervisor to ensure smooth handover between Controllers.

2.3.6 Recording of all information available to the workstation starts when the Controller “Signs On”; this is relevant both to the Incident and the recovery.

2.3.7 A further mode called “Watching Mode” allows a workstation to display a full copy of the data from another workstation. “Watching Mode” is entered by selecting sectors on a workstation that is not “Signed On” (whereas a “Signed On” workstation would move into “Elected Mode”). Normally all the workstations in the LAC Operations room are “Signed On” – even when unattended – so that they are readily available for use. However, if a Controller presses the “Select Sectors” button when the workstation is not “Signed On” it will enter “Watching Mode” – until 12 December 2014 this did not lead to service disruption, despite it occurring many times because a work station had been inadvertently left in a “Signed Off” state.

2.3.8 When the SFS receives a command to enter “Watching Mode” it generates an internal table to hold a copy of some system data (it is a copy, as it is just “Watching” not in control). The Atomic Functions are used to index (access) the table. Due to the inclusion of the military Controller roles on 11 December and the way in which the airspace was being managed on December 12, an unintentional request to enter “Watching Mode” led to the construction of a table of Atomic Functions representing all the Controller and Supervisor roles. This table had 153 entries, which failed the check on the size of the table (as the erroneous limit of 151 was used, not 193). This led to an (internal) “error”, known as an “exception”, being raised within the execution of the software.

2.3.9 As indicated above, discrepancies between the state of the workstation and the SFS are potentially unsafe as Controllers may be presented with the wrong data. Thus SFS was designed and programmed to shut down the primary SFS in response to this “exception”. However, it is important to preserve availability of the SFS, so control is transferred to the secondary. The secondary reprocesses the commands from the workstation (from a “retained commands” list) as would be entirely appropriate if a hardware fault had precipitated failure of the first SFS. In this case, as the fault was in the SFS code, and it was triggered by the command from the workstation, the same exception was raised in the software in the secondary SFS, and that too shut down. This “double failure” of the SFS occurred at 1444 and led to the imposition of limitations on aircraft movement (regulations) as discussed in subsequent chapters.

2.4 Recovery Process

2.4.1 NATS was able to recover the SFS quickly, and also to determine the cause of the problem (the software fault) within a matter of hours. Whilst, in some respects, these are separate issues they are treated together as knowledge of the proximate cause, i.e. where the fault lay in the software, helped the NATS team in bringing systems back on-line with confidence that they would then operate correctly. More detail on recovery timelines is given in chapter 3.

2.4.2 The redundancy of the SFS (the primary and secondary structure) is intended to cope with hardware (rather than software) failures. As many hardware failures are transient (i.e. not requiring repair of the hardware) part of the recovery philosophy is for the software to automatically restart the affected server following a failure. As a consequence the SFS servers (A and B) restarted automatically within a few minutes of the failure, however this did not lead to an immediate provision of full SFS capability as it is necessary to restore the flight data used by the systems as well. It is important to ensure that the data used by the systems are correct and consistent, e.g. that the NAS and SFS have the same data on aircraft flight plans, so the data has to be restored with care.

2.4.3 The standard practice in NATS is that engineering recovery is coordinated through a small group of designated engineers from those available in the Systems Control centre known as the Engineering Technical Incident Cell (ETIC). Whilst some recovery actions are automated, ETIC manually control all key recovery actions, e.g. the restoration of data, to ensure that decisions were made with due and careful deliberation; this is important, as the wrong decisions could have further downgraded the ATC performance.

2.4.4 ETIC asked for a further manual restart of the SFS to increase their confidence in the state of the hardware. At the time ETIC did this, they were not fully aware of the proximate cause of the failure (the software fault), hence this precautionary measure was taken.

2.4.5 ETIC decided that only one SFS (Server B) should be updated with flight data, as this gave more options should an error be made in the process of returning the system to service. The key decision was to update SFS with the information held in NAS. It is also possible to update NAS from the SFS; if this had been done under the prevailing circumstances then there would have been serious degradation of ATC capability. The decision to update SFS from NAS, not vice versa, took some time because of some confusion about the nature of the failure amongst the operational staff (see, for example, 3.7.3). ETIC took special care to prevent further escalation of the Incident, by ensuring that the data in NAS remained available and therefore usable by Controllers.

2.4.6 Following this process, SFS Server B was available to support ATC operations almost exactly an hour after the failure. Server A was not made available as a back-up until about five hours after the failure. This was done when ETIC believed that the risk of restoring the redundancy (operating both servers) was sufficiently low as the proximate cause of the failure was by then understood (it had been localised to a module known as `Waafu28` and supporting code, see Annex G for details), and the system had been operating stably for over four hours.

2.4.7 There is a critical factor that enabled the effective management of the recovery: the operation of ETIC. ETIC was populated with highly competent engineers, who had a very deep understanding of the systems and had the right attitude in terms of assessing and managing risk. They also operated a consensual decision making process which helped to ensure that appropriate decisions were made. Whilst it would, in principle, have been possible to restart the systems sooner, a prudent approach was taken, minimising risk.

2.4.8 Identifying a software fault (with confidence) in such a large system (the total application exceeds 2 million Source Lines of Code (SLoC)), within only a few working hours is a surprising and impressive achievement. This was made possible because of the way logs are kept and analysed.

2.4.9 The logs contain details of the interactions at the workstations, for example selections of the “soft keys” on the displays, messages sent between the workstation and SFS, records of significant events, e.g. system failures, and software exceptions. These logs are voluminous and they would be hard to analyse purely “by hand” although they are in a human readable form. NATS however have a set of standard functions used to analyse the logs.

2.4.10 Some of these functions run continuously, and produce error reports automatically – for example the responsible engineers and managers received email notification of the SFS failures. There are also queries and other functions that can be applied to the logs by the engineers on duty, enabling particular circumstances to be analysed. These enabled the events – pressing the “Watching Mode” soft key and the replay of the commands to the secondary SFS – that triggered the failure to be quickly identified.

2.4.11 The logs are also linked to the running code (strictly the source programs written to produce the running code). Thus it was possible to “trace back” from the record of the exceptions in the logs to the module in the SFS software that was judged to be the location of the fault; the fault was then detected by manual inspection of the module source code. The module is written in the Ada computer language and it includes a simple specification as well as the code. The specification refers to all functions (strictly it says “ANY”), but the body of the code uses data that is only capable of storing the civil functions; this discrepancy was the root cause of using the 151 “limit” rather than the correct 193. (See Annex G for more detail.)

2.4.12 There are two critical factors that enabled this rapid fault detection. First, the Lockheed Martin (LM) engineers (who had played a major role in the development of the code) in the UK and USA were able to obtain secure real-time access to the necessary data, e.g. logs, and thus contribute fully to the diagnosis of the Incident. Second, the New En Route Centre (NERC) team at Swanwick operates a collaborative culture, and their working is not hindered by organisational or commercial boundaries.

2.5 Overview of the Development and Verification Processes

2.5.1 This section presents a qualitative overview of the most relevant aspects of the software development process used by NATS (and its suppliers, particularly LM) for LAC. The intent is to clarify the extent to which the process gave opportunities for detecting and removing the software fault (proximate cause of the failure) prior to the events on 12 December 2014. An assessment of the processes for managing known faults is presented in section 2.6, together with a quantitative analysis of the processes, so far as this is practicable.

2.5.2 The LAC software development is governed by a standard known as POD SW01[9] (the current version is dated 2012, but it incorporates LM standards in force at the time the

[9] NERC Software Standards and Procedures – POD SW01, SO518/SW01, Issue 9, July 2012.

relevant modules were written, so it can be taken as representative of the standard in force at the time of initial development and subsequently).

2.5.3 POD SW01 includes two mechanisms for detecting faults in programs: reviews/inspections and testing. Reviews and inspections involve systematic reading of the code (or specification) to identify faults; it is done manually and without executing the program. Testing, on the other hand, involves executing the program with particular data and checking that the results meet expectations; if they do not, a fault has been found (although it may be in the test, the code or the specification, and subsequent manual effort is required to determine which). Reviews and inspections are done on code modules and other items, e.g. requirements. Testing is done at multiple levels, from individual modules, through subsystems up to the complete running system. (The technical aspects of inspection and testing, including coverage of the code, are discussed in more detail in Annex G.)

2.5.4 Not all software is equally critical; POD SW01 varies the requirements for review/inspection and testing with the assessed criticality of the component, according to the "category" of the software. Category 1 and 2 software, which can be viewed as having a potential safety impact, is treated more rigorously than software that is in category 3 and 4, which generally can only impact availability, not safety. The SFS and the modules that were significant to the Incident are in category 2. The enquiry has reviewed categorisation of the software (see Annex G) and concluded it is broadly correct; therefore the remainder of this analysis is presented on the basis that this is the correct categorisation.

2.5.5 According to NATS standards, the software development process artefacts, e.g. results of reviews and inspections, test results, etc. should be preserved for software in the operational system. The information for `Waafu28` was not available at the start of the Enquiry but has since been retrieved from the USA. It is possible from the information available to make a partial assessment of the effectiveness of the module testing that was carried out (see Annex G for details). In summary, all the tests passed which means that the developers would have had no prima facie reason to investigate the module in any more detail.

2.5.6 Reviews and inspections are done at two levels: a mini-inspection and a formal inspection. Individual reviewers do mini-inspections, and results are circulated. A formal inspection involves a meeting of all the reviewers. An inspection is done on a module, e.g. `Waafu28`. If a mini-inspection is done and reviewers are concerned, e.g. due to the level of faults, then a formal inspection can be called. POD SW01 requires formal inspections to be done on some development artefacts, e.g. unit (module) test plans, but it is optional for other artefacts, including code. However, a formal inspection was carried out on `Waafu28` and some faults were identified and later rectified (see Annex G for details).

2.5.7 There are other opportunities to find problems, as the software is integrated through subsystems and up to system test. Reviews are not carried out on integrated code (although the interfaces are subject to review). It is unlikely that the opportunities to identify this specific problem through test would be materially different at subsystem level than at module level, so this level of testing is not discussed in any more detail. The focus is on what might have been achieved at system-level testing, especially using the extensive test labs operated by NATS.

2.5.8 The failure involved interaction between the workstations and the SFS. The system went through formal acceptance tests, carried out by the supplier, witnessed by NATS[10] within an independence regime authorised by the CAA. The system acceptance test results that relate to the coordination between the workstations and SFS[11] indicate that there would have been

[10] NATS was a wholly owned subsidiary of the CAA until the public private partnership (PPP) in 2001, but the test independence requirements were still observed.

[11] New En Route Centre System Acceptance Test Report, CDRL Number: 051 8/ACC 06/004 Volume 93, FDP001 Coordination, Lockheed Martin Air Traffic Management, June 30,1999

over 130 Atomic Functions during this test process (the report doesn't give this figure directly but it can be inferred from the set of workstations used, etc.) but less than 151. There was some use of "Watching" in the test set; although some of the tests initially failed none gave a forewarning of the failure on 12 December 2014 (there was no loss of SFS).

2.5.9 NATS test labs have extensive representative hardware, including a full SFS set-up, a total of 78 Controller workstations and 5 Supervisor workstations (about half as many as in the operational set up, but with proportionately fewer Supervisor workstations). With the 83 workstations it is possible to get to the maximum number of Atomic Functions (i.e. 193) but this would be "artificial", i.e. not representative of operations, as it could not use fully representative traffic patterns and workloads that are manageable for Controllers. Prior to the Incident the test team preferred to use a configuration of 141 Controlling Atomic Functions and 5 out of 28 available Supervisory Atomic Functions, which was more representative of the operational set up. Thus, even with the extensive test facilities, it is not simple to get to a situation (the 153 Atomic Functions) that would have triggered the failure seen on 12 December.

2.5.10 The Enquiry team was shown a simulation of the failure that occurred on 12 December 2014. This enabled the team to see both how the failure would have been manifest to the Controllers, and how the logs can be used to investigate failures and to find root causes. This was done with hindsight and it does not mean it was realistic to have detected the problem beforehand.

2.5.11 System testing of every combination of parameters in a system of this size is not practical. There are many influences on the system, including airspace sector structures, and it is simply impossible to test all the feasible configurations. Each workstation has five modes. To test every combination of workstation modes at 1 second per test would take of the order of 100 years, without considering all the other parameters, e.g. set of flights and flight plans, aircraft emergencies, etc. Many of the tests involve flying an aircraft through a sector, and this takes many minutes, not a second. Therefore it is not sensible to expect that system testing will be able to reveal faults at such a level of detail, and it is more appropriate and more effective to look for such faults at the software module level. However there may still be merit in having a test facility that is truly representative of the operational system to reduce the risk of the eventual transition to operational use.

2.6 Overall Assessment of the Effectiveness of the Development and Verification Processes

2.6.1 Design of software of any significant complexity is difficult. It is unrealistic to expect that software faults will not be introduced in development (fault injection rates are typically between one per 100 and one per 30 SLoC for professionally developed software). Section 2.5 considered the means of detecting faults; this section considers how detected faults are managed, and the level of faults in the LAC software.

2.6.2 First, in1996, NATS found and fixed a similar issue in the `Waafu` package that caused a dual SFS failure, but the fault was in a different function to the one that caused the failure on 12 December 2014. This had no operational effect as in 1996 the system was still in the development phase. This could be viewed as an opportunity to have considered a wider review of the design and implementation of the SFS package. There is no evidence that this was done. However, at this stage in development, there will have been many failures and problem reports every day and it would be impractical to investigate every one to a greater extent than was necessary to be sure that the underlying fault has been identified and removed.

2.6.3 Second, there is the issue of whether or not the software development processes used by NATS and NATS' suppliers, conforming to POD SW01, are sufficiently rigorous for this class of software. This judgment is informed by comparison with other processes and standards, and then by considering metrics.

2.6.4 There are many standards for safety related software development. The CAA produces regulatory requirements, specifically CAA CAP 670[12], that constrains the standards and approach used by NATS. Part B, Section 3 of CAP 670 is known as SW01, and sets assurance requirements for software. SW01 is goal-based, that is it says what is to be achieved, not how to achieve it (although there is some guidance, e.g. on Commercial Off-the-Shelf (COTS) equipment[13]). Although SW01 post-dates the development of the faulty software, NATS has subsequently demonstrated the adequacy of its processes against the requirements of SW01 and continues to do so on a build-by-build basis.

2.6.5 NATS' Safety Case is at the system level, and is underpinned by a Software Safety Assurance Argument[14]. SW01 has five key goals, e.g. requirements validity (showing that requirements are appropriate for the intended use of the software). The argument (rationale) in NATS Assurance Argument shows explicitly how these goals are met by the results of the POD SW01 process, and assesses the work of the LAC suppliers in those terms. For example, for requirements validity, it considers the tracing of the software safety requirements to the system safety requirements, the allocation of the appropriate assurance level (criticality of the software), and the ability both to implement the requirement and to verify that it has been met in the final code. On industry norms, this argument is of a very high standard. It is also commendable in that it looks for counter-evidence – specific issues that might undermine the argument.

2.6.6 Relatively recently, the international airworthiness authorities have produced a variant of the software guidelines for aircraft for the development of air traffic management software, known as ED109[15]. Although not strictly required by the CAA, NATS have assessed conformance to ED109 in their assurance argument, showing that their processes and their application conform both to the UK requirements (SW01) and European, or International standards (ED109).

2.6.7 The software processes and assurance arguments have evolved and improved over time, and the versions for the current build (known as N38) show a mature approach that compares well with the practices in other high-hazard industries.

2.6.8 Although the processes and assurance arguments are important, what ultimately matters is the integrity of (or the level of faults in) the software that is deployed. Software fault density is measured in terms of faults per thousand SLoC, often written kSLoC. When figures are given for fault density they may represent known faults that were not removed as they were deemed not to be critical, faults found in service, or a combination of the two. It is hard to get definitive data, but figures from the more mature industries suggest that fault density for the best in class safety critical code has improved over the lifetime of NERC from about 1 per kSLoC to 1 per 10kSLoC, for faults found in service.

2.6.9 To appreciate how such figures are derived, it is important to understand the way in which known faults are managed. NATS and their suppliers use a process based around so-called Problem Trouble Reports (PTRs). When an issue is found in system test or in operation, a PTR is "opened" and the issue is assessed for severity, i.e. the degree of impact on operations. Severity 1 (the most critical) includes those with safety impact. Severity 1 PTRs are always resolved, and usually rectified during the current build. Lower severity issues are assessed and they may be rectified, or may be mitigated in other ways. This form of process is widely used in high-integrity software development.

[12] CAP 670 Air Traffic Services Safety Requirements, Part B, Section 3, Systems Engineering, SW01: Regulatory Objectives for Software Safety Assurance in ATS Equipment, CAA, May 2014.

[13] Acceptable Means of Compliance to CAP 670 SW 01, Guidance for Producing SW 01 Safety Arguments for COTS Equipment, CAA, March 2010

[14] Software Safety Assurance Argument –SO518/SAF/02 Issue 11, October 2014.

[15] ED109A Software Integrity Assurance Considerations for Communication, Navigation, Surveillance and Air Traffic Management (CNS/ATM) Systems, EUROCAE, January 2012

2.6.10 On this basis the fault injection rate over the last five builds of SFS is under ten per kSLoC, so below the lower end of the range quoted at 2.6.9 above. Based on PTRs raised for in service issues, and only considering the code developed or changed during the last five builds, the fault density is around 2 per 10 kSLoC, so comparable with good safety critical code. There is an important caveat; these are faults known to date, and there may be unknown residual faults that will be discovered during the on-going usage of the system and software. However there is no indication here that the software being deployed is unsatisfactory in terms of expected levels of integrity.

2.6.11 The final aspect of how faults (and integrity of the software more generally) are managed is the software release process. There is a formal process of assessing the status of the software, and a range of other factors (see below) to establish whether the software is fit for release to operations. As with the software development process, procedures have been in place for some time, and have evolved over the years. The faulty module, `Waafu28`, was signed off into service in 1997, using a fairly simple "cover sheet" for the technical documentation. (It was done in 1997, as the software that went live in 2002 was built up incrementally, and there were no changes to the part of code covered by this release sheet, before the system went live.)

2.6.12 The release process has matured over time and there is now a robust and comprehensive process. A collection of about 80 documents is brought together to support the release, including the system safety case, the software assurance argument, and human factors analyses. This collation of evidence is then formally approved for release using a set of certificates, with signatories including engineering, safety and ATC (to attest that the relevant training has been done). To indicate the robustness of the process, the N38 build was not released on the original schedule, as not all of the certification and underlying material was sufficiently mature. There are "streamlined" versions of the process for rapid rectification of problems; this was used to correct the software fault in `Waafu28` (the fault was rectified on 18 December 2014 and deployed to the operational system overnight on 7-8 January 2015).

2.6.13 The release for N38 covered the operational changes to support the military task that was relocating. It addressed ATC changes (including military issues), human error assessment, software changes and safety assessment (amongst other things). Thus it addressed the key changes, and associated potential risks, that led to the Incident on 12 December 2014. The coverage of these issues was appropriate, and it is unrealistic to expect that a release process could have prevented the Incident as the proximate cause of the Incident was not known in advance.

2.6.14 In summary, the NATS processes, as defined in POD SW01, are thorough and professional and were ahead of their time, when first developed. They meet the requirements of SW01 and the newer ED109 and have generally withstood the test of many releases over a long period. The resultant fault density appears better than would be expected for software of this criticality (with the caveat that there may be other undiscovered faults).

2.7 Software System Design

2.7.1 The Watching Mode is not needed operationally. It was initially used for training and familiarisation (in system "work up"), and was retained for presentation purposes, but cannot be used for controlling aircraft. This leads to the question of why the Watching functionality was implemented in such a way that there was the potential for complete loss of the SFS. The analysis starts by considering the system architecture and its approach to managing redundancy, the workstation interface, and then considers the possibility of removing the Watching functions.

2.7.2 The system architecture is designed to manage and recover from hardware failures, including loss of communication, loss of workstations and loss of servers, and it employs hardware redundancy to do so, see figure 2.1. This redundancy strategy has been a long-standing aspect of the system design, and effectively accommodates hardware failures. However a problem could arise if the SFS and workstation were "out of step". As the workstation originates commands, it has to "replay" (resend) these in the case of an SFS failure. To implement this, the workstations hold a "retained commands" list; this is automatically resent if there is a switch over between primary and secondary SFS.

2.7.3 Normally this strategy allows the system to maintain availability through failures of individual hardware elements of the system. The success of this strategy is shown by the fact that this is the first double SFS failure since the system went live. However, because the fault was in the SFS software, repeating the "Watching" command led to the double SFS failure. As "Watching" is not needed operationally it is important to consider ways in which the problem might have been avoided.

2.7.4 Not all commands go into the "retained commands" list. It seems obvious that the command to enter "Watching" should not go into the "retained commands" list. However, "Watching" is not a separate command, but is represented by elements of a message that transmits many (although not all) of the commands. To remove "Watching" from the "retained commands" list requires writing additional code to identify the "Watching" commands (analysing the message format). With hindsight, this might have been effective, but the importance of removing the "Watching" commands was not understood prior to the Incident on 12 December 2014, see below.

2.7.5 Before considering the advantages and disadvantages of modifying the design more fully, it is appropriate to consider the Human Machine Interface (HMI), recalling that the command to enter "Watching" state was given inadvertently. It is understood that a full HMI design evaluation was carried out at the time of the initial design, including the use of Human Factors (HF) experts. This is good practice, but it should be noted that standards for HMI design have evolved, and this is relevant for future developments.

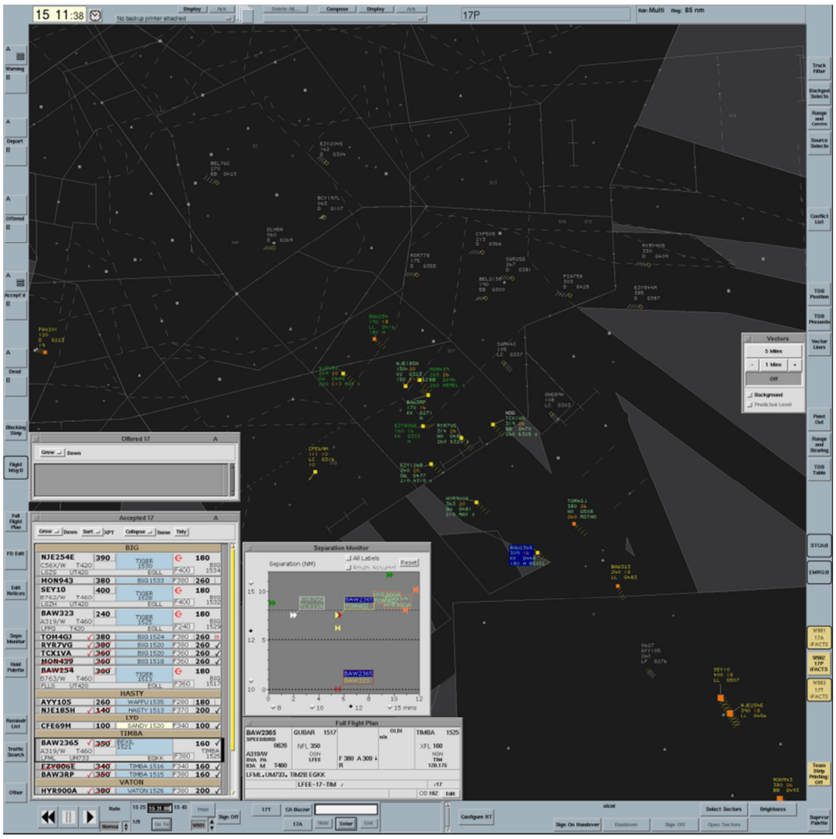

Figure 2.2: General Screen Layout

2.7.6 Figure 2.2 shows the general screen layout at a Controller workstation; the aircraft tracks and other data are shown on the main screen. The border contains "soft keys" and displays status information, including turning from grey to brown to indicate system problems. The overlays show additional information, e.g. conflict alerts from the software package known as iFACTS (interim Future Area Control Tools Support) which is designed to assist Controllers by providing predictive information 18 minutes ahead of an aircraft's current position. The "soft keys" are on a button palette at the bottom right of the screen; this is shown in more detail in figure 2.3.

2.7.7 When the Controller comes to a workstation that is available for use, the button palette would be as shown at the top of Figure 2.3. If the Controller presses "Sign Off" the workstation tries to enter "Watching" mode. The middle button palette shows how the screen would have appeared after selecting (or pressing) "Sign Off"; the bottom palette shows the result of selecting "Open Sectors" i.e. what the Controller should have selected in order to take over control of a sector.

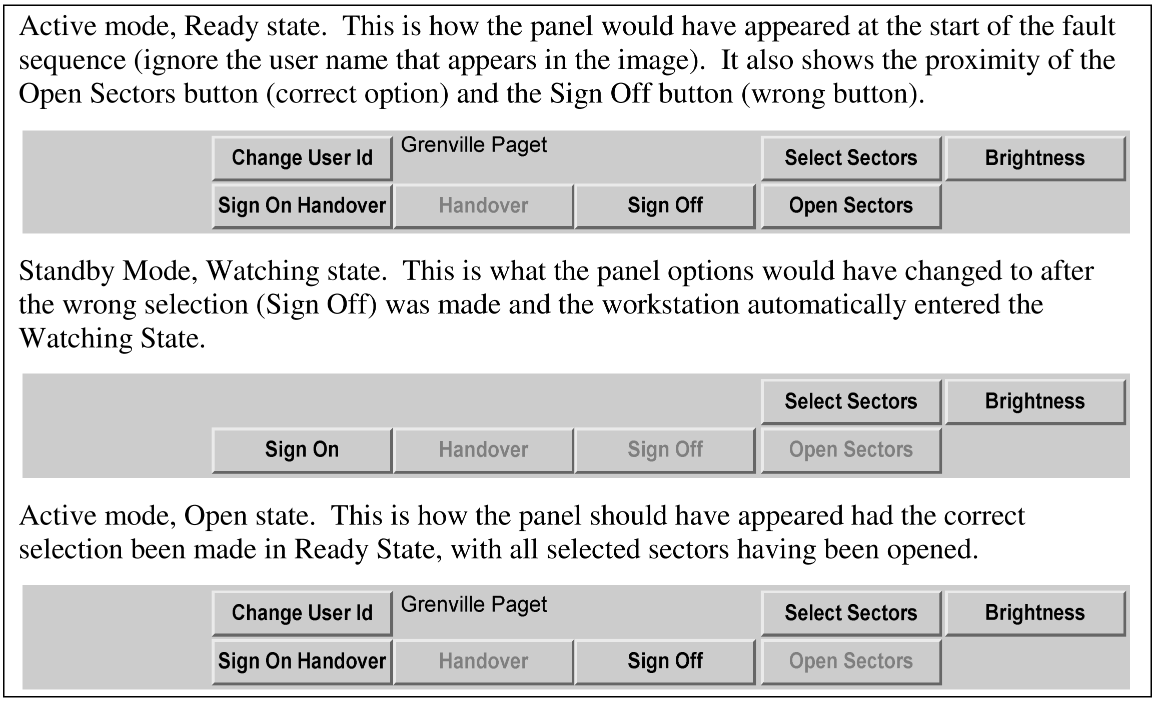

Active mode, Ready state. This is how the panel would have appeared at the start of the fault sequence (ignore the user name that appears in the image). It also shows the proximity of the Open Sectors button (correct option) and the Sign Off button (wrong button).

Standby Mode, Watching state. This is what the panel options would have changed to after the wrong selection (Sign Off) was made and the workstation automatically entered the Watching State.

Active mode, Open state. This is how the panel should have appeared had the correct selection been made in Ready State, with all selected sectors having been opened.

Figure 2.3: Different Configurations of the Button Palette

2.7.8 When it was decided that "Watching" was not required operationally (in 1998) the procedures were altered so that Controllers knew not to use this mode. However, because the soft "Sign Off" button on a Controller's workstation screen is adjacent to the "Select Sectors" and "Open Sectors" buttons (see figure 2.3) the "Sign Off" button can easily be pressed by mistake. NATS data (analysis of logs) now suggests that this error occurs a few times per week, and it had occurred earlier in the day on 12 December 2014. However, it was not generally known that this was occurring until analysis of the logs had been carried out (see below). Again, this emphasises the fact that a combination of factors was needed to cause the failure; the software fault was not sufficient in itself to give rise to the unavailability of the SFS.

2.7.9 Whilst it may seem obvious to remove unwanted functionality it is not so straightforward, and it is instructive to consider the options. The interface could be altered so that the "Sign Off" button was moved to reduce the chance of error. This would be of limited benefit, as Controllers could still make mistakes, and it is possible that such a change would simply make some other mistake more likely (and the consequences might be worse). Thus such a change would have to be investigated, and it is also worth noting that the screen technology and layout used limits the design options.

2.7.10 In principle, it would be better to modify the interface to prevent the workstation from issuing the "Watching" command, but without removing the code in the SFS that implements the command. This would be the simplest change that would "disable" the event sequence that occurred on 12 December. However the software development standards say that the programs should not contain any unreachable code – but there would be unreachable code in the SFS, so this option conflicts with good practice/standards. It might be possible to argue (in the system safety case) that the code had been reachable, so there was not a problem in leaving it in the system but, at minimum, there are challenges about following this strategy (e.g. defining a good rationale, and getting it accepted by the CAA).

2.7.11 The more "complete" strategy is to remove all the code that implements the move into "Watching" mode. This amounts to about 2-3kSLoC, distributed amongst many modules in the SFS and workstation code. Removing this code would "disable" the event sequence that occurred on 12 December 2014, but removing it is not without risks. Even though removing code ought to be easier than writing it, it has to be assumed that new faults could be introduced, and there would be a significant testing burden in checking that the change had been made effectively. POD SW01 requires re-inspection of changed code, and testing of both the modified and unmodified parts of the changed modules in the software. Thus removing the code implementing the "Watching" mode would be likely to involve several person-years of effort.

2.7.12 So far as can be ascertained there was no explicit evaluation of the cost of removing the code implementing the "Watching" function, but the decision in 1998 that Watching Mode was not required and only to change the procedures (paragraph 2.7.8), implicitly viewed software change as not being cost-effective. It is understood that Controllers reported that "Watching" was entered occasionally but it was not until after the event on 12 December 2014 that NATS inspected the logs and discovered that this happened several times a week (roughly every other day, see 3.6.2 for a discussion of error reporting). It is possible that identifying this fact earlier would have led to a different decision about disabling "Watching". However it is equally possible that an alternative view would have been taken; the event is occurring quite frequently without deleterious effect, so there is no need to alter the design to prevent the unintentional triggering of "Watching" mode.

2.7.13 In summary, it would have been possible to modify the system to disable or remove "Watching" mode. However NATS were not aware, prior to the Incident on 12 December 2014, of the consequences of leaving this capability in the system. Further, removing the capability is not straightforward, and there has to be an evaluation of the costs and risks of making such a change, as it is likely that there will be unintended consequences, in a system of such complexity. The decision to make a procedural change to limit the use of "Watching" is defensible, given the potential complexity and cost of changing the software, that no harmful effects had been observed and the facts that NATS has an obligation to keep downward pressure on its costs.

2.7.14 Finally, NATS have now fully resolved the problem by rectifying the proximate cause of the failure (by correction of the code implementing the check of the maximum number of permitted Atomic Functions), rather than making a more major change to disable "Watching". This is the cost-effective and timely solution to the problem; it was possible to make this correction in days, rather than the months of design and testing work required to disable "Watching" altogether.

2.8 Conclusions

There are ten key conclusions relating to the SFS failure and the recovery, from the software and systems engineering perspectives.

2.8.1 The proximate cause of the Incident on 12 December 2014 was a fault in the SFS software, specifically the use of an incorrect check on the maximum permitted number of Controller and Supervisor roles. This fault alone would not have led to the loss of availability of SFS. Other factors were:

- The inclusion on 11 December 2014 of additional (military) Controller roles;
- The Human Machine Interface (HMI) design;
- A Controller inadvertently selecting "Watching" mode;
- The system architecture, especially the approach to handling failures.

2.8.2 The software development process is both professional and in tune with good practice – particularly at the date it was initially defined. The low level of faults in the code (fault density) reinforces this positive impression; based on known faults it is better than industry norms for software of this nature. However the difficulty of finding records of both the initial review and testing of the relevant code (they had to be retrieved from the USA, and this took several weeks) tempers this generally favourable conclusion.

2.8.3 The system and software development processes offered a number of opportunities to find the fault. The most likely places for detecting the fault were in manual module inspections or module (unit) test, with inspection perhaps providing the better opportunity. The extensive test and development facilities run by NATS would potentially be a means of detecting such faults, but this would require the facilities to be modified to be fully representative of operations, and to be capable of "stressing" the system at its boundary conditions.

2.8.4 In principle, the fault could therefore have been avoided; in practice it is infeasible to remove faults entirely in systems of this size and complexity. It is difficult to analyse and test such systems as they are initially developed, as not all possible situations can be addressed in any realistic timeframe. This problem is compounded by the need to re-assess the system as the environment, or its usage, changes, as it did in this case. The release process draws together all the information to support a new release; the process is effective and it is unrealistic to expect that any such process could have prevented the Incident on 12 December 2014, as the proximate cause of the failure was not known at the time.

2.8.5 The Watching Mode was used early in the system life, but in 1998 (well before the system entered service in 2002) was determined as not being needed to support operations. However, the decision to retain Watching Mode was not unreasonable, given the information available prior to the Incident on 12 December 2014. NATS was not aware of the software fault that was the proximate cause of the Incident, nor did they know that the inclusion of military Controller roles on 11 December 2014 resulting in a total of controller and supervisor functions greater than 151 during the afternoon of 12 December, together with the erroneous activation of Watching Mode, would trigger an SFS failure.

2.8.6 The HMI design contributed to the failure. Some reasonably simple changes, e.g. rearranging the soft keys, and/or using different names for keys in different modes, could have reduced the likelihood of the inadvertent selection of Watching Mode. However, any change would require careful evaluation to make sure there were no undesirable unintended consequences – a far from trivial task. Further, changing the HMI would require software modification, retesting of the software, and additional Controller training.

2.8.7 The system architecture was intended to provide resilience against hardware failures. There are two SFS systems, and the workstations (reasonably) replay unexecuted commands to the second SFS server if it takes over duty when the first one fails. In this case, as the fault was in the SFS software, replaying the commands led to the second SFS system failing in exactly the same way as the first, resulting in the complete loss of SFS. Although the architecture did not help in the case of the software fault it is very effective in dealing with hardware failures.

2.8.8 The automatic system logging, and the ability for the logs to be made available to software suppliers in real-time, is a very valuable feature of the system design that enabled the fault to be located very rapidly. This was also supported by the highly collaborative culture between NATS and its key suppliers who worked as an effective team despite the geographical spread (UK and USA).

2.8.9 The Engineering Technical Incident Cell (ETIC) was a key factor in the effective management of the technical recovery. It was populated with highly competent engineers, who had a very deep understanding of the systems and had the right attitude in terms of assessing and managing risk. They operated a consensual decision making process which helped to ensure that appropriate decisions were made and they were careful in the course of action taken whilst the proximate cause of the failure remained to be identified.

2.8.10 Finally, and perhaps most importantly, it should not be presumed that any complex software system will be fault-free, and care is needed in the design of the broader system (procedures, training, etc.) to make it resilient to latent faults. In other words, for the system overall, it is important to understand how unknown faults are mitigated, as well as considering how fault density can be reduced, or minimised.

2.8.11 In summary, NATS' processes are thorough and professional and are believed to have been ahead of their time when first developed and they meet the requirements of the CAA. There is a strong and effective process for controlling software releases, requiring signatures from key stakeholders. The resultant integrity appears better than would be expected for software of this importance. However, there will be major changes in NATS systems over the next few years with the deployment of new pan-European systems (known collectively as SESAR). The current processes will need updating to address new factors that will arise in this international collaboration, including: achieving a stable requirement agreed with other European Air Navigation Service Providers (ANSPs); a new focus on NATS capabilities in a system and software programme management role; and significant changes to suppliers.

2.9 Recommendations

NATS are upgrading their capabilities to support new developments in airspace management, including a European programme known as SESAR (see chapter 4 for an overview). The following recommendations cover things that are being done well at present that should not be lost in SESAR, and some further actions that might help reduce the likelihood of an Incident such as occurred on 12 December 2014. Additional detail to help implement the recommendations is presented in Annex G.

R1. NATS should retain, for their deployment of SESAR:

- The system architecture approach, including hardware redundancy and the fault management capabilities to provide resilience in the presence of hardware and software failures and operator errors associated with configuring the system;
- The automatic logging of system behaviour, including Controller commands, software failure conditions ("exceptions") and hardware failures;
- The provision of real-time access to these logs and other system data by software development and support staff;
- The ETIC and its important role in managing technical failures and their recovery;
- The collaborative culture between NATS and its suppliers, integrating new suppliers into established information sharing mechanisms. (Para 2.8.2, 2.8.8)

R2. NATS should ensure that the approach to HMI for SESAR draws on modern design principles, enabled by the use of new display technology, to reduce the possibility of operator error. (Para 2.8.6)

R3. NATS should ensure the effectiveness of the software development process for SESAR, in particular how they specify and manage contracts on their suppliers, incorporating methods:

- Based on experience from the current software development, e.g. well established code inspection techniques;
- Reflecting advances in software technology. (Para 2.8.3)

R4. NATS should ensure that contracts or other suitable arrangements provide a complete, continuing evidence base for the current operational software. This should be demonstrated throughout the remaining life of the system by audits of the software development records and NATS should ensure that identified discrepancies or omissions are resolved. The reviews should be timed to support the five-year planning cycle, and instigated as necessary to respond to a perceived risk, e.g. an accumulation of change. (Para 2.8.2)

R5. NATS should, so far as is practicable, avoid including non-essential functions in SESAR at all, or separate them out from the rest of the code so that they can be "shut down" independently of the operational (main) functionality, or removed entirely, in the event of problems. (Para 2.8.5)

R6. NATS should consider the costs and benefits of developing a facility that could test the SESAR system at full capacity, and at the boundary conditions. (Para 2.8.3)

R7. Since SESAR is a collaborative programme with other Air Navigation Service Providers (ANSPs), involves other Regulators and introduces a new set of suppliers, NATS should ensure that through effective consultation and enlisting assistance from the CAA as necessary, it can achieve to the maximum extent possible the functional and non-functional requirements as well as the sound development practices that would be considered necessary in a national procurement. (Para 2.8.11)

R8. NATS should consider the costs and benefits of adopting any of the identified modern testing and software assurance methods on a targeted basis for the current LAC software environment, including determining whether or not the return on investment, e.g. from using modern static analysis of software code, is likely to be worthwhile, and the benefits of de-risking new tools on the current software before using these tools within SESAR.[16] (Para 2.8.11)

[16] Assessing the potential benefit and costs of applying modern software engineering techniques to the suite of software that supports NATS operations is a complex task, beyond the scope of this enquiry. It is therefore appropriate that NATS, who have both access to the software and the domain expertise, assess where it is useful to apply such techniques, and to consider the comparative value of investment in the current systems, the development of SESAR, and other initiatives, e.g. improving operational response to events. This assessment should be subject to review by the CAA as part of their oversight responsibilities.

Chapter 3. Operational Response and Recovery

3.1 Operational Response

3.1.1 To ensure ATC operations can continue to be conducted safely and effectively in abnormal circumstances, NATS document reversionary procedures for all anticipated failure modes via a series of fallback checklists and periodically exercise them. In the event of a major system failure, the principles applied by NATS are to secure the safety of the operation through a reduction in traffic whilst considering the correct course of action to maintain the operation in a safe steady state, and then achieve full system and traffic recovery.

3.1.2 At 1444 on the afternoon of 12 December, LAC controllers were alerted to a system failure by a warning at their workstations. The London Area Control Supervisory team were directed subsequently by the Operations Supervisor (OS) to follow fallback Checklist 4 for 'SFS Unavailable'. Loss of SFS alone does not present an immediate safety risk. The fallback procedures allow controllers to continue controlling in a reversionary mode using still available, but degrading, flight data for up to an hour. However, during the Incident, controllers were required to rely more on voice communication between air traffic agencies and on their practical air traffic management skills to coordinate flight data manually across air traffic sector boundaries, significantly increasing controller workload.

3.1.3 In response to the failure indication, ATC tower supervisors at airports within the London Terminal Manoeuvring Area (TMA) were instructed to stop all departures immediately and a tannoy summoned all operational staff present in Swanwick Centre to the LAC and London Terminal Control (TC) Operations rooms. A number of additional experienced senior controllers and technical experts were available in the Centre at the time and responded to the tannoy. Their presence, experience and expertise were viewed as key to the subsequent speed of analysis and decision-making on the LAC Ops floor.

3.1.4 The NATS National Airspace System (NAS) provides a central air traffic data hub that supports operations across multiple NATS ATM services, including LTC (traffic generally below 21500ft), Prestwick Centre (PC) and other local NATS units at airports. Its shutdown would seriously impact operations across the UK FIR. At 1450, anticipating a shutdown of the NAS as part of the technical recovery actions, the LAC OS forewarned TC of the possible need to revert to full manual control. It is not clear why the OS was concerned about an impending shutdown of the NAS. The failure indication presented at the workstations at 1444 was 'SFS Unavailable' (fallback Checklist 4), followed 2 minutes later by 'NAS Link Unavailable' (fallback Checklist 7) [17].

3.1.5 In their feedback to NATS' internal investigation[18], both the OS and the Swanwick Airspace Capacity Manager (ACM) refer to the Checklist 4 fallback procedure including the shutdown of NAS. However, such a procedure is neither referenced in the UK Flow Management Position (FMP) Emergency and Fallback Guidance (EFG), nor in any of the applicable role-specific Checklist 4 versions[19]. Moreover, had there in fact been a NAS failure, as opposed to a NAS-SFS Link or SFS failure, the EFG states that the Data Systems Supervisor (DSS) should have been expected to confirm this within 3 minutes of the failure. Nonetheless, it appears that the perception of a NAS failure and/or of the recovery involving shutting down the NAS may have led the OS and ACM to implement more conservative restrictions than those required by the EFG. This will be discussed in more detail below, but of note is the fact that the immediate actions for a NAS failure differ from those for NAS Link or SFS failures, in that the former include a direction to stop "all UK departures" and the latter refer to stopping "departures entering the London FIR[20]."

[17] During Nov 14, a 'NAS/SFS link failure' scenario requiring action of Checklist 7 was exercised at Swanwick.

[18] NATS SP301 Major Incident Investigation – SFS Failure 12th Dec 2014 – Preliminary Report V2.0, Jan 15.

[19] Tailored role-specific versions of Checklist 4 are provided, such as for OS, Local Area Supervisor, Tactical Controller, Planner etc.

[20] Encompassing the airspace over Wales and the southern half of England.

3.1.6 In addition to the immediate suspension of departures at London TMA airports, and in accordance with the EFG "London Area in Manual Mode (LAIMM)" checklist, the ACM took action to reduce traffic levels by instructing Eurocontrol's Network Management Operations Centre (NMOC) in Brussels to implement a pre-determined set of 4 air traffic regulations[21], the first effective from 1500 and the remainder from 1530. As a result, a zero flow rate restriction (colloquially referred to as a 'zero rate regulation' (ZRR)) was applied to the London FIR from surface to unlimited altitude. The ZRR was initially prescribed for a 4-hour period, with the intention of ensuring that it immediately stopped all departures from across the wider European Civil Aviation Conference (ECAC) region that were destined to transit London airspace en-route or terminate at London FIR airports. This therefore included departures from UK airports, other than Heathrow, Gatwick and Manchester, which were specifically excluded from the broader ZRR and were dealt with individually, and those from more distant airports with longer flight durations. It did not apply to non-European/long haul departures destined to transit the London FIR, but had the aim of significantly reducing the volume of planned traffic arriving into London airspace for the stated period.

3.1.7 The 3 standard contingency restrictions that were applied in parallel to departures from Heathrow, Gatwick and Manchester are each defined as taking effect from +45 minutes from implementation and thence for 3 hours 15 minutes duration[22]. However, the initial suspension on departures communicated verbally direct to tower supervisors was still in place at this stage, so the period of grace built into the formal airport-specific regulations did not take effect on the day, the significance of which is examined further below. The ACM also requested NMOC to promulgate a pre-scripted Air Traffic Flow Capacity Manager Information Message (AIM), identified as "LAIMM1" in the NATS EFG. As defined in the EFG, LAIMM1 refers to LAC "operating at reduced capacity" and also reports that TC "… is operating but … restricting departures at this time". However, the LAIMM1 label was not immediately recognised by NMOC staff. Once they had identified the correct template, NMOC promulgated an interpretation of it based in part on their understanding of the regulations that were by then in place, initially reporting the airspace as closed.

3.1.8 In the meantime, as the initial fault diagnosis proceeded on the LAC Ops floor, uncertainty as to whether the failure had occurred in the NAS or the SFS, and indeed whether or not a Checklist 4-led fallback was the correct option, continued. ETIC was convened at 1505 and shortly afterwards, at 1522, ETIC concluded that the NAS was in fact serviceable; that it should be left in operation; and that the SFS could be repopulated with NAS data. Amongst other things, this had the benefit of confirming the failure's implications and that they were confined to LAC operations, thereby removing the perceived threat of LTC, PC and other NATS ATM services having to revert to degraded manual control. At 1530, the ACM was advised by ETIC to prepare for a full flight database download from NAS to SFS Server B. Only after ETIC had finally confirmed the nature of the failure to the OS and ACM, around 40 minutes after the initial warning indication at the workstations, did they feel that the situation was well enough understood to consider acting on a lesson learned from the December 2013 system failure and manage LAC traffic tactically via LTC and PC's respective airspace areas of responsibility. By this stage, the regulations requested previously by the ACM had been advertised and implemented by NMOC and were taking effect across the European Network.

3.2 Operational Recovery

3.2.1 Following technical recovery of SFS Server B, the system was refreshed with flight data from the NAS. Full LAC ATM functionality, albeit limited initially to the configuration of

[21] The restrictions (regulations) requested were EGT1ACC, effective 1500; EGLLDEP (Heathrow), effective 1530; EGKKDEP (Gatwick), effective 1530; and EGCCDEP (Manchester), effective 1530.

[22] On the day, the ACM instructed NMOC to activate the 3 major airport restrictions from +30 mins, in a modification of the standard contingency regulations.

controller sectors extant at the time, was restored one hour after the initial failure. A 10-minute stabilisation period confirmed that all systems were functioning correctly, following which the ACM began to manage recuperation of traffic flows within LAC, TC and TMA airports.

3.2.2 At 1555, regulated departures were authorised by Swanwick from Heathrow, Gatwick and Manchester airports, initially at an extended Minimum Departure Interval (MDI) of one every 5 minutes (both northbound and southbound for Heathrow). However, acting on Engineering advice not to risk further failure by changing the configuration of the system, traffic capacity increases were capped to ensure they remained within the capability of the in-use configuration of controller sectors. During an Air Traffic Incident Coordination and Communication Cell (ATICCC) customer conference call at 1605, the NMOC was informed that the "EGT1ACC" ZRR and London airports' departure suspensions could be cancelled. The Swanwick ACM continued to restore traffic flows, regulated initially at 75% capacity[23] of the in-use sectors, striking a balance between a graduated increase and re-introducing traffic into the system as expeditiously as possible. By 1610, traffic was permitted to depart from the rest of the London TMA airports with MDIs applied. To assist with clearing the backlog of aircraft waiting to depart, the ACM negotiated slot extensions for Heathrow, Gatwick and Manchester departures with NMOC.

3.2.3 As many of the aircraft already airborne when the failure happened had continued to arrive at Heathrow, parking and stand availability became a critical issue and, following a request from the Heathrow Operational Efficiency Cell (HOEC), the Heathrow arrival rate was reduced to 20 per hour by NMOC at 1626. At 1645, Swanwick increased the Heathrow departure rate to one aircraft per 5 minutes on each of the six standard departure routes and at 1650, reducing congestion on the ground at Heathrow allowed the arrival rate to be increased to 40 per hour. By 1700, the Heathrow departure rate had been improved to one per 4 minutes per standard route and by 1730 all departure restrictions were cancelled for Heathrow, Gatwick and Manchester. Once Engineering had rescinded the instruction not to reconfigure sectors, controllers were authorised to split sectors as required from 1720 and by 1750 all en-route restrictions had been cancelled, with the exception of the Dover and Clacton sectors. It was not until 1935 that the remaining Heathrow arrival restriction was cancelled and the final flow restriction, Clacton, was removed at 2030.

3.3 Operational Control and Communication

3.3.1 Swanwick Silver (operational level) command team and the ATICCC were invoked at 1500. During the first internal ATICCC call at 1545, Silver confirmed that the secondary SFS was recovered and stable and that Ops would commence a graduated lifting of restrictions, including the ZRR, at 1605. However, at this stage, the root cause of the double SFS failure had not been identified. The 1600 ATICCC conference call was the first formal status report to customers, during which they were advised that the system had been restored and was stable, and that the recovery was underway, initially increasing to 60% of capacity, with significant delays expected over the next couple of hours. In addition, Swanwick requested NMOC to implement a series of flow rate regulations to support the graduated recovery. NMOC cancelled the initial 4 contingency regulations at approximately 1610.

3.3.2 Subsequently, at 1619, Silver team were apprised of the possibility that exceeding the permissible number of Atomic Functions was related to the system failure. At the Engineering Bronze (tactical command level) teleconference at 1630, the risk of recurrence of the failure was assessed as High, as the root cause was not fully understood and the engineering design team could not yet be specific with recommendations. By 1715, after further analysis by the engineering design team, the risk of recurrence was reassessed as Low, as the role that Watching Mode had played in the failure had been identified and it had since

[23] The same flow rate as used routinely in poor weather.

been confirmed that all workstations were signed on. Nevertheless, the root cause had still not been identified at this stage. Silver Chair directed that a temporary OPNOT[24] be rolled out for the weekend, that additional signage be provided to the LAC and military sectors and that the 'First Brief' system for oncoming controllers be updated to remind them of an existing Temporary Engineering Instruction (TEI) not to select Watching Mode.

3.3.3 At 1936, Silver Chair asked for engineering advice on the number of Atomic Functions in use and how he could ensure that headroom was maintained in the system. He was briefed on the maximum number of available Atomic Functions and the latent defect that reduced this to 151 if a terminal is in Watching Mode. At the time of this brief, it was reported that 142 Atomic Functions were active on the system.

3.3.4 Final restrictions on Heathrow arrivals were removed at 1955 and ATICCC was deactivated at 1956. During an Engineering teleconference at 2045, the number of Atomic Functions active was revised up to 153 and the fact that the system was therefore still exposed to the risk of SFS server shut-down if Watching Mode was entered resulted in Silver Chair direction to immediately remove 7 of the currently operating Atomic Functions. He also requested assurance that the 151 (Watching Mode) Atomic Function limit would not be exceeded. Arrangements were then put in place for the Service Manager to monitor the number of Atomic Functions in use in near real time, two sector terminals (OV2 and OV3) were cordoned off and additional signage was erected in the Ops Room. By 2124, the number of active Atomic Functions in AC had been reduced to 147 and was being actively monitored. Silver team stood down at 2140.

3.4 Effect

3.4.1 Acknowledging the commendable speed of the operational and technical responses, five specific aspects of Swanwick Centre's actions appear to have had consequences for the scope and severity of the Incident's impact, and the ease of the recovery. These were:

(1) Stopping departures at Heathrow, Gatwick and Manchester airports.

(2) Perceiving a need to conduct a NAS recovery from SFS data.

(3) Instituting a comprehensive ZRR for all London airspace.

(4) Initially applying all contingency regulations for 4 hours.

(5) The NATS-led 'generic' recovery.

3.4.2 **Suspension of Departures**. Almost immediately after the failure was indicated, Swanwick supervisors communicated directly with NATS Services ATC tower supervisors at London TMA airports by telephone and informed them that departures within and destined for the London area were to be suspended with immediate effect, although the anticipated duration of the suspension was not specified. At the London airports, there was understandable uncertainty as aircraft that were taxying were denied take-off clearance and instructed to hold their positions for an indeterminate period. The logic of stopping departures at all London airports for a brief period whilst the failure was positively identified can be acknowledged readily, but a sustained suspension is not a requirement of the fallback procedure applicable to and adopted for the actual failure on the day. Moreover, the EFG states that any potential confusion as to whether the NAS or SFS systems have failed should be cleared by the DSS within 3 minutes. Notwithstanding, the initial verbal suspension of departures was not rescinded when the 4 formal contingency regulations were applied and, consequently, the window in the latter designed to accommodate continued departures at Heathrow, Gatwick and Manchester for the first 45 minutes was not available to those airports and their operators. Moreover, the initial confusion over the status of the NAS and the perception that

[24] Operational Notice - OPNOTs may contain information and/or guidance relating to ATC procedures, but must not contain instructions.

all UK ATM operations might have had to revert to manual control appears to have distracted Swanwick supervisors from substantive consideration of implementing level capping procedures quickly, in order to exploit LTC and PC capabilities to help sustain a departure flow from the major London TMA airports. This was a procedure envisaged under the "lessons identified" from the 7 December 2013 communications failure incident, albeit in the context of longer term disruption, and is considered further below. Once preparations began for the recovery of full LAC functionality that commenced at 1605, it became the focus rather than any attempt to employ level capping in support of continuing departures. In all, therefore, up to 1 hour 15 minutes of potential departures from Heathrow, Gatwick and Manchester were lost.

3.4.3 Applying an instant halt to all departures from the major London airports accelerated congestion significantly, most notably at Heathrow, both for aircraft, airports and passengers. It is likely to have made the recovery more challenging than it could have been.

(1) Heathrow is the UK's major hub airport, and routinely operates at around 90% of aircraft stand and parking capacity[25] with up to 88 movements per hour. Between 1450 and 1635 there were no pushbacks from Heathrow terminals as taxiways were congested with aircraft that had taxied just before the system failed, parked aircraft awaiting stands and arriving aircraft being distributed to off-stand parking across the airfield, affecting some 120 aircraft during the period. Around half of these were departing aircraft that then had to be sequenced into the staged recovery when it was initiated at 1605. Even in a relatively short period of time, congestion at the airport became severe, passenger check-in was suspended and airport security areas were closed to assist in managing the growing volume of airside passengers.

(2) At Gatwick, the Panel was informed that the airport had been unaware that the contingency restrictions were initially specified for 4 hours. The airport's management stated that had they known this, they would have moved to de-plane passengers as they employ a 2-hour cut-off for keeping passengers on their aircraft awaiting departure. In the circumstances, this was an unintended benefit of the lack of understanding at Gatwick, since deplaning would have greatly increased the impact of the event locally and would have placed severe stress on the airport's ability to manage passenger volumes, both on-site and those arriving to fly.

3.4.4 The effect on airlines was of course immediate, but varied across the type of operation, with it being more acute for short-haul/regional operators than medium-to-long haul who are, comparatively, better able to absorb some element of delay. As a result of their frequent short-sector operations, high aircraft utilisation and critical dependence on their aircraft and crews being in the correct location to support their high tempo schedules, short-haul/regional airline operations were both affected more widely by the contingency regulations imposed and also faced strategic decisions more quickly about whether and when to suspend their overall operation to avoid a widely dispersed and incoherent fleet. For example, one major regional operator reported that they cancelled their third of four scheduled waves on 12 December 2014 in response to the Incident, placed their fourth wave on indefinite delay and were forced to contemplate cancelling their following day's complete schedule, a decision that they would have had to make by 2100, only 6 hours after the SFS failure. In the event, and even though the recovery was underway, by the end of Friday 50% of that operator's aircraft and crews were out of position.

3.4.5 Further afield, airports and aircraft operators had to rely on Eurocontrol's Network Operations Portal (NOP) Network News bulletins and AIMs for information and responded according to what they were being told. The Panel have not examined in detail the impact on

[25] By way of example, British Airways' Terminal 5 has 60 stands and operates each of these with a 15-minute interval between departing and arriving aircraft.

affected airlines and airports outside of the UK, but evidence from NMOC suggests that significant delays were experienced by a large number of aircraft throughout the European network.

3.4.6 **Status of London Airspace**. NATS maintain that London airspace was not closed during the Incident and NATS Gold (strategic level) command went to considerable efforts to rebut the suggestion in the media at the time. It is also true that whilst the SFS was down, most already airborne inbound traffic from adjacent ATM Area Control Centres (ACC) was indeed accepted via manual coordination between controllers. However, EGT1ACC, the standard contingency ZRR requested by the Swanwick ACM, is defined as applying to all Swanwick airspace from surface to unlimited. When the SFS fault occurred, LAC controllers on duty transitioned seamlessly to manual controlling of existing and arriving traffic. Nonetheless, in response to NMOC's direct enquiries at 1508 and 1530, it was not until 1535 that the ACM confirmed formally that Swanwick ACC was able and willing to continue accepting arriving traffic[26]. It may therefore be argued that, until then, it was not an unreasonable inference by NMOC that this, coupled with an immediate and enduring suspension of all London departures, and a zero traffic rate applicable to all London FIR airspace, effectively amounted to closure of that airspace. Following receipt of the ACM's confirmation, the NMOC NOP Network News headline message was updated at 1541 to advise that Swanwick would accept already airborne inbound traffic. Nevertheless, during the initial 45 minutes of the Incident, it is likely that ambiguity in Swanwick's communications and actions reinforced perceptions that London airspace was closed and resulted in up to 20 aircraft being diverted pre-emptively to alternative airports and around 150 flights being cancelled.

3.4.7 **Duration of the Regulations Applied**. NMOC's initial perception of the airspace being closed and the fact that the contingency regulations had been applied for 4 hours duration triggered a number of immediate procedural responses by NMOC, including:

(1) An attempt to communicate in plain English the 'facts', as they were understood, as expeditiously as possible to users, as it was critical information that would affect management of the wider Network and the operations of individual operators across the ECAC region, at the very least. This was achieved by:

(a) A NOP Network Headline News bulletin, posted at 1503 and updated at 1507, that reported, "ALL LONDON Airspace closed due to computer failure" and gave the duration as 4 hours. This message was then picked up and reported by media, including the BBC.

(b) Release of an AIM at 1517 that reported the computer failure and that "until the system recovers, area control units in the London does not accept traffic until 1900 UTC initially" (sic).

(2) Immediately suspending the Flight Plans of all affected flights – i.e. London-bound European departures scheduled to depart during the forthcoming 4-hour period - as the duration of the restrictions exceeded a predefined 1-hour threshold[27]. For durations below this threshold, Flight Plans remain live as it is assumed that the overall demand will remain more or less static and that the system will be able to accommodate the relatively modest recovery surge. For longer durations, affected Flight Plans are shifted to the end of the disruption window and aircraft operators are required to submit a Flight Confirmation Message (FCM) to reactivate them, once airspace flow rates are re-established and they have decided whether to operate delayed flights or cancel them. This provides a more accurate picture of actual demand to be managed as a recovery gets underway.

[26] His inability to do so before then was presumably a consequence of the internal NATS debate about whether or not the NAS would have to be shutdown and the potential effect that would have had across the UK FIR.

[27] As defined in the Eurocontrol Network Manager Air Traffic Flow and Capacity Management (ATFCM) Users Manual.

(3) The above action should have been accompanied by NMOC releasing an associated AIM to inform operators but, unhelpfully, this was omitted as the NMOC supervisor on duty felt that sufficient information on the outage had already been promulgated.

3.4.8 **Designing and Controlling the Recovery**. Once the system had been restored, Swanwick initiated and controlled the incremental recovery, initially defining both flow rates for London sectors and complementary departure rates from the London airports. However, the extent to which this was a fully informed and collaborative activity with customers is not clear. The ATICCC Home Page defines ATICCC's roles as, *inter alia*:

(1) Collate information regarding airspace capacity, airport infrastructure and airline demand;

(2) Agree on a strategy for allocating capacity and ensure that capacity is effectively utilised.

3.4.9 Accepting that the situation was fast moving, the four customer calls until ATICCC was deactivated at 1956 were predominantly structured around 'push' communications informing customers of actions taken, or planned, by NATS. There does not appear to have been a formal process to receive, triage and prioritize customer information and requests. It is NATS' view that attempting to respond to individual requests through the ATICCC call would have been unmanageable, given the 150 or so parties who participate. Customers were asked to make any requests for support to specific flights to the FMP. It may therefore be concluded that the NATS-led recovery was largely generic in nature and focussed on re-establishing LAC operations in the round. Indeed, in response to a question from Gold, Silver confirmed at 1727 that there was no prioritising of airports. Yet, whilst Gatwick and the regional airports were undoubtedly under stress, the situation at the UK hub at Heathrow was bordering on critical, with the potential for the airport having to suspend arrivals[28] and the very significant knock-on effect this would have had across the UK and wider aviation system.

3.4.10 Whilst acknowledging the logic of concentrating on recuperating the LAC system as a key component of the recovery, there is perhaps more that could have been done to enhance NATS' understanding of the actual situation on the ground and to permit that understanding to inform a more precisely targeted or tailored recovery. A coherent and shared picture of: the developing situation at the London airports; more widely across the European Network and in the airlines' operations; where pressure points existed or were building; how they might best have been alleviated and in what priority they should have been tackled, may have enabled a more nuanced approach to the application of limited ATM capacity and resources. Certainly, were serious and more prolonged disruption to be experienced in future, such collaborative decision making and crisis management would seem to be essential, not least to the effectiveness of any mitigating actions and the recovery.

3.4.11 The speed of the recovery on the ground at the airports was directly related to how quickly delayed departures could be completed and congestion, affecting both passengers and aircraft, relieved. Gatwick and the regional airports benefitted from a reasonable amount of manoeuvring space headroom to enable aircraft to be sequenced and positioned for take-off as required, but Heathrow faced a much greater challenge, effectively limiting initial movements to 'departure-in-turn' from where aircraft had ended up as the airport came under pressure to accommodate continuing arrivals in parallel with no departures. To compound matters at Heathrow, acting on the initial advice that the restrictions would be in effect for 4 hours, some aircraft had shut down engines on the taxiway to conserve fuel and required a finite time to restart, some required top-up refuelling and there were a limited number of fuel

[28] As it was, arrival rates at Heathrow had to be reduced during the recovery as a result of congestion on the ground at the airport.

bowsers available, whilst others were unaware that their Flight Plans had been suspended and would need to be re-filed, all of which placed an additional drag on the hub's recovery.

3.4.12 The recent introduction of the Eurocontrol Airport Collaborative Decision Making (A-CDM) system at Heathrow and Gatwick should have helped in creating shared understanding across stakeholders – airports, airlines, ATC and the NMOC – of movement priorities, aircraft readiness for departure and Network slot times. Ideally, informed in real time by progress with the London airspace system's recovery, the situation on the ground at the airports and FCMs, the NMOC would then have been able to allocate new Calculated Take Off Times (CTOT) intelligently for departing aircraft from the London TMA airports and contribute more effectively to managing congestion hot spots and Network performance more widely.

3.4.13 However, the initial suspension of all affected Flight Plans by the NMOC in response to the 4-hour ZRR prompted mass cancellations of CTOTs in the A-CDM system and, to some airports and operators at least, gave the appearance of A-CDM not being able to 'keep up' with the crisis. Added to this was the apparent difficulty in re-establishing CTOTs, as not all those affected understood the reason for them having lapsed. In some frustration, both Heathrow and Gatwick Airport Ops Cells dispensed unilaterally with the A-CDM system and resorted to managing departures locally within the MDI rates set by Swanwick and extensions to slot times agreed by the NMOC. This was less efficient than it may have been and had the effect of removing key data and communication pathways to and from the NMOC and the consequential need for increased telephone coordination. Moreover, such a fallback scenario had not been practised and, on the day, 2-way communications between HOEC and the NMOC proved to be severely limited and ineffectual.

3.4.14 As a direct result of the contingency regulations invoked by NATS in response to the failure, and in addition to those flights diverted or cancelled, some 353 flights were delayed[29]. It is estimated that the number of passengers impacted by these initial delays, diversions and cancellations is around 65,000. However, further delays continued during the recovery and into the evening. In total, it is estimated that in the order of 1900 flights were affected by the failure during the afternoon and evening of 12 December, impacting some 230,000 passengers. Additionally, several airlines reported cancellations and flight disruption the following day, with approximately 60 aircraft and 6000 passengers affected.

3.5 Previous Lessons

3.5.1 A previous NATS' investigation into a serious communications system failure that occurred on 7 December 2013 identified a number of lessons and prompted associated recommendations by NATS and the CAA most of which were reported as closed off and in place ahead of this most recent incident. However, amongst these recommendations were three of particular note in the context of the 12 December 2014 failure. The first was to review with stakeholders the industry's ability to respond to service failures and identify required changes to NATS' crisis management capabilities, resilience of systems, procedures and service continuity plans. Implementation of this recommendation was declared complete by NATS on 13 October 2014 and actions agreed included publishing pre-defined contingency route scenarios and addressing how revised routing options would be promulgated. On 8 December 2014, NATS further confirmed to the CAA that existing scenarios had been reviewed, shortfalls identified and new scenarios created where appropriate. In the event, none of the contingency routing scenarios were implemented on 12 December. NATS reported afterwards that, firstly, such scenarios are only appropriate if the incident is long term and, secondly, that the speed with which full technical operation was restored and restrictions were gradually lifted, coupled with the perception at the time that the

[29] According to the agreed process for Control Period 3 of their Licence, NATS reported 14863 minutes of 'Licensee Attributable En route ATFM Delay' for the subject period. This measure is arrived at via a formula defined in the NERL plc Licence, but does not relate directly to the delay experienced by passengers.

NAS would have to be restarted and the consequential effect that would have had on controlling capabilities, militated against activating any of the contingency routing scenarios, such as level capping departures within LTC airspace. However, the legitimacy of the assumption that such contingency routing options only have merit in longer term incidents is challenged by the acute effect departure disruption has on the London TMA airfields and the hub at Heathrow in particular.

3.5.2 The second relevant recommendation, made by the CAA, encouraged NATS to make best use of all means by which a crisis can be handled from an operational standpoint, including exploring the more effective use of and interactions with the Eurocontrol Network Manager (NM). This was also reported as complete by NATS on 13 October 2014, with no further action required. However, the evidence submitted by NATS to demonstrate completion was confined to defining means by which NATS did and would communicate with the NM, but fell short of any proposals to include the NM in informing options analysis or decision making during a crisis and its subsequent recovery. It is therefore evident that the intention of these two recommendations had not been addressed.

3.5.3 Finally, a review of the wider industry crisis response and resilience arrangements was recommended. Invitations to participate in an "industry crisis exercise" were extended by NATS to major stakeholders in May 2013 and the event was anticipated to take place in February / March 2015, although that date has now been postponed until after this Enquiry reports. This is entirely sensible but it will be important that the exercise includes Eurocontrol; given the complexity of planning and exercising such an exercise it may also be worth considering the introduction of external assistance from organisations familiar with such a role.

3.6 Safety

3.6.1 There were no safety events recorded within LAC and LTC during the period of fallback operations or during the recovery phase. Moreover, NATS have a mature Safety Management System (SMS) that is subject to assurance and continual improvement activities. Whilst specific review of the SMS is outside the scope of this inquiry, the improvement activities have recently included an internal safety culture assessment[30] and an independent review of safety governance and oversight[31] that made observations relevant in the context of this incident. These were related to the formality of, and responsibility for, closing out actions arising from safety review and incident investigations; the independence of internal incident investigators from line management in the associated units; and role-specific training for those involved in safety management activities. NATS have assigned responsibilities and initiated work to address the relevant observations.

3.6.2 Notwithstanding, post-incident technical analysis revealed that Watching Mode had been selected accidentally by LAC controlling staff multiple times a week on average in the months leading up to the 12 December Incident, despite the existence of a TEI stating that Watching Mode should not be selected. NATS' SMS includes a facility[32] to lodge a 'Safety Incident' Mandatory Occurrence Report[33] for significant safety occurrences and the reporting of lower level safety events is also encouraged. However, there is no distinct Error Management System (EMS) of the sort employed in other high-hazard industries, whereby occurrences that do not cross the above thresholds are nevertheless captured and the data used to inform independent trend analysis and risk management. It is possible that the existence of such a system would have highlighted the relative propensity for staff to miss-select Watching Mode and that this may have prompted earlier action to mitigate the hazard more effectively (however, see the discussion at 2.7.12).

[30] Safety Culture Assessment – NATS dated 26 Sep 14.
[31] NATS Internal Audit – Review of Safety Governance & Oversight dated Jul 14.
[32] The Safety Tracking and Reporting System (STAR).
[33] As defined in CAA CAP 382 "The Mandatory Occurrence Reporting (MOR) Scheme".

3.7 Conclusions

3.7.1 The timeliness of the response to the failure by Swanwick Centre and NATS staff was impressive and comprehensive crisis management capabilities were mobilised quickly, including support from the contractor engineering design team, some of whom were based in the US, and therefore were in the middle of their normal working day. NATS have a well-established process that aims to ensure staffing levels always meet routine roles and supervisory requirements in the direct conduct of their Operations, and the Swanwick Operational Resource Team was effective in ensuring that appropriate controller cover was in place throughout the recovery phase. However, it is clear that the presence of several additional senior, qualified and experienced personnel who actively contributed to the initial failure diagnosis, supervision and management of the operational recovery in LAC was key, but possibly owed more to accident than design. Ensuring that such experienced staff are available and ready to respond in a timely fashion in the event of an incident is perhaps an area that would benefit from more formal attention as NATS' business continuity plans are reviewed going forward.

3.7.2 The usefulness of aide-memoires and checklists to standardise and guide immediate actions in crises is well established. However, on the day, multiple role versions of Checklist 4 and their subtle incoherencies with the EFG appear to have contributed a degree of uncertainty amongst those leading the operational response in terms of what exactly was still available, how best to manage the failure's consequences and what operational recovery options may have been available. The checklists also give guidance on steps that would be necessary to enable a recovery of the failed systems, but are silent on when such preparations need to be made and what flexibility exists to defer the more acute actions, such as clearing sectors of traffic, until more is known about the timeline for the technical recovery. They also lack guidance on the likely effect actions taken 'locally' may have on the wider aviation system and any options for tailoring responses to the conditions to minimise adverse impacts.

3.7.3 There may therefore be merit in examining critically the scope and content of fallback procedures and checklists to ensure that they are both unambiguous and also useful in helping determine how to tailor responses, where and when appropriate, to help minimise friction. They must never of course become so complicated that they defeat their object, but there may be potential to capture within them opportunities to align responses and recovery profiles more closely with more refined scenarios and an appropriate assess/think/decide/act cycle, guided perhaps by a flow-chart of conditions-based options[34]. Amongst other things, this could include responses better tailored to expected or actual traffic flows, seasonal conditions and time of day, recognising that this cannot replace the expertise of the staff managing an incident on the day. Seeking assurance of operations staff's levels of systems knowledge, their understanding of the relationship between failure indications/warning messages and their associated checklists, and their proficiency in implementing them would also seem worthwhile. On the day, the workstation warnings presented to the LAC staff portrayed accurately the system failures, i.e. a failure of the SFS, followed shortly thereafter by a consequential failure of the link between the SFS and the NAS[35], yet it took some 40 minutes and the assistance of ETIC before AC supervisors were confident that they understood the nature of the failure with which they were dealing[36].

3.7.4 The standard contingency flow rate and London airport departure regulations are intended to be effects-based. However, they appear to be blunt tools in practice. Their definitions are open to interpretation and, consequently, there was not shared understanding amongst stakeholders of their application and intended effects. Moreover, initially applying them to

[34] Perhaps starting from a better perspective of what capabilities remain available and what fallback operations they can support, rather than what has been lost.

[35] The failure of the connection to NAS is specifically mentioned as an effect in Checklist 4 – SFS Unavailable.

[36] Not helped by the intervention of a Data Systems Supervisor DSS who initially misidentified the failure as a NAS-SFS Link Failure.

be effective for 4 hours as a means of stopping departures at more distant airports was both esoteric and triggered a number of second and third order effects that probably exacerbated the impact of the failure and made the subsequent recovery more cumbersome. It may be that risk-based judgement would allow a shorter duration to be selected initially, on the assumption that the nature of the failure, its consequences and the capacity of the system in reversionary mode could be assessed in sufficient time to make alternative arrangements for aircraft that proceeded to get airborne around the time of the failure, but by virtue of their flight times would not arrive in affected airspace in, say, the first couple of hours. It may therefore be worth reviewing the relevant regulations with Eurocontrol to ensure they are both clearly defined and strike an appropriate balance between maintaining safe operations, containing the failure's effects and minimising disruption to the wider aviation system. The absence of contingency route scenarios that could be employed quickly, particularly for departures from the major London airports, and are not reserved just for ameliorating longer-term disruption, would also appear to be an omission.

3.7.5 Notwithstanding, the 4 flow rate regulations applied initially were to a degree successful in achieving their aim – stopping inbound traffic to London airspace that wasn't already airborne. However, the fact that traffic already en-route would be handled to the best of Swanwick's ability was lost in translation, which resulted in confused messages about the actual status of the airspace and, probably, some unnecessary diversions and cancellations. Moreover, NMOC's instinctive suspension of Flight Plans filed for the period and the lack of awareness amongst operators of this standard procedure, exacerbated by NMOC's decision not to promulgate a related AIM notifying operators of their action, added significant friction to the recovery when it was initiated, prompted the unilateral withdrawal of two key stakeholders from the only obvious collaborative decision-making system and, in aggregate, made the recovery more clunky and probably prolonged it.

3.7.6 NATS' role as 'first responder' when the failure occurred was entirely appropriate, commendably quick and effective. However, immediately the failed system was recovered, the focus transitioned, and appears to have been confined largely to, leading the recuperation of NATS' services quickly. Indeed, the recovery initiated at 1605 was well before the root cause of the failure had been determined. It may be unreasonable to expect operations to have remained in reversionary modes and severely constrained until the root cause had been positively identified, which could of course have taken much longer than it did in this instance. Nevertheless, the safety governance – who was responsible, consulted, accountable, informed – surrounding decisions to return to 'normal' operations once SFS-B had been recovered and to subsequently declare full operational capability via ATICCC at 1845, some 1 hour 30 minutes before Engineering confirmed to Silver that the root cause had been identified and a further 30 minutes before the number of Atomic Functions had been positively confirmed and then reduced below 151, is unclear. At that time, a serious system failure necessitating significant crisis response actions had been experienced, part of the failed system had been restored, but the cause of the failure was unknown and there was no estimate of how long it would take to identify it – yet normal operations were declared and resumed, apparently with few, if any, substantive additional measures being implemented to prepare for any subsequent recurrence of the failure.

3.7.7 Some of the dialogue recorded in the logs of Gold and Silver commands indicates that considerable attention was being paid to NATS' standard (commercial) performance metrics – operating capacity, flow rates and en-route delay. Apart from one question relating to prioritising airports, there appears to have been little discussion about the impact on the aviation system more widely and how best to manage it, although there is no evidence to suggest that this was a conscious decision/choice. Accordingly, the recovery perhaps erred towards being dictated to customers, rather than being informed by and accomplished with them. This is likely to have resulted in a less expeditious 'enterprise' recovery. The critical requirement to maintain throughput at Heathrow hub could possibly have been acknowledged from the beginning and accorded a higher priority – perhaps by more determined and effective use of LTC when the service was eventually confirmed as available, and/or by prioritising the Heathrow hub recovery for a defined period to pump-prime or accelerate the wider network recovery[37]. The Panel was advised by both CAA and NATS that there is a "no preferment" clause in NATS' licence, but strict adherence to this in crisis and safety management may be inappropriate and, indeed, the relevant section does seem to allow some flexibility:

> *"7. In providing services under paragraph 1 the Licensee shall not unduly prefer or discriminate against any person or class of person in respect of the operation of the Licensee's systems, after taking into account the need to maintain the most expeditious flow of air traffic as a whole without unreasonably delaying or diverting individual aircraft or such other criteria as the Licensee may apply from time to time with the approval of the CAA.*[38]

As previously identified in the aftermath of the 7 December 2013 incident, NMOC's broader view of network capacity, demand and operators' requirements via its oversight of Flight Plans and FCMs could also perhaps have been utilised better in informing, coordinating and shaping the overall recovery sequencing.

3.7.8 Intuitively, the A-CDM system has significant potential to streamline and improve European aviation system stakeholders' understanding and collaborative decision-making – that is presumably why considerable investment is being made in it. However, its introduction into service and confidence in its benefits were undoubtedly dealt a blow by the decision of two of the UK's leading A-CDM airports to dispense with it early on in a crisis management situation. Notwithstanding what has already been undertaken in A-CDM education, training and resilience planning, it appears that more needs to be done to ensure that operators understand the tools' characteristics, strengths and limitations, are more proficient in their operation and, thereby, develop greater confidence in their use. The consequences of dispensing with the system should also be assessed carefully and any assumed fallback procedures should be tested and rehearsed. Nevertheless, assuming it lives up to its billing, A-CDM promises to greatly enhance shared situational awareness and, therefore, to offer real value in crisis management.

3.8 Recommendations

It is recommended that:

R9. NATS should examine the use, recording and governance of informal communications during crisis response to ensure consistency and minimise the risk of contradictory and/or ambiguous instructions. (Para 3.4.2)

R10. NATS should enlist appropriate expert support and expedite arrangements to conduct an industry-wide review of crisis response and resilience arrangements without delay. (3.5.3)

R11. NATS should consider introducing a formal Error Management System (EMS) to capture anomalous occurrences that fall below the safety event threshold, but which may indicate where changes in systems, procedures or training would benefit the management of risk. (Para 3.6.2)

R12. NATS should review routine availability of staff to support Swanwick and Prestwick Centres (and other sites), to ensure that there are always sufficient qualified and experienced personnel to support incident analysis and crisis management. (Para 3.7.1)

[37] To a degree, this option is already reflected in the EFG, with the recommended recovery departure rates weighted in Heathrow's favour, but these do not appear to have been employed on the day and, in any case, more could possibly be done in future.

[38] Air Traffic Services Licence for NATS En Route Limited plc, Jan 15.

R13. NATS should review their hierarchy of fallback procedure checklists for completeness, coherence and consistency, so that they support controllers via an intelligent checklist architecture that leads intuitively through conditions-based options, including making clear where the controller has discretion to adjust and refine responses as circumstances dictate or allow. (Paras 3.7.2 & 3.7.3)

R14. NATS should review their arrangements for continuation training of their operational staff in systems knowledge, failure identification and response. (Para 3.7.3)

R15. NATS, in conjunction with the Eurocontrol NM and key customers, should review their ‘standard’ contingency routing, flow rate and departure regulations to ensure they are suitably responsive, precise, effective and sensitive to their impact on the wider aviation system. (Para 3.7.4)

R16. The CAA should request a review by Eurocontrol of the means by which Eurocontrol defines, communicates and assures understanding by ANSPs and operators of critical network management actions and their implications. (Para 3.7.5)

R17. NATS should review their safety governance and assurance of operational decisions during crisis response and recovery phases. (Para 3.7.6)

R18. The CAA should facilitate engagement by NATS with the Eurocontrol NM, airports and airline customers to review roles, responsibilities and priorities in ATM crisis management and recovery. (Para 3.7.7)

R19. The CAA should engage with relevant UK airports and Eurocontrol to assure appropriate A-CDM system education and training, the effectiveness of A-CDM operation and that of any fallback modes. (Para 3.7.8)

Chapter 4. NATS Systems: Requirements, Management and Delivery

4.1 Introduction

4.1.1 Sustainable growth requires the evolution of NATS ATM capability to deal with inefficiencies in the current system, thereby generating significant benefits for passengers, industry and the environment. This will also provide opportunities to enhance aviation safety with the advent of new technologies and operational procedures that could reduce or remove safety risk factors from the current ATM system. For example, the introduction of Queue Management procedures, using specialist ATM tools to stream traffic using speed controls, will reduce reliance on airborne stack holding that carries the inherent risk of "level busts"[39] as well as the associated negative environmental effects.

4.2 Future ATM Requirements

Airspace Capacity

4.2.1 The evolution of ATM capability to generate future airspace capacity needs is principally based on a move, for an aircraft's knowledge of its own position, from ground-based navigational aids to the greater precision offered by satellite systems such as the Global Positioning System (GPS). Using the aircraft on-board flight management computers and satellites for positioning accuracy, aircraft are able to fly and maintain their position on pre-assigned routes to an extremely high degree of accuracy. Controllers can therefore have a much greater confidence that an aircraft will accurately follow its flight plan. As air traffic management is an international business that has operated to internationally agreed standards over many years, the evolution of ATM capability is required to align with the International Civil Aviation Organisation (ICAO)-led Aviation System Block Upgrade (ASBU) programme and US 'Next Gen' Programme. It also includes European mandated standards, designed to harmonise a collaborative global approach without requiring one single (monopoly) solution. In the UK, this work has been supplemented by additional rules and guidance published by the CAA. This capability is focused in three areas:

(1) Implementing a fundamentally more efficient route network in the busy terminal environment designed to exploit what are known as Performance Based Navigation (PBN) standards.

(2) Removing some of the fixed airways in the upper airspace and enabling more direct routeing in the cruise phase of flight.

(3) Streaming traffic through speed control to manage queuing and reduce stack holding thereby improving arrival punctuality.

4.2.2 In the past, although there were defined routes and procedures, Controllers could not assume that the aircraft would fly them very precisely. Hence separation standards and separation between defined routes had to be quite wide. With Performance Based Navigation, the aircraft asserts the accuracy that it will fly to, and ATC can have confidence in this. Hence defined routes can be closer together and there is confidence that the aircraft will fly them. Additionally aircraft can be given time base clearances to support queue management and can be expected to adhere to them.

Single European Sky ATM Research (SESAR)

4.2.3 ATM capability in the UK is not being developed in isolation. The Single European Sky (SES) initiative was established to tackle fragmented ATM arrangements and to deliver interoperability across Europe. The development of NATS ATM capability makes a significant contribution to the implementation of SES objectives. In particular, by coordinating UK deployment of solutions developed in the technology component of SES, known as SESAR.

[39] Level bust; when an aircraft does not level off at its assigned level but climbs or descends further than cleared

4.2.4 In the past, NATS had considerable flexibility in how these requirements were met. Since the passing of Single European Sky legislation in 2004 and its adoption into national rulemaking, the primary regulation of ATM has been from Europe through the European Commission (EC). The EC has introduced the performance regime and the concept of Functional Airspace Blocks (FABs – see below). Through SESAR, the European ATM industry has developed a European ATM Master Plan that defines capabilities and concepts that should be deployed across Europe to deliver seamless interoperability and to meet challenging performance standards.

4.2.5 The EC is driving deployment of these capabilities through the creation of the SESAR Deployment Manager (SDM) and by issuing mandates for key capabilities to be deployed by ANSPs, Airports and Airlines. The intention is to ensure that common European solutions are coordinated to deliver maximum benefit for airspace users. This European regulatory environment provides the context for NATS' future systems development and deployment. This includes rulemaking emanating directly from the European Commission, and regulations developed for the European Commission by the European Aviation Safety Agency (EASA), primarily Implementing Rules. European regulations are binding on the UK (and NATS in particular) and take precedence over UK National Legislation.

4.2.6 The high level SESAR concepts are well defined and the Panel was informed by NATS that there is a concept of operations describing the intent, operational use, the introduction of the new systems and the expected benefits. While the Panel did not have the opportunity to examine this concept of operations, their firm opinion is that successful projects require such a concept of operations as the fundamental first step before proceeding to draw up the requirements. Specifically, the first set of capabilities to be delivered has been published within the Pilot Common Project (PCP). The PCP identifies 6 ATM Functionalities to be delivered during the time period up to 2024:

- Extended Arrival Management and Performance Based Navigation in the High Density Terminal Manoeuvring Areas;
- Airport Integration and Throughput;
- Flexible Airspace Management and Free Route;
- Network Collaborative Management;
- Initial System Wide Information Management;
- Initial Trajectory Information Sharing.

4.2.7 While the final two functionalities are less mature, they have all been subject to extensive definition and validation and form part of the deployment planning of the SDM that will help to coordinate their implementation across Europe.

4.2.8 NATS remains responsible for their delivery, but is required to comply with the capabilities mandated by SES. For this arrangement to be properly effective, the requirements and performance standards levied on ANSPs must avoid variations in interpretation and application since these can drive significant differences in solutions and ultimately cost: this point was firmly made to the Panel by the CAA. NATS, in common with the rest of the European service providers, plans for delivery of an agreed programme of work over 5 years known as a Reference Period; the current period (RP2) covers 2015 to 2019. This regulation is coordinated in the UK by the CAA leading to the definition of a financial settlement within which the regulated monopoly manages operations.

UK and Ireland Functional Airspace Block

4.2.9 Underneath the SES and SESAR plans, there are Functional Airspace Blocks (FABs) comprising groups of European countries that collaborate to implement sizeable parts of the European ATM plan. The UK and Ireland Functional Airspace Block is driving the modernisation of the en-route airspace shared between the UK and Ireland. The associated plan is subject to consultation with stakeholders, such as airlines and airports, and there is Government and Regulator oversight. As part of RP2, States have been required to submit Performance Plans at the FAB level to contribute to the Europe-wide Performance targets, covering safety, cost efficiency, capacity and environment. The UK/Ireland FAB Performance Plan covering 2015-2019 was submitted to the European Commission in 2014. The Panel was informed that the EC consider this plan to be one of only two making sufficient contribution to European targets. NATS now tracks the common lines of action across the aviation sector in order to address inter-dependencies in the timescales required.

Contingency, Resilience and Business Continuity

4.2.10 The regulatory requirements associated with resilience and contingency for ATC systems and procedures flow from the international level (ICAO and EASA) and national obligations. However, the terminology used to describe these requirements is inconsistent and imprecise, and, as a consequence there is considerable latitude in interpretation. The CAA considers that the consistent application of sharper definitions could make an important contribution to clarity of expectations in terms of future performance.

4.2.11 From an ATM systems perspective, “resilience” describes the system and facility design features that prevent total system failure or severe degradation. Contingency tends to be used to describe the plans that are in place to provide control arrangements once a system failure, including the loss of a centre, has occurred. These resilience features, in some cases through secondary or tertiary systems, can be expected to provide a service capacity at a reduced level in order to maintain safety from both the system and the human performance dimension.

4.2.12 Currently, resilience is principally seen as reflecting the minutes of attributable delay that are incurred as a result of an incident: more resilience delivers less delay. Reducing traffic volumes to the safe level of degraded capacity that the system is deemed to be able to cope with, is the current mechanism for safely managing incidents. Whether this will remain the case in the SESAR environment is worthy of further investigation with the potential for inclusion of new approaches in the future concept of operations.

4.2.13 These mechanisms for managing a serious failure or degradation in the performance of ATM systems and tools, or other unusual events are extremely important. Emerging resilience risks, such as prolonged failure of power supplies or disruption of operations through a cyber-attack have the potential to disrupt or degrade service provision. Resistance to cyber attack and resilience against system failures form an integral component of all major new systems.

4.2.14 Neither the SES regulations nor the SESAR deployment plans specify details on resilience requirements or how resilience should be measured. It would be useful if NATS and CAA worked together to agree this concept of operations. Contingency, resilience and business continuity performance requirements can then be articulated so that stakeholder expectations are aligned. Ideally these would also be aligned within Europe to avoid driving varying requirements and cost across the network.

4.3 Investment Planning

Timelines and Sequencing of Deployment

4.3.1 There are a number of drivers that may require NATS to modify/upgrade its systems, but few are simple and most require a long lead-time and appropriate resource allocation. Existing equipment or upgrade plans are factored into the Service and Investment Plan (SIP) as part of routine replacement of systems against NATS own engineering and operational risk assessments. These are developed internally and may be discussed at a high level with customers and the regulator. Some of these upgrades will be of sufficient significance that they appear as individual line items within the SIP. Major systems replacement will receive much greater prominence and will be the topic of discussion with customers in the development of the SIP.

4.3.2 The UK Government made it clear in the 2013 Aviation Policy Framework[40] that the Future Airspace Strategy (FAS) Deployment Plan is a key driver in delivering part of the UK's contribution to SESAR. The underpinning rationale for the FAS programme has been to align investment plans between the key industry stakeholders (including airlines, airports, ANSPs and the regulator) in order to deliver performance improvements. It would be useful to investigate what flexibility the UK can accept in its requirements so that the opportunity to promote agreements with European partners does not founder on this familiar rock.

Managing the Scale of Change

4.3.3 Major evolutions of NATS ATM capability will always be deterred by their cost, risks and complex interdependencies and be subject to the willingness of the airlines to absorb the resulting charges. While incremental improvements have progressed satisfactorily, there is also a limited amount of change that NATS and the airlines are able to absorb in a given timeframe, owing to the impact of any change on operations. The amount of change that is possible is also constrained by the volume of aeronautical data that must be incorporated into one of the co-ordinated dates for changes to aeronautical data (which occur on an internationally agreed timetable every 28 days). This is so that changes can be made on a global scale, with every operator and air traffic service provider at the same time. Multiple successive changes introduce further risks associated with the bedding down period for one change and the scale of rolling changes that pilots and controllers can accommodate. Many changes may require further modifications once all of the operating characteristics have been understood.

Long Deployment Lead Times

4.3.4 The management of new systems, tools and procedures that have been designed and tested years prior to their deployment creates challenges. NATS use a process based on a defined project lifecycle to move from strategy to operational intent to delivery. Gateway points are set and must be achieved before the project can advance to the next stage. Major projects often span multiple financial years or SIP periods (refreshed every year) and may also straddle Reference Periods. Tracking the specific contribution of a multi-year project to particular Performance Plan targets in a particular year can prove difficult. Drawing a clear line of sight between the key initiatives and expected performance improvements is intended to ensure that implementation targets are stretching, but achievable, and ensuring that deployment remains performance driven.

Investment Trade-offs

4.3.5 The CAA informed the Panel that European legislation requires, and airline and passenger experience across Europe point towards, the need for increasing performance in terms of safety, service and resilience. The RP2 settlement is intended to be the framework for trading performance against affordability and with incentives and penalties to drive the right behaviours. NATS also strives to ensure operational personnel are engaged, trained and certified as the ATM capability evolves and this takes place alongside the continued provision of service. Taking a large cadre of controllers through a major operational change is recognised as a both a significant task and vital to the outcome; it can take months and may involve repeat sessions with individuals. Both fast-time and real-time simulations are used to confirm that the changes deliver the required operational outcome. The impact of evolving ATM concepts and of the greater capability of the new systems on the controllers' role, workload and culture will need to be carefully managed.

[40] Aviation Policy Framework, Department for Transport, March 2013.

4.4 Delivering Change

Management of the Residual Risks

4.4.1 Recent failures[41] indicate that the complexity of NATS Systems, which are continually adapted to deliver essential changes, require persistent vigilance in technical management and rigour in applying the change process. Regardless of the timescales with which new ATM systems are deployed, the approach to managing the residual risks associated with legacy systems should be maintained. This may come to require extending the funding and the people who support the existing systems as NATS works to introduce the new systems (including major testing) and to manage both the new systems and the old during transition Furthermore, the introduction of new systems can bring a flurry of important suggestions for their improvement over the first year or so in operation and this needs to be resourced and managed carefully.

Baseline for Safety Assurance

4.4.2 The introduction of more advanced ATM systems will come without the benefit of experience in using the systems in an operational capacity anywhere else previously (either in the UK or globally). One of the main challenges associated with delivering performance targets, including that of improved safety, is the ability to gather sufficient baseline information against which to assess the impact of the various components of the new system. Legacy systems have many years of accumulated safety assurance information from which to track the impact of changes. NATS will need to ensure that the approach to safety assurance of changes to new ATM systems is not exposed due to a lack of accumulated assurance information.

4.4.3 Controlled trials may be necessary to deliver part of the evidence to support the safety assurance of changes. Major updates to introduce advanced ATM systems are likely to require an appropriate architecture with the individual components and interdependencies clearly mapped, such that safety assurance can be delivered. The present architecture allows individual component safety cases to be linked so that in future, changes to the systems can be undertaken with an improved awareness of testing requirements for both the relevant component and the overall system. The same approach should be used for control of changes to the requirements for new systems. Overall, the work to develop and refine a SESAR ATM Master Plan has led to a reasonably stable and well understood definition of the system requirements.

Delivering Low Level Improvements

4.4.4 As a result of regular operations, incidents and dialogue with customer and other stakeholders, NATS often identifies low level improvements that be made to the overall ATM system (the combination of the people, procedures and technical systems). These potential improvements can be raised as specific change requests (CRs) or PTRs on systems where appropriate, or simply raised by staff as observations within the safety tracking and reporting system (STAR). All of these CR/PTR/STAR reports are tracked and assessed so that the requirement for change, importance and urgency can be assessed. Responses are:

- Implemented through an urgent system change;
- Implemented as a small scale procedure change;
- Scheduled for a future system build;
- Aligned with a future airspace change;
- Avoided by creation of a workaround;
- Deferred as not being sufficiently urgent or important.

[41] Notably the December 2013 system failure as well as the more recent December 2014 failure.

4.4.5 All of these approaches are tracked with appropriate approval and sign off for decisions taken to ensure that potential incremental improvements are tackled proportionately and judiciously. The Panel accepted the CAA view that the process is both appropriate and well executed.

4.5 NATS Capability: Supporting NERC

NERC Overview

4.5.1 The systems supporting the NATS operational activities at Swanwick, continues to be known as the New En Route Centre or NERC. NERC is a fully integrated system covering a range of capabilities including radar and flight data processing; voice communications and support information as well as simulator capabilities to support testing, development and training.

4.5.2 The original contract to develop the computer systems for NERC was awarded to IBM Federal Systems. IBM Federal Systems were acquired by Loral in 1994 which was, in turn, acquired by the US aerospace and defence contractor Lockheed Martin in 1996. The system went live in 2002 and is now supported by personnel from NATS and the main suppliers, with NATS having overall project accountability. Other UK and international companies have also been engaged on a subcontractor basis.

4.5.3 The software is written in a high level programming language called Ada, which was developed in the 1980's primarily for military Information Technology (IT) systems. The NERC System is a necessarily large and complex computer system. At its inception it was considered to be 'leading edge' in its use of technology and conformed to best practice as it was at the time.

4.5.4 Figure 4.1 (a repeat of Figure 2.1) provides a simplified view of the architecture of the NERC system, with the System Flight Servers at its heart, receiving Flight Data as key input from the NAS and supporting the operation of the workstations used by the controllers. Architecturally the NERC system has 5 independent networks for the workstations, with approximately 30 workstations on each network. There are 2 sets of redundant servers and the workstation networks can be connected to either set of servers to provide resilience and to support the software cutover process required for system updates.

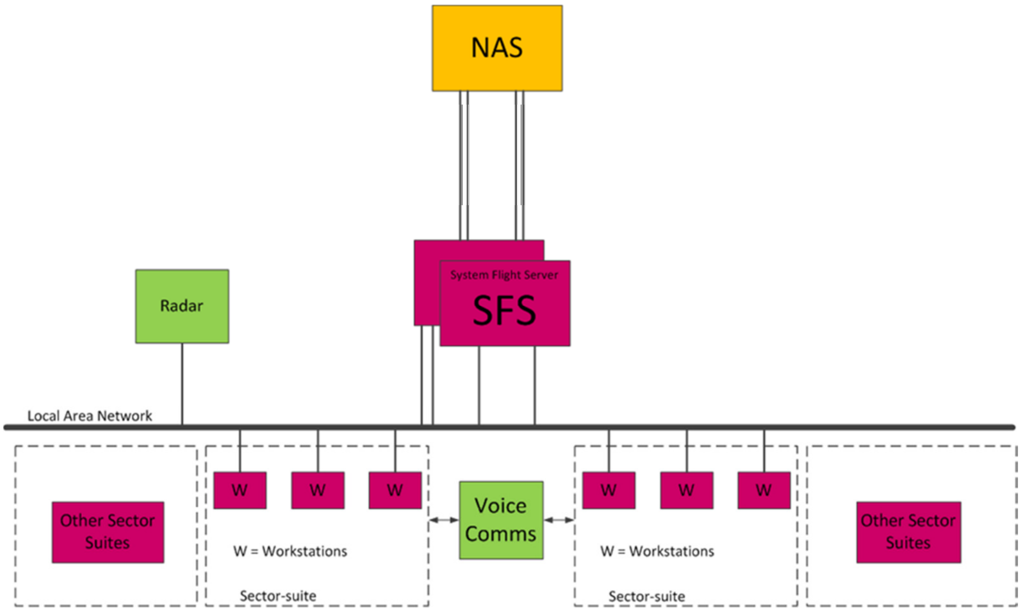

Figure 4.1: Major Systems supporting LAC (Simplified)

4.5.5 There are multiple layers of physical security between the outside world and the SFS processors. The SFS have no direct external data connections and 3 internal connections to the NATS systems:

(1) NAS.

(2) Datalink Front End Processor (FEP).

(3) Communications Gateway (CGW).

4.5.6 All the communications paths/systems have been risk-assessed under the NATS security risk assessment process with any identified risks mitigated or managed. There are also mechanisms in place to guard against 'denial of service' attacks.

NERC Change Lifecycle

4.5.7 The change lifecycle for the NERC system follows established, best practice processes and procedures for a system of this complexity. Changes are categorised by complexity using a structured Analysis Options and Design methodology (AOD). There are 5 levels of ascending detail:

(1) Cat 1 & 2 will consider the design options available to deliver a solution based on a request for a change.

(2) Cat 3 & 4 will require user involvement in the form of a user and system requirement documents to further formalise the requirement, with acceptance criteria specified and a design validation process undertaken satisfactorily against the defined options.

(3) Cat 5 would necessitate a formal baseline definition of the preferred solution to be undertaken.

4.5.8 System integration, verification and validation follows and includes the change or changes being put through a series of tests to ensure that they have been implemented so as to meet their design intent and that there has been no adverse effect on any part of the overall system.

NERC Verification, Validation and Deployment

4.5.9 The NERC software test strategy is subdivided into 4 parts:

(1) Coverage testing.

(2) Computer System Configuration Item (CSCI) build testing.

(3) Software regression testing.

(4) Sub-system acceptance testing.

4.5.10 Coverage testing is performed on each build on all units that have logical changes. Coverage testing is performed using appropriate test tools. These tools provide condition coverage for logical decisions within the code (see Annex G for an explanation and discussion of coverage criteria).

4.5.11 CSCI build testing is used to show that the changed software performs its function correctly. The tests are written by the individual developers. The tests are peer reviewed within the team, and the tests are run during development and again prior to handover.

4.5.12 Software regression testing whereby a set of regression tests are run on each package build for each sub-system. This includes a base set of tests covering core system functionality. A set of tests targeted at areas changed in the build are added to the base set.

4.5.13 Sub-system acceptance testing. These tests are written by independent testers rather than developers. They are used to prove that each sub-system meets its requirements and HMI (Human Machine Interface) specifications. The tests are peer reviewed by developers and the person responsible for the AOD under test. The tests are run during development and then a final QA witnessed run is carried out prior to build hand over.

4.5.14 This test and validation cycle typically takes about 6 months and would involve about 27 staff, of whom, 30% have more than 10 years' experience of the system. Following successful validation the changes will become operational through a Software cutover process. This process progressively migrates the servers and workstations to the new software in a manpower intensive overnight activity which typically takes around 3 hours.

Configuration Management

4.5.15 Configuration management of NERC is split into 3 functional areas: change management; document management and software configuration management.

4.5.16 Change Management (CM). Changes to the NERC system baseline are a result of strategy, CRs or PTRs. The configuration control board manages CRs into the build. Large CRs become AODs. Once an AOD is Category 5 approved, programme lead manages the build contents during implementation and verifies content of the build. CM ensures additional CRs linked to the build are closed.

4.5.17 Document management. Documents are stored on a shared network to allow access to the approved and relevant parties. When ready for publication, new and updated documents are received from the relevant parties, usually by email. Following quality checks, documents are stored on a separate area of the document management system. A document tracker system, maintained by document controller, ensures documents identified in AODs are delivered or appropriate concessions raised. Master records are then created for each build with links to the relevant documents, and the hard copies of master records are retained with appropriate backup.

4.5.18 Software configuration management. The approved contractors (LM, Altran and others) deliver source code and other relevant software files into NATS via their own configuration management systems and version control tools. These are then input into the NATS version control tool (IBM Rational Apex). The code is then compiled, links are produced to the appropriate routines and executable code is produced. Build releases and baselines are produced using the LM system support and change control system. File information is stored on a separate database. The configuration management team will then build the new software, verifying that it has built correctly based on the right content. The software is then released into operational service. The Panel considers that the overall change and software configuration management process conforms to best practice and that, in so far as the Panel was able to observe, it is properly executed.

Quality Assurance

4.5.19 There are two key roles in the quality assurance organisation

(1) Quality assurance manager. The QA manager is responsible for overall QA planning and governance. There are monthly quality risk assessments. The QA manager is a member of the project boards and critical projects reviews. The manager is also responsible for acceptance to provide customers with data delivery packs, which contain all necessary Deployment Assurance.

(2) Quality Work Package Manager. The QWPM is responsible for the routine delivery of quality services to the NERC programme, for example: design and code inspections; test witnessing and concession and defect prevention process management.

4.5.20 The Panel concluded that QA organisation adheres to best practice including a clear differentiation between quality control and quality assurance.

4.6 Required Capability: Deploying SESAR

4.6.1 This section covers the future capability that NATS will require to support its “Deploying SESAR” Programme which is a business transformation programme to deliver SESAR solutions into operation within NATS and to transform their operational and technical capabilities. The deployment of SESAR is a major programme which will progressively replace most of NATS legacy systems over a period of 5 years from 2015 to 2020. The programme is based around a transition approach which plans to deliver change to the different operations at a number of points in time called Deployment Points, in accordance with the chosen transition sequence.

4.6.2 Deployment should ensure that the entire operational concept; requirement and service design at each deployment point is captured. Deployment planning must disaggregate the necessary supporting services, ensuring that all elements deliver their contribution. In particular the plan must:

- Integrate, verify, validate and assure the entire change, made out of many "applications" running on a common infrastructure for each Deployment Point.
- Realise the benefits for each of the relevant project or projects.

4.6.3 For each Deployment Point there will be a hierarchy of activities that must be undertaken to define and implement the required solution. Each of these activities will be matched by a corresponding Verification, Validation and Assurance process as shown in figure 4.2 below.

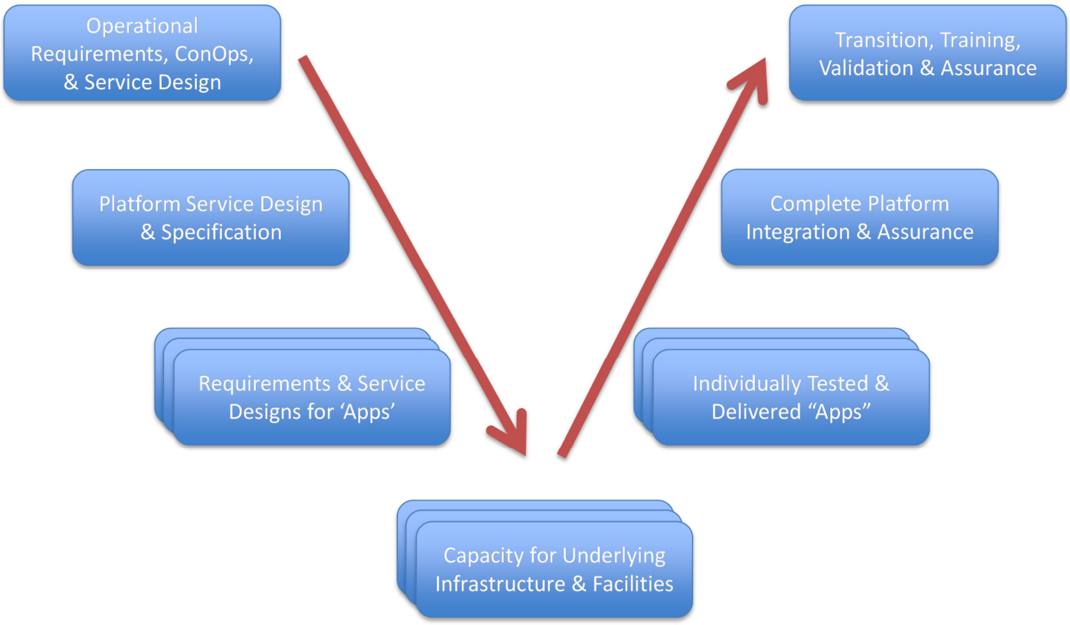
Operational Requirements, ConOps, & Service Design
Transition, Training, Validation & Assurance
Platform Service Design & Specification
Complete Platform Integration & Assurance
Requirements & Service Designs for 'Apps'
Individually Tested & Delivered "Apps"
Capacity for Underlying Infrastructure & Facilities

Figure 4.2: High Level Validation Approach

4.6.4 The phased approach inherent in the deployment points is consistent with best practice for major transformations. This must be accompanied by appropriate disciplines:

(1) The requirements and definition phases of the project must be completed in detail before any stage of implementation is completed.

(2) Clear exit criteria for each phase should be defined in advance and adhered to rigorously.

(3) Strong governance should be applied for all aspects of initial approval, management of change and phase approval.

4.6.5 Deploying SESAR has tight timescales and challenging collaboration aspects. It would not be sensible for NATS to attempt to accelerate this programme beyond the currently defined plan as this is likely to lead to shortcuts being taken in the early specification and requirement phases and will only lead to increased risk of late change and delay.

SESAR: Differences from NERC

4.6.6 The Deploying SESAR programme will be very different in nature from the NERC programme and NATS must understand the implications of these differences as they put in place the plans, expertise and governance necessary to deliver a successful programme.

4.6.7 The origins of the NERC system were in the United States FAA led system development called Advanced Automation System (AAS). The FAA subsequently changed their development programme so that the NERC development became a bespoke system for NATS alone. This allows development to continue with full control of the requirement by NATS and using suppliers selected by NATS: enormously simpler than a collaborative programme with an international supply chain.

4.6.8 SESAR is defining common solutions planned for deployment across Europe. Many of these capabilities are being developed as multi-national programs on a collaborative basis with other European ANSPs. For example a key element of SESAR deployment for NATS is the iTEC Flight Data processing system (FDP) which is being developed by the substantial Spanish information technology and defence systems company Indra Sistemas for 4 ANSPs: NATS; ENAIRE (Spain); DFS (Germany); and LVNL (Netherlands). The programme has operational and procedural requirements that are common to more than one ANSP and operational and procedural requirements that are particular to individual ANSPs. Pure UK requirements can stimulate similar unique needs from our partners; they are inherently more expensive to deliver (because the costs are not shared) and can be costly to maintain through life if further UK specific changes are introduced.

4.6.9 The design, implementation and system testing is being undertaken by a series of companies who are under contract to NATS to deliver specific parts of its overall SESAR deployment programme. NATS role in the SESAR development differs fundamentally from its role for NERC. NATS will be in a programme management capacity, whose prime responsibility will be to ensure that the SESAR objectives are fully met and are delivered on time and within budget rather than in a much more hands on development programme with an autonomous decision making ability. NATS will have responsibility to define its own specific requirements and ensure that these requirements are delivered by the suppliers. NATS will have final acceptance testing responsibility for the system before it enters operational service. NATS must therefore ensure it has the range of technical, business and managerial skills to undertake its SESAR role.

4.7 Conclusions

4.7.1 The international nature of the evolving ATM capability means that NATS must ensure that all of its future plans and projects conform to internationally agreed standards and harmonised timescales. This requires clear, consistent, terminology in both setting targets and ensuring common standards for requirements such as contingency, resilience and business continuity as part of the network design. A documented concept of operations is the essential foundation for articulating requirements and their subsequent specification – a project without a defined scope is difficult to assess in either timescale or cost.

4.7.2 Rulemaking activity by the European Commission and its agencies, and the framework for change provided by SES, means that the evolving capability must meet the requirements of SESAR. Any plan must also take account of the risks and impacts of multiple change cycles over successive Reference Periods. A rigorous approach to assurance is the necessary complement to introducing these new capabilities into service in a phased manner and in collaboration with partners.

4.7.3 The Panel accepted that the current process for incremental changes, implemented at the operational level, prioritising safety risks and impacts, is both appropriate and well executed. Throughout the Deploying SESAR programme of change, the role of the Controller will evolve, as emerging technology provides more procedural "systemisation" across the ATM network. The way in which failures are managed may also have to change, as applying traffic volume reductions as the principal control technique may no longer be viable in highly systemised airspace.

NERC

4.7.4 The NERC Operational System has been in service for many years. It has been upgraded over its lifetime to update hardware, support operational changes and to implement a range of system enhancements and problem resolutions. The technology has become dated although it was "leading edge" in its time. It now requires more "hands on" involvement to address changes than would be required for modern systems and the various change processes are manpower intensive.

4.7.5 There is an extensive range of hardware and test software available for use in the Verification, Validation and Simulation of test cases. Software testing has become more intensive since the NERC system was developed. Systems were tested to establish, beyond reasonable doubt, that they performed the tasks they were intended to perform, and delivered the required functionality. There was little or no testing to confirm the system did not perform unintended actions.

4.7.6 The NATS support team has benefitted from low staff turnover, with over 30% of staff having more than 10 years' experience of the system. This gives rise to a comprehensive understanding of how the system operates, while also introducing some fresh experience.

SESAR

4.7.7 The approach NATS is taking to SESAR Deployment is a step change from earlier developments. This European initiative, which is based on political directives, needs very different management skills, principally because of its collaborative nature.

4.7.8 Collaborative programmes (and there are not many of this size) require a wide range of skills to ensure a common understanding of the end products from the outset. The concept of operations (to be supplied by the overall NATS project sponsor) is the prerequisite for articulation of the requirements; their subsequent specification in a contract is a key step. Achieving the desired results also depends on a rigorous approach to acceptance and approval at every stage of the project. Deploying SESAR is a NATS wide development and it will have to ensure that it can attract and retain the right level of programme management and engineering skills needed to deliver this programme, including bringing forward experience from NATS previous programmes.

4.8 Recommendations

R20. CAA and NATS, in consultation with other stakeholders, should agree national definitions and requirements for contingency, resilience and business continuity. (Para 4.7.1)

R21. NATS and the CAA should agree on how to provide assurance that the evolving capability meets the functional and non-functional requirements of SESAR while complying with the performance regime of the Single European Sky regulations. (Para 4.7.2)

R22. NATS to investigate the availability of other techniques beyond traffic volume reduction as the principal means for managing degraded service incidents. (Para 4.7.3)

R23. CAA and NATS to assess jointly, before the end of RP2, the skills and expertise required to fulfil the role of Air Traffic Controller in the SESAR era. (Para 4.7.3)

R24. NATS should consider staff rotation within the teams responsible for testing, verification and validation of NATS existing systems to maintain freshness and rigour in long established processes. (Para 4.7.6)

R25. NATS should include, within their phased approach to SESAR deployment, scrutiny and control of the concept of operations, clear requirements and exit criteria for each phase defined in advance with strong governance of initial approval, management of change and phase completion. (Para 4.7.8)

R26. NATS should re-deploy experienced engineering staff from NATS existing systems to support the requirements capture and specification of the SESAR systems currently being planned. (Para 4.7.8)

Chapter 5. The CAA NATS Relationship

5.1 Introduction

5.1.1 Previous chapters have addressed the technical, operational and business management issues bearing directly on the Incident of 12 December. This chapter considers the effectiveness of the oversight of NATS En Route Limited (NERL) by the Civil Aviation Authority (CAA). This has addressed aspects of the statutory framework, governance, organisation, policies, processes and resources relevant to the Incident. The enquiry has addressed the question of whether there was a failure of the CAA's oversight in relation to the Incident of 12 December. The Enquiry has also considered the question of whether any shortcoming in current oversight arrangements might be expected to affect the likelihood of future incidents.

5.2 Background

5.2.1 The CAA was established in 1972 as the primary regulator of civil aviation in the UK. It is an independent body, with its own board of executive and non-executive directors. The chair and non-executive directors are appointed by the Secretary of State. The Chief Executive Officer is appointed by the non-executive directors subject to the approval of the Secretary of State. Other executive directors are appointed by the Chief Executive with the approval of the chair and at least one other non-executive director. The Board is advised on appointments by a Nominations Committee. The Board's policies and day-to-day actions are not subject to approval by government, except where specifically provided for in the legislation.

5.2.2 The CAA's responsibilities and powers in the licensing and regulation of NERL are set out in section 2 of the Transport Act 2000. The primary duty when discharging its licensing and economic duties is to maintain a high standard of safety. The secondary duties include furthering the interests of both aircraft operators and passengers, promoting efficiency and economy in the provision of services and ensuring that NERL is able to finance its operations. These duties are exercised within a broader regulatory framework established under the SES regulations.

5.2.3 The main tool for the discharge of these duties is the licence to provide air traffic services, required under sections 5 to 7 of the Act, which is held by NERL. The licence was issued by the Secretary of State in 2001 and included a number of terms and conditions. The terms of the licence are matters reserved to the Secretary of State. They include the duration of the licence and the circumstances under which it can be revoked. Other than under specified circumstances the licence cannot be revoked before 2031. The CAA describes this licence as an economic or operating licence, to distinguish it from its safety functions, which are implemented separately, but not in isolation from the licence issues. The CAA regulates NERL through the enforcement of the conditions in the licence and by modifying them from time to time whether by changing existing conditions, removing conditions or introducing new ones. The current conditions cover matters including:

- The obligation to provide services;
- Accounting requirements;
- Financial resources and ring-fencing;
- Production of a five year business plan, an annual Service and Investment Plan (SIP), and periodic reports;
- Service indicators, measure and standards;
- Setting the maximum level of NERL's charges to airlines for its businesses in regulated airspace.

5.2.4 With effect from 2012, charging regulation for the regulated airspace part of the business has to be set for fixed periods as part of a performance scheme specified under SES legislation.

These fixed reference periods are expected to be for five years going forward although the first period was for the three years 2012 – 2014. (For this first European period (RP1), the CA applied the final three years of the domestic price control [CP3]). The SES legislation requires the CAA as National Supervisory Authority to propose a cost efficiency target for the UK. From 2015 this forms part of a performance plan for the UK-Ireland FAB along with targets for the FAB as a whole, for safety, capacity (specified in terms of flow management delay) and environment (flight efficiency). There is then a process for this FAB plan to be approved by the European Commission.

5.2.5 In setting the cost efficiency targets under this plan the CAA has adopted an approach close to the standard UK utility regulation 'building blocks' methodology (modifying this where necessary to be consistent with EU legislative requirements). Under this, NERL submits a multi year business plan setting out its capital and operating expenditure plans, its cost of capital, and expected levels of service. The CAA critically reviews these plans and determines the level of charges needed to deliver the finally accepted business plan, together with the service standards to be met. These costs and charges then form the basis of cost efficiency targets submitted as part of the FAB plan for approval by the European Commission.

5.2.6 For NERL's Eurocontrol business, this business plan is then implemented through setting a maximum NERL element of the overall charge to airlines (known as the unit rate). It applies for the full four- or five-year period of the plan and within that is adjusted each year for the level of inflation in the UK. It is also subject to upward or downward adjustment in the event that the service standards are exceeded or missed ('the service bonus and penalties'). This adjustment resulted in NERL foregoing £7.3 million of revenue as a result of the system failure of December 2013 and £0.5 million following the December 2014 failure. When the maximum prices have been set, NERL is left to run its business in the way it sees fit, until the next price review. It may pay dividends, subject to restrictions on maximum gearing and minimum credit rating. So far there have been three price control periods: CP1 lasted from 2001 to 2005, CP2 lasted from 2006 to 2010 and CP3 which covered 2011 to 2014 (and included the first EU reference period 2012 -2104).. A new price control has been set for RP2 (2015-2019).

5.2.7 In addition to its general duties in relation to the economic licensing of NERL, the CAA has specific duties in relation to regulating the safety of NATS's operations. These derive from a number of sources, including the SES regulations, the Transport Act 2000, and Air Navigation Directions and Orders. They include responsibility both for the safety oversight of the NATS operations and for the licencing of individual traffic controllers.

5.3 The organisation of the CAA for the regulation of NERL

5.3.1 All major regulatory decisions by the CAA are matters for the board. The board includes executive directors with expertise in both economic and safety areas (respectively the Group Director for Regulatory Policy and the Director of Safety and Airspace Regulation). It also includes non-executive directors with specific experience in regulation, safety, finance and commercial aviation. Within the executive management team, oversight of NATS is coordinated by the NATS Licence Management Coordination Committee (NLMCC), which is chaired by the Head of Economic Regulation. The NLMCC meets regularly and its discussions are minuted. Around ten individuals, including amongst others economists and lawyers are engaged in one way or another in the oversight of the NERL licence. The actual number at any one time will fluctuate with more involved during a price control review. The CAA also makes use of expert consultants from time to time. Around 100 people are employed in the safety function at CAA, of whom around ten are dedicated to the oversight of NERL. Both the regulatory and safety functions can call on advice from a small number of computer systems experts at the CAA.

5.4 The CAA's approach to regulation

5.4.1 The CAA's approach to the economic regulation of NERL is performance based. It is broadly consistent with that of other UK regulators. It relies on three elements:

- The imposition of a structure of financial and other incentives, which encourages NERL to behave in ways which are expected to lead to outcomes which are in the best interests of consumers.
- The setting of performance targets and the publication of NERL's performance against those targets
- The requirement for NERL to inform and consult airlines on its major plans and its performance, and to take account of the views of airlines in making its final decisions. Airlines are assumed to represent the interests of passengers and other end users. The CAA also has a consumers' panel, which is now becoming increasingly involved in licensing matters that affect consumers.

5.4.2 It follows from this approach that the CAA licence team does not supervise or intervene in the day-to-day operating or investment decisions of NERL, and it is not currently organised or resourced to do so. The licence team does not have detailed knowledge of NATS systems and is not, for instance, familiar with SFS.

5.4.3 The safety regulation team has much greater familiarity with the details of NERL's operations. The CAA discharges its responsibility for safety by oversight of the NATS Safety Management System, people, systems, operational procedures and safety cases for changes. It receives and reviews a large number of NERL safety-related documents, and it undertakes continuous assessment of performance at Swanwick through dedicated inspectors. It does not intervene as a matter of course in the day-to-day operations of NERL or the engineering of systems. On 12 December it did not interfere in the operational response by NERL to the system failure, although some of the contingency arrangements in the event of the loss of an air traffic control centre would require specific approval by the CAA.

5.5 Findings

5.5.1 The governance structures of the CAA are clear. They are appropriate for the oversight of both safety and the economic activities of NATS, and there are structures in place to ensure that these two activities are coordinated. The Enquiry has been provided with, relevant papers on regulation by the CAA, in addition to the extensive material on the CAA's website. Panel members have met with the relevant executive directors and senior executives of the CAA on a number of occasions, and have heard points made by airlines. Meetings have been held with senior executives and the chair of NATS and the Director of Aviation at the Department for Transport.

5.5.2 The CAA is actively and extensively engaged in the oversight of the safety of NERL's operations, and no evidence has been found or suggested of any particular failure of oversight of safety in relation to the causes of the Incident. The organisation has staff with high levels of experience and expertise in this area; it also needs to keep its safety assurance processes under continual review, taking into account evolving best practice in other safety critical industries. It will be important for this to be built into oversight of the SESAR programme.

5.5.3 Four questions have emerged in relation to the broader regulation of NERL's activities through its licence:

- Whether the CAA's approach allows or encourages NERL to invest too little or too late in the maintenance and development of ATC systems
- Whether the design of the incentive-based system provides adequate financial penalties and/or rewards to ensure that NERL takes adequate measures to avoid failures

- Whether the CAA can or should increase its oversight of the levels of resilience to system failure provided by NERL
- Whether the focus of the CAA on NERL and its airline customers gives appropriate weight to the interests of air passengers in avoiding the disruption caused by major incidents

5.6 The CAA's approach to investment

5.6.1 The challenge of ensuring that regulated companies invest both at an appropriate level and cost effectively is a continuing issue for UK regulators. The CAA has adopted a variant of the standard regulatory approach to the question. This is to set the maximum prices that NERL is allowed to charge at a level under which an efficient company would make a 'normal' return its assets including the specified level of planned investment, adjusted for risk. These charges would then apply for the price control period (typically five years), and would then be reset. Allowing a normal return on planned investment should give the company a neutral overall investment incentive – that is, it should not have a commercial reason to skimp on investment, or indeed to invest excessively. In its raw form, this approach would leave a tactical risk that NERL could apply for a large investment programme to be included in the allowed level of charges for the next five-year period, but then deliberately avoid or delay investment, thereby taking the return on capital without actually spending the money. The CAA has recognised this and incorporated a mechanism that claws back any returns on planned investment which was not actually undertaken, through a downward revision to the allowed level of charges in the next control period.

5.6.2 This approach is consistent with regulatory best practice for other industries in the UK.

5.6.3 However, by making NERL theoretically financially neutral to the level of investment that the company makes, the CAA provides no incentive structure for NATS to make decisions on how much it should invest in any given period. In pure economic and financial terms, NERL should be 'indifferent' to how much it invests and when to invest. However, shareholders in NERL may have reasons other than the normal commercial focus on returns for wishing to underinvest. These can include giving high priority to minimising the level of ATC charges that they pay as customers during difficult times. This may outweigh the prospect of long-term returns on capital for them as shareholders in NATS. The CAA shields against such an outcome by linking the permitted level of charges to the achievement of specified levels of performance. This device is of limited effect when investment is made in one control period but its benefits are not felt until a later period.

5.6.4 Under this structure, at every price control review, regulators have to form a view on what constitutes the right level of investment as part of the process of setting the allowable level of charges for the next period.

5.6.5 The CAA does this primarily by requiring NATS to set out its capital programme along with its reasoning, and then consults airlines on this programme. Airlines then form their own judgement on what is an appropriate level of spending. Some of the airlines that participate in this process are also shareholders in NATS, and not all airlines participate. NATS is obliged to take the views expressed by airlines into account when finalising the capital programme. The CAA then subjects this programme to a review by expert consultants. This review is at a high level. It does not examine individual components of the programme in any depth; neither does it consider the risks associated with the delivery of the programme or attempt any independent cost–benefit analysis or programme optimisation. It primarily answers the questions of whether the programme was arrived at by an effective process, whether the cost levels for projects appear to be efficient, and whether the overall level of spend appears reasonable.

5.6.6 IATA regards this process as one of the best in the industry. It does, however, remain true that it excludes any mechanism for taking account of the interests of the ultimate customers – i.e. passengers – as distinct from those of airlines. In this respect it differs from other utility regulators who build the interests of end-customers formally into their processes. At the end of the process the CAA itself takes no responsibility for forming an independent informed judgement on the level of investment that should be undertaken by NERL.

5.7 The investment record

5.7.1 With this framework in mind, the Panel has examined the NERL record. Figure 5.1 shows the level of capital expenditure ('capex') approved by the CAA for each year since the Public-Private Partnership (PPP) was formed, and the level of spend actually undertaken, divided into the three price control periods. Over the 14-year period NERL has invested somewhat less than had been planned overall, but there have been distinct variations from year to year. In some years NATS invested more than had been assumed; in others, less.

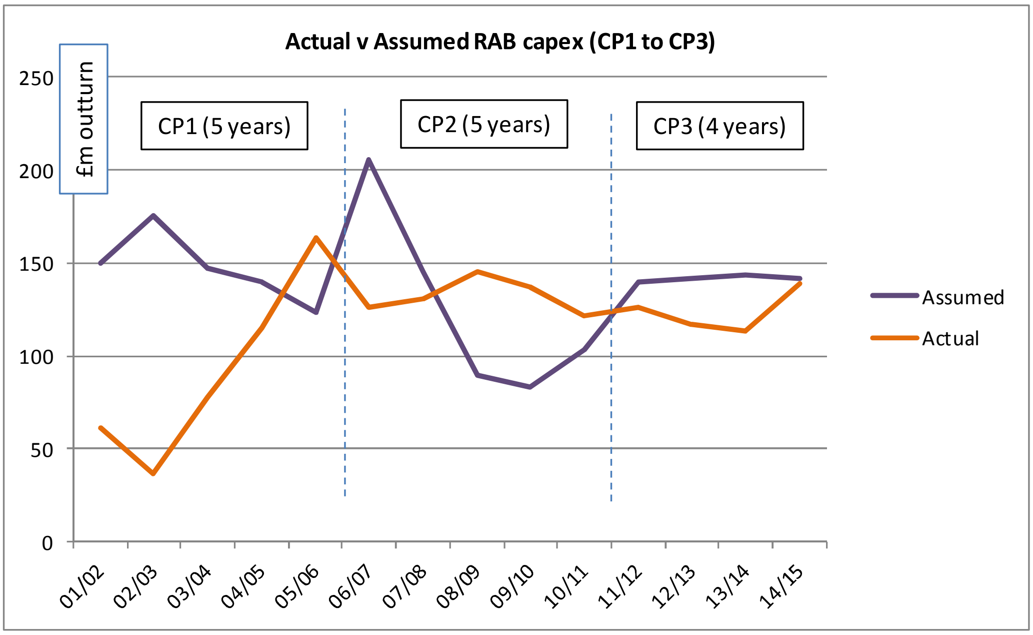

Figure 5.1: Actual v assumed capex; CP1 to CP3 (source: NATS)

5.7.2 For the first five-year period (CP1), actual expenditure was some 38% below the level planned by NERL. The difference is explained by the collapse in traffic and forecast demand growth after 9/11 that resulted in the financial crisis which overtook NATS in 2002. This made it difficult for the company to invest heavily while its financial condition was fragile. Expenditure was also reduced by the decision to delay the new Prestwick control centre to CP3.

5.7.3 In CP2, after a slow start, investment exceeded plan despite the fall in expected growth in demand from 2008 onwards, and the difficult financial conditions facing airline customers. Overall investment over the five-year period was approximately 5% above the planned amount.

5.7.4 Before agreeing the investment programme for CP3, the CAA subjected it to independent review by consultants Logica. However, in the event, the investment actually undertaken over the four years was approximately 12% below the agreed plan. This was partly as a result of the re-phasing of the iTEC-FDP programme and partly in response to airline requests to reduce costs.

5.7.5 Capital expenditure by NERL since the PPP has coincided with an impressive improvement in delay performance. Figure 5.2 shows that average delays have fallen from 60 seconds to around four seconds over the last 12 years. The bulk of this improvement took place in the early part of the period.

5.7.6 Eurocontrol's Network Operations Report shows that the current level of delays compares favourably with those of other ANSPs in Europe of a similar size.

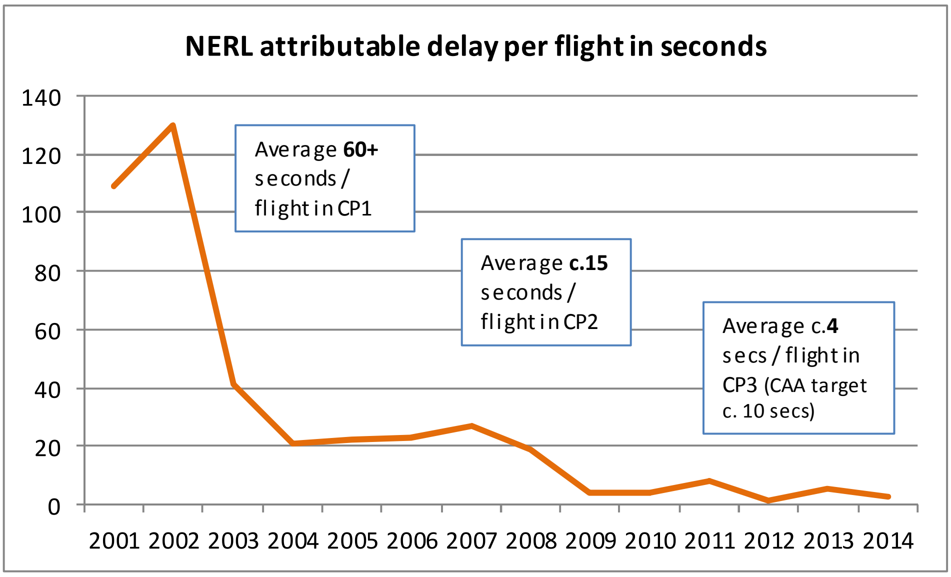

Figure 5.2: NERL attributable delay (source: NATS)

5.8 Current plans

5.8.1 In 2013, NATS submitted its five-year business plan as part of the process of setting the price control for the five-year period from 2015 to 2019. This followed a process imposed by the CAA. NATS initially produced two options: a larger capex plan of £653m that was intended to deliver service enhancement to airlines and capacity quickly; and a smaller plan of £603m in which a greater priority was to minimise charges to airlines at the cost of later delivery of service improvements. A number of airlines engaged in this consultation, some of which were also NATS shareholders. A number of other airlines chose not to engage in any depth. In the end, NATS settled on a programme towards the lower end of the range at £618m, but including most of the planned service enhancements. Projects relating to airspace development and resilience, including London Airspace Management Programme (LAMP) and Northern Terminal Control Airspace (NTCA), were prioritised at the cost of further delay to ITEC-FDP. The CAA subjected the final business plan capital programme to a review by consultants Arup and Helios, using documents and data provided by NERL. This review took place over a period of three months at a cost of £116,000, and its conclusions were published in January 2014. Arup/Helios generally reported positively on the planning process and consultation with users, but made some recommendations for improvement. The low cost of this review of a £600+ million IT investment over 5 years is indicative of its focus on process rather than the justification of the planned investments and their estimated costs.

5.8.2 At each price control review, the range of possible outcomes was framed by the initial business plan programme and options put forward by NATS. The panel has not found evidence that the CAA had undertaken any analysis or formed any judgement of its own as to whether the options put forward by NATS management fully represented the range of investment possibilities available to NATS, or whether an alternative or higher capital programme might be possible that would deliver greater capacity or resilience more quickly. The CAA was entirely dependent on NERL management, the airlines and its consultants for validating the programme.

5.8.3 In between the five-yearly reviews, NERL provides airlines and the CAA with its annual updated SIP, including changes to the size, shape and pace of the capital programme from year to year. It consults airlines on this plan in some detail. The panel has heard that as provided for in the licence, the CAA approves the form, scope and level of detail of each annual update, but within the current licence framework does not challenge the substance of the plan, including changes to the size, shape and pace of the capital programme from year to year. The CAA has told the panel that over the last year it has decided to scrutinise the annual SIP more closely and is able to have an influence on its content, but that it does not instruct NERL on any substantive changes to its annual plans. The CAA has regarded the SIP primarily as a vehicle for effective consultation between NERL and airlines on its plans.

5.9 Financial incentives

5.9.1 The CAA's approach to the regulation of NATS is based on the principle of giving the company strong incentives to make decisions that are in the best interests of the users of NERL's services. As described above, one of these incentives is the system of service quality bonuses and penalties. This is set out in the UK-Ireland FAB RP2 Performance Plan[42]. Under this approach, the level of charges that NERL is entitled to levy on airlines is reduced in the event of failure to meet specific targets. The relevant targets include ATC delay and environmental impact.

5.9.2 The maximum amount of the penalty for delay in any one year is set under European legislation at 1% of total airspace charges, broadly equivalent to £6m at current rates. It comprises two elements. 0.75% is made up of a measure of total delays, weighted towards longer delays and delays in peak operating periods. The remaining 0.25% is an excess daily score which penalises 'bad days' when there are high levels of delay, including those arising from system failures.

5.9.3 In addition, the CAA has created a penalty regime of up to 1% for failure to meet the environmental target. The target is known as 3Di. It is a complex measure of the vertical and horizontal optimality of flight paths and is a proxy for fuel efficiency. It is noteworthy that CAA has applied these regimes separately for a total value of 2% whereas other EU states have limited the total incentive regime to 1%.

5.9.4 As a comparison, The CAA has also established a bonus and penalty regime for a package of service measures at Heathrow airport. The maximum penalty for all measures at Heathrow taken together and published in the airport's licence is set at 7% of charges..

5.9.5 NATS accounts for penalties in the year in which they are imposed, by making a provision in its profit and loss account. However, for reasons concerning the technical working of the price control algebra, there is a time lag between incidents that incur penalties (or bonuses) and the compensating reduction (or increase) in unit charges to airlines. In practice, the change in charges applies in the second financial year commencing after the event.

5.9.6 The CAA is considering seeking primary legislation to extend its range of enforcement powers to include the ability to take action against past breaches of the licence and levy financial penalties ('fines') on NERL for significant failures of service amounting to an identifiable breach of its service obligations as set out in the licence, or of its statutory duties. The grant of such a power would not by itself guarantee that there will be no further system failures, but the expectation that the CAA will exercise sufficient oversight to minimise the risks of reoccurrence does require it to have adequate power to enforce its oversight. Powers to impose penalties were granted to the CAA in relation to airport regulation in 2012. Other regulators have similar fining powers that can in most cases extend to penalties of up to 10% of turnover. The have been used by the Office of Rail and Road (ORR)[43], The Water Services Regulatory Authority (Ofwat), Office of Gas and Electricity Markets (Ofgem), Office of Communications (Ofcom) and the Financial Services Authority (FSA) to incentivise

[42] FAB Performance Plan, UK-Ireland FAB, Second Reference Period (2015-2019), June 2014.

[43] Formerly Office of Rail Regulation

compliance, deter future non-compliance, punish serious transgressions and, where possible, remedy the consequences of a breach.

5.9.7 The CAA points to a number of potential benefits of such an approach. In particular the ability to impose financial penalties has an incentive effect on the owners of the company in that at least some of the money paid out in penalties would otherwise have been available for dividends. Furthermore, the publicity surrounding financial penalties can be a deterrent. Such publicity can lead to reputational damage to a company which multiplies the effect of the direct financial loss, and which management and the company will be expected to take all reasonable steps to avoid.

5.9.8 NATS has expressed concerns about such an extension of the CAA's powers, particularly if it were to be used as an additional mechanism to magnify the bonus and penalty system already in place for delays at the expense of the organisation's focus on safety. NATS is also concerned that the imposition of fines may have unintended and undesirable consequences, including a rise in the company's cost of capital which might cause an increase in charges to airlines. The enquiry does not believe that there would be a significant effect on the cost of capital unless NATS conducts itself in such a way as to incur repeated penalties. NATS is also concerned that the basis of penalties does not become unduly prescriptive. The enquiry understands that this is not the CAA's intention and that the CAA, like NATS is fully appreciative of the fact that the primary duties of both organisations concern the safety of operations. As in the case of other regulators, fines should only be imposed after full investigation of the circumstances of each breach of the licence and having taken account of all possible consequences.

5.9.9 The Department for Transport has informed the inquiry that an opportunity to legislate for the necessary change to the CAA's enforcement powers may arise in the near future as part of a broader Aviation Bill.

5.9.10 The remuneration arrangements for senior executives of public companies are generally designed to reward and incentivise performance. The current NATS executive remuneration scheme, as set out in the Annual Report for the year ending 31 March 2014[44], includes an annual bonus (up to 70% of salary for the CEO) and a long-term incentive plan (LTIP) rewarding longer-term achievements and vesting over a three-year period. (This was paid out at 78.45% for the CEO in 2014.) Both the annual bonus and the LTIP payment are made in cash. The annual bonus is partly attributable to the achievement of a profit target, partly based on personal objectives, and partly (30%) based on group non-financial objectives including delays; the Annual Report does not contain details of the LTIP scheme. The Remuneration Committee of the NATS Board has an additional discretion to recommend to the NATS non-executive directors that they withhold part or all of the bonus if they believe that management has fallen short in its stewardship of the company. In 2014, this reduced the potential bonus of the CEO by 12% to reflect the communications failure of 13 December 2013 that led to an interruption to normal service.

5.9.11 Unlike in many other commercial companies, the top executives do not benefit from a share-based LTIP that is designed to align the interests of management with the long-term development of the business.

5.9.12 The CAA has confirmed that, whilst in theory it may have the power to regulate executive pay though conditions in the NERL licence, it has not done so. Such a move would be an unprecedented level of intervention by a UK utility regulator. The CAA also accepts that there would be scope for the exercise of 'soft power' though discussions about remuneration with the board of NATS.

[44] Annual Report and Accounts 2014, NATS Holdings Limited, June 2014

5.9.13 In the course of this review, one airline has suggested that NERL might be required to pay compensation to airlines for any direct costs to them arising from incidents. At present, NERL is shielded from a requirement to pay this kind of compensation by an express exclusion in section 10(1) of the Transport Act. A change to the legislation to allow compensation would have significant precedential implications for airports and other transport operators, as well as for NATS. If the legislation were changed, mechanisms would have to be put in place to determine the methods used to calculate and audit compensation claims. These mechanisms would probably be expensive to operate and lead to arguments about whether consequential costs were incurred economically (see also 5.9.6)..

5.10 Passengers

5.10.1 The CAA's primary duty under the Transport Act 2000 is to maintain high standards of safety but its other duties include an express duty to further passenger interests.

5.10.2 Currently, NERL reports performance to the CAA largely on the basis of seconds of delay per aircraft, and assesses the costs of such delay in terms of the costs borne by airlines in line with the European performance regime. The CAA does not require NERL to report any information on the number of passengers who suffer ATC-related delays, diversions or cancellations; the length of any delays; or the costs of those delays and cancellations to passengers in terms of inconvenience and loss of time. NATS does not produce information in this form.

5.10.3 The total aggregate delay to aircraft attributed to NATS on 12 December was calculated by NATS as 14,863 minutes applied across 353 flights. This does not include flights cancelled as a result of the disruption, of which NATS estimate there to have been approximately 150 with up to a further 20 being diverted away from UK airspace. The European Network Manager has recorded a total delay of 18,433 minutes caused by the NATS Systems failure of which 3,983 minutes was recorded by neighbouring Air Traffic Control Centres (mainly Paris and Brest). The figures calculated by NATS and the NM are therefore broadly consistent. According to the European Network Manager, the aggregate delay reported by airlines was, however, more than double these ATC figures. He suggests that an explanation for this divergence may derive from the European Network Manager calculating individual flight delays from the datum of a new estimated (and delayed) departure time, whereas the airlines would continue to use the original departure time.

5.10.4 No passenger-related targets are set, and no bonuses or penalties are paid in relation to passenger delays, diversions or cancellations. The enquiry has experienced difficulty in establishing the number of passengers affected by the Incident of 12 December or the scale of lost time, inconvenience and distress they experienced. However NATS has assessed that around 65000 passengers were impacted directly by the flights delayed, cancelled or diverted as a result of this failure. As set out in Chapter 3, NATS has also estimated that the total number of passengers delayed on 12 December as a result of knock-on effects was some 230,000 with up to a further 6,000 passengers delayed on the following day.

5.10.5 Cost–benefit appraisals undertaken by NATS do not include the impacts on passengers from aircraft delays separately from the costs borne by airlines, and the CAA's assessment of the NATS investment programme does not take account of the additional benefits to passengers of early investment to reduce delays, diversions and cancellations or add to resilience.

5.10.6 The enquiry was told by NATS that the airlines regard the interests of passengers as a matter for airlines and this had conditioned the CAA's approach based on the proposition that, unless proved otherwise, the interests of passengers will be aligned with those of airlines. The CAA accepts that this a working assumption that has not been verified, and that it may not always hold true. From the enquiry viewpoint and because passenger information is neither collected nor contributes to the performance management arrangements, the assumption that airlines are an adequate proxy is unlikely to hold on all occasions. In exercising similar regulatory duties in relation to Heathrow and Gatwick airports the CAA adopts the same proposition that the interests of airlines passengers are normally the same. However, it does require the airport to measure a number of aspects of passenger service and satisfaction and incorporates measures of the passenger experience in the airport bonus and penalty regime.

5.10.7 The CAA and NATS also argue that estimating passenger delays and the associated costs and benefits may be very difficult in practice. Estimates of the amount of passenger delay on flights directly affected by events should be obtainable, given that estimates of delay to aircraft are available (but noting that cancelled flights do not contribute to the delay 'score') and there is data on passenger numbers on all flights. There are a number of established mechanisms for estimating the impact of delays, based on (for instance) cost–benefit analysis and customer preference studies, but each has its limitations and none can be relied on for precision. The enquiry was informed by the CAA that no analysis of this kind has been undertaken.

5.10.8 The CAA has told the enquiry that it is currently contemplating the integration of its consumer functions with its licensing functions. The enquiry welcomes this as a step towards the better recognition of the importance of the interests of passengers in NATS oversight.

5.11 Resilience

5.11.1 The CAA currently monitors the resilience of NERL's operations primarily through the three measures of delay described in paragraphs 5.9.2 above. The CAA has told the enquiry that, following the events of December 2013, it has given thought to how it might improve the oversight of resilience in the air services system. It has considered a number of questions of principle, including the definition and measurement of resilience, the linkage between investment and resilience, the setting of appropriate standards, and the need to ensure that safety is not compromised by those standards. So far it has not come to firm conclusions. In the meantime, in resetting the licence conditions for Heathrow airport after a number of major service disruptions such as that resulting from snow in 2010, it has imposed a specific condition on the airports to produce Operational Resilience Plans. The enquiry understands that such a requirement in relation to NERL is amongst the options being considered by the CAA for NATS.

5.12 Conclusions

5.12.1 The CAA has an appropriate governance and organisational structure to discharge its regulatory functions in relation to NERL. It has a great deal of operational expertise and clear regulatory processes in place. The enquiry has found no evidence or suggestion that any failure of oversight by the CAA contributed to the cause of the Incident on 12 December, or that lack of oversight led to any failure in the process of recovery. Neither has it found any failure in oversight that put safety at risk on that day. It will be important for the current high quality processes and staff experience to be built into oversight of the SESAR programme and, noting the experience described in the "Previous Lessons" section of Chapter 3, to ensure an appropriate degree of rigour is applied in closing off recommendations.

5.12.2 However, in the course of its review the enquiry has noted five respects in which the regulation of NERL might be refined and/or extended to reduce the likelihood of further such incidents occurring, and/or which might assist in the most effective recovery from such incidents.

5.12.3 Specifically, it concludes that the CAA should not depend on consultation between NERL and airlines to validate the ten-year capital programme, but should develop its own capacity to form judgements on this important subject.

5.12.4 The CAA should be able to call on enforcement powers in the event of major service failures similar to those it already holds in relation to airports and those held by other UK regulators.

5.12.5 In as much as these powers include the right to impose financial penalties, such penalties should not form part of the normal mechanisms of oversight, but should be deployed only as a measure of last resort, to be enforced where breaches if the NATS licence or statutory duties are so severe as to be beyond other remedies and after the full consequences of the penalty have been evaluated.

5.12.6 The remuneration policies of NATS currently offer limited incentives to executives to align themselves with the long-term interests of airlines and passengers. The CAA has scope to influence the development of these policies in the interests of NERL's customers.

5.12.7 The panel has found a narrow focus by both NATS and the CAA on the impact of service disruptions on airlines. The needs of air passengers need to be built more formally into the oversight of NERL. A precondition for this is the collection of data on the direct impacts of events on delays, cancellations and inconvenience to passengers.

5.12.8 The publication of a resilience plan by NATS would help the CAA to monitor and test the effectiveness of arrangements.

5.13 Recommendations

R27. The CAA should ensure that they have sufficient internal expertise to enable them to complement, select and manage external consultants in analysing and assuring the NATS capital programme, and overseeing its evolution through the annual Service and Investment Plan (SIP) (Para 5.12.3)

R28. The CAA should pursue the inclusion in any forthcoming Aviation Bill of powers to enforce appropriate levels of service by NATS, through, the grant of a power to levy fines for serious or repeated breaches of its licence. Such powers should only be invoked as a measure of last resort and having given full consideration to their possible implications for all aspects of NATS's culture and operations (Para 5.12.4; 5.12.5)

R29. The CAA and NATS should develop systems to estimate, monitor and publish the scale and direct impact to passengers of serious events causing air traffic control disruption (Para 5.12.7)

R30. The CAA should require NERL to submit a resilience plan for approval by the CAA as a condition of its licence (Para 5.12.8)

R31. NATS should review its remuneration policy in order to better align the incentives for management with the achievement of long-term objectives for delays and resilience. (Para 5.12.6)

Annex A to Independent Enquiry Report dated 13 May 2014

Annex A. Enquiry Terms of Reference

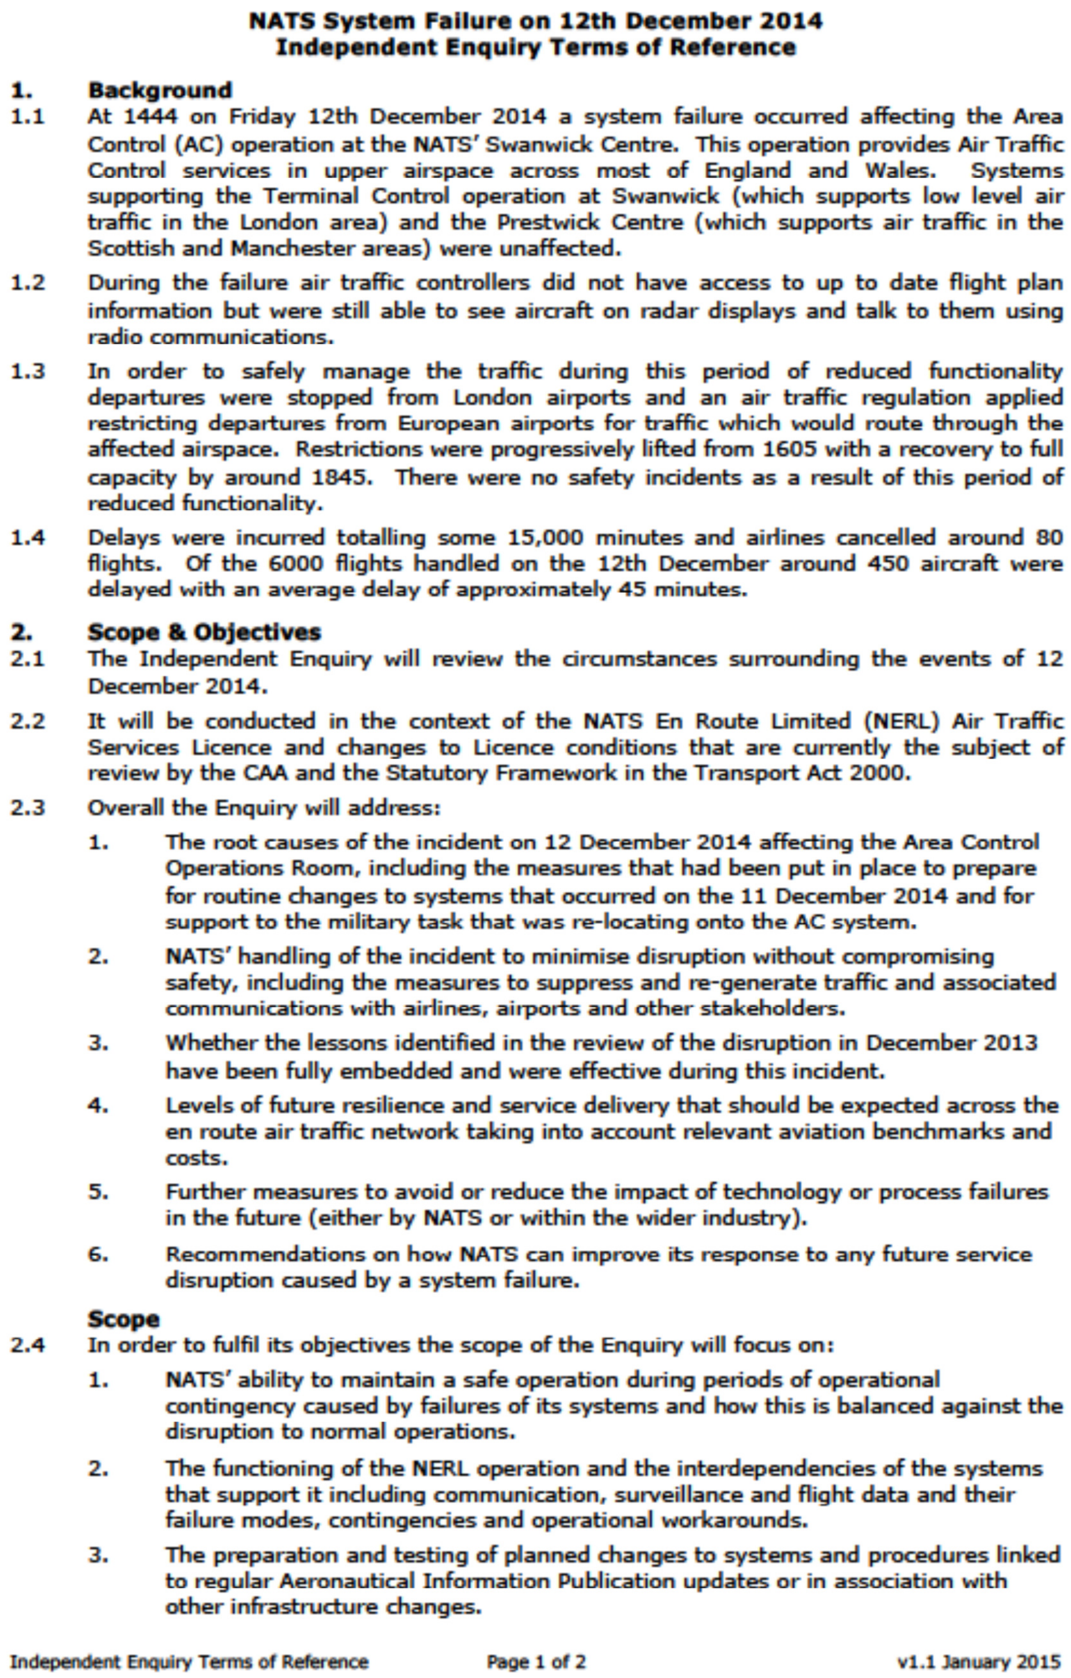

NATS System Failure on 12th December 2014
Independent Enquiry Terms of Reference

1. Background

1.1 At 1444 on Friday 12th December 2014 a system failure occurred affecting the Area Control (AC) operation at the NATS' Swanwick Centre. This operation provides Air Traffic Control services in upper airspace across most of England and Wales. Systems supporting the Terminal Control operation at Swanwick (which supports low level air traffic in the London area) and the Prestwick Centre (which supports air traffic in the Scottish and Manchester areas) were unaffected.

1.2 During the failure air traffic controllers did not have access to up to date flight plan information but were still able to see aircraft on radar displays and talk to them using radio communications.

1.3 In order to safely manage the traffic during this period of reduced functionality departures were stopped from London airports and an air traffic regulation applied restricting departures from European airports for traffic which would route through the affected airspace. Restrictions were progressively lifted from 1605 with a recovery to full capacity by around 1845. There were no safety incidents as a result of this period of reduced functionality.

1.4 Delays were incurred totalling some 15,000 minutes and airlines cancelled around 80 flights. Of the 6000 flights handled on the 12th December around 450 aircraft were delayed with an average delay of approximately 45 minutes.

2. Scope & Objectives

2.1 The Independent Enquiry will review the circumstances surrounding the events of 12 December 2014.

2.2 It will be conducted in the context of the NATS En Route Limited (NERL) Air Traffic Services Licence and changes to Licence conditions that are currently the subject of review by the CAA and the Statutory Framework in the Transport Act 2000.

2.3 Overall the Enquiry will address:

1. The root causes of the incident on 12 December 2014 affecting the Area Control Operations Room, including the measures that had been put in place to prepare for routine changes to systems that occurred on the 11 December 2014 and for support to the military task that was re-locating onto the AC system.
2. NATS' handling of the incident to minimise disruption without compromising safety, including the measures to suppress and re-generate traffic and associated communications with airlines, airports and other stakeholders.
3. Whether the lessons identified in the review of the disruption in December 2013 have been fully embedded and were effective during this incident.
4. Levels of future resilience and service delivery that should be expected across the en route air traffic network taking into account relevant aviation benchmarks and costs.
5. Further measures to avoid or reduce the impact of technology or process failures in the future (either by NATS or within the wider industry).
6. Recommendations on how NATS can improve its response to any future service disruption caused by a system failure.

Scope

2.4 In order to fulfil its objectives the scope of the Enquiry will focus on:

1. NATS' ability to maintain a safe operation during periods of operational contingency caused by failures of its systems and how this is balanced against the disruption to normal operations.
2. The functioning of the NERL operation and the interdependencies of the systems that support it including communication, surveillance and flight data and their failure modes, contingencies and operational workarounds.
3. The preparation and testing of planned changes to systems and procedures linked to regular Aeronautical Information Publication updates or in association with other infrastructure changes.

Independent Enquiry Terms of Reference Page 1 of 2 v1.1 January 2015

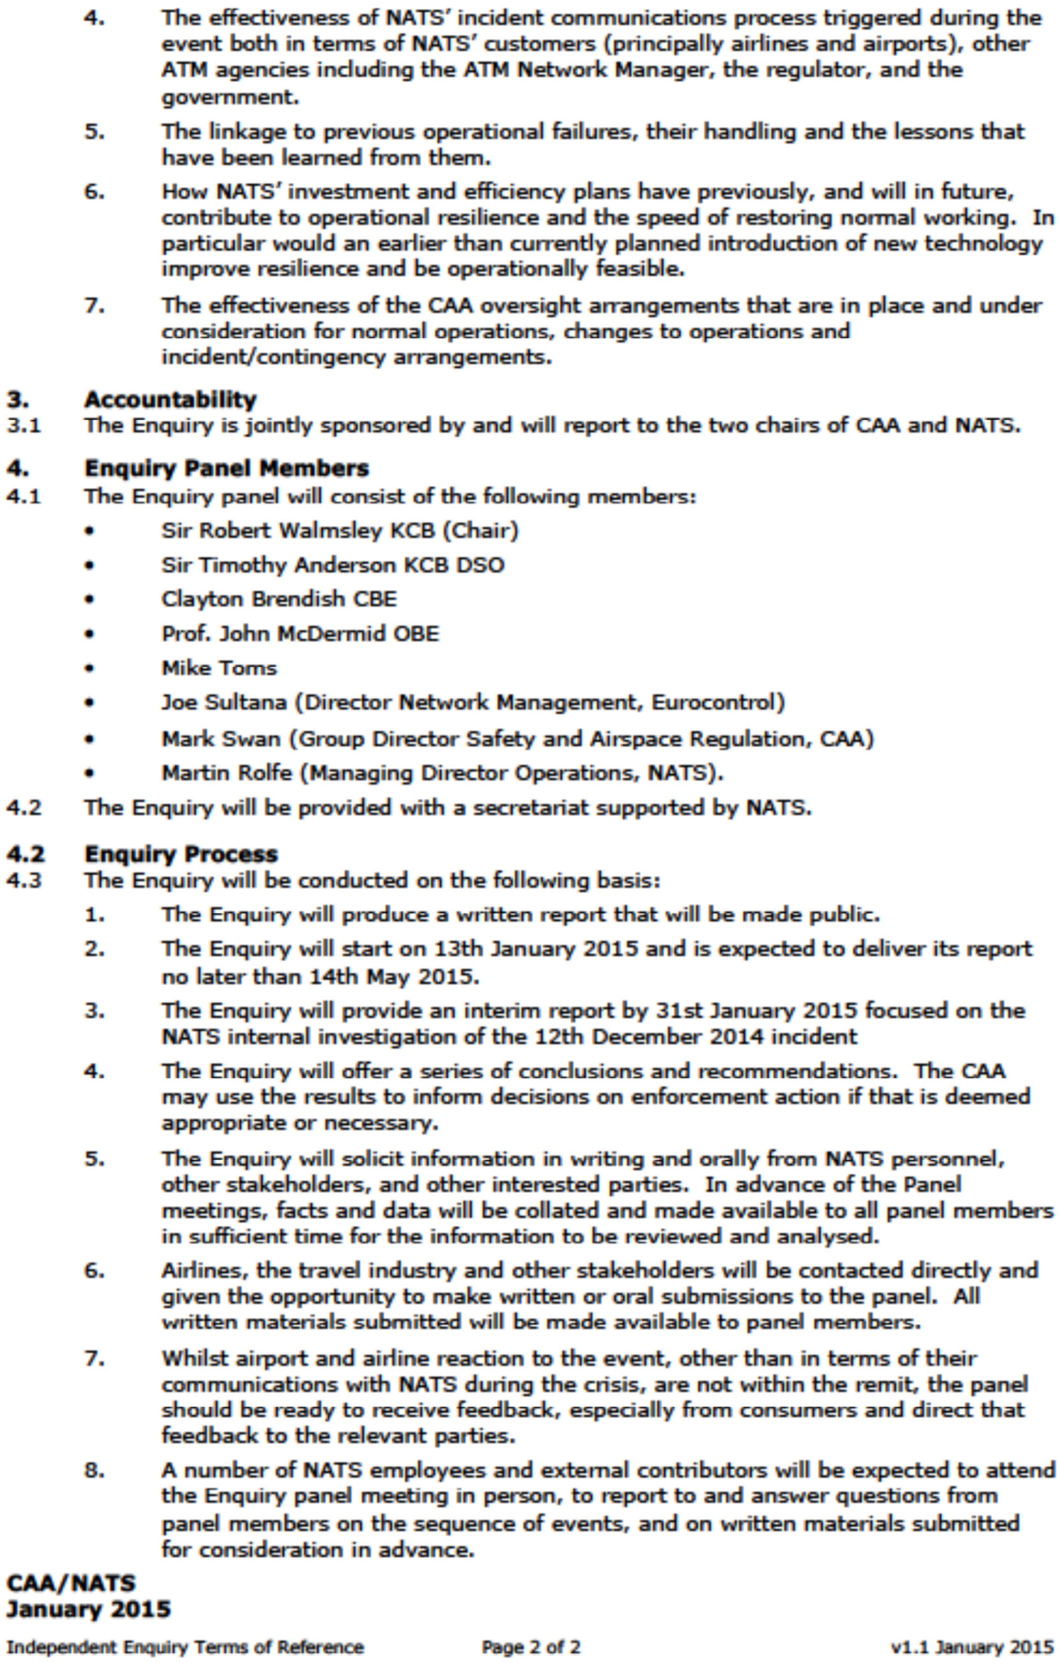

4. The effectiveness of NATS' incident communications process triggered during the event both in terms of NATS' customers (principally airlines and airports), other ATM agencies including the ATM Network Manager, the regulator, and the government.
5. The linkage to previous operational failures, their handling and the lessons that have been learned from them.
6. How NATS' investment and efficiency plans have previously, and will in future, contribute to operational resilience and the speed of restoring normal working. In particular would an earlier than currently planned introduction of new technology improve resilience and be operationally feasible.
7. The effectiveness of the CAA oversight arrangements that are in place and under consideration for normal operations, changes to operations and incident/contingency arrangements.

3. Accountability

3.1 The Enquiry is jointly sponsored by and will report to the two chairs of CAA and NATS.

4. Enquiry Panel Members

4.1 The Enquiry panel will consist of the following members:

- Sir Robert Walmsley KCB (Chair)
- Sir Timothy Anderson KCB DSO
- Clayton Brendish CBE
- Prof. John McDermid OBE
- Mike Toms
- Joe Sultana (Director Network Management, Eurocontrol)
- Mark Swan (Group Director Safety and Airspace Regulation, CAA)
- Martin Rolfe (Managing Director Operations, NATS).

4.2 The Enquiry will be provided with a secretariat supported by NATS.

4.2 Enquiry Process

4.3 The Enquiry will be conducted on the following basis:

1. The Enquiry will produce a written report that will be made public.
2. The Enquiry will start on 13th January 2015 and is expected to deliver its report no later than 14th May 2015.
3. The Enquiry will provide an interim report by 31st January 2015 focused on the NATS internal investigation of the 12th December 2014 incident
4. The Enquiry will offer a series of conclusions and recommendations. The CAA may use the results to inform decisions on enforcement action if that is deemed appropriate or necessary.
5. The Enquiry will solicit information in writing and orally from NATS personnel, other stakeholders, and other interested parties. In advance of the Panel meetings, facts and data will be collated and made available to all panel members in sufficient time for the information to be reviewed and analysed.
6. Airlines, the travel industry and other stakeholders will be contacted directly and given the opportunity to make written or oral submissions to the panel. All written materials submitted will be made available to panel members.
7. Whilst airport and airline reaction to the event, other than in terms of their communications with NATS during the crisis, are not within the remit, the panel should be ready to receive feedback, especially from consumers and direct that feedback to the relevant parties.
8. A number of NATS employees and external contributors will be expected to attend the Enquiry panel meeting in person, to report to and answer questions from panel members on the sequence of events, and on written materials submitted for consideration in advance.

CAA/NATS
January 2015

Independent Enquiry Terms of Reference Page 2 of 2 v1.1 January 2015

Annex B. Panel Details and CVs

Sir Robert Walmsley KCB FREng, Panel Chairman

Entered the Royal Navy as a Dartmouth Cadet and served a full career as an engineering officer, mainly in submarines and specialising in nuclear propulsion and acquisition. Retired as a Vice Admiral in 1996 to take up the MoD post of Chief of Defence Procurement where he served for 7 years. Since then he has been on six public company boards equally spread between the UK and US; currently he sits as a Director on the boards of Cohort plc and of Ultra Electronics plc; in the US, he was an independent Director of the General Dynamics Corporation from 2004 until May 2015. Since 2013, he has been Non-Executive Chairman of the Programme Board for Universal Credit in the Department for Work and Pensions. He was Chairman of the Major Projects Association from 2004 until 2013 and holds a number of advisory positions in the fields of energy, aerospace and defence.

Sir Timothy Anderson KCB DSO FRAeS

A military pilot by profession, he retired recently from the Royal Air Force after a full career including senior command and staff appointments spanning operations, acquisition and policy. Whilst Assistant Chief of the Air Staff, amongst other things he was responsible for Defence airspace policy and a non-executive director on the CAA Main Board. He established the UK Military Aviation Authority in 2010, the world's first fully integrated military aviation regulator. As its first Director General, he engendered a step change in military air safety culture, whilst overseeing design, operation and maintenance activity across the Defence air environment. He is currently a non-executive director on the Flybe regional airline Gp Board and advisor to a number of clients in the aerospace, defence and security sectors.

Clay Brendish CBE

Clay Brendish CBE is Non-Executive Chairman of Anite Group plc and SThree plc. Clay is also a Strategic Security Consultant with Compagnie Financière Richemont SA and provides IT consultancy to BT.

Clay was Executive Chairman of Admiral plc that he co-founded in 1979. Admiral plc employed over 2500 people in 8 countries. Following Admiral's purchase by CMG plc in June 2000, he became Deputy Chairman of CMG plc, retiring in 2001. Clay was a Non-Executive Director of BT Group plc from September 2002 to August 2011.

Between 1993 and July 2000 he was an advisor to the Chancellor of the Duchy of Lancaster and the Parliamentary Secretary, Office of Public Services, on their Next Steps agencies. He played a prominent role in the privatisation of a number of Cabinet Office Agencies. He was Non-Executive Director of Ordnance Survey between 1993 and 1996, a Non-Executive Director of the Defence Logistics Organisation between February 2001 and July 2004 and an External Member of the UK Defence Meteorological Office Board between June 1995 and March 2003 before being appointed Director and External Chairman of the Met Office from July 2003 to October 2006.

Professor John A McDermid OBE FREng

John McDermid joined the MoD as a student engineer, spent several years at the Royal Signals and Radar Establishment and worked in a software house for five years. He has been Professor of Software Engineering at the University of York since 1987. He set up the High Integrity Systems Engineering (HISE) research group in the Department of Computer Science and was Head of the Department from 2006 to 2012. HISE studies a broad range of issues in systems, software and safety engineering, and works closely with government and industry, e.g. Airbus, BAE Systems, the CAA, the MoD, QinetiQ and Rolls-Royce. He is author or editor of six books and has published over 380 papers. He has advised companies and government departments on several continents, including advising the US Nuclear Regulatory Commission (NRC) on software safety. He has set up and run a number of small companies, and became Chairman of Rapita Systems in January 2014.

Martin Rolfe

Graduated with a Masters Degree in Aerospace Systems Engineering from the University of Southampton in 1994. Joined Logica, a software house, to work on high integrity real time systems in the manufacturing and Transport sector. Joined Lockheed Martin in 1998 as a software test and verification engineer, which included working on the Swanwick Centre ATM systems. Following the operational transition of Swanwick area control in 2002, moved to the USA to work on the Federal Aviation Administration's En-Route Modernisation programme (ERAM) - a replacement system for the 22 US En-Route ATM centres. Initially starting as Flight Data Processing lead engineer and then moving on to be Chief Engineer for the programme. On returning to the UK took up the position of Managing Director of Lockheed Martin UK for Civil Government and Director of International ATM. Appointed Managing Director, Operations for NATS in 2011 with responsibility for the NATS En-Route business.

Joe Sultana

Graduated with an Engineering Degree from the University of Malta and joined the Air Traffic Services Unit in Malta. Appointed as Head of Air Traffic Services in the Maltese Department of Civil Aviation in 1982 and subsequently appointed as Deputy Director of Civil Aviation responsible for the Air Traffic Services organisation in 1984. Joined EUROCONTROL in 1991 as an Airspace Management Expert coordinating Airspace and Navigation projects. Within EUROCONTROL took on roles as RVSM Programme Manager, Head of Network Capacity Business Division and Head of Airspace, Network Planning and Navigation Division. Since 2008, joined the Central Flow Management Unit as Head of Operations. Became Deputy Director CFMU responsible for Network Operations and Information Management in 2009. Within the Directorate Network Management, Joe was first promoted to Director as Chief Operating Officer in 2011. Since 2013, Joe is the Director Network Manager responsible to fulfil the role of the Network Manager established with the Single European Sky.

Mark Swan

Mark was appointed to the Board as Group Director Airspace Policy in March 2008. In July 2013 he was charged with merging the Airspace and Safety groups and re-structuring the combined group to focus on performance-based regulation. He is currently Director Safety and Airspace Regulation.

Mark previously held numerous appointments in the Royal Air Force since joining as a pilot in 1979 and was formerly Director of Operational Audit for the Ministry of Defence from 2006 to 2008.

Mike Toms MA, FRAeS, MRICS,MRTPI

Graduated from Durham and Nottingham universities. Joined the British Airports Authority (later BAA plc) in 1980 and held various positions including Chief Economist and Director of Strategy. Seconded to the Airports Council International for 1991/2 as its founding Chief Economist. Joined the board of BAA plc in 2002 as Planning And Regulation Director. Retired in 2006 and subsequently served as a director of the Viridian Group and UK Coal plc and chair of Northern Ireland Electricity plc. Currently a director of Birmingham Airport, Bellway plc and Oxera Consulting and Chair of the Connections and use of Service Panel for the National Grid. Adviser on economics and regulation to a number of major airport operators in the UK, Europe. Asia and Australasia.

Annex C. Timelines for the Incident, Response and Recovery

C.1.1 This section describes the sequence of events that took place during the event on 12 December 2014, looking at this from an operational, engineering and overall incident management perspective. Each entry describes a specific event during the incident response and recovery together with the time of occurrence and a brief explanation of the nature of the event.

Time	Key Event
1444	LAC Workstations show "SFS Unavailable" *Controller workstations in the London Area Control (LAC) Operations room display a "brown border" indicating an error status and an associated error message indicating that the System Flight Server (SFS) is unavailable. The error indicates that both primary and secondary servers have failed and requires operation in a reduced capability fallback mode.* *SFS is a core part of the system that supports Area Control at Swanwick. In simple terms it receives flight plans from the central flight planning system (National Airspace System, NAS), processes these, and ensures distribution of them to the right controller working positions when required.*
1445	Loss of NAS SFS Link reported *Engineers in System Control at Swanwick receive an indication from NAS and the LAC control and monitoring that the link between NAS and SFS was lost meaning that SFS is no longer able to receive and distribute up to date flight plan information.*
1455	Swanwick Silver and ATICCC activated. All London TMA departures stopped. *Crisis management capabilities activated including the Swanwick Silver Team to coordinate the overall response at Swanwick involving members of the Swanwick Leadership Team and the Air Traffic Incident Communication and Coordination Cell (ATICCC)which leads interaction with customers.* *All departures from London TMA airports stopped to allow a gradual reduction in traffic reducing controller workload.*
1500	Initial Zero rate regulation applied *An air traffic regulation applied to restrict departures from within Europe for aircraft which would be expected to route through the affected airspace again allowing a reduction in traffic. In order to ensure all European traffic is captured a period of operation of 4 hours is applied.*
1505	Engineering Technical Incident Cell Convened *Engineering Cell at Swanwick activated which allows a joined up response to assessing and recovering from the failure, taking the engineering conversation away from the ops room, and allowing coordinated involvement of other engineering experts.*
1515	Gold team and CMF activated *The Gold team which represents the senior level of crisis management response activated at the Corporate and Technical Centre (CTC) together with the Crisis Management Facility (CMF) which coordinates information gathering and command and control where necessary on behalf of the Gold Team*
1522	SFS Server B Reset *SFS determined to be still running but disconnected from the workstations due to the error that head been detected. The B Server was reset to prepare it for repopulation with data and recovery to full operation.*
1525	SFS Server B Restored *Following the reset SFS Server B reports that it is restored and ready for service but as yet has not flight data available.*
1530	ATC advised to prepare for NAS flight data download *Preparation for the download of flight data from the central flight data processing system (NAS) to SFS to repopulate its flight plan database.*
1536	ATICCC activation messages sent to customers

	Emails and text messages sent to registered customers to formally notify them of ATICCC activation and to provide details of first ATICCC teleconference scheduled for 1600.
1541	NAS to SFS data download recovery commenced *NAS to SFS down load formally activated by Engineering.*
1543	NAS to SFS recovery complete *SFS now restored, repopulated with data from NAS and ready to resume normal service once ATC were ready for this to take place.*
1545	First internal ATICCC call *Internal teleconference within NATS between affected parties to coordinate status, ensure a common understanding of the situation and confirm key messages for customer communication.*
1549	ATC service resumed – electronic coordination enabled *Normal operation of the ATC systems at Swanwick was resumed (with reduced redundancy) allowing the ATC operation to return to normal procedures rather than the reduced capability fallback mode.*
1555	Departure regulations at Heathrow, Gatwick and Manchester cancelled. MDIs applied: 1 every 5 minutes northbound and southbound. *First removal of some of the air traffic restrictions allowing air traffic levels to begin to increase again. Minimum Departure Intervals (MDIs) applied to allow a managed flow of traffic into the surrounding airspace, with one departure every 5 minutes allowed for northbound aircraft and similarly for southbound aircraft..*
1600	First customer ATICCC call *Teleconference with customers to brief them on the status of systems, impact on traffic levels and anticipated recovery timescales.* Confirmed to customers restoration of ATC service and signalled commencement of a progression removal of ATC restrictions
1605	Zero regulation lifted, capacity raised to 75% *The regulations applied at 1500 were relaxed allowing operation at up to 75% of normal capacity.*
1610	Departure regulations lifted at the other London airports *Removal of departure restrictions.*
1630	Heathrow arrival rate set at 20 due to ground / stand congestion *Heathrow unable to operate at full capacity due to ground / stand congestion so a limit of 20 arrivals per hour set so as to manage the traffic flow in the adjacent airspace.*
1645	Heathrow departure rate increased to 1 every 5 minutes per SID *Heathrow departure restrictions significantly relaxed allowing 1 departure every 5 minutes for each of 6 Standard Instrument Departures operated from Heathrow..*
1655	Heathrow arrival rate increased to 40 per hour *Rate of 40 per hour is close to normal operations.*
1730	Departure restrictions cancelled *All departure restrictions from UK airports cancelled.*
1750	En route regulations except Dover and Clacton had been cancelled *The Dover and Clacton restriction were in place due to normal capacity constraints in those sectors rather than as a direct result of the failure. Note however, that the traffic may not have at such a high level in those sectors if the failure had not occurred*
1935	Heathrow arrival restrictions cancelled *Restrictions due to Heathrow ground congestion removed.*
2006	SFS Server A made available as standby – full redundancy returned *Although the second server had been recovered it had not been fully restored as the standby until engineering was sure that there was not a risk of both servers becoming unavailable as had happened at the original failure. At this point they were confident to re-enable the automatic standby capability.*
2030	Final regulation cancelled

Annex D. UK ATM Description

D.1 Background

D.1.1 Air Traffic Management in the UK is carried out in two adjoining regions, The Scottish Flight Information Region (FIR) and the London FIR. The London FIR is divided into:

(1) London Area Control (LAC), which handles civil aircraft over England and Wales in flight at high level.

(2) London Terminal Control (LTC) which is a smaller area, including the five main London airports, and covers aircraft generally flying below 21,500 feet, with the precise height demarcation with LAC depending on the location.

D.1.2 These areas are shown diagrammatically in Figure D.1 below. Aircraft passing through UK airspace (principally between Europe and North America) transit LAC en route; aircraft destined for the London Airports transfer from LAC to LTC as they descend and vice-versa for departing aircraft.

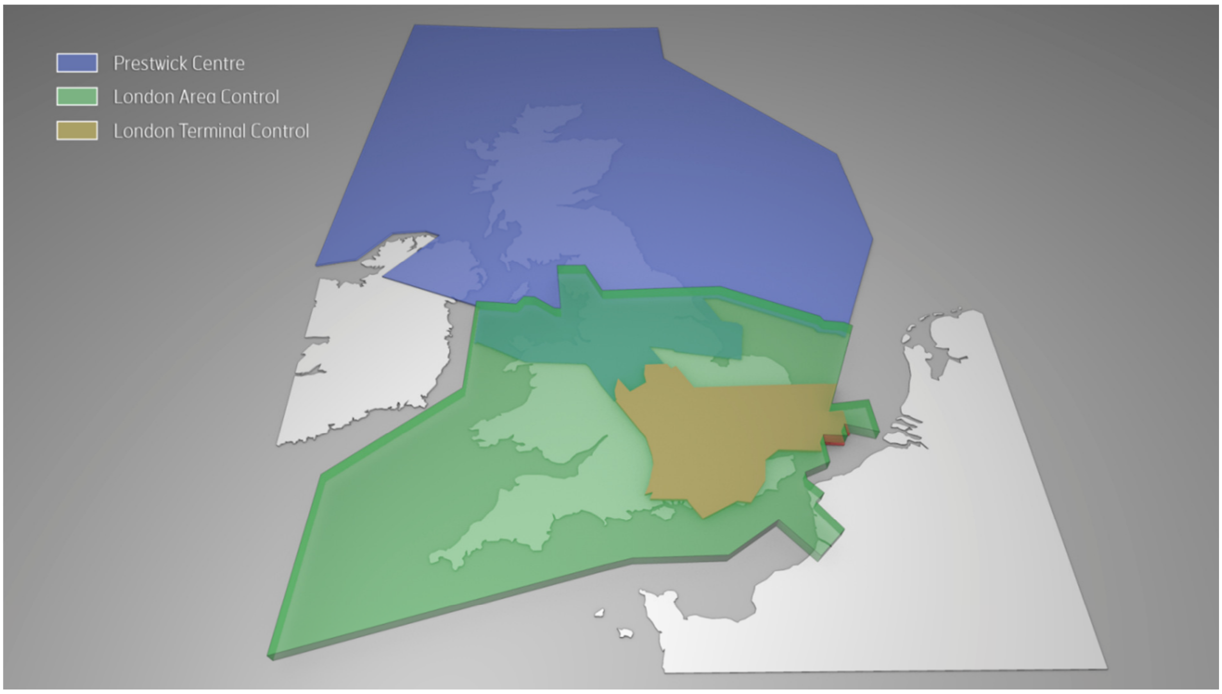

Figure D.1: UK Airspace by Control Centre

D.1.3 The Incident on 12 December abruptly affected ATC throughout London Area Control. Air traffic services for both LAC and LTC are operated by NATS and, together with military aircraft services for the UK, are provided from separate control rooms within the same building at Swanwick, some 5 miles South-East of Southampton Airport.

D.2 The LAC Operation

D.2.1 LAC is divided into a maximum of 32 sectors that can be combined ("band-boxed") at times of light traffic or separated or sub-divided ("split") when the traffic is heavier. The number of staff varies through the day, week and season but broadly depends on the number of aircraft expected to be flying in or through the London FIR. There are five "watches" of Controllers to manage the Operations Room on a continuous basis.

D.2.2 Each Controller can operate for up to 90 minutes without a break and controllers are rostered throughout the day to meet this requirement. When staff are not required at a workstation because of lighter traffic conditions, they are encouraged to leave the Operations Room, partly so as not to distract those engaged in operational duties.

D.2.3 Controllers normally work in pairs: a Tactical Controller who communicates with the aircraft under control and a Planning Controller who manages the flow of traffic into and out of their area of responsibility through liaison with adjoining NATS or other national ATC areas. An Air Traffic Services Assistant provides support to the controllers when required. The primary safety objective of these arrangements is to ensure a height separation of at least 1000' between aircraft or, where aircraft are within this limit, to maintain a lateral separation of at least 5 miles.

D.2.4 Each pair of Controllers is assigned to a particular sector or combination of (band-boxed) sectors. They are supervised in groups of 5-8 sectors by Local Area Supervisors. An Airspace Capacity Manager is focused on the overall flow of traffic in the LAC and supports the Local Area Supervisors in managing the band-boxing or splitting of sectors. The Operations Room as a whole comes under the charge of the Operations Supervisor. Both the Operations Supervisor and the Airspace Capacity Manager have designated Assistants.

D.2.5 NATS operates a network of radar stations that provide the position and height of all aircraft flying in the LAC. A data fusion system determines the best estimated position when an aircraft is detected by more than one radar so that the aircraft appears only once on the workstation screen; a label adjacent to the aircraft icon gives its height and can give the heading and other related information.

D.2.6 The Controller can call up all other necessary data associated with a particular aircraft, derived from its flight plan information. The flight information derives from a flight data processing system, also operated by NATS and known as NAS (or National Airspace System), and this is routed to a System Flight Server (SFS) that delivers the right information to each workstation.

D.2.7 In order for the workstations to ensure that the controllers have the right information available to them, the system must be aware of the role that each controller is fulfilling, and to support the transition of responsibilities as the process of "band-boxing" and "splitting" of sectors as described in paragraph D.2.1.

D.2.8 When a Controller signs on to a workstation in its initial powered state, it changes from "Base Mode" to "Prepare Mode" and recording to archive starts of all information available to the workstation; but the workstation cannot be used to control air traffic. The Controller then selects their designated sector thereby notifying the System Flight Server of the aircraft data required by the workstation; the workstation moves into "Elected Mode" and displays a copy of the data being used at that time to control the selected sector. If the Controller then selects "Open Sectors", a workstation goes into "Controlling Mode" and becomes fully operational while the workstation previously controlling that sector moves into the Elected State; this transfer of responsibility is managed by the Local Area Supervisor.

D.2.9 All of the responsibilities of the controllers and supervisors, including the processes and procedures they must adhere to in discharging their responsibilities are well defined and documented in the Manual of Air Traffic Services (MATS). MATS Part 1 is published by the Civil Aviation Authority (CAA) and applies for all air traffic services in the UK. Individual MATS Part 2 documents are published for each operational unit (LAC in this case) and define the specific responsibilities and procedures for that unit. These documents describe all aspects of the roles, including local activities, interactions with other units and procedures and procedures for abnormal situations including system failures.

D.3 Flow Management

D.3.1 In order to ensure safety of Air Traffic Management (ATM) operations the levels of traffic need to be managed to stay within the available capacity of the centres and sectors. In Europe this process of Air Traffic Flow and Capacity Management (ATFCM) is coordinated by the Network Manager, which is part of EUROCONTROL, the European Organisation for the Safety of Air Navigation.

D.3.2 Through this process each flight is allocated a Calculated Take-Off Time (CTOT) often known as a slot which defines a 15 minute window within which the aircraft can take off. The slots are calculated so as to ensure that no sectors along the flights route would exceed their available capacity and where necessary flights are delayed from the preferred take-off time to achieve this.

D.3.3 Available airspace capacity varies over time depending on a range of factors including weather, staff and system availability and inherent airspace design. ANSPs and Eurocontrol can apply "Regulations" which restrict capacity in an area of airspace when necessary to maintain safety, for example when a system failure occurs.

Annex E. NATS Organisation and Budgets

E.1.1 This Annex provides a high level view of NATS organisation and finances, including ownership structure and executive team..

E.1.2 NATS was partially privatised by the UK government in 2001, entering into a Public Private Partnership (PPP) with the Airline Group (AG) and retaining 49% ownership. The rest of the ownership now rests with AG, employees and with Heathrow Airport Holdings Ltd, as illustrated in Figure E.1 below.

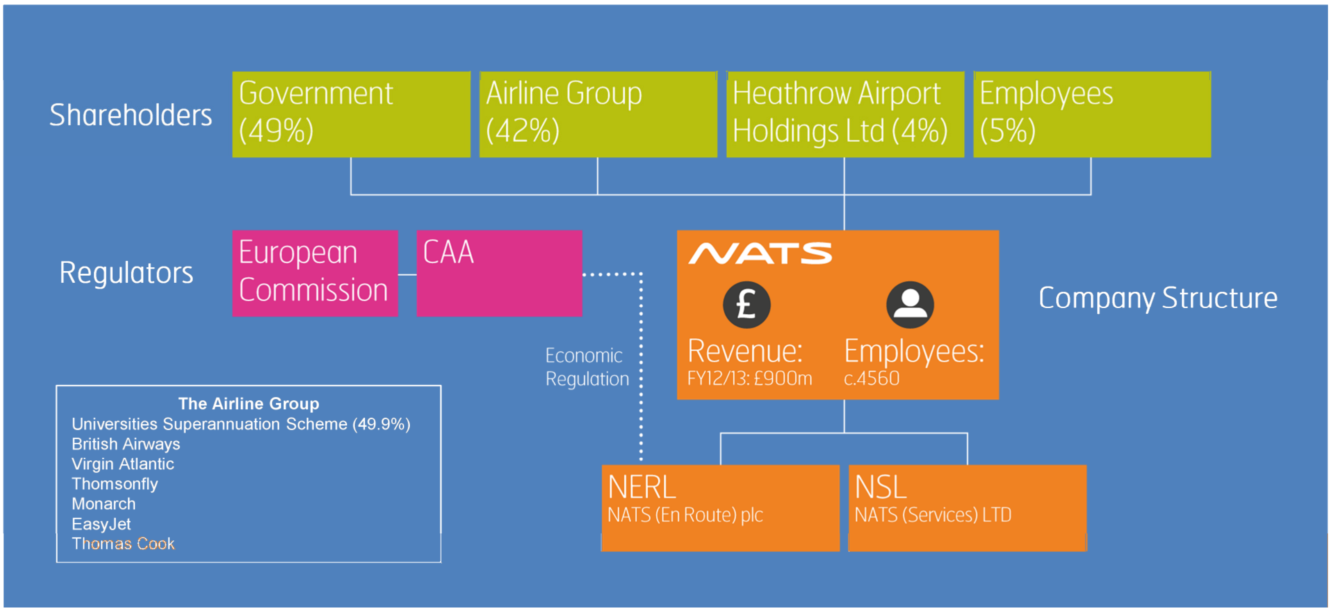

The Airline Group
Universities Superannuation Scheme (49.9%)
British Airways
Virgin Atlantic
Thomsonfly
Monarch
EasyJet
Thomas Cook

Figure E.1: NATS Ownership Structure

E.1.3 As illustrated in the Figure, NATS consists of three primary businesses, each of which is described below:

NATS Holdings Limited

- NATS Holdings Limited is an air traffic control services provider which owns two principal operating subsidiaries: NATS (En Route) plc and NATS (Services) Limited.

NATS (En Route) plc

- NERL (formerly NATS En Route Limited) is the sole provider of air traffic control services for aircraft flying 'en route' in UK airspace and the eastern part of the North Atlantic. It is economically regulated by the Civil Aviation Authority (CAA) within the regulatory framework of the European Commission's (EC) Single European Sky (SES) and operates under licence from the Secretary of State for Transport. It operates from two air traffic control centres, at Swanwick in Hampshire and Prestwick in Ayrshire. Note specifically that it is NERL that has responsibility for the London Area Control operation within which the system failure occurred on 12 December 2014.

NATS (Services) Limited (NSL)

- NATS Services provides air traffic services at 14 of the UK's airports and at Gibraltar: 8 of these contracts are subject to economic regulatory oversight by the CAA and the EC. It also provides engineering, consultancy (including training), defence and aviation information management services to UK and overseas customers. FerroNATS, a joint venture with Ferrovial, provides air traffic services at 9 airports in Spain.

E.1.2 Richard Deakin is the Chief Executive Officer of NATS and leads an Executive team who manage the business as a whole, supported by Operations and Services leadership teams who lead NERL and NSL respectively. The organisational structures and associated leadership teams for NATS and Operations are shown in Figure E.2 and Figure E.3 overleaf.

Chief Executive Officer
Richard Deakin

Manging Director Operations
Martin Rolfe

Managing Director Services
Catherine Mason

Finance Director
Nigel Fotherby

Safety Director
David Harrison

General Cousel and Company Secretary
Richard Churchill-Coleman

Human Resources Director
Gerry Skelton

Communications Director
Simon Warr

Figure E.2: NATS Executive Team

MD Operations
Martin Rolfe

Director Operations, Strategy
Simon Hocquard

Director Operations, Safety
Richard Schofield

Director Operations, Swanwick
Juliet Kennedy

Director Operations, Prestwick
Alastair Muir

Director Operations, Programmes
Garry Jackson

Director Operations, Engineering
David Hawken

Director International & Customer Affairs
Jonathan Astill

Director Supply Chain
Tim Bullock

Chief Information Officer
Gavin Walker

Figure E.3: Operations Leadership Team

E.1.3 The table below provides a high level summary of key financial data for NATS for financial year 2013-14 and with comparison to the previous year, while a similar table beneath provides the same information for NERL.

NATS (£m unless otherwise specified)	**2014**	2013(Restated)	Change	%
Revenue	**917.6**	899.6	+18.0	+2.0
Operating profit before exceptional items	**240.3**	215.1	+25.2	+11.7
Operating profit	**167.5**	215.1	-47.6	-22.1
Profit before tax	**157.5**	160.8	-3.3	-2.1
Capital expenditure	**129.7**	128.9	+0.8	+0.6
Net debt	**407.9**	457.1	-49.2	-10.8
Gearing[45]	**54.0%**	54.5%	-0.5%	-0.9
Dividends	**62.0**	40.0	+22.0	+55.0

Figure E.4: NATS High Level Financial Performance

NERL (£m unless otherwise specified)	**2014**	2013(Restated)	Change	%
Revenue	**742.5**	713.6	28.9	4.0
Operating profit before exceptional items	**217.5**	183.3	34.2	18.7
Operating profit	**144.8**	183.3	-38.5	-21.0
Profit before tax	**133.9**	129.3	4.6	3.6
Capital expenditure	**125.5**	126.8	-1.3	-1.0
Net debt	**570.5**	605.5	-35.0	-5.8
Gearing[46]	**54.0%**	54.5%	-0.5%	-0.9
Dividends	**57.0**	28.5	28.5	100.0

Figure E.5: NERL High Level Financial Performance

[45] Ratio of net debt to regulatory assets of the economically regulated business (NERL)

[46] Ratio of net debt to regulatory assets

Annex F. CAA Organisation and Budgets

F.1.1 This Annex provides a high level view of the CAA organisation and finances.

F.1.2 The CAA is the UK aviation regulator and exists to protect the interests of the consumer. This is central to all of its work; in safety, market regulation, consumer protection and in terms of the impact of aviation on the environment. The CAA is a public corporation and its core responsibilities are outlined in primary legislation (principally the Civil Aviation Act 1982, the Transport Act 2000 and the Civil Aviation Act 2012), European legislation, and in secondary legislation (notably the Air Navigation Order 2009).

F.1.3 Dame Deidre Hutton is the Chair of the CAA while Andrew Haines is Chief Executive and leads the organisation on a day-to-day basis supported by the executive team. The organisation structure of the CAA is provided in Figure F.1 below.

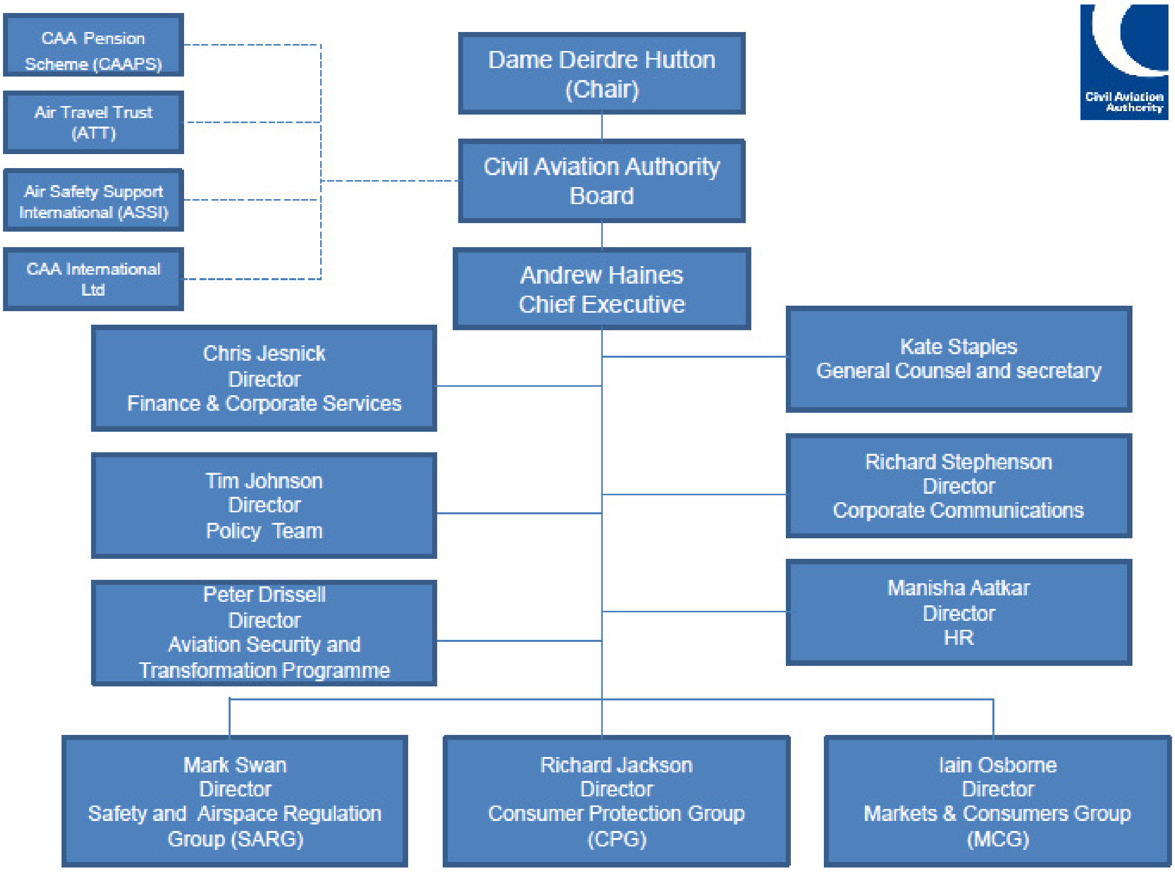
CAA Pension Scheme (CAAPS)
Air Travel Trust (ATT)
Air Safety Support International (ASSI)
CAA International Ltd
Dame Deirdre Hutton (Chair)
Civil Aviation Authority Board
Andrew Haines Chief Executive
Civil Aviation Authority
Chris Jesnick Director Finance & Corporate Services
Tim Johnson Director Policy Team
Peter Drissell Director Aviation Security and Transformation Programme
Kate Staples General Counsel and secretary
Richard Stephenson Director Corporate Communications
Manisha Aatkar Director HR
Mark Swan Director Safety and Airspace Regulation Group (SARG)
Richard Jackson Director Consumer Protection Group (CPG)
Iain Osborne Director Markets & Consumers Group (MCG)

Figure F.1: CAA Organisation

F.1.4 For the year ended 31 March 2014 our total Group income was £133.1m of which £75.1m relates to the statutory income we generate within the regulatory sector of our business and a summary of CAA finances for the year ended March 2014 is provided in Figure F.2.

F.1.5 The CAA is funded by those it regulates in accordance with the Statutory Charges Schemes published annually. The aviation industry and consumers expect the CAA to use the statutory income it receives efficiently and effectively. Our challenge is to ensure that the CAA is efficient without jeopardising the role it undertakes.

F.1.6 CAA International Limited, a subsidiary company of the CAA, provides consultancy services to a number of national and international aviation authorities to promote improved aviation safety standards worldwide. The company earned £19.8m of income for the year ended 31 March 2014.

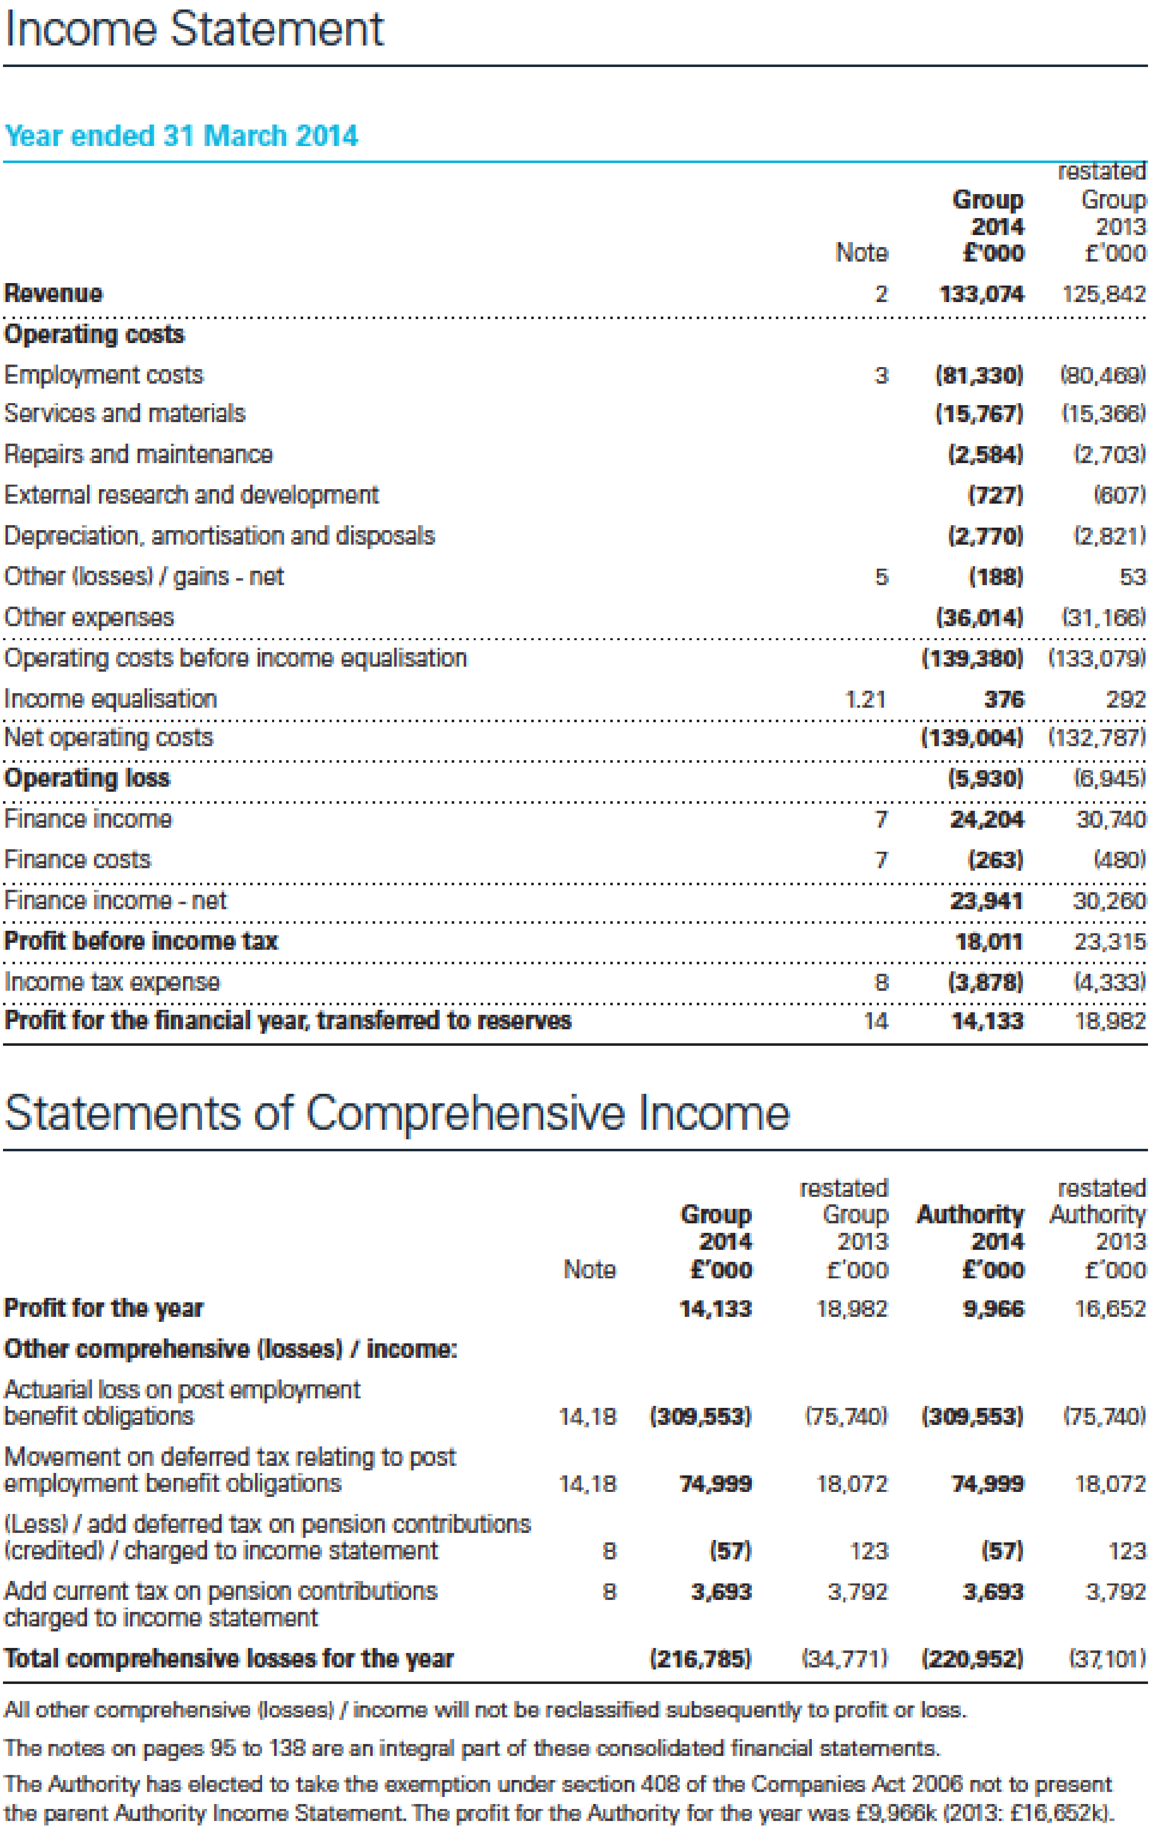

Income Statement

Year ended 31 March 2014

	Note	**Group 2014 £'000**	restated Group 2013 £'000
Revenue	2	**133,074**	125,842
Operating costs			
Employment costs	3	**(81,330)**	(80,469)
Services and materials		**(15,767)**	(15,366)
Repairs and maintenance		**(2,584)**	(2,703)
External research and development		**(727)**	(607)
Depreciation, amortisation and disposals		**(2,770)**	(2,821)
Other (losses) / gains - net	5	**(188)**	53
Other expenses		**(36,014)**	(31,166)
Operating costs before income equalisation		**(139,380)**	(133,079)
Income equalisation	1.21	**376**	292
Net operating costs		**(139,004)**	(132,787)
Operating loss		**(5,930)**	(6,945)
Finance income	7	**24,204**	30,740
Finance costs	7	**(263)**	(480)
Finance income - net		**23,941**	30,260
Profit before income tax		**18,011**	23,315
Income tax expense	8	**(3,878)**	(4,333)
Profit for the financial year, transferred to reserves	14	**14,133**	18,982

Statements of Comprehensive Income

	Note	**Group 2014 £'000**	restated Group 2013 £'000	**Authority 2014 £'000**	restated Authority 2013 £'000
Profit for the year		**14,133**	18,982	**9,966**	16,652
Other comprehensive (losses) / income:					
Actuarial loss on post employment benefit obligations	14,18	**(309,553)**	(75,740)	**(309,553)**	(75,740)
Movement on deferred tax relating to post employment benefit obligations	14,18	**74,999**	18,072	**74,999**	18,072
(Less) / add deferred tax on pension contributions (credited) / charged to income statement	8	**(57)**	123	**(57)**	123
Add current tax on pension contributions charged to income statement	8	**3,693**	3,792	**3,693**	3,792
Total comprehensive losses for the year		**(216,785)**	(34,771)	**(220,952)**	(37,101)

All other comprehensive (losses) / income will not be reclassified subsequently to profit or loss.

The notes on pages 95 to 138 are an integral part of these consolidated financial statements.

The Authority has elected to take the exemption under section 408 of the Companies Act 2006 not to present the parent Authority Income Statement. The profit for the Authority for the year was £9,966k (2013: £16,652k).

Figure F.2: CAA Income Statement for Year Ending March 2014

Annex G. Software and System Engineering Detail

G.1 Introduction

G.1.1 This appendix contains four components, complementing or amplifying elements of the main body. The four elements are:

- Categorisation of the software criticality that determines how rigorous the development and verification process will be;
- A source-code level description of the nature of the fault;
- A more detailed assessment of the testing and inspection of the key software modules;
- Guidance material to help interpret recommendations R3 and R4.

G.2 Software Categorisation

G.2.1 Section 2.5 briefly assesses the software development for the SFS code, on the assumption that the software category was correct. This is determined in the Hazard Analysis Report[47] (HAR), which assesses the severity of hazards and allocates categories to software, based on the potential contribution of the software to the hazard. The analysis here (Annex G) focuses on the SFS, and is not a full review of the hazard analysis process or report.

G.2.2 The HAR defines hazards in a non-standard way as "loss, corruption, delay, or misdirection of data necessary for a function at a sector suite or at a System Control workstation", rather than the more conventional definitions based on physical harm. However this is both appropriate for NERC, as it is an information processing system, and conservative, in that it is likely to over-state the risks associated with the system not underestimate them.

G.2.3 The HAR defines a systematic process starting with hazard identification, and using classical techniques such as fault trees to assess contributory causes to hazards. It focuses on "loss of" and "plausible corruption of" information in identifying hazards; in practice plausible corruption is most likely to be hazardous for SFS, as Controllers would be expected to quickly identify and respond to loss of data. The process provides a systematic way of identifying software contribution to hazards and its application covers software tools, as well as operational software (although this Annex focuses on SFS, not tools).

G.2.4 Hazards are categorised by severity; an extract from the table that defines the hazard classes is set out in Figure G.1 below.

Class	Definition	Rationale
Class 1	Inability to provide any degree of Air Traffic Control in one or more airspace sectors for a significant period of time.	Covers the situation where controllers have no possible means of controlling aircraft and separation will probably be eroded.
Class 2	Ability to maintain Air Traffic Control is severely compromised within one or more airspace sectors for a significant period of time.	Covers the situation where planned separation may not be maintained. Contingency Separation Measures can be applied, but the risk of infringing safe separation is extremely high until traffic has been curtailed to lower levels.
Class 3	Ability to maintain Air Traffic Control is impaired within one or more airspace sectors for a significant period of time.	Covers the situation where the ability to maintain planned separation is impaired and increased separation may be necessary. ATC procedures are able to compensate for the loss of function, but controller workload is likely to be high or the overall system capacity is affected.

Figure G.1: Extract from Hazard Severity Class Definition

[47] Hazard Analysis Report, NATS, NS--/--0518/SAF12, Issue 26, October 2014

G.2.5 As well as giving a clear description of hazard classes, the rationale makes clear why the air traffic regulation put in place on 12 December 2014 was appropriate.

G.2.6 Having assigned a severity class to each hazard, each Configuration Item is assigned a category according to the severity of the hazards resulting from its functional failure modes. The software category matches that of the hazard – for example if software can cause a Class 2 hazard, then it is assigned to Category 2.

G.2.7 The HAR says, referring to SFS, that: "Errors in this function could cause corruption of the track pairings, Actual Flight Level (AFL), or callsign (Flight Plan or Tactical Data Line (TDL)) which is severity Class 2." It also discusses loss of SFS capabilities including common mode failures and inability to revert to the back up system (although the HAR uses somewhat different terminology). In short, the analysis is thorough and does consider dual SFS failure.

G.2.8 Considering SFS, corruption is not a Class 1 event; the Controllers still have the primary radar and are able to communicate with all aircraft. Separation could be compromised if, for example, the flight data was corrupted to show the wrong flight levels (i.e. corruption of AFL). Thus SFS is categorised as Category 2, because such misleading data could contribute to a loss of separation.

G.2.9 Similarly, the loss of SFS is not a Class 1 event. The loss of SFS compromises (or impairs) Air Traffic Control, as the Controllers do not have the additional information they are used to receiving from SFS (and the iFACTS information on predictions of violations of minimum separation is also lost as it depends on SFS, even though it does not go through SFS). This issue is categorised as Class 3, because the problem is increased workload, rather than misleading data that could contribute to a hazard.

G.2.10 The software Category should reflect the worst-case event (hazard) associated with an item. Thus, SFS is Category 2 overall, as plausible corruption is Class 2 and loss is Class 3.

G.2.11 The events of 12 December 2014 were no worse than Class 3. Thus the SFS software has not been under-categorised, in regard to that event. Thus it is reasonable to assess the SFS software against the requirements for Category 2 software in POD SW01 as has been done in section 2.5, and is done in more detail in section G.3 below.

G.3 The Fault

G.3.1 This section amplifies the discussion in section 2.3, and provides more detail on the causes of the SFS failure on 12 December 2014. It also amplifies on some of the issues addressed in section 2.7, particularly on some of the difficulties in managing the exception in a more sophisticated manner. The issue of testing for the fault and the effectiveness of review and testing are discussed in section G.4.

G.3.2 The relevant code is written in Ada. The following description includes extracts from the relevant Ada source code modules; brief explanations of the program are given, but it is assumed that the reader is familiar enough with programming concepts and the language itself not to need a tutorial description of the language or the program fragments.

G.3.3 As indicated in chapter 2, the proximate cause (in the software) of the Incident on 12 December 2014 was a discrepancy between the size of a table constructed by the SFS software and the check on the size. The table was constructed in a utility module known as `Waafu28`.

G.3.4 The specification (Ada header) for `Waafu28` includes the statement:

```
 --> Return all workstations controlling ANY atomic function
```

Figure G.2: Waafu28 Header (Extract)

G.3.5 It is clear that this includes all Atomic Functions. The code (Ada body) builds a table, known as `To_List` that is intended to cover civil and military Atomic Functions. It does so taking records from a data item known as `From_Table` (it is a parameter, see figure G.5 below). The scope of `Waafu28` is illustrated in the following code fragments:

```
Workstation_List.Add(Element => From_Table.Atomic_Function_Record.
Civil_Airspace_Controlling_Sectors
…
To_List => Workstations);

…

Workstation_List.Add(Element => From_Table.Atomic_Function_Record.
Military_Airspace_Controlling_Sectors
…
To_List => Workstations);
```

Figure G.3: Extracts from `Waafu28` Body

G.3.6 Clearly this does cover military and civil sectors; although not shown, it covers the Supervisor roles as well, and hence all Atomic Functions. However the workstations list is of the wrong type, it is declared as:

```
Workstations: Workstation_List_T
```

Figure G.4: Definition of Type of Workstation

G.3.7 This type (which defines the structure of the data) can only contain civil functions (covering Controller and Supervisor), but that is far from obvious due to the naming of the type. The type is defined in another module, `Waafu00_Atomic_Function_Table_Utilities`, so the error is not immediately apparent on reading `Waafu28`. Section G.4 discusses opportunities for finding the fault, amplifying on the discussion in section 2.6.

G.3.8 The module `Waafu28` contains a single, function:

```
function Waafu28_Controlling_Workstations
(From_Table : in Wsaft.Atomic_Function_Table_T)
 return Workstation_List_T is …
```

Figure G.5: Function in `Waafu28`

G.3.9 The parameter to `Waafu28` is `From_Table` of type `Atomic_Function_Table_T` (defined in `Wsaft`; the exact definition of the type is not material). The way the function works, it goes through the input parameter (`From_Table`) and returns the result when it reaches the end of the input – the boundary condition.

G.3.10 As stated above, the type `Workstation_List_T` is defined in module `Waafu00_Atomic_Function_Table_Utilities`. The critical part is a definition that contains the size[48] that is then checked, viz:

```
new Wssil00_Simple_List (Element_T => Wsprc.Processor_Id_T,
Max_Size => Wsaft.Max_Civil_Atomic_Functions,…
```

Figure G.6: Type Definition in Waafu00_Atomic_Function_Table_Utilities

G.3.11 The names are not identical to those used in `Waafu28` but this is not an issue, due to the way the modules relate to one another (the calling structure). This is the point in the code at which the exception was raised around 1444 on 12 December 2014.

G.3.12 Ada allows exceptions to be handled in different parts of the code. Ada includes both generic exception handlers ("catch all"), and specifically programmed exception handlers. One other

[48] Note: here "=>" assigns a value to the element to the left, i.e. Max_Size and Element_T; it is not "greater than or equal to".

way of dealing with the “Watching” problem might have been to “catch” the exception more locally to where the exception is raised. To explore this possibility involves understanding the “call hierarchy”, i.e. the set of functions and other parts of the code that interact to implement the management of the “Watching” request. The relevant call hierarchy is shown in Figure G.7, which also shows where the faults that caused the failure on 12 December are located.

```
task body Fsea099_Sfs_Principle_Func_Task_T – Where Overflow exception is
  handled by generic exception handler leading to SFS failure
  procedure Fsea098_Process_Event
    procedure Fsafo010_Process_Incoming_Message
      procedure Fsafo050_Process_Aft_Update
        procedure Fsafo053_Process_This_Request
          procedure Fsafo062_Watch
            procedure Fsafo122_Check_Subset_For_Watching
              function Fsgst030_Is_A_Subset_Of_Existing
                function Waafu28_Controlling_Workstations – Where
                  Workstations array is defined as Workstation_List_T
                  (in package body) but then adds items to for
                  Civil, Military and Supervisor atomic functions.
                  The type Workstation_List_T is defined in
                  Waafu00_Atomic_Function_Table_Utilities
                    Function Wssil00_Simple_List.Add generic package –
                      where the exception is raised when the array
                      Overflows
```

Figure G.7: SFS Call Hierarchy

G.3.13 The top element of the hierarchy is a “task”; tasks are the main units of running software. The SFS contains several tasks that operate independently. The task is also where the generic exception handler (the “catch all”) resides. The task calls `Fsea098_Process_Event` on receiving a command from the workstation, and calls propagate down the hierarchy; exceptions propagate back up.

G.3.14 It should be noted that those functions starting with `W` are generic, and those starting with `Fs` are specific to the SFS. In order to make “intelligent” decisions in processing exceptions it is important to handle them close to where they are raised. However both of the lowest items `Waafu28_Controlling_Workstations` and `Wssil00_Simple_List.Add` are generic, so they cannot carry out sophisticated exception handling; the programmers who wrote this code would not have known the context of use.

G.3.15 The most likely place to “catch” the problem and handle it in a more subtle way than closing down the SFS is in the procedure `Fsafo122_Check_Subset_For_Watching`. At this point in the call hierarchy it is known that a “Watching” command is being processed, and that the exception raised by `Wssil00_Simple_List.Add` could be resolved by discarding the command. As commands should not be rejected “silently” the correct behaviour would have been to send back an error message, e.g. “ Watching request rejected, too many Atomic Functions”. The second part of such a message would not be very clear to a Controller, although the first part would be. Such an exception would be logged and brought to the attention of the engineers for exploration. If this had been done, then the double SFS failure could have been avoided.

G.3.16 In the design that was current on 12 December 2014, the exception “reached” the task `Fsea099_Sfs_Principle_Func_Task_T`. The task is not “aware” of the command that is being processed, so took the safe course of shutting down the SFS. Arguably the task could have made a more subtle discrimination based on the command, but it was intended mainly to deal with hardware problems, so it is unlikely that the designers would have considered this possibility, when the software was being designed.

G.3.17 In summary the proximate cause of the failure on 12 December 2014 was a discrepancy between the size of a data structure and the check on the maximum permitted size of that data structure. This arose because of the use of an incorrect type in the utility code, by the main SFS code, and the failure to find the fault prior to the change in circumstances (the inclusion of military Controller roles) and the unintentional invocation of "Watching" by one of the Controllers that triggered the fault which resulted in the SFS failure.

G.3.18 The specifications at code level are simple; see for example Figure G.2 above. There are some higher-level requirements to which the software has to conform. The requirements are structured in a hierarchy, with Level A at the top, supported by Level B. The requirement most relevant to the Incident on 12 December 2014, is at the B Level:

> B059480 SFS shall not permit the watching of an atomic function unless it is a subset of a set of atomic functions currently being controlled at a single workstation.

Figure G.8: B Level Requirement Relating to Watching

G.3.19 Implicitly, this means that the software has to construct the table of Atomic Functions to check the command against the current set of Atomic Functions, and the criterion for allowing "Watching" to go ahead is that the request relates to existing Atomic Functions (the intent is not fully captured in the text so, arguably, the quality of requirements is a contributory cause to the Incident). However there are no "clues" in the requirement about the maximum number of Atomic Functions to guide the developer or the tester – both would have had to use other sources of information. The software had its origins in an earlier development in the USA that did not support military Controllers, and this might help to explain the original program design, although it is unlikely that the underlying cause for the software fault can be found at this time.

G.3.20 It is not practicable to investigate further this potential underlying cause of the fault given the time elapsed since the software was developed. The form of requirements was quite typical of systems developed at that time and the fact that some creativity was required in developing the software is unsurprising. However, specification techniques have improved (although there are still questions about scalability) and there is, therefore, an opportunity to update the specification approaches used by NATS as they move towards the deployment of SESAR (see recommendation R3).

G.4 Testing and Inspection

G.4.1 Section 2.5 discussed the effectiveness of testing and inspection; this section adds some detail, including test coverage criteria, and discusses the available inspection records and test results from the initial development of the SFS and utility code.

G.4.2 One of the key opportunities for finding software faults is code inspection. POD SW01 includes mini-inspections and "formal" inspections. The latter are more likely to find faults, all other things being equal, but at a greater cost than mini-inspections. The procedures in POD SW01 allow an informed trade-off of cost against risk to be made, and code might only be subject to mini-inspections. The code of interest in the case of the SFS failure was subject to formal inspection, although the rationale for this choice has not been identified.

G.4.3 The inspection records (from 19 October 1994) show that `Waafu28` used to be part of a much bigger module known as `Waafu00` and that the bigger module was split into smaller modules that were easier to understand (`Waafu28` is still part of the `Waafu00` group of modules). The records also show that problems were identified with `Waafu00` and with `Waafu28` itself, and that these were corrected. The changes included correcting return parameters from the module.

G.4.4 Clearly, however, the inspections did not find the fault. Whilst it is not possible to be certain, the naming chosen seems to be a likely contributory cause (it is not obvious immediately that

Workstation_List_T is limited to civil functions), and this is only likely to be resolved by doing “cross-module” reviews (see recommendation R3 and the amplification at G.5 below).

G.4.5 Note that utilities, such as Wssil00_Simple_List are small parts of the program that will be used many times, elsewhere in the program. Thus, when reviewing and testing such parts of the SFS, it would not have been done in the context of the call from Waafu28, or any other module that uses the code. Thus neither testing nor inspection of the utility functions is likely to find the problem (again, see recommendation R3 and the amplification at G.5 below).

G.4.6 For category 2 code, POD SW01 requires condition coverage and boundary value analysis in testing. In other words the maximum size of data structures should be tested, and each condition should be made to take the value true or false – for example in an if … then … else … end if code structure both the then and else branches would have to be taken. Boundary values are considered first.

G.4.7 There is nothing in the code (or specification) of Waafu28 that indicates how big the parameter From_Table should be when testing this module. As it is necessary to write code to generate (legal or “valid”) input parameters, a normal approach to testing would be to keep the input parameter small, but to make sure that the return value is correct – thus meaning that the boundary (reaching the end of the input table) had been handled correctly. Thus it is possible (indeed likely) for the tests to be passed, the boundary value criterion to be met, but the specific fault of interest in this enquiry not to be discovered.

G.4.8 There is no reason to believe that testing Waafu28 with the full size of 193 (or any value over 151) would not have revealed the problem and raised the exception. Sometimes modules are “stubbed out” in testing, but this does not seem plausible for Wssil00_Simple_List as Waafu28 depends on types defined in that module. However it is not clear whether or not the current limits (151 and 193) were known at the time Waafu28 was tested (14 November 1994, after the code inspections) nor whether or not “test limits” were used, as discussed above.

G.4.9 The module Wssil00_Simple_List has a clear boundary condition, specifically: Max_Size => Wsaft.Max_Civil_Atomic_Function. Boundary value testing would yield the exception (the one that occurred on 12 December 2014) on exceeding Max_Size but this would be seen as correct, not a problem. As the data structure is in a utility function, the testers would have no knowledge of the context of use to determine that the Max_Size was inappropriate (clearly there would be other contexts when it would be appropriate e.g. checking only civil Atomic Functions).

G.4.10 The discussion of coverage focuses on Waafu28 as Wssil00_Simple_List is used just for the data type so boundary value testing assesses the relevant property.

G.4.11 The test results show six tests having been passed, with no failures. The tests group together sets of tests to achieve some overall testing goal. For example the test for decision coverage ensures that each condition in the program takes at least one true and one false value during the test. The test summary incudes results of checks, such as shown in Figure G.9:

```
CHECK_ANALYSIS ( WAAFU00_ATOMIC_FUNCTION_TABLE_UTILITIES.
                   WAAFU28_CONTROLLING_WORKSTATIONS,
                   DECISION_COVERAGE,
                   Lower Limit   100.00,
                   Upper Limit   100.01 );  PASSED
                   Value      100.00 %
```

Figure G.9: Coverage Test Check

G.4.12 Thus there is evidence the test coverage criterion was fully satisfied. However, as with boundary value analysis, there is no reason why the data structure should have been anything like the 151 limit. Indeed a `From_Table` of seven entries (for the three civil, three military and one supervisor roles) would have been sufficient to conduct this test. Thus, even though the test criterion is met, it is very unlikely that this form of coverage testing would have found the problem.

G.4.13 In summary, at the level of individual modules, there was an opportunity to find the problem when reviewing and testing `Waafu28`. It is more likely that it would have been found in review/inspection than by test, as the programmers will have understood the code, and should have known how `Workstation_List_T` was defined.

G.4.14 Given the structure of the software, it is not realistic to think that the problem would have been found by considering `Wssil00_Simple_List`, as it is in a utility module.

G.4.15 It is worthwhile highlighting the difference between fault finding in development, and in operation. The fault was found quickly after the events of 12 December, which perhaps suggests that it should have been found in initial review or testing. However, the situation is quite different and it is perhaps helpful to use a "needle in a haystack" analogy. In initial testing, the NATS and LM staff did not know that this particular fault (needle) existed, nor where to look for faults (needles) of any type in hundreds of modules and over two million SLoC. In contrast, on 12 December 2014, the LM and NATS staff knew there was a needle, and the sort of needle (the logs said that there was, and identified the type of exception). The logs also narrowed down the search to a few straws (counting modules as straws) or a few hundred straws, if we consider each SLoC as a straw. Thus the speed in finding the fault in operation does not mean that it was an "obvious" fault, rather that the failure and the logging gave NATS and LM staff some good "clues". Therefore it cannot reasonably be concluded that there was a shortcoming in the conduct of this aspect of the development process.

G.5 Additional Guidance on Recommendations

G.5.1 Some of the recommendations in section 2.9 are fairly self-explanatory; however others are more open to interpretation. The aim here is to amplify two recommendations that are quite broad in scope in order to help NATS implement them. The primary focus here is on the third recommendation (R3), which gives the broad software engineering guidance; it starts by amplifying the recommendation, then considering the points in turn.

G.5.2 NATS should ensure the effectiveness of the software development process, making both pragmatic changes, and reflecting new technology; pragmatic changes include:

- Improve naming conventions (for data, types and functions) to make their intended use more obvious;
- Extend the guidelines for inspections to look more carefully at type definitions in external modules and calls on utility functions to check their correct usage, in each context of use;
- Introduce periodic audits of verification and validation evidence, repeating tasks to regenerate evidence, where it has been lost.

Updates to reflect changes in software technology, since the NERC was developed, include:

- Model-based development, including automatic code generation;
- Static analysis tools, for Ada and for other languages;
- Testing tools, including tools that automate test coverage analysis.

In revising processes, thought needs to be given on how to manage suppliers, e.g. by ensuring that they understand and respect NATS requirements for processes and standards, and on integrating and testing software that is likely to be developed in a range of different software (programming) languages, in different companies and in different countries.

G.5.3 The guidance covers each of the bullet points in turn. It should be noted that the guidance is only intended to help guide interpretation, and deliberately does not mention particular methods or tools.

G.5.4 Naming conventions – names of data are very important in program comprehension. The aim should be to ensure that the name makes the use of the data clear. Where two or more data structures have very similar meanings, care should be taken to ensure that the names distinguish the uses, ideally with the discriminant near the beginning of the name, not at the end, especially if the names are long. Similar guidance applies to the functional aspects of the code, for example, what does `Waafu28` signify? The same applies to specifications.

G.5.5 Review and inspection guidelines – reviews and inspections are focused on individual modules, and need to be bounded. However, mistakes can be made by inappropriate use of data structures or functions defined in other modules; this risk can be addressed by requiring that the source definitions of the types used be considered explicitly when reviewing the using module. Similarly, use of utility functions should be reviewed to ensure they are appropriate in the context of use. Guidelines could be developed to support these inter-module checks.

G.5.6 Model-based development – the use of computer-based modelling tools is now becoming mainstream, with some significant uses in critical applications. Generally the approaches use graphical models showing the structure and (to an extent) function of programs. The models are often analysable to show, say, completeness, e.g. that outputs are defined for every combination of inputs. This helps ensure validity of the requirements (an SW01 goal). Also, many tools are capable of automatic code generation, although errors in tools can introduce faults in the executable program. It is thus possible to remove cost and error from the development process, albeit with a "burden" to demonstrate the soundness of the code generators for more critical applications.

G.5.7 Static analysis – the available analysis tools have improved significantly in recent years, and can deal with relatively large programs, in a range of languages. As well as finding errors, e.g. undefined variables, they can demonstrate important properties of programs, e.g. exception freeness. (This might have found the proximate cause of the Incident on 12 December, although it depends how well the tools deal with inter-module dependencies.) These tools are becoming more cost-effective, and are a useful complement to (and perhaps substitute for part of) manual code reviews and inspections.

G.5.8 Testing tools – NATS use what are often viewed as industry standard tools, but there are now alternatives that may be more cost-effective, and that may help automate some of the tasks that are particularly effort intensive. There is also some evidence that the use of coverage criteria is not very effective as a way of finding faults (despite their prevalence in standards) and that testing based on "fault hypotheses" may be more useful.

G.5.9 The final point in G.5.2 is particularly important. NATS are largely dependent on suppliers for their software, and there is a need to manage software quality in the supply chain. There is little point in NATS defining processes that suppliers do not or cannot use. Thus NATS needs to consider: what it believes are good processes; how it can mandate these on the supply chain; what it can assess retrospectively (e.g. with static analysis); what it can encourage suppliers to adopt and how it ensures in both quality of execution and in contractual terms that NATS' conditions are respected. As for encouragement, the strongest argument is likely to be economic; if it can be shown that faults are detected and removed more cost-effectively with new methods and tools, then it is much more likely that companies will accept the technology. The contractual terms and conditions can also play a role. In the SESAR environment where NATS are only one of several customers, economic and contractual incentives are important tools. Also, the engagement of other ANSPs and regulators should help to encourage the supply chain to adopt modern practices.

G.5.10 Recommendation R4 refers to a complete, continuing evidence base (CCEB), and covers audits of the CCEB. The CCEB is all the information necessary to evolve the software in the light of change requests, and to carry out forensic analysis in the events of problems. Thus it covers specifications, architectural descriptions, verification results, etc.

G.5.11 Although the current NATS configuration control rules mean that evidence, e.g. test results, relating to each build should be retained, it is possible in long-lived, geographically distributed developments for material to be mislaid, so there is not a CCEB. (In the current case, the code in question has been the responsibility of three different companies in its life.) Whilst good processes will reduce the likelihood of problems, they will not eliminate them. Periodic configuration audits give a means of identifying problems, thus enabling them to be rectified. There is a cost benefit trade-off: the more frequent the audits, the sooner problems will be found, and can be corrected, but the greater the cost. Recommendation R4 links the audits to NATS' planning cycle, but there may be merit in carrying out more frequent audits. Note that the audit process should cover the complete CCEB, not just verification evidence.

Annex H to Independent Enquiry Report dated 13 May 2015

Annex H. List of Meetings

H.1.1 Table of meeting held by the Enquiry Panel and with relevant stakeholders.

Date	Location	Meeting
Fri 9 Jan	CAA House	Enquiry pre-meeting
Tue 13 Jan	CAA House	Panel kick off meeting
Fri 16 Jan	CAA House	CAA Economic regulation initial briefing
Mon 19 Jan	CAA House	CAA Safety Regulation initial briefing
Tue 20 Jan	CAA House	Panel Root causes meeting
Fri 23 Jan	Swanwick	Panel visit to Swanwick Centre
Mon 26 Jan	Brettenham House	NATS Regulation initial briefing
Mon 26 Jan	CAA House	Panel meeting
Mon 2 Feb	CTC	NATS Engineering
Tue 3 Feb	CAA House	Panel Regulatory discussion
Tue 10 Feb	Teleconference	Panel teleconference
Thu 12 Feb	Waterside	British Airways
Thu 12 Feb	Heathrow Tower	Heathrow Airport and NATS Tower Heathrow
Thu 19 Feb	Gatwick	Gatwick Airport and NATS Tower Gatwick
Tue 24 Feb	CAA House	Flybe
Wed 4 Mar	CTC	NATS Engineering
Fri 6 Mar	Luton Airport	Luton Airport and NATS Tower at Luton
Wed 11 Mar	CAA House	Regulatory discussion with CAA
Wed 11 Mar	CAA House	IATA, Ryanair and EasyJet
Thu 12 Mar	CTC	NATS Engineering
Fri 13 Mar	CTC	Panel visit to NATS Corporate & Technical Centre
Tue 17 Mar	CAA House	Director Aviation; Department for Transport
Tue 17 Mar	CAA House	Panel Meeting
Tue 17 Mar	Brettenham House	Regulatory discussion with NATS
Mon 23 Mar	Brettenham House	Panel Chair meets with NATS Chair
Fri 26 Mar	Swanwick	NATS Operations
Tue 31 Mar	CAA House	Panel Meeting
Mon 13 Apr	CTC	NATS Engineering
Tue 14 Apr	CAA House	Panel Meeting
Mon 20 April	Brettenham House	Panel Meeting on Regulatory aspects
Tue 12 May	CAA House	Panel final review Meeting

Annex I. Glossary

I.1.1 Glossary of key acronyms and technical terms used within the report.

AAS	Advanced Automation System	US FAA led ATM automation project that was the forerunner to the UK's NERC project and the foundation of the ATM system for LAC.
A-CDM	Airport Collaborative Decision Making	Concept and supporting systems which aims at improving operational efficiency at airports by sharing key information between partners including airport operators aircraft, operators/ground handlers, ATC and the Network Operations
ACM	Airspace Capacity Manager	The Airspace Capacity Manager acts as the focus for Air Traffic Flow and Capacity Management (ATFCM) within NATS, fulfilling the requirements and responsibilities of the Flow Management Position (FMP) within both the London and Scottish Flight Information Regions (FIRs).
AFL	Actual Flight Level	Flight Levels are a measure of altitude expressed in hundreds of feet based on a standard sea-level pressure. Actual Flight :Level is the current altitude of an aircraft expressed as a Flight level based on this standard pressure setting.
AIM	ATFM Information Message	A message transmitted by the Network Manager Operations Centre (NMOC) to provide information, advice and to promulgate instructions relating to the application of current Air Traffic Flow & Capacity Management (ATFCM) measures. It is also used for the initial publication of the Network Operations ATFCM operating procedures that affect all users
AIRAC	Aeronautical Information Regulation And Control	Process whereby changes to aeronautical information (e.g. airspace routes and procedures) are coordinated internationally and published according to a regular 28 day cycle.
ANSP	Air Navigation Service Provider	Standard term for the organisations that provide ATC services within a state or region.
AOD	Analysis, Options and Design	NATS process for carrying out structured analysis of potential changes to the NERC system in order to assess their complexity, consider options to deliver a solution and ultimately to create a formal baseline for the preferred solution suitable for incorporating in a build.
ASBU	Aviation System Block Upgrade	The Aviation System Block Upgrades (ASBUs) are the International Civil Aviation Organization (ICAO) defined framework for harmonising avionics capabilities and the required air traffic management (ATM) ground infrastructure as well as automation. An ASBU is a package of capabilities (modules) which has essential qualities of: • Clearly defined measurable operational improvements with appropriate metrics to determine success • Necessary equipment and/or systems in aircraft and on the ground along with an operational approved or certification plan • Standards and procedures for airborne and ground systems • Positive business case over a clearly defined period of time
ATC	Air Traffic Control	A service operated by appropriate authority to promote the safe, orderly and expeditious flow of air traffic.
ATFCM	Air Traffic Flow & Capacity Management	ATFM extended to include the optimisation of traffic patterns and capacity management. Through managing the balance of Capacity and Demand the aim of ATFCM is to enable flight punctuality and efficiency, according to the available resources with the emphasis on optimising the network capacity through the collaborative decision making process.
ATFM	Air Traffic Flow Management	A service established with the objective of contributing to a safe, orderly and expeditious flow of air traffic by ensuring that air traffic control capacity is utilised to the maximum extent possible, and that the traffic volume is compatible with the capacities declared by the appropriate air traffic services authority
ATICCC	Air Traffic Incident Co-ordination and Communication Cell	The NATS Air Traffic Incident Co-ordination and Communication Cell (ATICCC) has been established to provide a management focus to co-ordinate the post incident business recovery process following the loss of any substantive parts of the aviation industry support infrastructure, particularly Air Traffic Control. A team of experienced operational managers comprising representatives from NATS, Airports, Airlines and other appropriate organisations will man the ATICCC to co-

		ordinate and manage the recovery process.
ATM	Air Traffic Management	The aggregation of the airborne and ground-based functions (air traffic services, airspace management and air traffic flow management) required to ensure the safe and efficient movement of aircraft during all phases of operations.
	Atomic Function	The NERC system allows any controller or supervisor role to be present on any workstation. Each ATC function has a unique identifier known as an Atomic Function which allows these to be managed. "Sector 7 Planner", "Local Area Supervisor 5", "Mil East" are all examples of Atomic functions.
	Band-box	LAC is divided into a maximum of 32 sectors that can be operated independently. When traffic is light it is efficient to combine sectors together to be operated by one sector team. This process of combining sectors is known as band-boxing and as a result the sectors are said to be band-boxed.
	Bronze Team	BRONZE teams are the tactical level in the event of an incident. They can be any shape or size and are set up by SILVER to undertake specific tasks.
CAP 670	Civil Aviation Publication 670	CAA publication titled ATS Safety Requirements and which sets out the safety regulatory framework and requirements associated with the provision of an air traffic service.
CCEB	Complete Continuing Evidence Base	The CCEB is all the information necessary to evolve the software in the light of change requests, and to carry out forensic analysis in the events of problems. Thus it covers a range of key data including specifications, architectural descriptions, verification results.
CGW	Communications Gateway	A processor that interfaces the LAC data processing environment to wider NATS engineered environment.
CM	Configuration Management	Configuration management refers to a discipline for evaluating, coordinating, approving, and implementing changes in artefacts that are used to construct and maintain software systems. An artefact may be a piece of hardware or software or documentation
CMP	Configuration Management Plan	Plan for how configuration management will be undertaken for a specific organisation or system.
CP1, CP2, CP3	Control Period 1/2/3	Prior to the creation of Reference Periods under SES NATS was regulated by the CAA under a UK only scheme. Price controls were set for Control Periods commencing at the time of the PPP in 2001. CP1 and CP2 were 5 year periods while CP3 was a 4 year period from 2011 to 2014 to bring its conclusion into line with the European Reference Periods.
CR	Change Request	A document containing a call for a specific change to a systems or process.
CSCI	Computer Software Configuration Item	An aggregation of software that is treated as a single entity within the configuration management process.
CTC	Corporate and Technical Centre	NATS corporate headquarters and home to many key functions including engineering, programmes, HR, Communications, Training, Strategy, Supply Chain, Simulation, Safety and IS
CTOT	Calculated Take-Off Time	A time calculated and issued by the appropriate Central Management Unit, as a result of tactical slot allocation, at which a flight is expected to become airborne.
	Deploying SESAR	The Deploying SESAR programme is a significant part of NATS investment programme throughout RP2 and will transform its operations in support of Europe's Single European Sky. The three main outcomes of the programme will be the replacement of many ageing legacy ATM systems; deployment of a modern ATM platform and the introduction of the SESAR concepts of operation to enable increased flexibility, capacity and efficiency
DSS	Data Systems Specialist	Role with responsibility to lead and direct a watch providing 24 hour immediate support of NAS FDP computer systems and associated peripherals in support of for ATC operations
EASA	European Aviation Safety Agency	EASA is the European Union Authority in aviation safety.
EC	European Commission	The European Commission is the executive body of the European Union responsible for proposing legislation, implementing decisions, upholding the EU treaties and managing the day-to-day business of the EU.
EFG	Emergency and Fallback Guidance	Guidance on how to manage emergency and fallback situations for the UK Flow Management Position.

EMS	Error Management System	Error Management System (EMS) are processes / systems often used in in high-hazard industries, whereby occurrences that do not cross the above safety thresholds are nevertheless captured and the data used to inform independent trend analysis and risk management.
ETIC	Engineering Technical Incident Cell	ETIC is an engineering communication and coordination cell that can be convened following the occurrence of an engineering event. The ETIC will be the means of communication to upper management and external assistance and will provide an engineering focal point for incident management away from the ATC operational environment.
EU	European Union	Union of 28 European member states.
	Eurocontrol	European Organisation for the Safety of Air Navigation
FAB	Functional Airspace Block	An airspace block based on operational requirements and established regardless of State boundaries, where the provision of air navigation services and related functions are performance-driven and optimised with a view to introducing, in each functional airspace block, enhanced cooperation among air navigation service providers or, where appropriate, an integrated provider
FAS	Future Airspace Strategy	The Future Airspace Strategy is the CAA’s strategic framework for UK airspace. Its aim is to provide a policy structure to enable a modernised air traffic management system that provides safe, efficient airspace, that has the capacity to meet reasonable demand, balances the needs of all users and mitigates impact on environment.
FCM	Flight Conformation Message	A message to be sent to Enhanced Tactical Flow Management System confirming the operation of the flight.
FDP	Flight Data Processor	The Flight Data Processor is a core ATM system which coordinates flight plan and track data and distributes real time flight information to controller working positions. Modern FDP systems process flight plan, adaptation, manual input and other data to provide an accurate 4D trajectory calculation for a given flight.
FEP	Front End Processor	A data communication interface that enables NAS to exchange data with SFS
FIR	Flight Information Region	An airspace of defined dimensions within which flight information service and alerting service are provided.
FMP	Flow Management Position	A working position established in appropriate air traffic control units to ensure the necessary interface between local ATFCM partners (i.e. ATCs, AOs and Airports) and a central management unit on matters concerning the provision of the air traffic flow and capacity management service
	Gold Team	GOLD represents strategic level command in the event of an incident. It comprises the NATS Executives and focuses on strategic matters including corporate communications, interfacing with Government and customers and on continuing to run the business.
HAR	Hazard Analysis Report	The Hazard Analysis Report presents the results of the Hazard Analysis conducted on the operational equipment supporting a specific Air Traffic Service (ATS) The hazard analysis supports the safety case and specifically the assurance for the design in Safety Case Part 2.
HF	Human Factors	Human factors is the discipline of designing products, systems or processes to take proper account of the interaction between them and the people who use them so as to minimise the impact of human error and maximise safe human performance.
HMI	Human Machine Interface	An HMI is the software application which presents information to an operator and allows them to enact control operations in order to fulfil their task.
HOEC	Heathrow Operational Efficiency Cell	The HOEC provides an to collaboration and early decision making at Heathrow, by enabling key stakeholders to work together with access to common shared information.
IATA	International Air Transport Association	Trade association for the worlds airlines with some 250 airline members.
ICAO	International Civil Aviation Organisation	The International Civil Aviation Organization (ICAO) is a UN specialised agency, created in 1944 upon the signing of the Convention on International Civil Aviation (Chicago Convention). ICAO works with the Convention’s 191 Member States and global aviation organizations to develop international Standards and Recommended Practices (SARPs) which States reference when developing their legally-enforceable national civil aviation regulations.

iFACTS	Interim Future Area Control Tools Support	A Trajectory Prediction (TP) and Medium Term Conflict Detection (MTCD) system that identifies and display predicted conflict information to controllers to support decision making
iTEC	Interoperability Through European Collaboration	iTEC brings together the air navigation service providers of Spain (ENAIRE), Germany (DFS), the UK (NATS) and the Netherlands (LVNL) – alongside systems provider Indra. It was initially established in order to develop a next-generation Flight Data Processing (iTEC-FDP) system and to explore collaboration on a Controller Working Position (iTEC-CWP).
kSLoC	Kilo Source Lines of Code	Measure of the size of software system based on counting the lines of source code (in units of 1000 lines). For example, a 20 kSLoC software module contains 20,000 lines of software code.
LAC	London Area Control	London Area Control handles civil aircraft over England and Wales in flight at high level.
LAIMM	London Area In Manual Mode	LAIMM is a fallback mode for LAC which is entered into as a consequence of certain failure events, e.g. NAS failure, when some automation features become unavailable for controllers.
LAMP	London Airspace Management Programme	NATS programme to re-organise the operation of airspace around London airports to improve capacity, safety and environmental performance.
LM	Lockheed Martin	US aerospace and defence contractor and the prime contractor for the original NERC system. LM are one of a number of suppliers who continue to provide systems support to the LAC system under a single team managed by NATS.
LTC	London Terminal Control	London Terminal Control, including the five main London airports, and covers aircraft generally flying below 21,500 feet, with the precise height demarcation with LAC depending on the location
LTIP	Long Term Investment Programme	The LTIP is the name given to NATS capital investment plan which forms the underpinning for the Service and Investment Plan.
MATS	Manual of Air Traffic Services	The Manual of Air Traffic Services contains procedures, instructions and information, which are intended to form the basis of ATS within the UK. It is published for use by civil Air Traffic Controllers and may also be of general interest to others associated with civil aviation.
MDI	Minimum Departure Interval	A minimum time interval that is required between successive departures on the same Standard Instrument Departure.
MOR	Mandatory Occurrence Report	The objective of the MOR Scheme is to contribute to the improvement of flight safety by ensuring that relevant information on safety is reported, collected, stored, protected and disseminated. The sole objective of occurrence reporting is the prevention of accidents and incidents and not to attribute blame or liability. The MOR scheme is fully described in CAP 382 - The Mandatory Occurrence Reporting Scheme. This document collates the relevant rules and regulations and provides guidance on occurrence reporting, including examples of what should be reported and by whom.
MTCD	Medium Term Conflict Detection	Software algorithms that compare the predicted future trajectories of multiple aircraft in order to identify potential conflicts.
NAS	National Airspace System	Civil Flight Data Processing system operating centrally for the whole of the UK.
NERC	New En-Route Centre	The project name for the London Area Control computer systems.
NERL	NATS En route plc	NERL (formerly NATS En Route Limited) is the sole provider of air traffic control services for aircraft flying ‘en route’ in UK airspace and the eastern part of the North Atlantic. It is economically regulated by the Civil Aviation Authority (CAA) within the regulatory framework of the European Commission’s (EC) Single European Sky (SES) and operates under licence from the Secretary of State for Transport.
NLMCC	NATS Licence Management Coordination Committee	The CAA’s NERL Licence Management Coordination Committee
NM	Network Manager	Function provided by the Eurocontrol Network Manager Directorate (NMD) as described in the Network Manager Implementing Rule of the European Commission.
NMOC	Network Manager Operations Centre	The NMOC is the primary operational capability of the Network Manager, and delivers core operational services including flow and capacity management and flight planning operations.

NOP	Network Operations Portal	A set of information and actions derived and reached collaboratively both relevant to, and serving as a reference for, the management of the Pan-European network in different timeframes for all ATM stakeholders, which includes, but is not limited to, targets, objectives, how to achieve them and anticipated impact
NTCA	Northern Terminal Control Airspace	NATS programme to re-organise the operation of airspace primarily around Manchester airport to improve capacity, safety and environmental performance.
OPNOT	Operational Notice	Notice to disseminate information which, although significant, does not warrant the issue of a Temporary Operating Instruction. OPNOTs may contain information and/or guidance relating to ATC procedures, but must not contain instructions. OPNOTs exist to provide short term operational information, on a limited distribution basis
OS	Operations Supervisor	Key operational management role responsible for the provision of clear people leadership and direction of ATC operations to the Watch in the Operations Room ensuring a safe, efficient and effective service delivery
PBN	Performance Based Navigation	Area navigation based on performance requirements for aircraft operating along an ATS route, on an instrument approach procedure or in a designated airspace.
PC	Prestwick Centre	NATS control centre at Prestwick which provides ATC services for the Scottish FIR, part of the London FIR covering lower level airspace in the North of England and a large are of Oceanic airspace over the North Atlantic.
PCP	Pilot Common Project	The PCP contains the first set of ATM Functionalities that, having completed their research, development and validation cycle through the work of the SESAR Joint Undertaking, have demonstrated their readiness for deployment and their capability to produce benefits in particular if they are deployed in synchronisation
PPP	Public Private Partnership	NATS is a public private partnership between the Airline Group, which holds 42%, NATS staff who hold 5%, UK airport operator LHR Airports Limited with 4%, and the government which holds 49%, and a golden share.
PTR	Problem Trouble Report	For the NERC system problem reporting and defect tracking is carried out in accordance with the PTR process which includes the identification, assessment and 'tagging' of any PTR that impacts safety..
QA	Quality Assurance	Quality assurance (QA) is a system of checks designed to ensure that products are free of faults. A quality assurance system involves regular quality control inspections that test and monitor the quality, accuracy and fitness for purpose of the product, from the design stage through to manufacture
QWPM	Quality Work Package Manager	The QWPM is responsible for the routine delivery of quality services, for example: design and code inspections; test witnessing and concession and defect prevention process management.
RP1, RP2	Reference Period 1, Reference Period 2	The Performance scheme of the SES is one of the key pillars of the Single European Sky aiming at achieving improved safety performance and efficiency. The Performance scheme is organised around fixed Reference Periods (RPs) before which performance targets are set both at EU-wide level and National/FAB level. The first reference period (RP1) runs for three years from 2012 to 2014. The second reference period (RP2) will be from 2015-2019
SDM	SESAR Deployment Manager	The SESAR Deployment Manager (SDM) is the body that synchronises and coordinates the modernisation of Europe's air traffic management system under the political oversight of the European Commission.
SES	Single European Sky	Initiative launched by the European Commission in 2004 to reform the architecture of European air traffic management. It proposes a legislative approach to meet future capacity and safety needs at a European rather than a local level
SESAR	Single European Sky ATM Research	SESAR (Single European Sky ATM Research) is the technological pillar of the Single European Sky. It aims to improve Air Traffic Management (ATM) performance by modernising and harmonising ATM systems through the definition, development, validation and deployment of innovative technological and operational solutions. These innovative solutions constitute what is known as the SESAR concept of operations.
SFS	System Flight Server	Software that stores and distributes the next 4 hours of flight data in Swanwick Area Control and records which sector is being operated from which workstation
	Silver Team	SILVER is the operational command level during an incident. It comprises senior managers at individual sites and manages the response to an incident within the Site

SIP	Service and Investment Plan	NATS (En Route) plc (NERL) is required by Condition 10 of its licence to submit to the CAA each year a Service and Investment Plan (SIP). The purpose of the Plan is to provide an annual update of NERL's investment plans and to show whether there have been material changes to those plans.
SLoC	Source Line of Code	Measure of the size of software system based on counting the lines of source code
SMS	Safety Management System	A SMS is an organised approach to managing safety, including the necessary organisational structures, accountabilities, policies and procedures. Additionally it focuses on ensuring that safety management is integrated into the day to day activities of the organisation with an organisational culture that reflects the safety policy and objectives. At the core of the SMS is a formal Risk Management process that identifies hazards and assesses and mitigates risk.
STAR	Safety Tracking and Reporting (system)	The Safety Tracking and Reporting (STAR) system is a single, authoritative NATS wide safety data tracking and reporting system. It enable timely and accurate passage of safety data across NATS, and a completely electronic safety investigation process including an audit trail..
TDL	Tactical Data Line	Simple display of data for an aircraft providing the controller with key tactical information.
TEI	Temporary Engineering Instruction	TEIs are formal instructions raised for an operational asset or process, e.g. as given in a System File or System Management Manual, where it is necessary to: • Temporarily supplement the standard operating instructions • Temporarily vary the standard operating instructions • Temporarily add a new instruction in lieu of a formal procedure
TMA	Terminal Manoeuvring Area	TMA is a term used to describe the airspace around a major airport or group of airports and with airspace and procedures designed to manage the flow of traffic into and out of the airports.
TP	Trajectory Prediction	Software algorithms that predict the future position of aircraft over time based on their filed plan and clearances and taking into account a range of factors including aircraft performance, and weather conditions (wind).
UTC	Universal Time Coordinated	UTC, is the primary civil time standard by which the world regulates clocks and time and which is used throughout ATM
	Watching Mode	NERC workstations can operate in a number of workstation modes depending on the role they are fulfilling. Watching mode is where one workstation can display a full copy of the data from another workstation but with no control function.
ZRR	Zero Rate Regulation	Regulations are methods of matching traffic demand to available capacity by limiting the number of flights planned to enter an airspace or aerodrome, achieved by the issuing of departure slots. A Zero Rate Regulation which sets this limit for regulated traffic to zero are applied in circumstances e.g. of system failure when ANSPs need to severely constrain traffic in order to ensure safety.