Securing a Mosaic Infrastructure: Challenges in Airport IT Security

David C. Shields

University of Maryland University College

Table of Contents

Introduction 3
The Tiles: Common Airport Information Systems 3
The Strongbox: Physical Infrastructure Security 7
Unification of Chaos: Internal Information System Security 8
The Public Myriad: Securing Public Information Systems 11
Conclusion 13
Appendix A 14
Bibliography 15

Securing a Mosaic Infrastructure: Challenges in Airport IT Security

Introduction

From the moment people arrive at an airport for a flight, their activities are a frantic array of check-ins, security pat-downs, hurrying to locate departure gates, finding an eatery and even checking a few email messages on the free Wi-Fi before boarding the plane and waiting for the roar of the twin jet engines to fly them into the sky. Surprisingly, the frantic level of activity is usually much smoother than the average person may realize and few people have insight into the massive amount of systems, networks, and information databases that must work in harmony to produce the experience. Unlike most business IT infrastructures, an international airport has to offer a mosaic of different network connections, software/hardware combinations, international relay networks, and federal organization networks in an environment that is as equally exclusive as it is fragmented. Furthermore, international airports due to their expansive systems, high volume of passengers from varying cultures and very visible profile in the world, must be positioned in a way to protect their networks against multiple security threats from terrorists, foreign nationals, and internal threats. IT managers in an environment such as this must be able to generate functional infrastructure that supports a wide array of disparate information systems while still maintaining security between the 'borders' of the various virtual networks and still offering a smooth experience for each passenger and business that shares the grounds.

The Tiles: Common Airport Information Systems

Thanks to the multiple terrorist attacks on USA airports, most airports are quite low-key about the various systems their facility contains as offering such information to outside parties is a security issue in of itself.

Even still, there are a large number of obvious systems in use that any airport patron is capable of determining and a few that some information is known about. The first and most prominent of the systems at stake is a complex infrastructure software suite which is wholly operated by the Federal Aviation Administration (FAA) and truly is the backbone to link multiple international airports in the USA and the world. Although its actual name is not widely known, it is referred to as an Airport collaborative decision making system (ACDM) (Katsaros & Psaraki-Kalouptsidi, 2011, p. 351). This software functions as a way for air traffic controllers to determine what airplanes are in the sky at any given point in time, where the planes departed from and are arriving at, as well as helps manage the take-off, landing, and re-routing of planes. It is from this system that most airports are able to display the arrivals and departures messages passengers are so commonly used to seeing. There are other software applications used by the FAA but this information is classified for obvious reasons.

Once the software from the FAA has formed the 'bones' for the airport, the next likely systems for a passenger to see would be those used by the individual air travel companies such as Delta, American Airlines, etc. From the moment one arrives in the airport, they will likely be drawn to the particular kiosk of their chosen travel company for ticket check-in, updated flight information, baggage check in, and various other services. In fact, a large number of airliners now offer automated self-service kiosks which are, at their core, information systems (Abdelaziz, A. Hegazy, & Elabbassy, 2010, pp. 17-18).

After a traveler has checked in, got their ticket(s), and passed through security screening, they will find themselves in a concourse that will contain countless shops, eateries, and convenience areas. Almost all of these vendors will have some sort of POS system or kiosk that is used to make purchases.

While it is not likely that an attacker would want to exploit these machines, the fact remains that it is still very much possible to do so. Furthermore, these information systems would need to have a link to their home base for processing transactions which would require the already diverse network infrastructure to support and separate these items from the rest.

The traveler, now with food in hand (or whatever merchandise they purchased at the aforementioned kiosk) arrives at the gate where their plane is to depart and finally has about 30 minutes of time to relax before they must fly out to wherever their destination may be. It's likely that they would want to produce their laptop or smartphone and check flight schedules, weather in their destination, maybe catch up on those emails before they are required to silence all outgoing transmissions. How is this handled? Easy, hop on to the public Wi-Fi access system offered by the airport for its patrons – yet another of these diverse information systems which must be secured.

Eventually, the traveler receives the signal they have been waiting for, time to board their plane. Before this can occur, multiple operations such as flight pre-checks, fuel checks, and even visual security checks of the plane must occur (Transportation Safety Administration, 2001, pp. 335-338). All of this input must be documented in TSA and FAA approved systems prior to the actual boarding of the plane for the safety of those flying. Again, the systems of the TSA, the FAA, and the individual jetliner company must be in agreement before the boarding staff will receive a confirmation and approval.

Finally, the traveler has been approved to board the plane, found their seat, plane is preparing to take off and they are already being asked to turn off all electronic devices.

One must be prone to wonder why a request is made and that is linked to yet another information system that is resident on most modern planes such as the Boeing 777 and 787, Embraer jets, and Airbus jets known as “Fly by Wire” (FBW) (Brière, Favre, & Traverse, 2001, p. 305) .
The FBW system, although not entirely computerized, makes heavy use of computerized processors to adjust everything from yaw and pitch to fuel being transmitted at varying levels. Although planes have a backup manual control to use as a failover, the very thought of what could happen if a person with an electronic device happened to ‘hack’ the system in flight is a frightening consideration (Valacich & Schneider, 2010, p. 410).

Assuming all goes well in flight, there is one more tile in the mosaic infrastructure that is specific to international airports. So let’s say that the traveler in question travelled to the USA from a foreign country, even if they are a US citizen. Once they arrived at the international airport (or perhaps before departing their original airport), all of their luggage and even the person themselves would have to be verified by the United States Customs and Border Patrol (USCBP) team, the Department of Homeland Security (DHS), the Department of Immigration and Customs Enforcement (ICE), and countless others to ensure they are not a dangerous person or otherwise a danger to flight. Every one of these checks must be completed within a database that is stored on an information system connected to the already colorful mosaic infrastructure.

It is plain to see that the infrastructure required to support the myriad of disparate information systems used by an international airport will require a very ironclad security infrastructure. The key term to consider in the creation of security infrastructure for this environment is separation as almost all of the systems mentioned above have to be separate from each other.

Even so, the network needs to be secured in a way that will not hamper ongoing operations while still meeting all federal guidelines for security of the machines. Great care must be taken to ensure that the security safeguards in place are as extensible as they are secure. The airport infrastructure and its comparable security safeguards can be divided among three groups: physical security of facilities, securing internal information systems and securing public information systems.

The Strongbox: Physical Infrastructure Security

The ACDM system used by the FAA and the array of other FAA systems should already be quite secure against any cyber-attacks as their systems would likely be created under the jurisdiction of the Department of Defense. Even still, the responsibility of securing the facilities themselves and the network equipment to which they are attached is still an issue that must be handled by the airport itself. As with any physical hardware in a public environment, the rooms which house the network equipment will need to be secured through multi-factor authentication. On the physical level, all personnel that need access to this room should have a physical badge that not only authenticates the user visually by a photo of the user but also identifies the user in a monitoring system that will only unlock the door if the particular user is identified in a security database as being on the approved access list. This check can be made by use of a unique ID such as that assigned to an RFID. Furthermore, the room should be under surveillance 24 hours a day, seven days a week with all video archived for incident recreation.

The approach outlined above is the most secure choice within reasonable budget limitations. It offers the advantages of multifactor authentication, what the user has (their ID card), and without much effort, a password could be implemented for each user which would make a second layer: what the user knows.

For an additional cost, the airport might want to even consider a biometric system for all three factors, something the user is (Valacich & Schneider, 2010, p. 422). With the inclusion of a badging database that manages authentication through the door with the badge, the additional advantage is offered of an audit trail which would make it simple for the airport to determine who entered the secure room at the time the compromise occurred and the records of door access could be combined with video from surveillance to truly give a visual of who the person is. However, the disadvantages of such a system include the high level of initial investment required to create the facility and the level of complexity required to manage the various interconnected systems. Regardless, with the high level of security needed to ensure that these systems are not physically compromised coupled with the higher likelihood of them being targeted, the expense is certainly justified.

Unification of Chaos: Internal Information System Security

Once the physical facility which houses the network hardware has been created, the work to secure this infrastructure has only just begun. Due to the nature of cyber-attacks, the most secure physical room will be completely irrelevant if the contents within are not staunchly secured from attacks on the networking components themselves. The airport will want to create two completely separate networks to be used, neither of which is able to interact with the other unless secure VPN tunneling is used. The network holding the FAA hardware and connectivity to the ACDM should be independent from anything else in the room as these are the most critical systems to ensure safe landing and takeoff and proper routing of aircraft due to weather or other external circumstances. The network will need to contain a strong firewall system, a heavily scrutinized Intrusion Detection/Prevention System (IDPS) with constant monitoring by staff.

Additionally, all hardware and software in use on network devices needs to be updated as often as updates are available. It might be a wise idea to use this network to support the other government entities as well including the USCBP and ICE organizations but to prevent the chance of cross-stream attacks, each organization should be kept on its own private VLAN resident on the switch as this will either remove or minimize the chance of a compromised node bringing the other nodes in the other groups to compromise as well (Prowse, 2011, p. 113). In the event of an emergency, all of these devices will need to be protected by a UPS system with a minimum of 12 hours battery support in the event of a power loss or disaster.

The network used by the airport operations and its support personnel should be on a second network with similar controls as those on the FAA. However, due to the wide variety of technology ecosystems that this network needs to support, it will be much more fragmented into VLANs than the other two. The airport itself will have various business units such as marketing, compliance, vendor relations, facilities and executives which keep the airport itself running and these various units will be collected onto the same LAN fragment. These units will share resources such as servers, internet filtering systems, and the like and will work best if they share the same network ‘space’. These must be kept separate from the other groups as their systems may not be as tightly monitored as those of the FAA and other government entities and any cross-communication with the two units could be brought under great scrutiny in the event of a security breach.

Additionally, the operations infrastructure will house the independent systems of the various airlines (i.e. Delta, JetBlue, etc.) including self-service kiosks, baggage check systems, and passenger information systems.

For security and dependability reasons, each airline is likely to want its systems kept on separate VLANs with each subnet secured and firewalled to prevent data leaks. These micro-networks within the greater LAN will each share their own server resources, communications systems, and may even have independent staff on hand to assist with any outage issues. Depending on the policies outlined by each airline company, the infrastructure may require the data on these machines to be backed up locally on a shared server for easy recall and data recovery. Additionally, it would be wise to dedicate various IDPS with port sniffing capabilities (perhaps even set to promiscuous mode) to ensure that no data from the systems has a collision or corruption as it travels across the LAN (Prowse, 2011, pp. 234-235).

There is no doubt a great deal of complexity involved with the creation of these two separate networks (each of which are even further subdivided) to run the internal infrastructure of the international airport. Not only are the network switches, firewalls, and security appliances (IDPS) expensive but also the competence of the network engineers and support staff needed to create such safeguards and effectively secure them are expensive as well. There is also the requirement to provide separate WAN connectivity to each individual network via some sort of line provided by an ISP (Internet Service Provider) which comes with its own cost and monitoring fees. However, the advantages to security and stability of the network and its components far outweigh the costs. Each network can be kept separate from the other so that if an intruder were to somehow compromise the architecture of one network, the other network would be less likely to be compromised as it would require a completely separate attack to be launched (Prowse, 2011, p. 114). Furthermore, keeping the two LANS completely unique and the subnets unique would greatly reduce the time needed to locate a security issue as the varying IDPS and auditing systems would have less data to investigate in each network.

In fact, the separate transmission paths would also serve to boost transmission of data among each network segment so that high integrity systems such as those implemented by the government offices would be less likely to suffer signal degradation during high volume periods at the airport.

The Public Myriad: Securing Public Information Systems

The final segment in the mosaic of the network that needs to be secured is that of the public facing infrastructure for the airport. Much like the other networks before it, this network should also include a completely separate signal feed as it is the network most susceptible to attack and instability. The public network system requires some other considerations that are not common in the other two networks, however, which make it more difficult to manage. First and foremost, the public system must be accessible to virtually any device that the patrons and vendors (such as food vendors) might bring to it. It must be able to handle multiple different laptop devices from many manufacturers, virtually all smartphones and tablet devices in the world, and accept transmissions from multiple different manufacturers' network connectivity devices. Second, with so many people using the network at any given time, it is much more likely that this network could be compromised and, if connected to the other networks, could be used to infiltrate the more secure networks used by the airport operations and government entities. Even still, it would be unwise to have a public accessible network to run the vendor terminals as they would be processing financial data. For this reason, a VPN system across a segmented VLAN on this public network would be the most efficient way to create a secure network without adding additional network WAN connections to the environment.

On the public Wi-Fi, there would be a great amount of overhead, technical support, and bandwidth concerns that might be daunting to an already taxed IT infrastructure team.

For this reason, many public Wi-Fi hotspots have started to outsource their Wi-Fi services to external contractors such as Boingo.com. These companies require little more than a WAN signal and space to place their hardware and they can create an extensible Wi-Fi solution that has a dedicated task force for support issues instead of bothering harried IT departments (Platt, Carper, & McCool, 2010, pp. 18-19). Because of this, the airport should outsource this particular network function to allow their team to focus on the more important issues in the IT infrastructures elsewhere in the building.

As was the case in the previous network setups for the government and operations folks, the cost and complexity of managing yet another segmented LAN could be a challenge. However, the costs and issues linked to a computer failure on a critical system such as that being used by the FAA or the USCBP could lead to poor service and loss in confidence. The Los Angeles International Airport (LAX) witnessed something like this in 2007 when a USCBP computer system went down and caused delayed flights for over 20,000 international passengers (Brombacher, 2007, p. 1). A costly delay such as this could cause not only issues with passengers being literally stuck on the runway while the system is restored but also causes logistic issues because a crew must be kept onsite to keep the plane fueled and occupied while the issue is resolved. Therefore, it makes perfect sense to design the network in a way that would provide the highest quality of performance and dependability because the repercussions of being undependable will be remembered by the passengers and the news media for a very long time (Brombacher, 2007, p. 1) .

However, the advantages of having a dependable infrastructure are only the first of many advantages to the solution proposed. If the public Wi-Fi is outsourced to another company, the IT team has one less thing to concern themselves with as long as the network is secured.

Many IT projects, especially in airports, have often been negatively impacted by inflated requirements, lack of human resources and a constantly dwindling budget (Montealegre & Keil, 2000, pp. 418-419). By outsourcing the public Wi-Fi, the IT Team can focus their resources on handling the other issues in the rest of the infrastructure as well as the projects that need to be completed and hopefully devote their funding to other projects. The segmented LAN for the vendors will ensure that their POS systems and other information systems are less likely to be impacted by a high amount of public LAN traffic and are less likely to be compromised by attacks on the public LAN. Lastly, the dedicated IDPS scouring the networks will be more likely to catch any attempts at malicious hacking that the public Wi-Fi vendor does not detect. This will ultimately result in a more secure network with better performance for all parties.

Conclusion

An international airport is a place that must make a secure ecosystem for all the wide variety of machines they house and support. It is highly recommended to keep the high-risk systems separated from the public systems to limit as much risk of system compromise as possible. The government network, the operations network and the public network should be kept physically separate and on completely separate connectivity networks. The operations network and airliner network could share a common network switch but be segmented into individual VLANs in order to prevent accidental data leakage or connectivity bottlenecks. The public vendor network and public Wi-Fi can also share a common router but should be segmented into separate VLANs for connectivity and security reasons and the public Wi-Fi should be supported by an external vendor. Each network will be scanned by individual IDPS systems in order to maintain integrity of data and to reduce the amount of traffic each individual system must monitor (For an outline of the network, see Appendix A).

If the mosaic tiles of an international airport are properly secured, managed with a focus on security and separation and all physical areas protected, the high amount of risk with such a target will be reduced by a significant amount. With properly trained IT staff, managing even as segmented a network as described is not as significant of an undertaking as it might seem and the resulting high security and high availability network will be worth the effort. Even the most fragmented mosaic, when meticulously fitted together, can create a wonderful work of art.

Appendix A

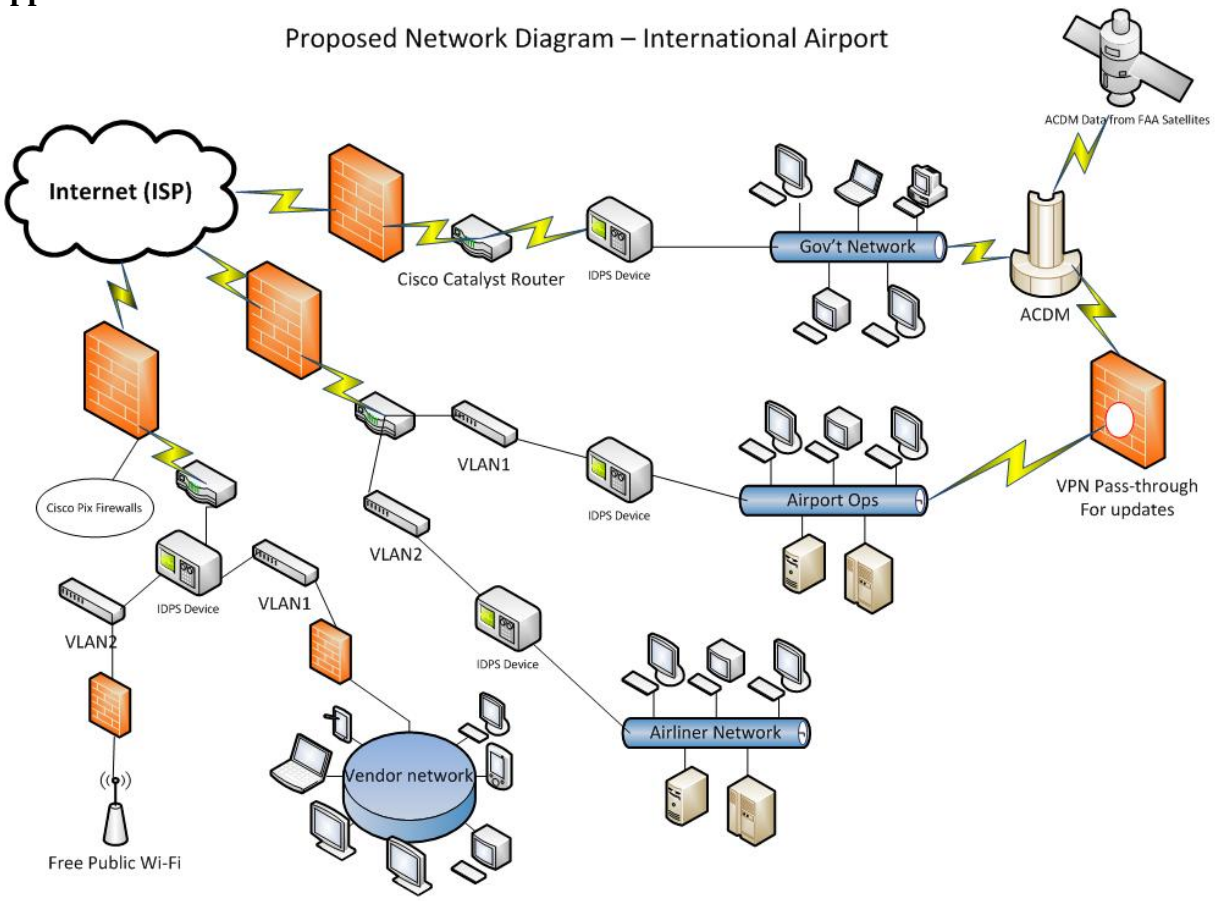

Bibliography

Abdelaziz, S. G., A. Hegazy, A., & Elabbassy, A. (2010, May). Study of Airport Self-service Technology within Experimental Research of Check-in Techniques Case Study and Concept. *Inernational Journal of Computer Science Issues, 7*(3), 17-26.

Brière, D., Favre, C., & Traverse, P. (2001). Electrical Flight Controls from Airbus A320/330/340. In C. Spitzer, *Avionics Handbook* (pp. 255-330). Boca Raton: CRC Press LLC.

Brombacher, A. C. (2007). Dependability... *Quality and Reliability Engineering International, 23*(767), 1. doi:10.1002/qre.894

Katsaros, A., & Psaraki-Kalouptsidi, V. (2011, July-September). Impact of collaborative decision-making mechanisms on operational efficiency of congested airports. *Airport Management, 5*(4), 351-367.

Montealegre, R., & Keil, M. (2000, September). De-escalating Information Technology Projects: Lessons from the Denver International Airport. *MIS Quarterly, 24*(3), 417-447.

Platt, R. G., Carper, W. B., & McCool, M. (2010, July). Outsourcing a High Speed Internet Access Project: An Information Technology Class Case Study in Three Parts. *Journal of Information Systems Education, 21*(1), 15-25.

Prowse, D. L. (2011). *CompTIA Security+ SYO-201 Cert Guide.* Indianapolis: Pearson Education.

Transportation Safety Administration. (2001). *Subtitle VII Aviation Programs.* Washington D.C.: Transportation Safety Administration.

Valacich, J., & Schneider, C. (2010). *Information Systems Today: Managing in the Digital World.* Upper Saddle River: Prentice Hall.