Fundamentals of Network Segmentation

Arthur Wyatt

ICTN 6880

Dr. Pickard

East Carolina University

Abstract

One of the most important things to consider when designing or creating a network architecture or infrastructure is what the security needs are and how best to achieve that security. When deciding the level of security needed there are several things that need to be taken into accounts such as who will need to access the information, how much of the information needs to be secured, Hardware consideration, and accountability. The next step to is to decide how best to address those needs. Many of those aforementioned considerations can at least in part be addressed with the implementation of network segmentation. four methods of segmentation will be the main focus of this paper. Through out this paper several strategies on how, where, and when to use the various methods will be discussed. Of those, the ones that will be discussed in detail are the use of physical segmentation, logical segmentation, virtual segmentation, and lastly a subset of physical called air gag. Lastly, the paper will end with a brief discussion of the main advantages of network segmentation in the form of increased performance and security, the quality of life improvement with ease of management and lastly how segmentation helps meet standard compliance.

Introduction

Designing a networking is not a simple task. The parts and processes to effectively do so may seem simple from a high-level view but the interweaving flow of how the various parts work in tandem together can be a difficult task. This paper focuses on some small subpart of the overarching processes in designing a network. This paper attempts to provide a fundamentally basic view and explanation of what network segmentation is and how it may make the lives of network administrators easier. There are multiple different methods to segment a network and several of them are discussed in this paper. Specifically, this paper discusses the segmentation method of physical segmentation, logical segmentation, virtual segmentation, and a subset of physical segmentation called air gap. The discussion attempts to describe how each of the methods could be implemented and how they would achieve network segmentation and several of the disadvantages of each. Lastly, the paper will include some overall advantages that network segmentation provides such as performance improvements, ease of management, meeting compliance, and the added security benefit.

What is Network segmentation

Network segmentation is the best practice policy of either physically, logically, or virtually separating parts of the network into smaller parts call segments. Furthermore, because it is "network" segmentation and not user segmentation, for example, it allows for the creation of policies and rules that define how each network segment interacts and communicates with each other. There are a limited number of techniques that can be used, and many of them will be discussed shortly in this paper, but the criteria used to determine where to break or segment the network into smaller part is almost entirely subjective. The reason for this is because just like almost everything in life, except hats and sweaters, and even then, it is one size fits most, there is

no one size fits all. There are plenty of guidelines and papers published that suggest several ways in which to design a network for compliance, security, ease of management, etc. but at the end of the day it is up to the person or persons that are tasked with designing the network to decide what the best metric is to create the network segments. Later in this paper, some loose examples of metrics that can be used to segment a network will be mentioned, but first, it is important to understand the different fundamental methods of segmenting a network. Those fundamental methods are physical, air gap, logical, and virtual segmentation.

Physical segmentation

Physical segmentation is arguably the most basic form of network segmentation. Physical segmentation is the process of using separate and discrete hardware for every segment created. For example, take two customers, customer A and customer B. Both customers have come to a service provider and would like to use their service but want their traffic isolated from each other for compliance or security reasons. One way to do this would be to give each their own switch in which they would be the sole user of. An example of this is shown in figure 1.

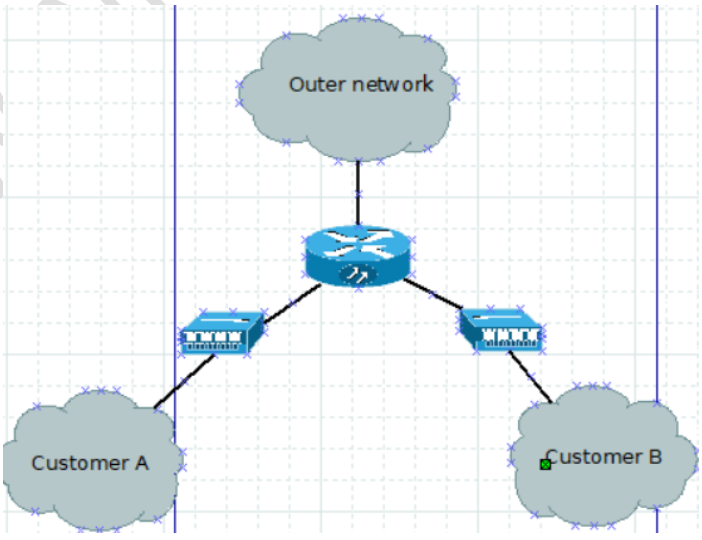

Figure 1. Physical Segmentation

In figure 1 both customers A and B get their own switch. This allows for those network segments to be isolated from each other until it gets to the router which should be configured, through the use of access control lists and other security methods, to prevent them from ever talking to one another. A more complex scenario would be to go even further with the device separation. If it was necessary or prudent a case could be made that each customer should even have independent routers for further isolation. This method is fairly simple and easy to understand and implement but with current technology, it is not widely used because of the major disadvantages it has.

Physical Segmentation Disadvantages

The disadvantages physical segmentation are that it is wasteful, inefficient, and can cost significantly more than the other options. In the example provided above in figure 1 each switch has more ports available than are actually being used. Imagine using only a single port on an entire twenty-four port switch. That would leave each customer 23 ports for expansion but if there is no intention of the expanding then those 23 ports would go unused and would ultimately be wasted. The next issue with physically segmenting each customer with individual devices is that it is inefficient when it comes to scaling. Again, referring to the example given in figure 1, if a third customer, customer C, was to be added to the topology then an entirely new device would need to be bought, added, and configured to accommodate their inclusion. This added complexity would be even more exaggerated if multiple devices, such as a switch, router, and server, would need to be added. Which leads to the last major issue with physical segmentation and that is the cost of using such a strategy. In the small scale adding cheap more cost effective devices may be feasible but if that is every scaled up to the need to accommodate thousands of isolated segments that would not only be expensive because of the additional hardware that would have to be bought, but the power and management hours needed to configure and

troubleshoot any issues. Even though physical segmentation has its drawbacks that does not mean that it should never be uses. The next section will discuss a sub section of physical segmentation referred to as air gap.

Air Gap

Air gap segmentation is a sub section of physical segmentation because the segments are physically isolated by hardware. However, air gap differs from physical segmentation because the air gap takes this separation to the extreme because there is literally no connection what so ever between the segments. Furthermore, there is only "air" between the air gapped network and other networks. Air gapped computers or networks are isolated away from any potentially unsafe, insecure, or uncontrolled network. Which means that it is not connected to the internet or connected to any computer that might connect to the internet and is even physically removed from the locations of such networks or computers (Nohe, P. (2018, March 13)). A simple example of an air gapped network is shown in figure 2.

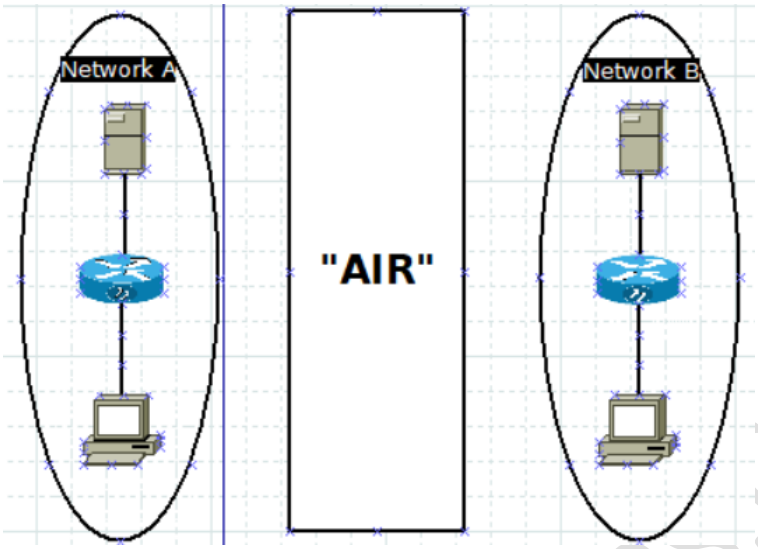

Figure 2. Air Gap example

Figure 2 shows what two networks that are air gapped from one another. In figure 2 both network A and network B are not connected to the internet, but they are also air gaped from each other. They are identical networks each with one computer, one router, and one file server but there are no connections between them. In fact, if it were not for removable portable media there would be no way to share anything between them. Because of the nature and level of isolation that air gap segmentation provides one might consider this to be extremely secure and that it would be virtually impossible for this network to be compromised. Due to this idea, that in a correctly designed and implemented air gap environment information has no way to traverse the gap between networks and, therefore, there is no way for malicious actors, such as hackers, worms, and malware, to cross the gap and access the network (Byres, E. (2011, August 17)). Some examples of systems that this type of security might be advantages for could be military or government computers, industrial control systems such as supervisory control and data acquisition (SCADA), and critically important systems like nuclear power plants and medical equipment (Nohe, P. (2018, March 13)). Because this is a sub set of physical segmentation it

obviously has the same disadvantages in regards to something like cost and an argument can be made about how it is wasteful and inefficient too, but in the circumstances in which air gaps would be used it is more than likely its important enough or the costs outweigh the protentional cost that this can be overlooked. However, due to the isolation, there are other disadvantages.

Air Gap Disadvantages

The additional disadvantages air gap has is due to the difficulty of managing such a system or environment. Unlike other connected systems administrators

, by design, cannot remote in via secure shell (SSH) or some other management software to apply updates or make configuration changes. Instead, someone, probably an administrator, has to determine that an update or change is necessary to go to a computer that is connected to the internet and download the patch or software change. Next, they have to load that on to a USB or laptop. Then they have to travel to air gapped site. In some cases, this may only be across the hall but it just also just as easily be a several hours drive from headquarters. Only after they get to the site can they affect any change to the air gapped environment. This may be worth entirely worth but what if they computer that was connect to the internet was compromised and then the compromise was transferred to the USB or what if the laptop used is infected? Then the issue is now the administrator has introduced the malicious actor or threat to the closed environment that, because of the difficultly in making changes, is not up to date patching wise and is very likely vulnerable to the thing being introduced (Byres, E. (2011, August 17)). In short, air gap segmentation is important and has its uses but it shouldn't be looked at the "Fort Knox" of networking because, just like many other instances, a careless human error can very easily introduce something that was previously thought impossible. A perfect example of this is that some researchers have a proof of concept that shows a way of using electromagnetism to

transmit binary data not only through the air gap but also through a Faraday cage, a cage designed to uses electrostatic and electromagnetism to prevent the transmission of data, to a receiver and thusly steal data that way (Khandelwal, S. (2018, February 08)).

Logical Segmentation

Logical segmentation is using software capable devices to create software isolated groups. This function similar to physical segmentation except it addresses several of the disadvantages that the physical method has, and these will be addressed a little in this paper. There may be more ways to employ logical segmentation, but this paper will only briefly discuss two of them which are Virtual Local Area Networks or VLANS, and Virtual Routing and Forwarding or VRF. An example of the former, VLAN, is shown in figure 3.

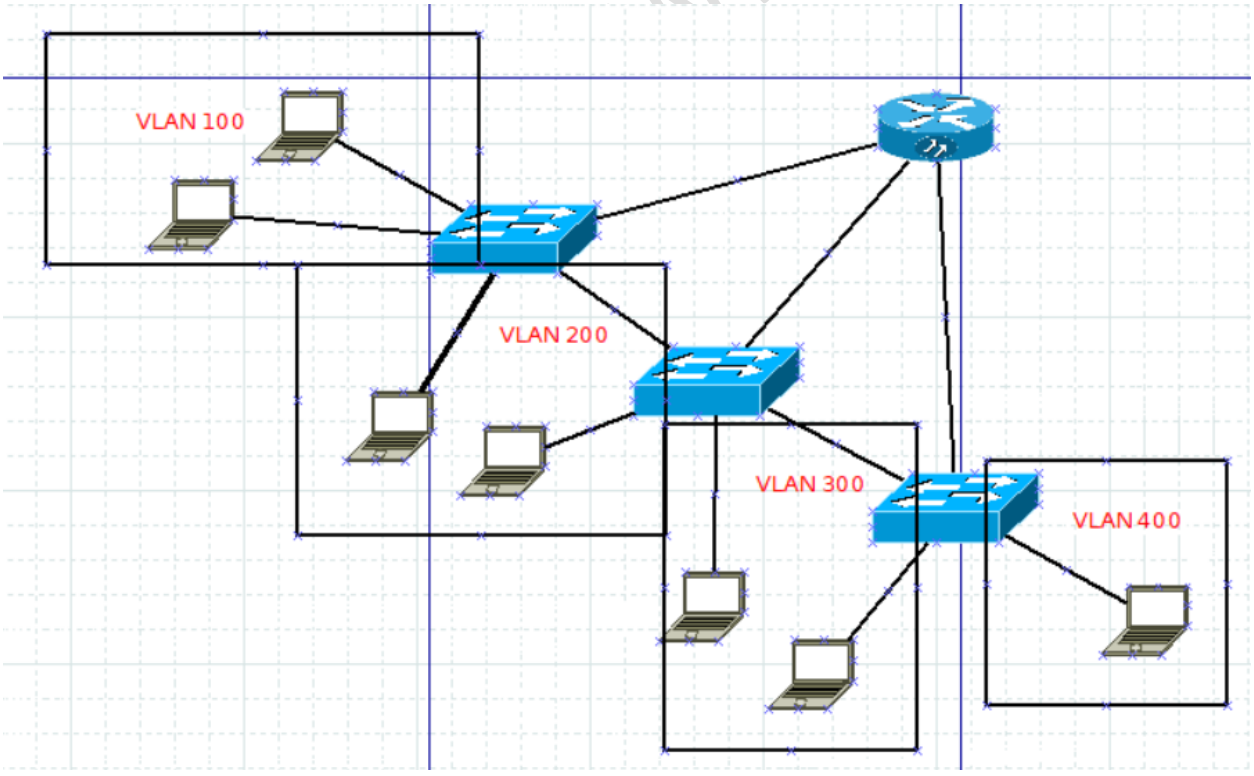

Figure 3. VLAN example

Figure 3 shows an example network with a single router, three switches, and four logically segmented networks. One thing that should be immediately apparent is that there are more logical segments that there are switches which would not be possible with physical segmentation. Furthermore, there are 2 VLANs that are not located on a single switch. That's because with VLANs the switch has a mac address routing table that allows them to route VLAN to other switches as long as the destination is on the same logical network segment. This is useful because it allows the administrator to configure remote resources like a printer or server to appear as if it is on the local area network. The reason to do so could be for security or for convenience but none the less it option is there.

Previously, it was stated that VLANs address some of the disadvantages that come with physical segmentation. The main disadvantages of wastefulness, inefficiency, and cost are all objectively lessened by using logical segmentation through VLANs. In physical segmentation, if there are unused ports on a switch they remain unused and wasted but with VLANs, those ports can be assigned to a separate segment and used there. Using this method, it may not even be necessary to configure a second switch altogether. In the original example shown in figure 1. Customer A and Customer B need to be segmented apart from one another and through physical segmentation, this is only possible through 2 switches. However, with VLANs, it would be possible to put both Customers, A and B, on separate VLANS but on the same physical switch. Doing so would make using a VLAN less wasteful, and cost less (fewer devices) than physical segmentation. Additionally, it would also be more efficient in regard to scalability because if another customer, customer C, needs to be segmented as well it would be possible to simply use any extra unused ports on the already installed switch to accommodate that. Also, with VLANs, it becomes possible to simply move, remove, or create network segments at will with some

change in the software configuration. This makes VLANs incredibly useful and more cost effective than the previously mentioned methods. VLANS are not perfect though and while there are fewer disadvantages than the other methods discussed up to this point they do suffer from a few.

The disadvantages of VLANs results from the fact that they are located on a limited number of devices and that making changes is incredibly easy. Traditional switches operate on layer 2 of the TCP/OSI model and as a result, can not route network or IP traffic. This means that they can not send traffic from one VLAN to another. Instead, the traffic will still need to be sent to a router to route it. There is one exception to this and that's layer 3 switches can route IP traffic, but that topic is not covered in this paper. This means that a router table entry needs to be made for each of the VLANS. On a small network with a limited number of VLANs, this is not an issue but as the network grows and the number of VLANS increases the routing table can become cluttered. This can increase latency and, in the extreme cases, can cause the router to malfunction. Similar to routing tables, switches create MAC tables to route traffic within the VLAN itself. this means that entries need to be made, and in the case of intra-VLAN communication the router need never be involved and if enough VLANs are created the same issues present. Additionally, there are attacks that are designed to exploit how the MAC table on a switch works. In the most basic of explanation these attacks work by forcing the switch to fill its MAC table to the point of using all of the available space and the resulting effect is that switch then acts similar to a hub and forwards traffic out of every port (Lauerman, K., & King, J. (2010)).

The next method the logical segmentation is VRF. VRF is extremely similar to VLAN because it allows for a single router to effectively have multiple routing tables. Because this is as similar as

it is to VLAN this topic will not be discussed in great detail. However, because the VRF configured router holds a separate routing table for each segment it effectively isolates the network segments by not providing, unless desired, a path to the adjacent networks. Routers already allow administrators to segment to the network via the use of the access control list and subnetting strategies. However, VRF takes this a step further but allowing the segments to be invisible to other network segments.

Segmentation Through Virtualization

The heading of this section may make it confusing because in the previous section VLAN and VRFs were discussed and both of those have virtual in the name. however, this section is different because those previous methods rely on vendor provided devices that include the functionality to use those strategies. This section is more about true virtualization of the entire network structure without the need for physical routers and switches. This is, of course, referring to Software Defined Networking or SDN. Software Defined Network has is a fairly recent development to the networking being a major topic of discussion in recent years.

SND is a more complex topic than this paper will delve in to but it is important none the less. SDN removes the reliance on hardware centric network configuration and provides a central place for everything. Literally everything. SDN works by virtualizing the computers, switches, routers, and most everything else. This makes it an incredibly powerful tool as it allows administrators to dynamically create, configure, move, remove, and revert to previous iterations with a few clicks and near instantly. This is a game changer because it makes managing a large network from a single simple, easy, and affordable. (What is Software-defined Networking?)

A basic example of a protentional use of SDN would be if in a large already created networking there were security policies that needed to be implemented, updated, or changed on multiple devices it can be done from a central control plane instead of allocating man hours for a human to physically go and make the configuration changes. To continue to put this into perspective imagine that the changes need to be made on physical devices that were geographically isolated from one another. That would mean that either one person or multiple people would have to travel to the location of the device and make the changes manually by connecting to the serial port of the device. This could present with issues because what if mistakes were made and the configuration change that was made was done incorrectly. This would have more hours to troubleshoot, find, and rectify the human made error. The error rate can be reduced with the use of automation like scripts of configuration files, but they would still need to gain access to the physical device. With SDN those issues seem to disappear. A human made error can still be introduced but the troubleshooting and the rectifying process can be completed centrally and completed much faster. With something that sounds as great and SDN does it may be hard to imagine that there are any disadvantages. However, even though SDN is still a relatively new technology there are some apparent disadvantages.

SDN Disadvantages

One of the disadvantages of SDN is that because it is so new there are few people who are already trained in it. This means that if a company wants to use SDN then the current staff will need to be trained on SDN and the management tools that are used with it. This can be costly in both time and monetary ways. The second disadvantage is that SDN is a completely different take on network infrastructure. This means that in order to implement SND into an existing network infrastructure major configuration changes will need to be made to accommodate it. The

last major disadvantage of SDN that will be discussed in this paper is that SDN introduces a single point of failure. If the SDN controller goes down then so does the network and the ability to make changes to it. (“Advantages of SDN”)

Advantages of Segmentation

Putting aside that segmentation has been considered a best practice policy for over the last decade, segmentation provides a lot of quality of life advantages that most everyone in the industry will agree are fundamentally necessary to meet with the demand and requirements of networking today. The main advantages that will be discussed in this paper will be performance, security, compliance, and management.

Segmentation provides performance advantages but allows the administrators to configure service only segments. For example, if there is a service or application that requires a high about of bandwidth then that can be addressed through segmentation by putting that service or application on its own the dedicated segment with a high bandwidth allocation reserved for it and isolated from other protentional bandwidth sinks.

Segmentation also provides an ease of management advantage. This is a simple but drastic advantage because it can often be taken for granted. If the network was on a giant conglomerate of network devices then any change required would have to be applied to every device. For example, if the network of one giant single network and it was necessary to add another set of devices then every device in the network would need to be updated to include a pathway to and from these added devices. This could be extremely cumbersome and time consuming. With network segmentation, the number of devices that would need to be updated can be drastically lowered. For example. If the network was segmented into 5 zones connected in a bus topology

and then a sixth zone was added on to the end only the directly connected router would need to be updated and it is the only one affected by the change. The other zones would function normally and send traffic out of the default gateway and on to the next router to reach the newly added zone.

the next major advantage of segmentation is the effect it has on compliance. For example, to be PCI DSS compliant the network would have to meet some pretty stringent security needs for everything in scope. However, with the use of network segmentation that scope can be reduced and thus make it easier to become compliant. In fact, PCI themselves have made available a paper outlining how to use network segmentation to reduce the scope and help with meeting compliance. P. (2016, December)

the last major advantage of network segmentation is that it increases security by helping enforce security policies. This works by using the least privilege or access policy which basically means that those who do not need access to a resource should not be given access to that resource. Segmentation makes the process of preventing that access similar and easier by limiting the access points to something which can them be secure with firewalls, access control lists, or even intrusion detection and prevention systems.

Conclusion

in conclusion network segmentation is a very important part of designing a network. This paper has discussed how physical segmentation works in that it separates networks by using separate physical devices and that doing so has the disadvantages of being wasteful, inefficient, and costly. It also has discussed that air gap segmentation is an extreme version of physical segmentation and comes with its own risks even though by design it should be completely

isolated from the rest of the world. Also, how one can use VLANs to achieve logical segmentation was talked about and how there are a few but limited disadvantages to doing so. The final method discussed was to use SDN to virtualization of the networked which boasts some impressive benefits when it comes to managing, configuring, and changing the network on the fly. And last the paper briefly discusses some of the major advantages provided through network segmentation in the forms of added security, ease of management, standard compliance, and performance improvements. In short, every network administrator should know at a fundamental level how the various types of segmentation work along with the advantages and disadvantages of each. Without a strong basic understanding of segmentation is may not be possible to create what is best needed for the given situation. If extreme security needs like some government instances in the form of air gaps? Or is ease of management most needed in the way that SDNs provide? The only way to make the best decision is to be armed with the knowledge of the various methods so that one can make the most informed choice.

References

About Security Zones. (n.d.). Retrieved from https://www.cisco.com/assets/sol/sb/isa500_emulator/help/guide/ag1463340.html

Advantages of SDN | disadvantages of SDN. (n.d.). Retrieved from http://www.rfwireless-world.com/Terminology/Advantages-and-Disadvantages-of-SDN.html

Buns, T. (2017, December). Designing a Network with Segmentation. Retrieved October 12, 2018, from https://infosecwriters.com/Papers/tbuns_NetworkSegment.pdf

Byres, E. (2011, August 17). SCADA Security's Air Gap Fairy Tale. Retrieved October 12, 2018, from https://www.automation.com/automation-news/article/scada-securitys-air-gap-fairy-tale

Fort, J. (2018, July 09). The advantages of network segmentation. Retrieved November 7, 2018, from https://community.jisc.ac.uk/blogs/csirt/article/advantages-network-segmentation

Khandelwal, S. (2018, February 08). Hackers Can Now Steal Data Even From Faraday Cage Air-Gapped Computers. Retrieved October 10, 2018, from https://thehackernews.com/2018/02/airgap-computer-hacking.html

Kirkpatrick, K. (2013, September). Software-Defined Networking. Retrieved October 12, 2018, from http://delivery.acm.org/10.1145/2510000/2500473/p16-kirkpatrick.pdf?ip=150.216.255.46&id=2500473&acc=ACTIVE SERVICE&key=A79D83B43E50B5B8.D2E2D13E69DBEDD9.4D4702B0C3E38B35.4D4702B0C3E38B35&__acm__=1541988178_5a5c5d6b1e7139b8015f89141e3fbad4Lauerman, K., & King, J. (2010). Layer 2 Attacks and Mitigation Techniques for the Cisco Catalyst 6500 Series Switches Running Cisco IOS Software. Retrieved October 20, 2018, from https://www.ndm.net/ips/pdf/cisco/Catalyst-6500/white_paper_c11_603836.pdf

M. (2009, December 10). What is VRF: Virtual Routing and Forwarding. Retrieved November 5, 2018, from https://www.plixer.com/blog/netflow/what-is-vrf-virtual-routing-and-forwarding/

Nohe, P. (2018, March 13). What is an Air Gapped Computer? Retrieved November 5, 2018, from https://www.thesslstore.com/blog/air-gapped-computer/

P. (2016, December). Guidance for PCI DSS Scoping and Network Segmentation. Retrieved October 12, 2018, from https://www.pcisecuritystandards.org/documents/Guidance-PCI-DSS-Scoping-and-Segmentation_v1.pdf

Piyevsky, S. (2015, October). Logical Segmentation and VLANs – An Overview. Retrieved October 16, 2018, from http://www.industrial-ip.org/en/industrial-ip/convergence/logical-segmentation-and-vlans-overview

The Business Case for Network Segmentation Modern network segmentation to reduce risk and cost. (n.d.). Retrieved October 12, 2018, from https://assets.extrahop.com/whitepapers/ExtraHop-Modern-Network-Segmentation.pdf

What is a Virtual Local Area Network (VLAN)? - Definition from Techopedia. (n.d.). Retrieved November 12, 2018, from https://www.techopedia.com/definition/4804/virtual-local-area-network-vlan

What is Software-defined Networking? (n.d.). Retrieved November 5, 2018, from https://www.citrix.com/products/citrix-adc/resources/sdn-101.html