U.S. Department of Transportation

Office of the Secretary
of Transportation

Inspector General

Office of Inspector General
Washington, DC 20590

June 17, 2010

The Honorable James L. Oberstar
Chairman, Committee on Transportation
and Infrastructure
United States House of Representatives
Washington, DC 20515

The Honorable Jerry F. Costello
Chairman, Subcommittee on Aviation
Committee on Transportation
and Infrastructure
United States House of Representatives
Washington, DC 20515

Dear Chairmen Oberstar and Costello:

On November 19, 2009, 129 of the Nation's air traffic facilities experienced an outage of the Federal Aviation Administration's (FAA) telecommunications that delayed thousands of travelers and grounded hundreds of flights nationwide. As you requested, we conducted a review to (1) identify the cause of the FAA Telecommunications Infrastructure (FTI) outage, (2) review FAA's corrective action plan to prevent future critical outages, (3) examine FAA's ability to oversee FTI and the contractor, and (4) identify oversight vulnerabilities or best practices of other critical systems in the National Airspace System (NAS) owned or operated by the private sector.

Summary

The November 2009 FTI outage raised questions about FAA's and the prime contractor's (Harris Corporation) ability to effectively manage FTI as well as the integrity of the network design and whether it can support initiatives for the Next Generation Air Transportation System (NextGen). Specifically, we found:

- A Harris technician incorrectly configured an FTI router (which directs air traffic data, such as flight plans, through the network) at Los Angeles Center. The error caused the FTI network to send air traffic data on the wrong routes, which blocked

approximately 75 percent of the routes across the FTI fiber optic network. Service restoration was delayed for 5 hours because an automatic tool that alerts technicians to network failures and their locations did not work as intended. Therefore, Harris technicians could not readily identify the source of the problem, which could have minimized the impact of the error on the NAS.

- FAA and Harris have taken corrective actions to prevent another critical outage. These include deactivating the FTI legacy network and router configuration design (which are now obsolete) that caused the outage and fixing the automatic alert tool. However, both FAA and Harris officials acknowledged that there is still a risk of critical outages as new NextGen services are added to FTI's new fiber optic network.

- FAA's oversight of the FTI contractor could have been more effective. FAA was unaware that Harris officials had configured the network in error and made other procedural errors. In 2008, we recommended that FAA develop improved controls over the contractor's FTI equipment configuration and take steps to prevent unscheduled outages and restore them on time to improve service reliability.[1] While FAA agreed to take action, we found it still has problems ensuring FTI services are restored within contractual requirements. To its credit, FAA plans to address this and other FTI issues, in response to the findings of an independent review panel convened to investigate the November outage.

- FAA has not developed best practices to oversee NAS systems not owned or operated by the Government even though it is increasingly shifting more acquisitions and services to the private sector to reduce costs. Moreover, FAA's internal reports have concluded that FAA and Harris need to identify FTI network vulnerabilities—a critical step highlighted by the fact that the program office for the Automatic Dependent Surveillance-Broadcast (ADS-B), a key NextGen system, has decided against using FTI due to network reliability and security concerns.

Background

The FTI network provides voice, data, and video communications that support operations at more than 4,000 FAA facilities and remote sites nationwide, as well as some Department of Defense facilities. The network provides more than 25,000 telecommunications circuits and service delivery points, upgraded switching and routing services, and centralized network monitoring and control. As a major contributing system to NextGen, initiatives are underway to further upgrade the FTI fiber optic network to increase capacity, or bandwidth;[2] provide greater flexibility;

[1] OIG Report Number AV-2008-089, "FAA's Progress and Challenges in Meeting FTI Transition Goals," September 30, 2008. OIG reports are available on our website: www.oig.dot.gov.

[2] *Bandwidth* refers to a data rate measured in bits per second, such as network throughput (i.e., the average data rate of successful data transfer through a communications path).

and continue to reduce latency.[3] FAA's mission for FTI is to achieve an integrated suite of products, services, and business practices to better meet the telecommunications needs of the NAS. With FTI, FAA will transition from traditional dedicated circuits to on-demand service where appropriate. According to FAA, these services will provide lower unit cost, more efficient bandwidth utilization, improved information security, and state-of-the-art business processes and technology.

In July 2002, FAA awarded an Indefinite Delivery Indefinite Quantity contract to Harris to begin transitioning FTI into the NAS and to provide management and support functions for the FTI network. FAA does not own the network, and its contract with Harris is essentially a 15-year lease that expires in 2017 and covers the cost of acquiring, operating, and maintaining the FTI network. The contract has a current maximum value of $1.4 billion and a ceiling amount of $3.5 billion, with no limits on quantities, meaning there is almost no limit on Harris' ability to sell additional services to FAA until the ceiling is reached. FAA has expended about $1.2 billion on the contract and currently spends about $146 million annually on the FTI program.

Harris Configuration and Procedural Errors Caused the FTI Outage and Delayed Service Restoration

On November 19, 2009, FAA's FTI system experienced a NAS-wide outage while Harris was transitioning FTI Internet Protocol (IP) services from a legacy network to the new FTI Operations Internet Protocol (OPIP) fiber optic network.[4] Our review found the root cause of the outage occurred when a Harris engineer incorrectly configured one of several temporary routers (known as Logical Transition Bridges) between the old and new networks. These "bridges" were installed at the 26 core FTI sites with different route maps to allow continued flow and separation of air traffic information (e.g., flight plans and weather data) across both networks and prevent routing problems during the network transition. The configuration error essentially went unnoticed and ultimately created a "domino effect" across the FTI network when all circuits on the new fiber optic network failed, resulting in 820 flight delays.

Specifically, all Air Route Traffic Control Centers (ARTCC), Network Enterprise Management Centers, and the FAA Command Center did not have the data they rely on to manage flights when multiple FAA systems were affected by the FTI outage. These included the following:

[3] *Latency* expresses how much time it takes for data to get from one designated point to another. Excessive latency creates bottlenecks that prevent data from filling the network pipe, thus decreasing effective bandwidth.

[4] *IP* typically uses various routing and communications procedures and communicates with multiple sites simultaneously. IP services do not use the actual Internet; they just follow similar procedures. With FTI, IP services were previously carried over the Asynchronous Transfer Mode circuits. In the new fiber optic network, IP services are now carried over optical circuits, which will increase bandwidth and reduce data latency.

- *Enhanced Traffic Management Service*, which provides information on a national level to predict traffic surges, gaps, and volumes based on current and anticipated airborne aircraft.
- *National Airspace Data Interchange Network*, which distributes flight plan data, weather information, and other air traffic control messages within the NAS.
- *National Defense Program Surveillance*, which provides surveillance data from FAA long- and short-range radar to Department of Defense and Department of Homeland Security agencies.

The error that led to the outage and flight delays was first made on October 21, 2009, with the temporary router installed at Los Angeles Center, during the Salt Lake City Center's transition to the fiber optic network. The error essentially made the FTI network "believe" that air traffic information could be routed either locally within Los Angeles airspace or between Los Angeles and Salt Lake City airspace. When the link between the two facilities was reestablished after installation, the authorized routes between the networks and Los Angeles and Salt Lake City resumed operations, masking the error, which remained dormant in the system.

On November 19, 2009, a Harris engineer took down the link between Los Angeles and Salt Lake City Center to place a new router in service on the FTI fiber optic network at Los Angeles. The new router also contained the configuration error, since the route map (which failed to define the authorized traffic between the two networks) was copied from the replaced router. The error occurred because the engineer failed to append a needed "no" statement to the configuration file after replacing the router and before reestablishing the link between the two facilities. As a result, when the link went live, the network sent air traffic data on the wrong routes between the airspace locations and resulted in multiple outages. According to Harris officials, the failure to append the "no" statement caused the network to begin using the wrong routes and sending network traffic through Los Angeles, which impacted services at 129 facilities instead of the 21 that should have been impacted. The error also broke down the separation between the old and new networks, causing routing errors when the FTI network began sending all network traffic through Los Angeles Center to Salt Lake City Center. The link between the two facilities did not have the bandwidth to support all traffic, and approximately 75 percent of routes were blocked across the FTI fiber optic network for 5 hours while Harris tried to find the problem.

Harris eventually started restoring services after technicians discovered that taking a core router offline at Salt Lake City eased the problem. The FTI network had over-utilized core routers at that site when channeling all the additional traffic over from Los Angeles. Regardless of the initial error, the impact on the NAS could have been minimized if Harris had identified the source of the problem sooner. We found the following procedural breakdowns contributed to delayed recognition of the problem:

- An automated tool that alerts Harris personnel to network failures when any router central processing unit exceeds 60 percent of its utilization for 10 minutes did not work as intended. This was due to a configuration error that suppressed the wrong alarms. As a result, a filter meant to silence recurring alarms for only a specific router at FAA's Technical Center caused the alarms to be filtered for all routers on the FTI network.
- Once Harris engineers were made aware of the outages, they started pursuing the wrong problem. They initially looked at backing out configuration changes to the network routers at the Los Angeles ARTCC and the Herndon, Virginia, Command Center, since the two most recent maintenance actions had been at these sites.
- Harris mistakenly thought another router supporting Los Angeles ARTCC was the problem. Harris engineers were having trouble accessing the router remotely and sent a technician to reset the router manually. Harris later found that this router was not the problem, but this was a temporary distraction that further delayed service restoration.
- Harris could not initially determine whether FTI was being over-utilized. Spot checks of core routers at other sites (but not Salt Lake City Center) did not indicate a general high-utilization level on other routers.

FAA Has Taken Corrective Actions To Prevent a Recurrent Outage, but the Risk of Future Critical Outages Remains

FAA has eliminated the possibility that the events leading to the November 2009 critical outage could reoccur. As of December 13, 2009, all IP services were cutover from the legacy network to the new FTI fiber optic network, thereby eliminating the need for the temporary routers between the two networks. As a result, Harris has deactivated the legacy networks and the route map configuration design used during the transition. Harris officials state they have also corrected the problems with the automatic alert tool. In addition, Harris is working to require support personnel to actively monitor the FTI network during future maintenance releases that involve installing or replacing core routers.

However, both FAA and Harris officials acknowledge that an inherent risk of critical outages remains since Harris plans to transition more existing services and new NextGen services to the FTI fiber optic network. Additionally, Harris will face challenges and risks as it continues to design and build out the new fiber optic network, which is expected to support future services.

Risks of FTI Outages for Existing Operational Services

As of December 2009, Harris reported that there are a total of 20,982 services operating on the FTI network. While FAA has transitioned all 1,808 IP services supporting FAA flight plans to the new FTI fiber optic network, additional risks will

be introduced when the Agency begins transitioning other FAA services, such as En Route Automation Modernization (ERAM) and NextGen platforms, which also require IP technology to operate.[5] Moreover, there are 19,174 existing services that may be transitioned to the new FTI fiber optic network. According to FAA, transitioning these remaining services is important because FTI now connects older FAA systems that provide safety-critical voice and surveillance radar information by utilizing Time Division Multiplex (TDM) technology, instead of IP services.[6] Table 1 describes the current technologies supporting FTI services, and the percentage of services operating over the two FTI networks.

Table 1. FTI Services Operating on the Fiber Optic and Legacy Networks (as of December 2009)

Technology Supporting FTI Services	Backbone Infrastructure	Number of Services	Percent of FTI Services
IP	FTI Fiber Optic Network	1,808	8.6%
TDM Point-to-Point	FTI Legacy Network	19,174	91.4%
Total		**20,982**	**100%**

Source: OIG Analysis of Harris and FAA Briefings

FAA required Harris to begin transitioning about 1,492 of the remaining services to the FTI fiber optic network in April 2010; the transition will last several months. However, FAA has yet to determine whether the remaining 17,682 services will be transitioned to the new FTI fiber optic network but plans to conduct analyses to determine feasibility.

Moreover, critical voice and surveillance data communications services continuing to operate over the FTI legacy network (e.g., AT&T and Sprint networks are used to support the 17,682 FTI services) are still vulnerable to outages because Harris has little control over how these networks are managed. For example, according to a November 2009 internal FAA study, Harris confirmed either the complete lack or inadequate proof of diversity[7] between FTI primary and alternate network paths at several critical facilities—including FAA's Technical Center and Baltimore Air Traffic Control Tower. According to Harris officials, FTI sites that supposedly had diversity no longer had it after AT&T or Sprint made upgrades to their network. Therefore, FAA's installation of the new FTI fiber optic network was also an effort to

[5] The $2.1 billion ERAM program will replace the existing hardware and software at facilities that manage high-altitude traffic.

[6] FTI remaining services use Time Division Multiplex (TDM) technology. Specifically, they require a circuit connection between two end points that utilize TDM technology to transport voice and data. The TDM services are considered critical because they transport FAA's critical voice and radar data information in this manner throughout the NAS.

[7] For the purposes of this report, we refer to diversity as instances where there is not adequate separation between FTI primary and alternative paths. We did not examine FTI's overall architecture or design.

maintain the diversity and redundancy of critical FTI services through a dedicated infrastructure where only the Government (and its contractor) could provide and support services to avoid the diversity violations encountered with FTI services operated over the legacy network.

Risks Regarding FTI Support for NextGen Initiatives

According to FAA, the new FTI fiber optic network also establishes the foundation for the telecommunications architecture required to support NextGen initiatives. However, FTI program officials stated that they have yet to determine whether the fiber optic network can support NextGen and plan to further assess the network to determine its integrity and ability to support these services. As FAA continues to modernize the NAS, increased usage of IP services is expected. For example, emerging NextGen technologies, such as the System Wide Information Management and Data Communications, will be IP-based and will be implemented on FTI. However, FAA has yet to establish a timeframe for implementing these services.

Moreover, concerns about FTI have already caused the ADS-B program office to decide against using FTI to provide its telecommunications services. The ADS-B contractor (ITT Corporation) stated that ADS-B requires high service availability and low latency for services to be provided as proposed in the contract. For example, the ADS-B service requires that an ADS-B report be delivered to Air Traffic Control automation within 700 milliseconds of receipt at a radio station and that services can be restored within 6 seconds (e.g., safety-critical capability.) The loss of this capability raises to an unacceptable level the risk associated with providing safe and efficient local NAS operations. At the time of the contract award in 2007, ITT did not believe FTI could meet these requirements.

FAA also faces challenges as Harris continues to design and build out the FTI fiber optic network to make it more stable and capable of supporting advanced NextGen technologies. For example, Harris is upgrading several locations that support eight Centers (i.e., facilities that manage high-altitude traffic.) While the upgrades, slated for completion in fiscal year 2011, could improve the reliability and efficiency of network traffic, they could also introduce risks to network operations if not properly planned and managed. To address FAA's safety requirements and provide a back-up capability, Harris is also building out the FTI fiber optic network to support FAA's Business Continuity Plan (BCP) initiative. The BCP concept effectively creates a temporary Center at the William J. Hughes Technical Center in the event of a long-term Center outage due to natural disasters (e.g., storms, fires, etc.). Harris has installed the necessary equipment at the Technical Center to implement the spare Center concept. However, this is considered only an interim step for continuity planning as FAA must still transition thousands of small remote sites and about 300 larger sites to the new FTI fiber optic network. This will allow Harris to reroute

services more efficiently in the event of a disaster. However, FAA has yet to determine if or when these sites will be transitioned.

The new FTI fiber optic network also has vulnerabilities with potential outages and security risks that will require sustained management and oversight. The new network is designed to reduce the risks of future widespread failures because it is partitioned by en route airspace using Border Gateway Protocols (BGP), and the separate domain will assist in containing any network anomalies to specific delegated airspace. However, according to the National Institute of Standards and Technology, the use of BGP does not come without risks.[8] For example, if the BGP routing protocol fails to carry out the routing function, portions of the network may become unusable for periods of time—ranging from minutes to hours. While most of the risks to BGP come from accidental failures, there is also a security risk that attackers could disable part or the entire network. Therefore, it is imperative that FAA and Harris institute proper controls to ensure the safety and security of the FTI fiber optic network.

FAA's Oversight of FTI Vulnerabilities and the Contractor Should Have Been More Proactive

FAA's oversight of FTI and Harris was not as effective as it should have been. Although FAA has three representatives on-site at the Harris facility to monitor FTI outages and managers at the two FAA Network Enterprise Management Centers, they have limited ability to oversee FTI. FAA tends to have a reactive, rather than proactive approach to assessing network vulnerabilities. For example, at the time of the outage, neither FAA nor Harris could readily identify the root cause of the outage or what corrective actions were needed to resolve it.

In 2008, we recommended that FAA improve its processes and procedures for restoring FTI outages within contractually established timeframes to meet reliability, maintainability, and availability (RMA) standards. While FAA agreed to take action, it continues to have problems ensuring FTI services are restored within contract requirements. FTI services vary depending on the RMA levels. For example, RMA-1 services such as radar must be restored within 6 seconds. However, RMA-4 services, such as En Route Air to Ground Communications, account for about 80 percent of FTI services and must be restored within 3 hours. At the time of our audit, we found that an average of 7 percent of FTI services experienced outages and were not restored on time. While this may seem like a small percentage, the trend has not improved, with just over 8 percent of FTI services not meeting availability requirements as of December 2009 (see table 2).

[8] National Institute of Standards and Technology, Special Publication 800-54, Border Gateway Protocol Security, June 2007.

Table 2. Percent of Individual Services Not Meeting Minimum RMA Requirements

RMA Level*	Meeting Requirement	Not Meeting Requirement	Total	Percent Not Meeting Requirement
RMA 1	313	21	334	6.29 %
RMA 2	1,543	73	1,616	4.52 %
RMA 3	861	61	922	6.62 %
RMA 4	19,188	1,888	21,076	8.96 %
RMA 5	1,937	87	2,024	4.30 %
RMA 7	134	9	143	6.29 %
All Services	**23,976**	**2,139**	**26,115****	**8.19 %**

Source: Total of 26,115 FTI services included in the individual service information data in the December 2009 FTI Performance and Management Report. This is based on the past 12-month reporting period.

To its credit, after the November 2009 outage, FAA established several independent groups to review the cause of the outage. These groups include FAA's Safety Event Response Team, an internal group sponsored by the FTI program office, and another group of experts chartered by the FAA Administrator.[9] The Administrator's group was further tasked to review the integrity of the FTI architectural design, FTI's ability to support NextGen initiatives, and any potential threats of future, critical outages. This group was also asked to examine whether Harris has adequate personnel, processes, and technology deployed to provide a robust communications network with adequate security and backup capabilities to meet FAA's needs. The FAA Administrator's group issued its report on April 20, 2010, noting several steps FAA should take to improve oversight of FTI and its contractor. Many of these bolster our analysis of actions needed to improve oversight of NAS systems owned and operated by the private sector, which is further discussed below.

FAA Has Not Developed Best Practices To Oversee NAS Systems Not Owned or Operated by the Government

Over the last several years, FAA has sought to transition more acquisitions and services to the private sector to reduce cost. For example, FAA transitioned Flight Services Stations (FS-21) and FTI programs and is deploying the ADS-B infrastructure, which will be a service-based system owned and operated by the private sector. FAA plans to rely on a similar approach to develop and implement Data Communications, which will be another multibillion-dollar investment. Despite this shift in its implementation strategy, FAA has not assessed best practices for overseeing systems not owned or operated by the Government. The ATO's Chief

[9] The panel convened by the FAA Administrator was made up of the following participants: the Chief Information Officers from FAA, FAA's Air Traffic Organization, and DOT; FAA's Assistant Chief Counsel for Acquisition and Commercial Law; the Assistant to the President and U.S. Chief Technology Officer; the Chief; Executive Officer of Noblis, Incorporated; and the former Director of Command, Control, Communications, and Computer Systems.

Operating Officer has stated that a paradigm shift in FAA's oversight is needed for such systems.

Based on discussions with FAA officials and our review of the two independent reports recently issued, FAA needs to be more proactive in assessing and addressing FTI network vulnerabilities. FTI program officials acknowledged the need to develop an in-house capability to monitor FTI network performance. To address this vulnerability, FAA states it is developing a new automated toolset to monitor FAA systems operating on FTI's OPIP fiber optic network. While the toolset does not allow FAA to directly monitor the status of the FTI network, this is a good first step; however, it is too early to assess its adequacy. Ultimately, a new oversight approach for NAS systems provided and serviced by the private sector is needed. We identified the following actions FAA should consider for FTI and other systems it does not own or operate—many of which could also address several areas noted in the independent review groups' recommendations listed in the enclosure to this letter:

- Ensure sufficient in-house expertise by providing training and experienced staff to the FAA team charged with oversight of the contractor.
- Use modeling, simulation, and network monitoring tools to examine failure mode simulation, routing configuration changes, and alarms for unexpected and significant routing changes.
- Ensure the use of a quality management system that includes checklists, peer review, and validation/verification for system changes.
- Create a government/industry team responsible for identification of vulnerabilities, recommendations to improve survivability, and research into new and improved methods of building high-availability networks.
- Require independent periodic reviews of existing and proposed network architectures.

Conclusion

FTI plays a significant role in the U.S. aviation system—the largest and safest system in the world. Maintaining this safety record and transforming the system to meet future demand depends on reliable equipment and technology, which will become more complex as FAA continues the transition to NextGen. While the events that led to the November 2009 outage were legitimate causes for concern, FAA has taken steps to prevent a recurrence, and other actions are planned. As FAA continues to implement FTI, it is imperative that the Agency exercise due diligence and become more proactive in overseeing the contractor's performance and addressing FTI network vulnerabilities. It will be important for FAA to follow through with plans to review FTI architecture to assess whether it can support NextGen. At a minimum, FAA should use and document best practices to more effectively oversee Harris and

FTI vulnerabilities. Until FAA fully addresses this issue, the potential for oversight lapses and service outages remains.

We are encouraged by FAA's announcement on April 20, 2010, that it is accepting the findings and recommendations of the independent review panel assigned to investigate the November 2009 outage. FAA is still determining the best way to implement the recommendations, which are aimed at improving overall FTI reliability as well as FAA's internal procedures for dealing with outages. Another review on the reliability of FTI to carry critical navigation, communication, and other NextGen services is pending. Therefore, we are not making any formal recommendations at this time. However, we will continue to monitor FTI and report on FAA's progress in addressing these issues as necessary.

We discussed the results of our review with the Director of Air Traffic Control Communications Services and incorporated his comments where appropriate. If you have any questions, please contact me at (202) 366-1959 or Matthew E. Hampton, Deputy Assistant Inspector General for Aviation and Special Program Audits, at (202) 366-0500.

Sincerely,

Calvin L. Scovel III
Inspector General

Enclosure

cc: Secretary of Transportation
Federal Aviation Administrator

Recommendations of the FAA Administrator's Independent Review Panel

The FAA Administrator's independent review panel assessing the cause of the November 19, 2009, outage issued its report on April 20, 2010. The report detailed a number of recommendations to improve the reliability of FTI and FAA's internal communications and procedures for dealing with an FTI outage.[10] The panel's recommendations included the following:

1. Consider using automated tools to implement router configuration changes and to support independent verification procedures.

2. Review maintenance operations and associated checklist design from a human factors and risk reduction perspective to help minimize the potential for human errors. Consider using the FAA's Aviation Safety (AVS) and external experts in this review.

3. Implement end-to-end situational awareness of the network, both Local Area Networks (LANs) and the FTI, as well as including appropriate applications.

4. Implement a capability to report network and application service outages and describe the impact to FAA customers (internal and external) using a common language.

5. Consider developing a functional model of the FAA's FTI network to simulate and test configuration changes and upgrades.

6. Consider a needs assessment of the FTI workforce staffing and skill levels to ensure adequate levels of network technical support at all times.

7. Consider modifying the FTI contract award fee and/or performance incentive structure based on the observations in this report.

8. Provide an alternate means for rapid and standardized entry of flight plan information into the National Airspace System (NAS) to mitigate failures in the flight plan filing system.

9. Evaluate the ADS-B and FTI network architectures to determine the viability of using each as potential back-up for selected services of the other.

10. Perform a review of currently identified essential services and categorize them according to priorities in support of NAS safety and capacity.

[10] "FAA Telecommunication Infrastructure Review Panel Report on November 19, 2009, Outage," Federal Aviation Administration, Washington, D.C., issued April 20, 2010." The full report and recommendations can be found at: http://www.faa.gov/air_traffic/publications/media/FTI_Phase1.pdf.