AVIONICS COMMUNICATIONS SECURITY

2017 INDUSTRY REPORT

PRESENTED BY:

SMARTSKY™ NETWORKS

Avionics

COMMUNICATING IN THE SKY

SECURITY, RELIABILITY, LATENCY AND BANDWIDTH ARE KEYS

The market for in-flight wireless internet communications in the U.S. is booming and is expected to grow rapidly for some time. If the terrestrial consumer market is any indication, passengers' appetites for such services may well be insatiable, driven by new services and applications that satisfy their needs and interests with greater efficiency.

Business travelers, whether on airliners or corporate aircraft, increasingly expect to complete in the air the type of tasks that they perform in the office, whether reviewing and revising a PowerPoint presentation, editing a complex spreadsheet or joining a video conference.

Aircraft personnel themselves want greater access during a flight on the performance of their airplanes, its engines and onboard systems, and even with crewmembers to help better manage turnarounds on the ground, maintenance costs and in-flight services. They also want greater access to flight information to preclude emergencies or react to them on a more timely basis as the situation unfolds.

Just as that burgeoning demand promises increased use, it carries risks as well. An in-flight communications network must be robust enough to provide reliability, bandwidth and quality of service to satisfy Wi-Fi uses. There are few quicker ways to dissatisfy customers than to raise their expectations of fast, reliable Wi-Fi access and then fail to provide it.

In-flight networks must also be hardy enough to ensure the communications security of its users as greater amounts of sensitive data flow through it, especially with cyber threats being reported with increasing regularity. Each new channel for voice or data communications between the ground and aircraft, each additional individual user of those resources and every connected device mounted on an airplane represent new potential vulnerabilities.

Given this growth in demand, the average user may be surprised to know that there has been little progress with security technologies to protect shared-network users from hackers, scammers and spies. In 2015, the U.S. Government Accountability Office concluded the FAA lacked cybersecurity protection against the vulnerability of its own systems and devices connected to in-flight Wi-Fi. Even Gogo recommends that users "refrain from accessing or transmitting sensitive or private information." It tells customers, "Make sure your laptop has both a firewall and malware protection."

The challenge then becomes how to deliver desired resources securely without shutting down or tightly limiting the growth of these channels for communications. A standard tool of business communications with computers, the virtual private network (VPN) can help ensure such security, but using it in the in-flight environment presents challenges.

SECURITY CHALLENGES

Unless you have access to your own VPN, using Wi-Fi in the air can make your private information easily accessible to others. This is also true of most shared networks such as public Wi-Fi or airport networks. Tourist-frequented Wi-Fi networks are targeted by scammers and hackers more often than others.

"It's no different than being in a Starbucks. Somebody there could be sniffing your data," said Russ Syphert, a data communications security expert and chief technologist for Cyber Business Analytics. Just as with ground-based computer communications, he said, in-flight Wi-Fi likely will never be 100% secure. The idea, as with home security systems, is to make stealing your information so difficult that the hackers will choose easier targets.

Airlines and other aircraft face their own risks in securing air-ground communications. To understand the broad and complex nature of those risks, consider that today's aircraft don't have just one communications network on them. At a minimum, they have three separate networks which the International Civil Aviation Organization (ICAO) separates into three domains.

- The Aircraft Control Domain (ACD) consists of cockpit systems, the primary function of which is to support the safe operation of the aircraft.
- The Airline Information Services Domain (AISD) provides services and connectivity between independent aircraft elements such as avionics, in-flight entertainment and cabin distribution systems, as well as connections with off-board networks. The routing, computing, data storage and communications for non-essential aircraft applications happen inside these networks.
- The Passenger Information and Entertainment Services Domain (PIESD), the most dynamic onboard network, is better known as in-flight entertainment and connectivity (IFEC) services as featured in most modern cabins. It includes multiple systems from different vendors that may or may not be interconnected. According to a working paper on aircraft domains presented to ICAO by Honeywell Aerospace Senior Program Manager Aloke Roy, these can include passenger device connectivity systems, broadband television, seat actuators or messaging systems. Over the past 15 to 20 years, this has been where most of the growth in demand for in-flight communications capability has taken place.

> *"Devices on the different networks on the airplane sometimes do talk to each other, but only by design, and typically in a one-way direction," said John Craig, chief engineer of cabin and network systems at Boeing Commercial Airplanes. "Each device will accept data only from pre-approved sources or critical sensors."*

Awareness of cybersecurity risks is relatively new for civil aviation, Craig said. "We in aviation had a bad habit of thinking that we were special, that we weren't affected by all the cybersecurity stuff everybody else in the business world was having to deal with," he said. "But Stuxnet was the wake-up call."

Stuxnet is the famous computer worm believed to have been jointly developed by the United States and Israel and used by Israel to disrupt the operations of Iran's central nuclear-fuel enrichment plant in Natanz in 2009 and 2010.

In response to that event, the aviation industry formed A-ISAC, the Aviation Information Sharing and Analysis Center, as a nonprofit cybersecurity think tank and information-sharing group under the aegis of the U.S. Homeland Security Department's National Infrastructure Protection Plan. New standards for keeping air-ground communications secure were established, and team members from the industry continue to share information to keep advancing the state of the art of in-flight cybersecurity.

"We started with seven or eight member companies, but we're now up to about 33 members — airlines, suppliers, service providers, manufacturers," Craig said. "We even have an airport member. We've reached critical mass and have established a good infrastructure for information sharing."

An eccentric and well-known computer hacker, Chris Roberts, and others continue to remind the industry that it isn't immune from cyberattack. Roberts startled the global aviation industry in April 2015 with public claims — taken seriously by the FBI — that over the previous four years he had been able to access commercial jets' flight controls from his seat in the passenger cabin "11 to 15 times." His tale was shown to be highly exaggerated and, in all likelihood, untrue.

As airlines move in coming years from proprietary technologies for air-ground communications — such as the Aircraft Communications Addressing and Reporting System (ACARS) — to internet-based open systems, they are keenly aware that doing so increases the risk that outsiders could access onboard systems. So, as that shift is being made, the industry is also taking steps to build increased security layers around these more capable networks.

"It's a good thing to move away from legacy systems and in the internet protocol direction," said Franc Artes, architect of the security business group for tech giant Cisco Systems' chief technology officer. "But it brings extra risk. You have to think about it as a whole ecosystem in terms of security, from production supply chain to the consumer on the aircraft and the consumption of those communications services on the airplane."

PROTECTION FOR THE PASSENGER

Air travelers these days expect to be able to stay fully connected to the world even while they are in the sky. They want to be able to carry on proprietary business and personal conversations via email, texting and even with telephone calls in order to access information they deem important. Increasingly, passengers also want to be able to securely Skype on their personal mobile devices via streaming technology. Fulfilling this customer-focused need without creating substantial cybersecurity risks necessitates a communications link broad enough to ensure reliable VPN connectivity.

Most veteran business travelers are familiar with VPNs.

"If you're high enough up in your organization to be handling sensitive data, you've probably been read in on all the security protocols your company uses," said Syphert of Cyber Business Analytics. "You're either using a VPN and/or encryption or you're trained not to handle that kind of data in such places."

At Coca-Cola, for example, high-ranking executives who travel aboard the company's three Gulfstream jets are well versed in voice and data communications security protocols. They use them at all times, whether at home, in a hotel or elsewhere, said Jamie Buff, the company's aircraft program manager in Atlanta. "Everything they do is through a VPN, even at home," he said.

To better address modern cyber threats, Coca-Cola is also preparing to add another layer of voice and data communications security to its planes: a virtual firewall that will create a second VPN to protect the communications of executives on board in addition to the individual VPNs they will use when making calls or using the internet.

To understand the importance of using a VPN and the complexity of maintaining its connection within an airplane, it helps to have an understanding of such a network's architecture.

Microsoft's Technet web information portal defines a VPN as a private network that enable organizations to transmit data and information between two end points across the internet by emulating the properties of a point-to-point private link. When a passenger is using a VPN in an aircraft cabin, he or she is tapping into a private network across the internet using a remote-access connection.

A VPN uses a networking technology called tunneling that encapsulates one type of protocol packet within the "datagram" of a different protocol. This shields a data packet transmitted within a VPN from open access points as it moves across the internet and, in the case of in-flight connectivity (IFC), across the aircraft connectivity link to the internet. It does this by requiring the VPN's end points to agree constantly on what has been sent, what has been received and, quite importantly, over what time span, as the packets move through a series of networking protocol "handshakes" based on configuration variables, such as IP address assignment, encryption or compression parameters. A tunnel management protocol is the mechanism that creates, maintains and terminates such a tunnel.

Any disruption of the handshakes, such as taking too long to acknowledge that a transmitted data packet has been received (also known as a "time out"), will cause the VPN to shut down. In the case of IFC it is critical that the in-flight network not inhibit the VPN from completion of its task.

NETWORK'S IMPACT ON VPN

The effectiveness of VPNs in air-ground communications is impacted by three variables:

- Maintaining the link (reliability)
- Latency (avoiding security "time outs")
- Bi-directional connectivity (balanced uplink and downlink)

Reliability can obviously be affected by the quality of the equipment used in the network. For an in-flight communications network, reliability is also governed by the wireless link between the aircraft and the ground. The greater the distance that the packets have to travel along the wireless link, the greater the risk of a VPN time out.

Since electromagnetic waves such as radio signals travel at the speed of light, the delay between a signal's transmission and its arrival at the intended recipient over short distances is imperceptible in most cases to humans and easily manageable by communication systems. Over greater distances, of perhaps a thousand or more miles, the signal's delay becomes noticeable to us and more of a challenge for communication systems, especially when the signal travels through a VPN.

With a network that relies on a few large, widely separated gateways (such as on the U.S. East and West coasts) to transmit and capture signals, the distance between those gateways can cause timing problems. This can present an issue in linking to the internet, which operates at speeds for the exchange of messages on the order of 30 to 100 milliseconds (and often 30 to 50 milliseconds.)

Physical separation and the associated delays can be a much greater issue for networks that rely on satellites to provide communications. The beams of a geosynchronous communication satellite can cover large areas of the Earth's surface. But for such a spacecraft to stay in one position over the surface, it must match the Earth's rotation about its axis. That requires an orbital altitude of nearly 22,240 miles above mean sea level.

A radio signal can take 120 milliseconds or more to travel that distance to the satellite (depending on the slant angle between the aircraft and the satellite) and just as much time to travel back from the satellite to a network's ground station. Add the time travel from the ground station into the internet and back, and the message-exchange time can exceed 500 milliseconds.

A network using satellites in lower orbits of a few hundred miles above Earth can cut that transit time. But each satellite's beam covers less of the Earth's surface, so more "birds" are needed to cover an area the size of the continental United States. That means a signal may have to travel up and back to more than one satellite, or relay between satellites, in order to connect two distant points on the surface. That back-and-forth also takes time.

An option for cutting the travel time for a communication signal is a network that wirelessly communicates with the aircraft from ground stations that are much closer than satellites. Such regional nodes enable signals from the aircraft to be directed quickly into the internet, processed and returned to the in-flight customer in much less time.

Additional security of that VPN transmission can be augmented through the use of 4G LTE protocols commonly used in ground communication networks and available recently in air/ground environments. The LTE network authenticates the radio, and the radio authenticates the network credentials. The longer root key (128 bit vs CDMAs 64 bit) requires a greater effort to break through security. LTE's second security layer of encryption, Integrity Protection, verifies that the signaling has not been modified over the radio-access interface and that the origin of signaling data is accurate.

LATENCY AND VPNS

The round-trip time between sending out data and receiving an acknowledgment that it was received is called latency. It is a critical factor in effective VPN use.

The security of VPNs generally relies on use of the Transmission Control Protocol, or TCP. One of the main protocols used for internet communications, it provides reliable and orderly message exchanges, checking data packets for errors along the way. In addition, streaming services such as Netflix and Amazon Prime Video also depend on TCP.

An essential element of TCP is the acknowledgment of a data packet's receipt. It must be received before the protocol will allow transmission of the next data packet. If the acknowledgment does not come back and does so in a timely manner, TCP considers that packet missing and calls for retransmitting it. However, if the problem persists, the protocol will trigger the VPN to interpret that as fault in the "tunnel" and terminate the connection.

Clearly, a long transmission time (or "high latency") can be a problem for in-flight communications by causing the VPN to time out and break the link between the passenger in the sky and the system on the ground.

U.K.-based Bentley Walker, the largest European re-seller of satellite internet equipment, has detailed the challenges of maintaining a VPN connection over a satellite-based internet service. According to a 2012 Bentley Walker report, two-way satellite networks must employ special software to deal with the 44,500-mile-plus round-trip distance for satellite communications signals. One technique is a version of "spoofing" in which a satellite service provider's system tries with software on its spacecraft to fool each end of the VPN that data packets have been received and acknowledged. Without this software, the increased latency will result in TCP time outs that severely limit the communication link's performance.

Bentley Walker estimated that average latency for satellite networks is about 550 milliseconds, as compared to the low latency of 35 to 100 milliseconds for ground-based networks. Clearly, an air-to-ground network with low latency offers advantages in providing cybersecurity for in-flight communications.

BANDWIDTH AND THE IN-FLIGHT EXPERIENCE

Another factor that can affect a VPN in flight and also have great potential impact on the overall performance of communications between those in the aircraft and resources on the ground is the bandwidth available.

The demand for bandwidth continues to increase as passengers' appetites for greater access to high-capacity in-flight applications grow and as airlines make greater use of operational and systems data downloaded during flights.

Today's in-flight networks often boast of the bandwidth available to transmit to the aircraft, which can be on the order of multi-megabytes per second. But a major limitation in most in-flight communications is the return link (from the aircraft to the ground) that typically has a bandwidth of only 300 to 500 kilobytes per second.

That comparatively small bandwidth can hamper VPN performance by slowing transmission of data packets and message acknowledgments, especially in cases in which high latency is also an issue.

For the passenger, especially the business traveler, that imbalance between the forward link and return link can severely limit what can be accomplished on a flight.

"If you want to do a WebEx or a videoconference from an aircraft, data must be able to get off the aircraft just as quickly as it comes from the ground," said Darren Emery, Director, product support engineer for SmartSky Networks. "You need multi-megabit-per-second connections in both directions, otherwise you just can't do it. If you want to transmit a 20-megabyte PowerPoint off the aircraft, because you're late with something as we always seem to be, it is almost impossible do it on a 300-kilobit-per-second connection from the aircraft to the ground."

That forward and return link balance and the need to meet future demand for greater capacity are driving some in-flight communication service providers to build networks that can match return link bandwidth more closely to the forward link capacity.

BOTTOM LINE

In-flight communication services, both for the entertainment of the general passenger and the efficiency of the business traveler, face steady growth in demand that is certain to increase as passengers insist on the ability to do more and to do it faster while aloft. That demand will only increase further as airlines and other aircraft operators seek greater access to operational and systems data from their flights while they are in the air to improve the cost-effectiveness of their own operations as well as the in-flight experience of passengers.

In addition, passengers and aircraft operators will insist on even greater levels of security for the information they exchange and access through aircraft in the sky.

Satisfying those growing demands poses challenges for providers of in-flight communication services, some of which are constrained by the architecture and physical limitations of their networks. However, new air-to-ground services are coming to the in-flight market with technologies for more reliable, higher capacity network architectures that will enable customers to have both more robust data capabilities as well as the cybersecurity required for connected operations while in the air.