

Transport Committee

Oral evidence: [Air traffic control disruption](#), HC 1849

Wednesday 18 October 2023

Ordered by the House of Commons to be published on 18 October 2023.

[Watch the meeting](#)

Members present: Iain Stewart (Chair); Jack Brereton; Paul Howell; Grahame Morris; Gavin Newlands; Greg Smith.

Questions 57–90

Witnesses

[III](#): Martin Rolfe, Chief Executive, NATS; and Alistair Borthwick, Chief Financial Officer, NATS.

Examination of witnesses

Witnesses: Martin Rolfe and Alistair Borthwick.

Q57 **Chair:** I welcome our third panel. Could both of you state your name and position, please, for the purposes of our record?

Martin Rolfe: By all means. My name is Martin Rolfe and I am the CEO of NATS.

Alistair Borthwick: Good morning. I am Alistair Borthwick. I am the chief financial officer for NATS.

Q58 **Chair:** Thank you very much for your time this morning. I think you heard a fair chunk of the previous sessions and the questions that we will want to put to you. I begin by inviting you to give a very brief summary of what happened on 28 August and subsequent days. We have your preliminary report. What are the next steps to look at some of the issues raised in it?

Martin Rolfe: By all means. Thank you very much for inviting both of us to this panel. It is good to be able to give our input.

Our job, as the provider of the air traffic control service under licence by Government in the UK, is to make sure that we can safely control the flow of traffic. We have about two and a half million flights a year and hundreds of millions of passengers. We were expressly set up by Government and airlines many years ago to make sure that we did just that. In order to do it, we have to know the intention of every aircraft. We need to know where it is flying from, where it is flying to and what points it is flying over on its way. We call those "waypoints".

On 28 August, we received an incredibly unusual message in the form of a flight plan—a stated intention—that was unlike any we had seen in the last five years, since the system that we were using has been in operation. It was sufficiently unusual in that it did not just have duplicate waypoints, as you may have been informed earlier; it was the combination of them and the sequence of them. It was sufficiently different that the system decided that the safest course of action was to stop processing and, essentially, allow a human to intervene. That is the basis of our safety-critical systems. In the event that something happens that is unexplainable, it is passed on to humans because ultimately humans are better at interpreting confusing data than computers are. That is what happened.

As a result, we reduced the flow of traffic in the skies over the UK. That is simply to make sure that the air traffic controllers can safely handle what is coming when they do not have all the information that you might otherwise expect them to have. Regrettably, obviously, that results in cancellations and delays on the very rare occasions that it happens. This is the first time something like this has happened in over 10 years. We resolved it as quickly as we could. Obviously, we completely understand



HOUSE OF COMMONS

the distress and the upset that any delays cause. We do not underestimate that. Perhaps I can reiterate my apology to everyone who was involved in that on 28 August. We absolutely stand by our original apology, and I will continue to apologise.

Our engineers then worked as quickly as they could in circumstances where the problem was not fully understood. At this point we did not know exactly what the problem was. The on-site engineers called in back-up engineers. They called in the provider of the software in Austria. They did all of that over the course of that morning. Normally, most of the technical issues we see across the diverse NATS estate are solved within minutes. This was a very unusual incident where it took a long while to diagnose the problem and then to restore the system. Obviously, our priority was making sure that it was restored safely and in a way that was sustainable, so that it was not going to fail again.

One other thing to mention is that this was part of the system. There are lots of parts to an air traffic control system, the most important being the voice communications with pilots and aircraft. That is the safest and most important piece. That carried on working throughout. The radar systems carried on working throughout. This was around the flight planning system and was sufficiently unusual that it required human intervention.

We were back up and running with the system at about 2 o'clock or 2.30 in the afternoon, which is unusually long. Everything was then back to normal from a NATS perspective, recognising that it obviously took a lot longer for the passengers who were disrupted to be repatriated. We were back up and running without any regulation or delays to the operating airlines over the UK by about 4 or 5 o'clock.

Chair: Thank you. I want to take up some of those points in a minute, but Grahame has a question.

- Q59 **Grahame Morris:** You have given us a bit of an overview. I hope you heard some of the concerns and suggestions from, particularly, the previous panel. You explained what happens when there is one of the waypoint name anomalies, but, for my understanding, why would the entire system shut down? Why would it not flag up that flight plan to be corrected by a human, by someone plotting the course manually? Why was the back-up system similarly affected? Was a decision taken to shut down the back-up system as well because you were assuming that the whole system was hacked, infected or whatever the term is?

Martin Rolfe: I am happy to explain that further. We handle duplicate waypoints in flight plans all the time. We have handled 16 million or more and probably thousands of those have had duplicate waypoints.

- Q60 **Grahame Morris:** It must be very common.

Martin Rolfe: It is.

- Q61 **Grahame Morris:** There are at least six places called Easington in



England that I know of, so it must be—

Martin Rolfe: It is not that common.

Q62 **Grahame Morris:** We've got an airport at Horden.

Martin Rolfe: It is not that common in the global dataset for aviation waypoints, but it does exist. We handle them all the time. As we explained in our preliminary report, which was written immediately after the incident, with an awful lot of effort and diligence, it was not the fact that it just had duplicate waypoints. It was a very curious combination of the flight plan that came in from the airline, how that was translated by Eurocontrol, how that came into our systems, and how our systems then processed it.

Normally, you are absolutely right: a flight plan simply with duplicate waypoints would be processed perfectly well. Most flight plans are processed perfectly well. If we know of a circumstance where one is anomalous, we obviously programme the system to make sure that never happens again. In this case, what we have done is to resolve the system. We have put in a number of different fixes to make sure it can never happen again.

Moving on to your point about the back-up systems, they are generally, though not always—maybe I will come back to that—two separate physical systems. I think you heard earlier that they were on the same computer. Michael was probably a little confused there. They are actually two very separate systems. The systems are the same hardware, if you like, but they are there to make sure that if one system's power fails, if the communications link fails or if the system throws up a random error, as your iPad might, the back-up can take over. On four occasions over the last five years that back-up system has come into operation and carried on working. Nobody was affected and no flights were impacted.

In this case, the back-up system processed the same, very unusual, flight plan and came to the same conclusion, which was, "This is sufficiently unusual that it needs human intervention."

Q63 **Grahame Morris:** Is the back-up system replicated exactly the same as the primary system, or is it a different software package?

Martin Rolfe: It is the former. The voice communication system did not fail. It never fails because it is so safety critical. It has different software and different hardware. It has different communication lines and there are all sorts of differences to make sure that they can never fail. It is a bit like the flight systems on an aircraft, where you have different software, different control systems and no interaction between the parts.

For other parts of our system, which are safety related but not safety critical, we generally use the same software on different hardware with separate connectivity, separate power and so on. That is a deliberate decision. Obviously, as you change software or have different software,



HOUSE OF COMMONS

you make the systems increasingly complex. As you do so, you introduce more different failure modes. It is a balance between engineering complexity, the likelihood of failure and the actions you can take.

Obviously in this particular instance, although it is something we would never want to do, the actions we can take are to reduce the number of flights and keep everybody safe. If the voice system were to fail and you could not talk to aircraft, you could not do anything at all. It is a different level of safety assurance.

Q64 Chair: The previous panel suggested that other air traffic control systems would not have, to use their word, collapsed their own system, but yours did. How would you respond to that?

Martin Rolfe: I would start by saying that the system did not collapse. The system shut itself down for a set of reasons that it understood, in the sense that it knew there was something we needed to be worried about. Using terms like "collapse" is probably a little bit simplistic.

Other providers around Europe suffer from similar incidents. It is not necessarily with exactly the same software, but it is very common. Despite what you may have heard earlier, the French had a problem that was similar in nature. It was not duplicate waypoints, but it was similar in terms of the flight planning system. That put them out of action for about seven or eight hours across the entire region of France. In fact, the Italians had a similar problem on the same day as us, 28 August. It was not quite as severe as the one we had. These things are not uncommon. Normally they are fixed without anybody noticing and without delaying any passengers, but any large technical system will have days when things need to be fixed.

Q65 Chair: On your resources, some criticism was made in the previous panel of the fact that the back-up team were working remotely. It was a bank holiday and I think he alluded to them watching daytime television rather than monitoring. How would you respond to those criticisms? Are they fair?

Martin Rolfe: No. We are a very serious organisation with safety at the heart of the culture. That would never be the way we would operate. We do not treat a bank holiday differently from any other day in the context of operations. We operate 24/7. We operate 365 days of the year. It does not matter what the day is; we have the right number of air traffic controllers and the right number of engineers on site. In addition, we have the ability, as you would expect of any modern company, to have people dial in remotely to help and assist when things fail. That is exactly what we did. In fact, I would probably argue that because we were able to access people more quickly, including those overseas in Austria where the company that supplied this particular software was located, we were able to get them working on the problem within hours. Clearly, it would have taken days had we not had that ability. No, there is no truth in those assertions.



Q66 Chair: Before I turn to colleagues, I want to pick up on the duplicate waypoint issue. I appreciate that you say that the flight plan that caused the problem was more complex than just having duplicate waypoints, but the duplicate waypoints were clearly part of the problem. They have been in existence for decades. Why has that not been fixed if it causes problems?

Martin Rolfe: There are two parts to that question. In terms of the overall global datasets, there are actions in place with ICAO—the International Civil Aviation Organisation—globally to remove duplicate waypoints. That has been an ongoing effort over, as you say, probably decades, and certainly as air travellers become more global.

There aren't any duplicate waypoints within the UK. These were duplicate waypoints outside the UK. The systems are designed to handle them. In fact, they do handle them. We probably get them hundreds of times every day. Like any of these failures, it is only when you have an incredibly specific set of circumstances that has never been seen before that you end up in a situation where you have to resort to what I call a fail-safe mode of operation, where you involve humans in the loop. That is really the reason why this happened. The duplicate waypoints are a contributory factor, but it is considerably more complicated than that. It is laid out more in the report. I am sure that the CAA, in their independent review, will want to dig into that even more.

Q67 Chair: This is my last question before I hand over. Is a basic problem that you are integrating preliminary data input with the airspace management function in the hardware? Is that the basic problem, and do you need to have a better defensive data checking system to stop these issues arising?

Martin Rolfe: I will answer your question, although I am not 100% sure I understood it. If you are asking whether or not there should be better data checking for the data that comes in to avoid this happening again, there are two stages to that.

The first stage is whether it is something you recognise and therefore something you can throw away if it is bad. This is not a spam filter; this is something that is dealing with safety-related data. It has to know enough about it, and understand enough about it, to know whether it is safe to throw it away. It seems like it might be easy to throw something away, but actually what you do not want to do is throw away something that is incredibly important further down the line in these systems because you could end up in a much worse situation.

In our case, what happened was that the first barrier of, "Is this the wrong data?" was passed through, with it saying, "It's not wrong in a traditional sense," in that it is not gobbledygook, but when the system went to process it, it said, "Hold on a minute, this is so unusual that we need a human to intervene on the grounds of safety." It is not a failing of the system as such. I do not want to imply in any way that we would



HOUSE OF COMMONS

expect this to happen routinely, and it does not. We have already made sure that it cannot possibly happen again, but it is perhaps one of the artefacts of safety-critical engineered systems.

Q68 Paul Howell: I want to continue the point we are talking about before I move on to a specific question. I am still struggling to understand why somebody cannot just look at one rogue flight plan as opposed to entering the system and then everything shuts down. You can call it a controlled shutdown if you like, but it effectively carefully collapses the system. That is what happens.

If it was just because of two waypoints, surely a message should come up and say, "This has two same waypoints. Can somebody look at this?" That should not take five hours to sort out. What is different about this occurrence which means you did not get that message saying, "Can somebody look at this very specific flight plan and sort it out?"

Martin Rolfe: You are absolutely right. The system invariably takes duplicate waypoints. It does not even have to flag them to a user. It will process them. I cannot work out what 99.9% of the time is when it comes to one in 15 million, but on this particular occasion the combination of sequencing and the translation from what was the standard flight plan via the Eurocontrol system to ours, and then into ours, was sufficiently unusual that the system did not recognise it as a problem with duplicate waypoints. It said, "This is so unusual I need a human to intervene."

Q69 Paul Howell: But why would it not say, "There is a problem with this flight plan, so will somebody look at it?"

Martin Rolfe: Because in the way the system processed it, it came up with an answer that said, "The answer here is minus 10; I don't understand that." It did not then relate it back and say, "It must have been because of this flight plan." It is processing thousands of flight plans every minute. Effectively, the system was unable to understand that the issue was definitely and solely with that flight plan. What it came up with, effectively, was, "I don't understand this. It is sufficiently anomalous that I need to get a human involved." That is precisely why it did that.

It is a complex, technical system. Obviously, it is something that the CAA will also look at in their review. It is as I have explained. Effectively, it was sufficiently anomalous.

Paul Howell: At some point in the CAA report we need to get to the bottom of what is actually the cause of this. It sounds, for want of a better phrase, like a bug in the system.

Q70 Gavin Newlands: Why would the flight plan be compliant for IFPS Eurocontrol but not your system? What is the reason for that?

Martin Rolfe: Eurocontrol translates, to some degree, the flight plan that comes in, to say, "Do I need to send this to the UK? Do I need to send it to France? Do I need to send it to Italy?" It does some basic checking on



HOUSE OF COMMONS

that, but obviously what it does not try to do is convert it into a route, for example, for the entire region. Our systems then have to take that and translate it. There is quite a complex operation that goes on with every individual flight plan. As I say, the final report will obviously delve much more into that. The problem is fixed though, so it will not happen again.

Q71 Paul Howell: The fix time became the problem. It was the fact that it was about five and a half hours instead of the four hours that is the lead time on a flight plan submission. Should we therefore change that to five and a half, to six, to seven or to eight? Would that actually make a difference, or have you changed things so that something else has fixed that now?

Martin Rolfe: Yes. The answer to your question is both. What we have changed is the software, such that the problem will never happen again. The combination sequence that we have discussed will not matter.

Q72 Paul Howell: You say you have fixed it so that it will never happen again.

Martin Rolfe: Yes.

Q73 Paul Howell: Therefore, you know very specifically what happened.

Martin Rolfe: Yes, absolutely. We fixed the specific problem and that will not happen again. In terms of the four hours, it is a little bit of a balance. Historically, four hours has been more than enough time to fix any problems we have seen. The difficulty with going to more than four hours is really one of data stability. I think that comes back to the Chair's question a little bit.

If you imagine flight plans coming in from days to weeks to minutes ahead of time, the further you go back or the more data you store, the more out of date it becomes and the more it needs to be updated. There is a balance in the system design to say what is the right amount of data to store such that you can continue operating. Of course, remember that for those four hours we were still operating at full capacity. What is the right amount to store? You have to balance that with how much that data is likely to change in the intervening period.

In the preliminary report, in section 9, we identify whether looking at that four-hour window is the right answer, given the experience of 28 August. I am sure that is something the CAA will also look at in their review.

Q74 Paul Howell: You could push that forever and come up with a different conclusion. I understand there has to be some sort of balance in where things are.

The interim report listed that you had updates and improvements that you planned to put in by 4 September. Can you confirm that they happened, and that they delivered what you expected? What are you doing next in terms of any further investigations or further changes you might think about? Given the last few years, we are all, unfortunately,



very aware that something catastrophic can happen. Covid is a perfect example. We have fixed it in the sense that people know how to treat covid as a thing now, but what happens if there is another similar biochemical incident? You could have another similar thing. How are you broadening your consideration to make sure that you are not just fixing the very specific but fixing the concept of what went wrong?

Martin Rolfe: To deal with your first point, the items in the report that you referred to were all completed. They were all safety assured and all tested. They are in and working, and we are certain of that. That piece has been done. Because these sorts of things, fortunately, happen so infrequently, we spend a huge amount of time and effort looking at what we could do differently, particularly in resilience.

What we will do from this one is across a couple of different areas. They are outlined in section 9 of our preliminary report. It is around whether there are similar systems that work in similar ways. We would look across that. How do we communicate in the event of disruption? How do we make sure that, when we have systems that go into these kinds of modes, there are better and quicker ways of restoring them? We will undoubtedly get a lot of lessons from this. We will obviously co-operate fully with the CAA review as part of that. We always make sure that we try to learn everything we possibly can from these events. They are rare but they are incredibly serious, and we recognise just how incredibly serious they are.

Q75 **Paul Howell:** I have a final point before I hand back to the Chair. You referred to resilience in what you were just saying. Resilience also comes down to your staffing, numbers and people. I am sure you have heard that there has been quite a lot of criticism of you being understaffed. In the last panel, there were also questions raised about the challenges involved in actually training new people in the process.

There are two things. One is about your comment. Do you have the staff that you need on the ground? If not, what are you doing about it? Do you agree with the suggestions made by the previous panel about changing the protocols about who could be trained, how long they could be trained and where they could be trained? I must admit I find it bizarre that you cannot train people in a simulator for that type of work. That is just as an example.

Martin Rolfe: I absolutely stand by what I said at the start of the summer with the Aviation Minister and the Secretary of State for Transport. We are appropriately staffed, with probably one exception, which I will come back to in a moment, which is Gatwick. We have the right number of staff to do the job we need to do. Despite what you might have heard earlier, our delay level is incredibly low. I heard some of the earlier witnesses talk about how we compare to the rest of Europe and that it does not matter. I would argue that it is probably a sensible benchmark.



HOUSE OF COMMONS

We account for about a quarter of all European traffic. In NATS we handle about a quarter of all European traffic. We are about a 50th of the total delay. By any objective measure our teams do a very good job, which does not take away from the fact that obviously 28 August was something that we would never want to see happen or do not appreciate the impact of.

On our staffing at Gatwick, which maybe I should come back to, we took over that contract last year. We had it many years ago with much higher staffing. The airport then decided to award it to a smaller start-up subsidiary of the German national provider. We took it back because they were not able to perform the requirements of the contract. It was understaffed from a resilience point of view when we inherited it, back in October last year. We took over operations in October last year. We have been working with the airlines and with the airport to increase resilience. It takes a long while, and I will answer your question on the licensing piece. When we took it over, it is fair to say that the airlines were delighted that we were taking it back. Everybody understood at the time that it had been understaffed. We are doing everything to resolve that.

Alistair Borthwick: To clarify and be clear for the panel, there is a big distinction between the Gatwick airport contract, which sits within our services business and is commercially one contract that is available for anyone who provides tower air traffic control services to bid for—hence why the start-up that Martin referred to was able to do so previously—and, in relation to the technical incident of 28 August, the en route business which relates to our regulated monopoly business. That is the main provider of services, referred to as NERL. They are two very different businesses. One is the monopoly, and one is the commercial business that bids for contracts with towers.

Martin Rolfe: To go back to your point on licensing, I fully agree with what Sophie said in the earlier panel. We are somewhat unique in the world, certainly in Europe, in that we cannot take qualified controllers from elsewhere without them having to repeat all of their training in the UK. In fact, even more unusually, we cannot even take military air traffic controllers who have been licensed in the UK and work at military air traffic control towers. We cannot even use them without retraining. That is something we have been pushing Government to change policy on over the last five or six years. We are continuing to work with them to do that. It is a long-term challenge for staffing of air traffic controllers.

On the simulator piece, perhaps what was described earlier was not quite accurate. We use a lot of simulator training. In fact, most of the first year of training is simulator training. The second and third years are on-the-job training. It is a little bit like a specialism in medicine. You get your basic degree or competence in air traffic control. In the case of an airport, you then have to go and learn exactly how that airport works. You do that on the ground in the airport, whether that be Gatwick or Heathrow, or in our centres under the regulated business, as Alistair just



HOUSE OF COMMONS

mentioned. We use simulators a lot, but there is also a very strong requirement in the regulations, as currently set out, to do on-the-job training.

- Q76 **Paul Howell:** I thought I heard from one of the previous panels that the training time was nine months. You have now said year one, year two and then on the job. How long does it actually take from start to finish?

Martin Rolfe: Generally about three years.

- Q77 **Paul Howell:** Before they can even participate in any way, shape or form? For me, in my simple head, there is a very different challenge to working as an air traffic controller at Gatwick, in your monopolistic world, compared to Teesside airport or somewhere like that, where the traffic levels are different. Is there just a basic position to get an air traffic controller into the pot somewhere, or are you talking about high-level training for somebody at Gatwick airport or monopolistic?

Martin Rolfe: Everybody goes to the same first training. The first year is the same for everybody. That is at the college on simulators and teaches you to be an air traffic controller. It teaches you all the language and all of the procedures, and so on.

The second part is your specialism. You are absolutely right that it would probably take less to train at Cardiff airport, where there is much less traffic, than it would at Heathrow. Heathrow and Gatwick tend to be at the longer end. They probably take up to two years. It might be less. It depends a lot on the aptitude of the controllers. They are a very scarce resource. The aptitude required for controllers is very difficult to find.

On the nine months that you heard earlier, we have articulated nine months as being pretty much the shortest time at somewhere like Gatwick that you could go from being an experienced controller somewhere else, coming in and then effectively converting to be a Gatwick controller. The familiarisation piece—the specialty piece—takes about nine months.

- Q78 **Paul Howell:** For a final piece of clarity, whether it be from the military, from overseas or wherever, if they were coming in to join that training programme, are they joining it right at the start, at the end of year one or at the end of year two?

Martin Rolfe: Day one.

- Q79 **Paul Howell:** Day one?

Martin Rolfe: Yes, if they are to come back in.

- Q80 **Chair:** You say you are lobbying for a change to allow speeded-up training for military and others who have some knowledge. Who needs to change what? Is that for the CAA? Is it a legislative change that the Government need to do?



HOUSE OF COMMONS

Martin Rolfe: It is both. It is secondary legislation. It is also CAA regulation. It is both of those, and we are working with both of them to progress that.

Q81 **Jack Brereton:** The previous panel were pretty scathing of NATS, particularly around the lack of communication with airlines, which resulted in them having much less time to actually deal with the problem. It resulted in them not having any communication from yourselves until the next day, which they suggest was written communication. Why was there a total lack of communication with airlines about what was happening?

Martin Rolfe: You may not be surprised to hear that that is not quite how it worked on the day. You may have heard from the last panel that they heard through Eurocontrol. I think that was a phrase that was used quite frequently. That is exactly how the system is designed to work.

If you imagine a big distributed system like the aviation network, what you do not want to be doing is having an individual call from me to an airline CEO or an individual call from Alistair to an airline CFO. The idea is obviously to get coherent, accurate and safe information out incredibly quickly as soon as you know that something needs to happen. That mechanism is Eurocontrol, which has a central website and repository hub, whatever you want to call it. As soon as we know that something is happening, we report it there. It is then immediately accessible to everybody.

When everybody says they heard it from Eurocontrol, what they mean is that we provide it to Eurocontrol, and it then gets promulgated to absolutely everybody. That is obviously the focus of the operational team. At that point in a crisis, we are communicating with the operational teams of the airlines and the airports, not generally the CEOs. The main factor in all of this was to make sure that we kept everything safe and that everybody understood what was happening, and the regulations we were putting on to slow down aircraft. Once we had done that, we started to talk to airlines. In fact, I think I spoke to most of them; I certainly spoke to Ryanair that morning.

We probably did not do as good a job in following up with the trade associations. They are probably second tier on our communications because they are not running the operations. One of the things we have said as part of the preliminary report, and as part of the CAA review, is that we think there are lessons to be learnt there. You can always communicate better in these circumstances.

Q82 **Jack Brereton:** Do you think the message came from Eurocontrol soon enough? There were several hours before some airlines got that message. Some of the previous panel said that they did not get the message for a couple of hours after the actual incident had occurred.



HOUSE OF COMMONS

Martin Rolfe: It is instantly available. As soon as we do it, it is instantly available to everybody who needs to know. What you are hearing there is perhaps the difference between management engagement of airlines versus operational engagement of airlines. The operational people in the airlines would have known instantaneously. In fact, that is why you heard from some of the earlier witnesses that they heard it from Eurocontrol, and that was the first they had heard of it. That is exactly how the system is designed to operate.

Q83 **Jack Brereton:** I think we have some disagreement there between this panel and the previous panel. I want to ask about the issue around reimbursement of costs. A number of those on the previous panel were saying that, in their view, NATS should reimburse some of the costs around accommodation for passengers, and things like that, which was because of the disruption and was out of the airlines' control. Is it your view that there should be some compensation from NATS to address the costs that had to be met by airlines?

Martin Rolfe: I will start with that question, and maybe ask Alistair to chip in. I absolutely understand the frustration of the airlines in the context of the expenses they have to incur. This, effectively, goes back to the way NATS is constructed. NATS's primary purpose, and in fact the first line of its licence, is to prioritise safety over all other things. The idea is that in a situation like 28 August, or perhaps even worse—say something really terrible had happened—we are focused not on cost but on safely fixing and restoring traffic. That was set up specifically by Government and by the airlines in that way.

There is a separate piece in our regulatory settlement with the CAA whereby, if we do not meet certain service targets, separate from safety targets, we will reduce the charges in future years. The airlines will receive rebates in future charges if our services are not acceptable, as the CAA deems. The CAA have said as part of the terms of reference for their independent review that they will be explaining to the aviation system why it is structured as it is currently structured.

Q84 **Jack Brereton:** That is a completely convoluted system.

Martin Rolfe: I think it is a system that places safety at its complete heart.

Q85 **Jack Brereton:** You are not taking any responsibility, are you, for the actual impact?

Martin Rolfe: We do take responsibility. In fact, I have absolutely taken responsibility. It is our problem to solve. We solved it, and we solved it as quickly as we possibly could. The challenge is that the structure exists as to how this is supposed to work. It can clearly be changed by the CAA or by Government, should they choose to. All I am saying is that it was put in place for a reason. It was put in place to ensure that there were not unintended consequences. All things are possible, but that is why it is what it is.



HOUSE OF COMMONS

Q86 **Jack Brereton:** Who should ultimately be held accountable for what happened on that day?

Martin Rolfe: I am absolutely clear that I am accountable for what happened on that day, without the shadow of a doubt. I run the NATS operation. I run the NATS company, and it is my responsibility. I have never implied otherwise. I apologise for what happened on that day, but, ultimately, I am responsible for making sure that every single one of the 5,500 flights that did fly that day were operated safely. That is something that I will always prioritise over any other consideration, which is not to say that I do not take responsibility, but it is absolutely the priority of the company.

Q87 **Chair:** To pick that up, the mechanism is there such that airlines will be compensated in lower future charges for the care costs that they incurred. Our previous panel told us that the fees are going up, not down. How do you reconcile that?

Martin Rolfe: There are two separate things. I will ask Alistair to come in on this. The fees going up are largely as a result of covid. As you can imagine, we did not receive any income from anyone during covid, but the regulatory regime allows us to recoup costs during that period to keep the system running. Obviously, we were still operating medical flights, military flights and so on despite the airlines not flying. We wanted to support them. That is why the fees are going up in that context. The actual underlying cost is coming down.

On the second part, there isn't a direct relationship between compensation and reduced charges. Perhaps that is where I will hand over to Alistair, to give a bit more colour to that.

Alistair Borthwick: To answer on both aspects, I heard the challenge previously with regard to fees going up and not coming down. Since 2002, the overall costs of NATS have reduced by 31% in real terms over that time. The reference that was made in the letter is looking at that duration of time and the benefits to our customers and passengers with regard to the overall decrease in real-terms costs that have been seen over that time.

With regard to the point on compensation, Martin has outlined a penalty and bonus regime attached to that which serves to look at the overall delay that might be apparent within the system and how far beyond arrival times and other aspects flights might be. It also deals with aspects of driving sustainability and environmental performance. If either of those are missed over the period of our price control, which is a five-year window, the additional cost to NATS or NERL, our regulated business, could be up to an additional £75 million over the five years of a price control.

Mechanisms exist to deal with poor service performance, or lower than service performance targets that are established within the regulatory



regime. It is not necessarily directly linked to compensation for those events. If our service does not meet the standards required, penalties kick in as part of the regulatory framework.

Chair: It seems a very opaque system. I imagine those are questions that we will want to put to the CAA and to Government. Gavin, you wanted to come in.

Gavin Newlands: You have essentially asked the question I was going to ask. It seems a bit odd to put in a line about ever-reducing prices in the year you put prices up 29%.

Q88 **Chair:** I read your mind, Gavin. Time is against us. I want to cover two quick additional points. First, there seems to be some discrepancy over the number of delays and flight cancellations. Mr O'Leary in particular says that you have significantly under-reported the impact. How would you respond to that?

Martin Rolfe: We took great care in the preliminary report, which I think was what he was referring to, to make sure that we sourced all the information we had from an independent place. That was Eurocontrol, which is the repository for that kind of information. We obviously do not cancel flights, so we do not necessarily know how many flights are cancelled by the airlines. What we put in there was the data that was publicly available at the time from Eurocontrol. We even said there that we do not know the true extent. We were simply trying to characterise it as best we could at the time. We were producing the report very quickly, as the Secretary of State had quite rightly asked, and as the airlines had asked, to understand what the technical causes were. We were not trying in that report to determine the amount of impact it had. In fact, that is a big part of the CAA's independent review. They have far more ability to gather that data than we do. We reported what we knew from sources that are independent of us.

Q89 **Chair:** One last question from me. Looking back to the root problem, which was the exceptional data input, just clarify why you cannot design the system so that, rather than the whole thing shutting down, you simply block that flight plan.

Martin Rolfe: You can design a system to do that, but the way it generally happens is that you design it to all the known parameters and all the parameters that are established up front to say, "These are all the things that we can possibly think of that might go wrong." You design it with that in mind. Then you test it with that in mind. Then you run it with that in mind. Generally speaking, over a period of time, you find things that are not handled. You have an issue and you fix the issue, normally very quickly. You then add a new rule to say, "And in this case handle it like this." Now that we know this circumstance can occur, and it was a very convoluted set of circumstances, we have designed that into the system and will use it in any future design. Effectively, what you are trying to do is predict as many different obscure possibilities as you



possibly can. Invariably, you do not catch them all at the design phase of any complex system.

Q90 **Chair:** But why can't the system just say to that aircraft, "You are not flying until we resolve this," and let the rest of the system continue?

Martin Rolfe: I have not explained completely the way flight plans work. The aircraft can be in the air when we receive the flight plan. This aircraft was well on its way before we received the flight plan. I don't think I am betraying any confidence in saying that it was a long-haul flight. It was probably five or six hours into that flight because we only received the flight plans four hours before, for the reasons we have talked about.

You cannot stop an aircraft taking off. If it has a valid flight plan, it has to take off. You can imagine how complex it would be globally to say that nobody could take off from China until everybody on its entire route had processed the flight plan. That is just an artefact of how the global aviation system works.

Chair: Thank you. I wish we had more time to pursue some of these issues, but we are up against the clock a little. Thank you for your time this morning. I imagine that once the further reviews are done this is a subject we will want to return to. For this morning, thank you again for your time.