

SECURITY IMPLICATIONS FOR DATA COMM

Aniruddha Karmarkar

Lockheed Martin,

ani.r.karmarkar@lmco.com

Ravi Vaidyanathan

Telcordia Technologies Inc.

rvaidyan@telcordia.com

Abstract

The FAA Data Comm Integrated Services (DCIS) program is a key enabler of NextGen applications. The FAA plans to share the existing data radio currently used for supporting Airline operations, for Air Traffic Control (ATC) data communication between the controller and the pilot. The FAA plans to use existing commercial service provider (CSP) infrastructures to support the ATC data communication network service. This exposes the air traffic control message traffic to the CSP networks outside the FAA domain. Traditionally, FAA has operated all its communications within a relatively closed network environment. The DCIS program necessitates operation in a relatively open network environment where the FAA network is connected to the CSP networks which in turn are connected to US and international airline operation centers. While air to ground communication security is subject to collaborative standards evolution, the potential of distributed network attacks that may be launched from outside CONUS on the CSP networks needs to be addressed. This paper explores potential vulnerabilities of the ground network and possible techniques to mitigate attacks such as distributed denial of service attacks launched by malicious actors from CONUS or OCONUS locations on CSP network infrastructures.

Background

The FAA's Next Generation Air Transportation System (NextGen) will provide a more convenient and dependable air transportation system. As part of NextGen the FAA is planning to use VDL Mode 2 radios to support ATC services using data communications under the aegis of the Data Comm Integrated Services (DCIS) program. DCIS will supplement the currently used voice based communication system for ATC services. The CONUS-wide CSP air-ground network and the VDL Mode 2 radio on the aircraft is currently used for airlines operations. The FAA plans to share the

existing CSP communication infrastructure in CONUS for ATC traffic. The shared use of VDL Mode 2 radios in NAS opens up new cyber security vulnerabilities for ATC communications. The issue of network and cyber security using commercial infrastructure for data communications is a new paradigm for the FAA's current relatively closed networks and systems.

The FAA will use the current voice based communication system as a backup system for Data Comm. Data Comm supported FAA enterprise applications will enable trajectory based operations, resulting in increased use of automation for air traffic operations throughout the National Airspace System (NAS). Once Data Comm is widely deployed over next 10 to 15 years, the use of the backup voice network for air traffic operations will become infrequent.

The Data Comm program necessitates operation in a relatively open network environment where the FAA network is connected to the CSP networks which in turn are connected to US based and international airline operations. While air to ground communication security is subject to collaborative standards evolution, the potential of distributed network attacks that may be launched from inside or outside CONUS on the CSP ground networks must be considered.

The consequent cyber security vulnerabilities in this environment are not only new but were not previously anticipated. The existing code of federal regulations does not adequately address cyber security vulnerabilities for open collaborative networks. Consequently, there are no existing policies, certification criteria or procedures that provide assurances that cyber security vulnerabilities will not cause unsafe flight conditions. Unmitigated, these vulnerabilities could have a detrimental effect on flight safety.

In this paper we provide a brief overview of the Data Comm program in the context of potential security implications for the Data Comm supported

FAA applications, we then describe the potential network security and cyber security issues, and then provide potential solutions for addressing those issues.

Data Comm Overview

In this section, we provide an overview of the proposed FAA Data Comm network and outline some of the security considerations associated with the Data Comm network.

- Network components including A/G components such as the VHF Ground stations (VGS) and ground components such as the central message processor are shared between best effort AOC services and mission-critical ATC traffic. Such disparate network components are typically not certified (from a security perspective) to the same degree; potentially allowing malicious actors or misconfiguration in the airline

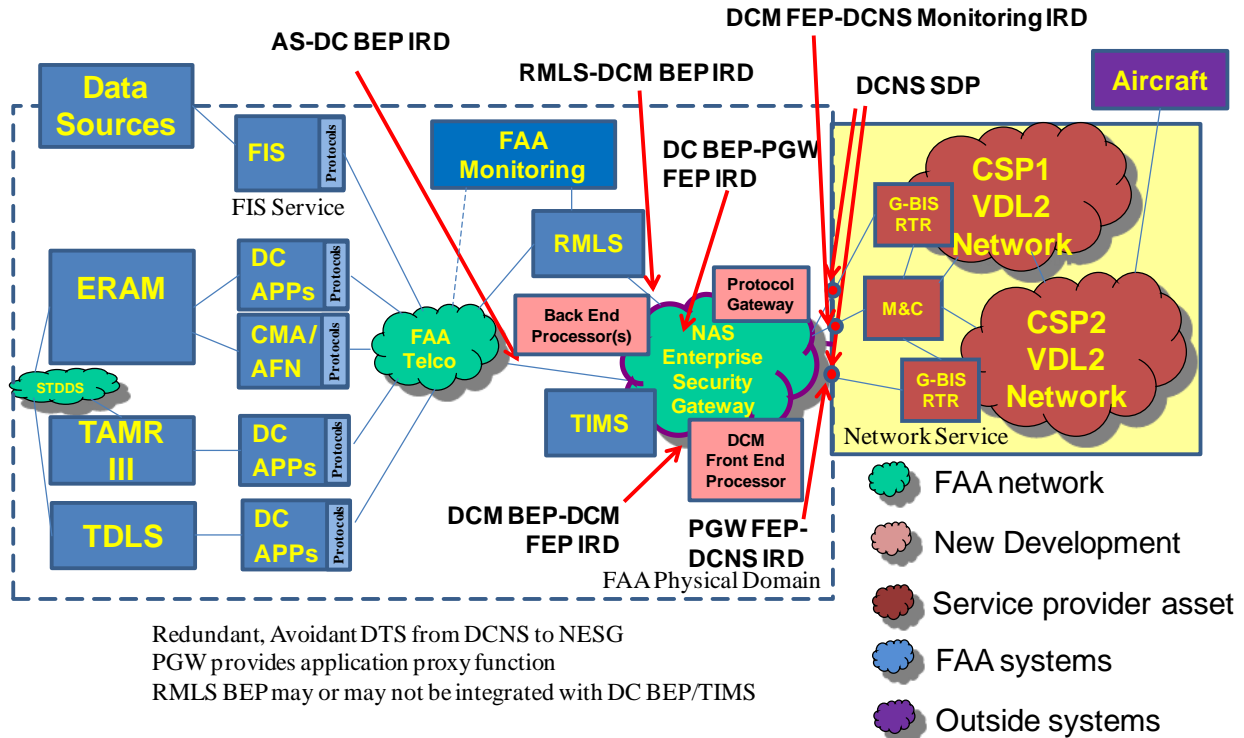


Figure 1 Data Comm Network High Level View

Figure 1 illustrates a high-level view of the Data Comm network showing the interconnections between the Airline Operations centers and the Air-Ground System (AGS) for transmitting airline operations (AOC) traffic. AGS traffic traverses a CSP central message processor, typically deployed in a redundant configuration. ATC messages traverse the CSP owned and operated AGS and then traverse the Ground Network System (GNS) before being interconnected to the FAA end systems over the FAA Telecommunications Infrastructure (FTI) network.

In general, we can make the following observations about the security considerations surrounding this network architecture:

operations centers to affect ATC services.

- The network architecture has certain critical component and capacity bottlenecks. In particular, the data rate supported by the VDL mode 2 air-ground networks is 31.5 kbps, which is smaller than the ground network circuits, which are typically, leased lines with aggregate bandwidths ranging from DS0-T1 (64 Kbps-1.5Mbps). This mismatch in bandwidth implies that a malicious or misconfigured end system may easily be able to flood a VGS subnet, resulting in a denial of service attack

affecting any aircraft serviced by that subnet.

- Average usage rates at a VGS station are typically an order of magnitude lower than the provisioned network circuits with bandwidths of (64 Kbps-1.5 Mbps). A successful network attacker can further exploit the mismatch between average usage and available capacity by opening up a small covert channel to infect network nodes and application systems within the CSP network and other adjacent networks with malware, and botnets.

Carefully considered usage of best-practice security recommendations can mitigate several of these threats and reduce the impact of compromised elements in commercial components such as the airline operations center.

Network and Cyber Security context for Data Comm

With the widespread use of networking, information technologies, and networking and mobile communication technologies, malicious actors are continually seeking out cyber vulnerabilities in systems and gaining access to sensitive and protected data.

Cybersecurity issues transcend the protection of personal data or networks from hackers or even organized crime. Cyber warfare is a major national security issue. Protecting the security and freedom of our networks is as critical as protecting freedom of the seas and space. It is difficult to establish evidence that proves beyond a reasonable doubt that a particular entity staged an attack. Not only is it difficult to identify and prove whether the attacker is sanctioned by a foreign government, it is also hard to distinguish between active direction by foreign officials and mere tolerance or lax enforcement. Consequently, accountability for cyberattacks is extremely difficult to determine. [1].

Recent cyber security incidents such as the widespread disclosure of highly classified diplomatic communication data by WikiLeaks, Stuxnet computer virus attack on Iran's Bushehr nuclear plant, the federal government Department of Energy (DOE) experiment to test the vulnerabilities of the

US power grids, and the reported vulnerabilities of the FAA air traffic control systems for a possible cyber attack indicates that the use of Data Comm for ATC communication should address cyber security vulnerabilities.

While most networks support a mix of packet and circuit services, the general trend is toward increasing use of packet networking. The most significant concerns with stable, secure packet based network operation are configuration or operator errors that open up backdoors for malicious hackers, and attacks against control plane protocols, including routing protocols (the easiest vulnerabilities to exploit).

The commercial internet consists of a large number of interconnected networks operated as different autonomous systems (AS). Border Gateway Protocol (BGP) is an Inter-domain Routing Protocol widely used for IP (Internet Protocol) networks. BGP is the IP stack equivalent of Inter-domain Routing Protocol (IDRP) protocol used for the ATN stack.

Routing protocols, generally the Border Gateway Protocol (BGP), determine forwarding paths between end hosts in the enterprise. BGP update messages are exchanged between peer AS networks. These update messages provide information on available paths and end destinations reachable through each AS. By design, each AS trusts the information it receives via BGP update messages. Spoofing refers to the forgery of identity information (e.g. IP addresses in packet headers), typically by malicious attackers in order to avoid detection, or prevent source identification. Spoofed or incorrect BGP update information from any network can easily propagate through multiple ASes over multiple paths to corrupt routing tables and can affect network denials of service.

Misconfigurations, Misdirections and spoofing

In 2008, Pakistan Telecom, under government order, tried to prevent in country access to YouTube. The telecom intentionally "blackholed" requests for YouTube videos coming from within Pakistan. In doing so, it mistakenly told the upstream global carrier that they should send all YouTube traffic to Pakistan Telecom [2]. The upstream carriers accepted the route update and sent it to other carriers in the world. Soon requests from around the world for YouTube videos were diverted to Pakistan,

specifically to a “dead end” or black hole. This resulted in partial service disruption of YouTube for several hours. The YouTube Pakistan incident highlights several points.

- (1) All routing protocols have inherent design weaknesses because of the implicit trust each node places in updates it receives and because of the transitive nature of that trust. Errors and misconfigurations can and do occur. They are the largest source of issues in network carriers.
- (2) Malicious adversaries have additional options that are more subtle than that used in the YouTube incident. Spoofing addresses, altering routing tables, DNS cache poisoning are some of the techniques commonly used for misdirection, and spoofing.

The use of routing protocols such as IDRP within the context of Data Comm implies potential for similar misconfiguration affecting ATC communications.

Distributed Denial of Service

A denial-of-service, or DoS, attack occurs when someone directs a large number of requests to a target server so quickly that the server can't respond, and the site becomes inaccessible. A distributed denial-of-service, or DDoS, attack occurs when hundreds or thousands of compromised computers are enlisted.

On April 27, 2007, the Estonian government moved a controversial Soviet-era World War II memorial from a square in the capital city of Tallin.[3] Weeks of cyber attacks followed, targeting government and private Web sites. Some attacks took the form of distributed denial of service (DDoS) attacks. Hackers used hundreds or thousands of computers and pelted Estonian Web sites with thousands of requests a second, boosting traffic far beyond normal levels. The Estonian series of well coordinated attacks became the largest cyber attack replacing Titan Rain[4] as the new largest cyber attack in the world launched by hackers on foreign computer systems. Till then, Titan Rain launched by foreign hackers, in 2003, on US computer systems; was the world's largest targeted cyber attack.

The Estonian cyber attacks were able to shut down some sites for a long time. The government didn't lose any infrastructure, but the events proved

extremely time consuming, expensive to combat and indicative of weaknesses in Estonia's cyber security.

The interconnection of international CSP and airline commercial networks to the Data Comm network exponentially increases the risk of potential DDoS attacks being launched against ATC infrastructure by compromised computers in the commercial domain.

Targeted Malware and Botnets

In many government agencies it is believed that the use of highly customized proprietary software application implementation prevents the systems from attacks normally observed in the public domain internet. The Stuxnet [5] computer virus attack demonstrated that this is a misplaced belief. Stuxnet is of particular interest in the context of the Data Comm program and NextGen applications. Stuxnet is part virus, part a botnet.

Botnets are a collection of compromised computer systems controlled by adversaries. Attacker's command and control typically uses multiple layers to obscure master sever identity. Each layer often includes fault tolerance capability to provide resiliency. Different layers and elements of attacker command and control are generally geographically distributed. Nodes in the command and control structure often rotate addresses, roles, domain names, etc. to hide their communications patterns. Adversaries use various techniques to avoid detection and attribution.

A major computer virus attack was launched on the Windows computers at Iran's Bushehr nuclear power plant. Stuxnet is not an ordinary, computer virus. It does not erase hard drives. It is not a typical botnet. It does not steal bank accounts.

Like any ordinary virus, Stuxnet installs itself in Windows computers as a “rootkit.” Once a virus gets installed, it has full control of the computer, and is completely invisible to the computer operator. But Stuxnet doesn't harm most computers. Stuxnet searches the computer for industrial control software (ICS). Each industrial device in an automated plant runs a special kind of ICS program on a specialized microprocessor known as a “programmable logic controller” (PLC). Stuxnet looks for Windows software that's communicating with a device running a PLC. Once the stuxnet virus identifies the right kind of factory or plant, it installs yet another virus

into the plant's PLC software. From that point on, the attacker can invisibly control the entire plant or factory from a remote location.

More sophisticated defenses are required to guard against targeted malware and botnet type attacks in Data Comm networks.

Recommended Approach for Data Comm

In this section, we delineate the broad, best practice based approach to securing the Data Comm network against the various types of threats previously described.

Carefully considered usage of best-practice security recommendations can mitigate several of these threats and reduce the impact of compromised elements in commercial components such as the airline operations center. We suggest the following measures to be deployed in the Data Comm network environment:

Metering and rate-limiting

Metering and rate-limiting are effective tactics for safeguarding against denial of service attacks, in particular excessive AOC traffic against a particular VGS station or subnet. Metering involves limiting the number of messages per second initiated by specific end-systems based on their normal usage profiles, i.e. source-based metering. Further, limits on normal AOC usage per VDL subnet can be enforced based on historical observation or traffic demand prediction. This constitutes destination-based metering, i.e. controlling the amount of traffic permitted per destination VDL subnet.

Pairwise communication using access lists

Access lists can be used to ensure that pairwise communications can only take place between authorized entities, e.g. filtering out source-destination address pairs that are not permitted; thus ensuring that ATC directives do not originate from the airline operations network. In conjunction, anti-source spoofing measures [6] can be used to ensure that sources do not misrepresent themselves by masquerading as other entities.

Preventing misconfigurations

Human error is a major cause of misconfiguration, network service disruption, and degradation. Automated routing policy management can significantly reduce configuration errors and mitigate service disruption risk. Routing policy management is best performed as part of the network management system and consists of route registration, configuration generation, and configuration deployment.

Route registration

During the provisioning process and subsequently when mission or network conditions require, the ATC, AOC and CSPs exchange lists of known, valid and invalid prefixes that might originate from the ATC, AOC and CSPs for each peering relationship. This information helps define transit (CSP) network import policies and user (AOC and ATC) network export policies. Other similar prefix lists help define transit (CSP) network export and user (AOC and ATC) network import policies. The network administrators store import and export policies in a "data base", the route registry, using Routing Policy Specification Language (RPSL).

Configuration generation and deployment

After changes to the route registry, the network administrators initiate a configuration generator that reads the route registry and generates device specific routing configuration for each border router in the network. After generating configurations, the network management system sends device specific configurations to each border router for deployment into operations using existing configuration interfaces.

Initially created by Merit Network, Inc. as the Routing Assets Database (RADB), Internet Route Registry (IRR) [7] [8] provides a route registry and configuration generation capability. The IRR database stores routing policies using RPSL [9] [10] a human friendly, device independent language for routing policy. The configuration generator, called rtConfig, converts RPSL to device specific configurations. IRR is freely available, community-tested by large commercial provider networks, and includes multi-vendor support.

In commercial use, IRR is best known as a public registry of routing information for networks. Organizations that operate networks (ISPs, universities, business enterprises) publish their

routing policy and route announcements in IRR to facilitate the operation of the Internet. They use information in IRR to troubleshoot routing problems, generate access lists, and automatically configure backbone routers.

Benefits of Route Registry

Minimization of Human Error

Device specific configurations are automatically generated based on objects and policies in the IRR. Manual effort is minimized.

Improved Cyber Security

When users register routes or prefixes, there is greater confidence in and better definition of Martian routes. Attempts to spoof routes in the registry are blocked at each border router based on a priori configuration thereby enhancing routing security.

Use of route registry reduces denial of service attack vectors and dependence on dynamic routing protocols without trust vectors. This could be a particular problem for ATN communications, due to the use of the IDRIP routing.

Route registry data should be protected from cyber attacks. A cyber attack against route registry can result in crippling the network and thus affecting the air traffic. Use of Byzantine algorithms [11] [12] for Registry distribution will ensure that compromise of registry database will not affect system availability.

Byzantine algorithms are a class of fault tolerant algorithms. The algorithms were proposed more than 25 years back by Leslie Lamport in his pioneering work on Byzantine General's Problem. The algorithms protect distributed systems against Byzantine failures. Byzantine failures occur when a system fails in an arbitrary manner i.e. provides corrupted information or corrupted local state. Signatures for Byzantine failures due to natural system failures and intentional malicious attacks are often similar making it harder to detect the original cause of the problem. This makes pinpointing the issue of attribution of cyber attack to an adversary very difficult. Byzantine failures are a major topic of research on cyber security.

Byzantine failures also open up back doors for the attacker to perform escalation of attacks at higher layer of the victim's network and application infrastructures.

Preventing malware and botnets

Often attackers initially probe for vulnerabilities at the network layer. Once they identify the vulnerabilities they launch a covert network attack at an appropriate time. A successful network attack often provides the attacker a covert channel. Using the covert channel the attacker can then probe potential vulnerabilities at higher layers and applications. The attacker can then introduce malware and botnets on the victim's network nodes and application systems.

The use of end-to-end encryption and authentication schemes at application layer provide good data integrity and can mitigate concerns over the validity of ATC directives between controller and pilot. While this is the most desirable end-state option for mitigation of malware and botnets, the use of such schemes requires standards evolution and mature implementation.

Commercial strength crypto devices supporting Advanced Encryption Standard (AES) are available for cipher blocks of 128 bits/192 bits/256 bits. Key sizes of 128 bits (AES-128), 192 bits (AES-192), or 256 bits (AES-256) are readily available. They may be adequate for Data Comm ground segment hop-by-hop link encryption.

The mechanism for key distribution between the pilot and the controller will have to be addressed. A X.509- like PKI infrastructure will be required for ensuring the identity of the two end users. VDL Mode 2 uses a narrowband channel of 25 KHz. Pilot to controller communications over a Carrier Sense Multiple Access (CSMA) Media Access Control (MAC) layer, and the movement of the aircraft across VGS boundaries requiring frequent link layer handover mechanisms in a crowded air space will require frequent synchronization between the encryption devices at the controller and the pilot. This means establishing the encrypted channel between the pilot and the controller will lead to longer handover times. The same procedure will have to be repeated if the air/ground communication drops a link layer frame over the air. This may make the air/ground communication secure but will not enhance the overall operational safety of the system. There are simplified approaches evolving for synchronization that should be carefully considered. We recognize that a lot of work is needed in this area before it will be practical to implement application

layer secure channel between the pilot and the controller.

In the meanwhile we recommend a more practical approach, a hop by hop encryption of links within the CSP infrastructure and on all external links. Use of relatively cheap commercial crypto devices supporting relatively low data rates in the ground segment may not be as expensive as installing and coordinating deployment of commercial grade crypto devices in the avionics at a low overall cost.

In addition to installing the crypto devices we recommend installing message guards performing deep packet inspection at VGS stations and the boundary between DCIS demarcation point between the FAA and the DCIS program i.e. the FANS gateway and the ATN G-BIS router.

Conclusion

Data Comm is a key enabler of NextGen applications, including trajectory based operations, and promises many benefits including reducing controller/pilot workload, fuel savings to airline operators etc. The use of data communications between controllers and pilots and the FAA's deployment plans for Data Comm that call for a shared RF network between airline operations and air traffic messages introduces new threat models that need to be carefully considered. In this paper, we have described some applicable threats from the perspective of the Data Comm network and outlined a best-practice based security approach that can be leveraged to mitigate or minimize the impact of potential threats and breaches to the Data Comm network. We believe that more detailed study is required in this critical area, to ensure smooth, secure transition to data communications as the primary channel between controllers and pilots.

References

[1] "Cyber ShockWave exposed missing links in US security" Michael Chertoff, Former DHS Chief in Federal Computer Week, March 11, 2010

[2] Pakistan Telecom hijacks Youtube,
http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml

[3] Estonia Cyber attack
http://en.wikipedia.org/wiki/2007_cyberattacks_on_Estonia

[4] Titan Rain– Time August 25, 2005
<http://www.time.com/time/nation/article/0,8599,1098371,00.html>

[5] Stuxnet <http://en.wikipedia.org/wiki/Stuxnet>

[6] Ravi Vaidyanathan et al, "On the use of BGP AS numbers to detect spoofing", Globecom WPS 2010

[7] Internet Routing Registry, Merit Network Inc.,
<http://www.irr.net/>

[8] Internet Routing Registry Daemon, Merit Network Inc., <http://www.irrd.net/>

[9] C. Alaettinoglu et. al., IETF RFC 2622, "Routing Policy Specification Language (RPSL)", June 1999

[10] Blunk et. al., IETF RFC 4012, "Routing Policy Specification Language next generation (RPSLNg)", March 2005

[11] M. Pease, R. Shostak, L. Lamport "Reaching agreement in the presence of faults". *ACM* April 1980

[12] L. Lamport, R. Shostak, M. Pease, The Byzantine Generals Problem, *ACM* July 1982.

Acknowledgements

The authors would like to acknowledge Joseph Pobiell, Fred Bay, and Steve Hansen for their insightful comments.