the Availability Digest

www.availabilitydigest.com

The FAA's Availability Woes

December 2009

Without a filed flight plan, a commercial airliner can't fly. This was painfully demonstrated on Thursday, November 19, 2009, when the FAA's[1] National Airspace Data Interchange Network (NADIN) failed at about 5:30 AM Eastern Time and was down for three hours.[2]

NADIN is used by the air carriers and general aviation to enter flight plans into the National Airspace System (NAS). These flight plans are routed to the appropriate air traffic control facilities, including those airport control towers and enroute traffic control centers that will handle the flight.

NADIN uses the services of the FAA Telecommunications Infrastructure (FTI) to communicate across the National Airspace System. It seems that on November 19th, a scheduled maintenance on FTI in Los Angeles corrupted a router in Salt Lake City. The router's backup failed to take over, and the FTI services upon which NADIN depended were lost (a classic failover fault).

With NADIN down, air carriers could no longer enter their flight plans into the FAA's systems. Instead, they had to fax or email their flight plans to the FAA so that controllers could enter the flight plans manually. With 50,000 flight plans filed per day, the backlog built quickly; and planes sat on the tarmac, often for hours, waiting for takeoff clearance. Often, a flight was simply cancelled by the air carrier.

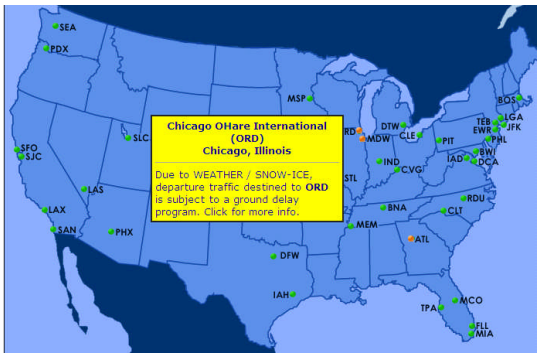

FAA Flight Status Web Site

http://www.fly.faa.gov/flyfaa/usmap.jsp

Fortunately, the failure occurred early in the morning. Affected most were flights originating on the East Coast of the U.S., though delays spread across the country as equipment needed for later flights never arrived on time. The delays were the longest at the busiest East Coast airports, including Atlanta, Washington, New York, and Chicago. Because planes could not take off, those airports were handling only 40% to 50% of the traffic that they normally handled – the rest were on the ground waiting.

It took well into the afternoon before air traffic returned to normal.

[1] The U.S.' Federal Aviation Agency.

[2] Fliers hit delays, cancellations after network outage, *Market Watch*; November 19, 2009.

The FTI failure took down not only flight planning but also Notices to Airmen (Notams) as well. Notams advise pilots of temporary problems such as airport runway closures and navigation aid outages. It also cut off communications between many different FAA computers that normally coordinate flight delays to minimize the time that airliners must wait. The relaying of weather information to pilots was affected. However, radar and air/ground communication services were not affected. There was no safety problem.

The FAA was forthright in its coverage of the outage. An FAA statement said:

> "At approximately 5:00 am EST, a router problem disrupted a number of air traffic management services including flight-plan processing. The problem was resolved at approximately 9:00 am EST. Air traffic control radar and communication with aircraft were not affected during this time, and critical safety systems remained up and running.
>
> The failure was attributed to a software configuration problem within the FAA Telecommunications Infrastructure (FTI) in Salt Lake City. As a result FAA services used primarily for traffic flow and flight planning were unavailable electronically.
>
> The National Airspace Data Interchange Network (NADIN), which processes flight planning, was affected because it relies on the FTI services. During the outage air traffic controllers managed flight plan data manually and safely according to FAA contingency plans.
>
> There is no indication the outage occurred as a result of a cyber attack.
>
> System wide delays and cancellations will continue to be assessed throughout the day.
>
> A team of FAA technical and safety experts is already investigating the outage. FAA Administrator Randy Babbitt is meeting with representatives from Harris Corporation, the company that manages the FTI, to discuss system corrections to prevent similar outages in the future."

The Flight Planning Process

A flight plan is filed with the FAA for every commercial flight and for any general aviation flight that will be flying in bad weather. It informs the air traffic control system of the type and speed of the aircraft, the requested route and altitude, the time of departure, the duration of the flight, the pilot, and the fuel and the number of persons on board.

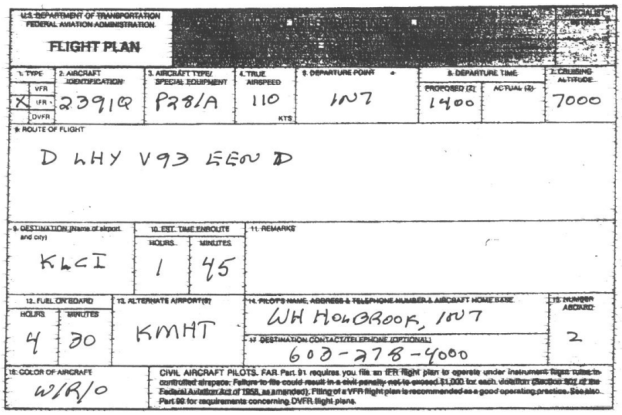
U.S. DEPARTMENT OF TRANSPORTATION
FEDERAL AVIATION ADMINISTRATION

FLIGHT PLAN

1. TYPE	2. AIRCRAFT IDENTIFICATION	3. AIRCRAFT TYPE/ SPECIAL EQUIPMENT	4. TRUE AIRSPEED	5. DEPARTURE POINT	6. DEPARTURE TIME PROPOSED (Z)	6. DEPARTURE TIME ACTUAL (Z)	7. CRUISING ALTITUDE
X IFR	23919	P28/A	110 KTS	1N7	1400		7000

8. ROUTE OF FLIGHT: D WHY V93 EEN D

9. DESTINATION (Name of airport and city)	10. EST. TIME ENROUTE HOURS	MINUTES	11. REMARKS
KLCI	1	45	

12. FUEL ON BOARD HOURS	MINUTES	13. ALTERNATE AIRPORT(S)	14. PILOT'S NAME, ADDRESS & TELEPHONE NUMBER & AIRCRAFT HOME BASE	15. NUMBER ABOARD
4	30	KMHT	WH HIGHBROOK, 1N7	2

16. COLOR OF AIRCRAFT: W/R/O

17. DESTINATION CONTACT/TELEPHONE (OPTIONAL): 603-278-4000

CIVIL AIRCRAFT PILOTS. FAR Part 91 requires you file an IFR flight plan to operate under instrument flight rules in controlled airspace. Failure to file could result in a civil penalty not to exceed $1,000 for each violation (Section 902 of the Federal Aviation Act of 1958, as amended). Filing of a VFR flight plan is recommended as a good operating practice. See also Part 99 for requirements concerning DVFR flight plans.

The FAA computers may modify the flight plan based on other traffic and preferred routes, and the approved flight plan is returned to the pilot.

NADIN

What is NADIN?

The FAA's NADIN system processes flight plans for every filed flight in the U.S. Air carriers enter flight plans directly into the NADIN system, which approves them and then forwards them over the FTI network to the appropriate air traffic control facilities. (General aviation pilots generally enter flight plans over the Web or phone them into a flight service specialist, who enters the flight plans for them.)

NADIN is a 24-year-old system managed by Harris Corp. It has redundant centers in Atlanta and Salt Lake City. These centers back each other up. Should one fail, the other is intended to carry the full flight-planning load.

However, with the FTI backbone network down, no flight plans could be communicated to either NADIN system. Consequently, air carriers were reduced to faxing or emailing their flight plans to FAA personnel, who then entered them manually into the NADIN system. With a volume of 50,000 flight plans per day, this was a cumbersome and time-consuming process. Flight plans could not be approved, and aircraft could not move.

NADIN's Failure History

There have been at least nine failures in flight communications systems in the past three years. Though the November 19th failure could not be attributed to the NADIN system (it was an FTI fault), NADIN has not been without its problems.

Friday, June 8, 2007

On the morning of Friday, June 8th, 2007, the Atlanta NADIN center failed. As planned, all flight planning activity was transferred to the Salt Lake City NADIN center. However, this overloaded the Salt Lake City Center; and it, too, failed. Though the centers were returned to service in late morning, delays ranging up to four hours lingered into the late afternoon. Tens of thousands of passengers were inconvenienced.

May, 2008

In May of 2008, the FAA's NOTAM system, related to NADIN, failed and was down for almost a day. The crash was caused by a disk failure in an end-of-life Sun system. Failover to a backup system was attempted, but it failed also. It turns out that the data corruption caused by the primary disk failure also corrupted the backup disk (RAID arrays were not being used). Embarrassingly, the FAA had a replacement system onsite; but it had not yet been installed.

Tuesday, August 26, 2008

On Tuesday, August 26, 2008, at about 1:30 pm, the Atlanta NADIN facility crashed. All flight-planning activity was routed to the Salt Lake City center. Flight plans that had been filed with Atlanta had to be refiled. However, the total flight-planning activity so overloaded the Salt Lake City center that it, too, once again crashed. As of mid-afternoon, the FAA stopped accepting new flight plans. The delays became so bad at all major airports that the FAA stopped predicting flight delays. Air travel did not return to normal until that evening. The Atlanta center was not returned to full service until 11 am the next day. Its failure was blamed on a failed software upgrade.

FAA Telecommunications Infrastructure (FTI)

What is FTI?

The FTI system is an IP backbone that provides communication services between all FAA facilities.[3] It replaces seven aging legacy telecommunications systems and provides routing, network monitoring, and security services for voice, data and video communications to more than 4,000 FAA and Defense Department facilities.

[3] FAA Telecommunications Infrastructure (FTI), Harris Corp. white paper.

FTI began phasing over in 2002, carrying a price tag of $2.4 billion. It currently supports 50,000 users using integrated satellite and microwave links over optical networks with over 100,000 manageable devices.

It was an FTI failure that took down the NADIN flight planning service on November 19th. It, too, has had a less-than-stellar failure history.

On May 24, 2005, an FTI failure at O'Hare Airport in Chicago delayed over 100 flights. The problem was caused by an incorrect configuration of communication equipment.

On September 21, 2007, all flights coming into Los Angeles were delayed due to a computer malfunction. All flights within a one-hour flight time of Los Angeles were put on ground hold until the problem was corrected.

On September 25, 2007, all FTI services were lost at the Memphis Air Route Traffic Control Center (ARTCC), one of 20 enroute control centers in the contiguous United States. 566 flights were delayed for up to several hours. The problem was in a failure of an optical ring network that is supposed to be fault tolerant (traffic can flow in either direction around the ring).

On November 9, 2007, all primary and alternate FTI services were lost at the Jacksonville ARTCC, delaying almost 100 flights.

NextGen – The FAA's Next Generation System

The U.S. air traffic control system is old. It is built on the radar and radio technologies of the '50s; and it is becoming increasingly difficult for it to handle today's air traffic, which is expected to double or triple over the next one to two decades. It is written in obsolete languages, it is difficult to upgrade, and hardware and software crashes are becoming increasingly common.

So what is the FAA doing about this? Its answer is NextGen – the next generation air traffic control system that will be phased in, with the final goal to be up and running in full by 2025.

What Is NextGen?

NextGen is a transformation of the entire U.S. national air transportation system.[4] It replaces legacy ground-based navigation (radar, navigation aids) and ground/air voice communication with satellite-based technology and digital communication. It comprises the following components:

ADS-B

Automatic Dependent Surveillance Broadcast (ADS-B) is the backbone of the NextGen system. It uses GPS satellite coordinates to provide controllers and pilots with much more accurate position information than is available today with radar surveillance.

Transponders located in every plane broadcast their locations and altitudes to other aircraft in the vicinity and to air traffic control facilities. Pilots are given a digital display of nearby traffic with collision warnings.

The greatly improved positional accuracy allows separation between planes to be substantially reduced, thus increasing the capacity of our airspace. In addition, GPS navigation allows planes to fly direct routes rather than having to follow predefined airways. This will not only save time but fuel as well.

[4] NextGen Fact Sheet, *FAA white paper*; February 14, 2007.

SWIM

SWIM, the System Wide Information Management system, will provide the communications infrastructure for NextGen. It uses an extension of FTI as its foundation.

Digital Communications

Digital communications will replace much of the voice traffic between pilots and controllers. Though voice communications will still be used for immediate requirements such as collision alerts, other information such as weather, neighboring traffic, and route changes will be communicated directly to cockpit displays.

Network Enabled Weather

Network Enabled Weather will provide a common digital weather picture across the National Airspace System (NAS). It will use tens of thousands of global weather observations and sensors from ground, airborne, and space sources. It is expected to cut weather delays in half.

NAS Voice Switch

The NAS Voice Switch will replace seventeen legacy voice switches that are over twenty years old. It will provide the foundation for all air/ground and ground/ground voice communications.

NextGen Advantages

Among the many advantages that NextGen will bring to the U.S. national aviation transportation system are:

- Flight safety will be improved due to increased pilot awareness of nearby traffic.
- Flights will be able to go direct rather than have to follow predefined airways (the current system, which is akin to an airborne interstate highway system).
- There will be fewer weather delays.
- There will be fewer in-flight delays due to congestion.
- Airport operations will be improved. The same satellite technology used in the air will be used at airports to reduce separation and improve safety.
- It will eliminate billions of dollars of annual costs to the U.S. economy due to gridlock in the skies.

Where Does NextGen Stand?

This is the disappointing part. NextGen is scheduled to be in place by 2025 at a cost of $22 billion. Many of its features will be showing up in the 2012 to 2018 time frame (ADS-B is already being tested in Alaska and by UPS).

However, the FAA's budget authorization expired in 2007. Since then, it has lived on a succession of temporary budget extensions and acting administrators.

In effect, the FAA has been a victim of its own success. As President Lyndon Johnson said, "If it ain't broke, don't fix it." Congress has unfortunately been following that dictum. Progress is being hampered by two main funding issues:

- The split of service fees between commercial and business aviation.
- The cost/benefit analysis of the airlines equipping their planes with the new technology.

Until NextGen gets fully funded, we will have to live with the U.S.' current tottering air traffic control system.

Lessons Learned

One obvious lesson to be learned from the FAA's experiences is that each node in a redundant system must be configured to handle the full application load should the other node fail. The repeated failure of the Salt Lake City NADIN Center to handle the full flight-planning load following an Atlanta Center failure is, on the face of it, inexcusable.

Another observation with questionable merit is that mission-critical, 24x7 systems that are decades old can be held together with baling wire and chewing gum. But for how much longer? The FAA touts 99.74% uptime. Do you want to trust your life to systems with less than three 9s of availability? If the U.S. Congress waits long enough to fund FAA future development, FAA system availability may approach two 9s.

Finally, the old tongue-in-cheek saying may be more true every day that Congress puts NextGen on hold:

> "If you have time to spare, go by air."

Acknowledgements

In addition to the references cited above, material for this article was taken from the following sources:

FAA glitch shines spotlight on troubled telco project, *Computer World*; November 19, 2009.
FAA hit with network glitch; Flight plans go manual, *ZDNet*; November 19, 2009.
Flight delays linger after FAA computer failure, *USA TODAY*; November 19, 2009.
F.A.A. Computer Problem Snarls Flights, *The New York Times*; November 20, 2009.
FAA Computers delay hundreds of flights, *CNN*; August 26, 2009.
F.A.A. Computer Trouble Causes Flight Delays, *The New York Times*; August 27, 2009.
Flight Delays Caused by Computer Failure, FAA Says, *The Washington Post*; August 26, 2009.
FAA computer glitch causes flight delays, *USA TODAY*; June 9, 2007.
FAA Computer Failure Disrupts Travel, *CBS News*; June 8, 2007.
FAA Continues to Ignore Flawed FTI Program, *PR Newswire*; May 24, 2005.
More Never Agains, *Availability Digest*; August 2008.
So You Think Your System is Robust?, *Availability Digest*; August 2007.
So You Think Your System is Reliable?, *Availability Digest*; January 2008.
FAA's Next Generation, *Air Transport World (pg. 49)*; February 2009.
FAA shifts focus for Next Generation air traffic control to near-term possibilities, *GCN*; March 19, 2009.