Designing a Network with Segmentation

Timothy Buns

ITEC 6880

December 2017

Abstract

Network design entails many different angles, and it should take some very careful planning when implementing a new network or configuring a redesign. In this paper we will be looking at the logical point of view of how to separate the different parts of a network. Almost any network is going to include a data center that includes servers. However, there may different use cases for those servers, such as web servers, that need to be internet accessible, or databases that should be kept secured. There are also different types of users and even devices that should be taken into consideration in terms of how to group different functional unites together. Being that there is no one size fits all blueprint, sometimes a design can get out of hand and too specific, and of course this will depend on the company. There are many different ways to design a network in terms of segmentation. Using Virtual Local Area Networks (VLANs) to do this is the most relevant avenue for doing this. However, one of the challenges that can be seen is how and where to use VLANS. In this paper, the different methods will be discussed and which is deemed to be best practice. We will look at some of the common methodology used and the pro's and con's of each. In addition to how to properly segment a network, we will also look at the how to into increase security by using techniques in conjunction with grouping different members of the network. The basic principal of least amount of access can assist in protecting against possible breaches and minimize damages if one were to occur.

Introduction

When looking at all the different types of variants that touch a network, it can be daunting to how to make sure each item should be treated. Many decisions will be based on the company and how that company operates. Emphasis will be placed on certain areas, based on function and what it relies upon to stay functional. For some, it could be people. Take, for example, a call

center or dispatch center for a police department. Phone lines are critical and the computers they use are key to helping them do their job. Manufacturing will take a completely different approach. While they may rely on people to perform the jobs, machine to machine communication may be the critical piece that enables the humans to do the jobs and producing.

In this discussion, we will look at the specifics of using VLANS and VRF's in an enterprise network. There will be certain use cases for having one or both methods configured. We will look at some examples of some common configurations that are used when implementing either technique and some of the advantages and disadvantages.

Segmentation

You will often see in today's network distinct functions separated. In order to achieve this, VLANS are put in place. In brief, VLANs provide a method to reduce the amount of broadcast domains and confine the hosts to only can easily communication with other hosts within that domain. If a host needs to communicate with a host in a different domain, there is only one entrance and exit to go through. This provides two advantages. Having multiple domains or VLANS keeps the conversations with that group. This also provides extra security again by limiting what VLANS can have communication with other VLANS.

Taking a step back to when networks were in their infancy, it was even suggested to design networks to have large VLANS for ease of management. What made this feasible was that there just weren't as many devices connected to the network at that time. While routing was used at this time, it was quite expensive to have layer 3 devices. Using layer 2 networks with wide broadcast domains made financial sense at that time as they were cheaper and faster. Networks did not play the roles that they do today and the technology certainly was not what we see now.

So what is the best way to segment the network?

When looking a network design and how devices and users should be separated, the real question needs to be asked of the business is how it functions. A manufacturing facility is going to have much different requirements then a healthcare facility. Also, are there any security policies that need to be designed around. PCI and HIPPA standards are very big hurdles to climb and making sure that the network meets compliancy. There are times when it makes logical sense to place an area of the network in the same VLAN, but due to standards it is a better practice to keep them separate. This is brought out by Gary Glover in his article on PCI where he highlights that (10) "non-segmented environments, or "flat" networks, have their card-processing systems mixed in with back-office systems. In these environments, the entire network is in scope for PCI DSS compliance. This can significantly increase the amount of work needed to secure your business's network."

VLANS

Let's briefly look at what a VLAN is and how it works which will help define how they are implemented. VLAN which stands for Virtual Local Area Network is essentially a group of network devices. We know of VLANS today as broadcast domains. When a computer is connected to the network, there is generally lots of conversation traffic that goes on behind the scenes that the user is not aware of. At times a device on the network will announce itself in form a broadcast. This is an easy way to let all of the other devices know that they are here and are active. These domains can be thought of as rooms in a building. It is much easier and faster to share information with someone that you are in the same room with. This is one of the benefits of a layer to network is there is no intermediary between two devices and simples the communication process to an extent. At least one of the downside of this is when there are too

many people in the same room trying to communicate. Ten people in a room may not be too strenuous, but as that number rises to 1,000 for example, the same conversations will be much more difficult. This is generally why it is best practice to try and keep these broadcast domains or VLANS smaller in size. When a person needs to communicate with someone from another room is where routing comes into the picture and essentially there is a device with routing functionality in the middle handling those types of conversations.

So, at this point we have discussed VLANS and have a good understanding what they are and do. Without getting too far down in the details of where to segment it would be equally as good to get a simple understanding of layer 3 segmentation.

VRFs

Once traffic tries to communicate with another device outside of its domain and routing takes over, this is another aspect of segmentation that can be designed to remain in place. One way of doing this is using a VRF which stands for Virtual Routing and Forwarding. (8) "VRFs allow for complete separation of different routing instances from one another. This simple and effective concept of hiding networks from each other and limiting the ability of devices from interacting outside of defined boundaries creates a more secure network."

This is a straightforward way of creating multiple routers on a single device. So, again if a person wants to talk to a person from another room it must talk to the router. Within a single instance, a router contains a table or a lit with directions or routes on how to reach each destination. At this point all routes are held in one table. Implementing VRF's allow for that separation. Using multiple routers keeps traffic separate and not able to easily communicate with hosts that reside in other routers as they also have separate routing tables.

What is the difference between VLANS and VRFs

After discussing VLANS and VRFs a question may arise as to what is the difference? Besides one of the obvious main points of VLANs operating at the layer 2 level and forwarding traffic based on mac address and VRF's operating at the layer 3 level routing traffic based on IP addresses, they are indeed very similar in they provide the means to segment. The function is the same and both will separate traffic and even provide the ability to implement security policies blocking traffic. Take for example and network that had a staff network and then a production network. It would very simple to create a VLAN for each area and then implement an ACL on each interface that would restrict both VLANS from communicating with each other. However, there may be a requirement to totally isolate the traffic from end to end. This is getting more into the security side, but even with ACL's in place to restrict access, they would eventually still follow the same path to a destination as they would share the same routing table. With VRF's implemented, true isolation could be achieved. Due to there being two routing tables, traffic can take their respective paths. This also makes it easier if there are two internet providers in place where staff traffic needs to go out one link and guest traffic needs to go out the other. With one routing table it makes it much more difficult to achieve this.

Another likely scenario where VRF's would be seen is for an Internet Service Provider(ISP). This is a common way for providers to use the same equipment while keep traffic protected and private from other customers. Having two routing tables provides the ability for customers use overlapping IP's for any specific case they may need it. This would definitely would not be possible using VLANS.

Another area of questioning is how big or small should VLANS or network segments be? Again, there is no firm right or wrong answer here and it may depend on the network. However, there

are best practices when it comes to this part of the design that should be taken into consideration in order to keep the network running at its best. This can relay into a more of a subnetting topic, but the general rule of thumb is to keep subnet size to a subnet size of 254 or 512 addresses. This does not mean that they cannot be any bigger and there is certainly nothing wrong with using smaller ones. When it comes to determining how big or small a VLAN or subnet should, it can be up to one's own interpretation or environment on what is a correct answer.

Advantages of VLAN segmentation

In the infancy of networks, to increase the size of a network, hubs were the quickest and cheapest way to do this. (3) "The network hub is the simplest device serving as a physical layer interconnect. "Due to how hubs work, sharing a single segment, it caused collisions on the network, which, when multiplied could cause majors issues for communication between hosts. (3) "The limitations described have made hubs virtually obsolete. They should not be used in modern substation installations. "Thankfully, switches have become more economical and are able to be used more often to remove this issue. However, even though switches remove collisions it still does not equate to being able create VLANs as big as you want. As touched on earlier devices still create a considerable amount of traffic, so best practice suggests there still be limitations. So, while there should be a max amount of devices in a VLAN, there are many advantages to have them grouped together.

Logically, it helps ease any pains of management. (9) "The ability to contain network problems, improve performance, and reduce congestion are all key benefits that come from a well-segmented, well-maintained network." Even, if the network is small enough to only need one VLAN per building or section of a building, it is nice to know when looking at addresses where clients should be located or what they should be doing. Which this brings out another great

aspect. VLANS can reside in multiple switches if need be. So, if there was a staff VLAN created, but there were staff located in two areas, the same staff VLAN could be added to the respective switch in the area to maintain the same grouping and separation from other VLANs. In terms of switching, this is very helpful with troubleshooting as well. (6)" By grouping our network users and resources into different VLANs, problems emanating in the network can easily be identified and fixed by mere tracing group such hosts belong to." For instance, if there is a client connectivity issue. Knowing where the device is located can help pin point where a network problem may be. Having a structured VLAN environment provides the tools to do this.

Touching on security, VLANs provide advantages in this arena as well. (7) "Because VLANs segment a network, creating multiple broadcast domains, they effectively allow traffic from the broadcast domains to remain isolated while increasing the network's bandwidth, availability and security." Due to clients not being able to communicate with other clients on different VLANS, it sets up a great way to maintain security boundaries between functions of the network.

VLAN design

Creating the correct VLAN design is not going to be exact for each environment, but there will be similarities. However, there could be a reasons to put different areas of the network into their own VLAN. What should be cautioned against here is not to get too overloaded with minor details and really focus on the functions that each VLAN should serve. (4) "The organization must be careful not to simply through the VLANs at the network and they must make sure that the solution does not generate more network administration then it saves."

For most modern day networks there will certain functions that should always be separated. Network management devices will generally be on their own management VLAN. There are a

number of validations for doing this. For one, if there is an issue going on within the network where clients were having connectivity issues, it would make troubleshooting that much more difficult if getting access to the switches were not possible, which would then, require physical access to the device.

Bringing up the security aspect, it is most definitely bad practice to co-mingle regular network traffic with management. This would essentially allow any user or device access to the administrative prompt. While they would still need the username and password, it is one less hurdle for someone to need to make.

For placement of users, here again it depends on the company. There are many examples of using departments as separate VLANs such as Finance, HR and the list could go on. This can be quite a bit of overkill. It may depend on the size of the company and even what the company does. If the finance department handles critical data such as social security numbers and other sensitive data, it may warrant creating a separate VLAN segmented from the rest of the user population. This could be crucial if a security breach happens. There are certain viruses, that once it is able to infect a computer on the network, it then starts probing on what else it can find and install on. Separate VLANs minimizes those types of risks of infecting other machines on the network. This is brought up here as the last thing a company wants to here is that their financial data could be at risk.

Getting back on topic though, in general users can be grouped together no matter what the job may entail. This can simplify and reduce the amount of unneeded VLANS. However, there is the exception of guests or non-employees will be allowed to use the network. Undoubtedly, guests should at the very least be segmented from other traffic. While, this isn't a security discussion, some companies even choose to take the segregation step further and put this type of traffic into

what known as a DMZ. In short, a DMZ is a network that sits on the other side of the firewall totally separate from the reset of the network traffic. It may depend on resources, but if a store for example wanted to offer free Wi-Fi to its customers, it may warrant purchasing a separate internet connection just for this purpose. While this goes above and beyond of network segmentation, it and leaves no question that those two networks will never even see each other.

Now that the data VLAN is covered, there are other functions that need to be addressed as well such as printing. Depending on the size of the network, it is not uncommon to see printers mixed with the general data VLAN. On the contrary there may be a need to separate out the printers into their own VLAN or at least in a different segment so that the users can not directly reach them. Many companies have print server's setup that they want to ensure users must use. In a university setting, the campus may have centrally located printers for all to use. Since these have consumables that must be maintained, users will have an account tied to their credentials and using a print server will allow their account to be charged. It would be too easy for users to find and print directly to the printers if they were on the same VLAN. Another issue that comes up with printers, is that the firmware is sometimes a forgotten piece of software. Just like anything else, they are just as susceptible to vulnerabilities just as much as computers. While keeping them in a separate VLAN is not a free pass to not update the software, it is equally not as common for manufactures to not be vigilant to make those updates available.

Another commonly segmented part of the network is for security devices. More and more companies have the need for security cameras. Keeping this devices separate serves many functions including the main one being that this would users should not or other devices should not be able to reach the cameras directly. In addition, some of the more tech savvy users may even be able to intercept video streams which would not be wanted. Keeping all camera related

activity may actually even help performance as traffic would not need to involve the need of routing and be able to directly reach the server it is sending the video feed to.

Voice over IP(VOIP) is also considered to be normal practices these days. This means that telephone conversations are going to be traveling the same lines that all of the other traffic is using as well. Leaving the security aspect aside, grouping all voice or even video traffic together does many things. The best reason is Quality of Service(Qos). (5)" The classification of the traffic can be based on the application(s) involved, the connection (source/destination addresses), or the protocol used. To be effective, this QoS must be provided "end-to-end"; that is, across various individual network elements and networks, which are between the "end" points." This a way to give this type of packets a special mark and the network can be configured to give any packet with that special mark priority over all other packets. To briefly explain, many applications today use the TCP protocol. These are where packets are guaranteed This is similar to what radio dispatchers use to confirm the other party received their message. When a dispatcher calls on a hand radio to a recipient, the receiving person says copy to relay to the dispatcher that they received the message. Voice and video applications generally use UDP protocols because it is much faster as reception of the packets do not have be confirmed. The downside is that the packets may never get delivered. This can sometimes happen during a phone call conversation that the callers may notice the call may have some jitter. While we don't necessarily appreciate the jitter it is probably a better option than having bits and pieces of the conversation showing back up a few seconds later. Being able mark these type of packets and configure them to have priority helps provides a better user experience. In some networks, bandwidth may not be a problem, but this may not be the case for all. Some sites may be limited in the bandwidth they are able to provide. For instance, say you have a remote site in which only

a 10MB service can be afforded. This would be an absolute true use case for enabling QoS ensuring that you want to make sure get to use that pipe if it is full.

A server VLAN is almost expected in today's network. Depending on the security needs there may even be multiple. One of the benefits to having the servers grouped together in one broadcast domain is the speed. When routing is introduced to a network path it can add extra hops and time to a network conversation. (1) "Compared to switches, routers require more processing of incoming traffic. As the volume of traffic passing through the routers increases, so does the latency in the routers, which results in reduced performance." Often times servers need to communicate with each other and being in the same VLAN creates a more direct conversation in addition to placing less stress on the network to not have to deal with the routing. This will depend on the severs being in the same physical location as some companies have data center spread out across the region. Furthering in the security discussion, it is much easier to keep a solid layer of protection for a group as opposed to having to care for individual machines throughout the network. In terms of access, access lists can be configured so that only certain segments of the network can contact the server network. This also suggests another best practice. As mentioned, there may be a need for multiple server VLANs. So, there should definitely be consideration for servers such as web hosting or applications that are reachable from the internet to be placed in their own segment. As these are more susceptible to attacks, it is best to keep them separate from not only the other servers but from any other part of the network as well.

Disadvantage of VLAN Segmentation

Whenever you are looking at the advantages of a concept, it is natural to look at the disadvantages. However, I don't necessarily see any downsides to VLANS as they are more a fact of life in a network and will be used. The disadvantage I see is all in the design and how a

network is configured to use them. We discussed different examples of how segmentation can happen where an administrator may assign certain functions. To reiterate, even though there are many advantages to keeping network segments small there is also the caution to not end up with too many VLANS as this is can be harder to maintain and can put more strain on the network sacrificing speed.

VRF Segmentation

Now that we have discussed segmenting the network using VLANS at the layer 2 level. There is the next discussion of how to segment the network at the layer 3 level using VRF's. A VRF is a logical router. There are two flavors to choose from, VRF which is generally seen in ISP environments and VRF lite which are generally seen in enterprise networks as it strips off some of the unneeded functionality that would not be used.

When thinking of VRF's and how to segment the network, the same thought process should be used as when dealing with VLANS. What function will the particular VRF be serving? More than likely there will be much of the same design and thinking that the VLAN structure is already using. While there may be a one for one function, this is another great avenue to be thinking about how you want the network to be separated. The one area that VRF's bring to the table is security.

VRF's can provide some flexibility of making segregation changes to the network while at the same being less disruptive. Let's look at one example deployment method of how VRFs can be used in an enterprise format. When looking at VLAN design we already looked at the different functions how what typical configuration would be in place. What's to be noted though is that larger networks may need more VLANS that fall into the same functionality. What this is

eluding to is that larger networks will have bigger populations that go over the best practice guidelines of how many devices should be in a subnet. Another use case scenario would be if there are is a large campus with multiple buildings it may not be best practice to use the same VLAN in every building even if there would not be enough clients to fill up the configured subnet. This can cause some management pain especially there are specific security standards in place. Configuring VRFs for this type of scenario is a great method of grouping multiple VLANS with the function into a group of their own. Any kind of security policy such as access-lists can be shared and implemented at on place at the router level. While this a great way to segment traffic, the amount of processing power needs to be kept in mind as well. When looking at where to place servers and storage, if they are placed in separate VRFs, that traffic will be going across routing interfaces. This may or may not be acceptable, but just another factor to keep in mind when thinking design.

As we have discussed VLANS and VRFs and how to use both methods for segmentation it may seem as though that there are many similarities. The question may arise that would it be necessary to add the complexity of VRF's into a network? The answer will be different for each specific network. More than likely, the smaller the size of the network the less likely VRF's will make sense. The larger and more wide spread networks are the use cases where they will make the most sense. The example was brought out of how a network may have multiple campus building with similar needs in terms of segmentation at each location. VRF's allow for a much more detailed containment of the VLANS. When multiple VLANs are needed that are functionally similar, VRF's enable the ability to easily create these for each location and at the same time keep the specific VLAN traffic local to that building or region.

Advantages of VRF environments

When thinking about widespread VLAN use this could potentially cause some network issues. One way to prevent network loops and broadcast storms on a layer 2 network is STP. (2) However, “the biggest issue with extending VLANS across multiple switches is spanning-tree loops.” This is a great protocol for what it can do. However, if a VLAN was configured for many locations, STP would be heavily relied upon to perform its regular function at multiple points on the network. Another downside in this area, is if a loop were top happen on the network, even though this should only effect the VLAN where this has occurred, it is still sharing the same medium or pipe as the rest of the network traffic. This could easily cause issues on other segments of network as they are now having to compete with extra bandwidth being consumed by the problematic VLAN. So, while this is a possible way to use STP in a widespread format and for the network to perform as normal, it is probably not the best use case and would be suitable for implementing VRF’s.

Utilizing VRFs provides the ability to isolate the layer 2 level of the network while at the same keeping the segmentation needed.

Security VRFs provide

Using VRFs provide many benefits in terms of segmentation, but also opens many avenues when looking at the security aspect. As we have been discussing it may be necessary to have different functional segmentation throughout a network. In a large network this can make it difficult to manage in to order to maintain any kind of security boundaries. In addition, when trying to secure multiple locations it can become very costly. In large campus type of network, there will be a need to have multiple network locations serving many of the different VLAN designs as previously discussed. If there was need to have firewall capabilities for the entire network, it is possible a firewall would need to be placed at each location.

Making use of VRFs can greatly simply the network in addition to reducing costs. A common way to look at this design is to think of a network in zones. Just as previously discussed there will be multiple zones or VRFs based on the function they serve. All zones will be configured to route back to the firewall in order for communication to take place between zones. With this type of scenario, the campus firewall can be utilized for implementing rules on how what communication can take place. This is a huge benefit from a management perspective in this can happing in one place opposed to dealing with multiple access-lists in a variety of locations. No matter where a server or device is located in the network, as long as it is placed in the correct zone it will be bound to those rules.

It can almost be thought by having VLANs in addition to VRF's in a network provides multiple levels of security.

Conclusion

Much has been discussed in terms of network segmentation. There are multiple different ways to think about the design of network and how to logically separate out the different functions and devices that are connected. We looked at doing this at no only the layer 2 level of the network, but also the layer 3.

VLANs and VRF's were discussed at a high level and showed how they could provide many avenues to think of on where to place network functions. There is management aspect to being able to have a good understanding of where the clients are at on the network and what they are doing. There is also, the health aspect of the network where VLANs provide the ability to limit the size of the broadcast domains creating a better environment for communication. While VRF's provide the additional segmentation abilities to keep device communication within the

zone or domain in which they are intended for. Having one or both of these methods in the network also provide layers of network security that can not only be required by the company but also by standards in which must be complied with due to the type of data that crossed the network.

It was stressed that knowing how and which methods to use in a network will highly depend on the use case of the network, how large it is and namely how much resources a company has to use. Having a good understanding of how the company works will provide the ground work and baseline for VLAN considerations and whether or not VRFs will assist or only add to the complexity of the network.

References

1 Varadarajan, Suba. (No Date). Virtual Local Area Networks. Retrieved from http://www.cse.wustl.edu/~jain/cis788-97/ftp/virtual_lans/index.html

2 Cisco press – Barnes, D. Sakandar, B. (2005) Cisco Lan Switching Fundamentals. Cisco Press. Retrieved from https://books.google.com

3 Leischner, G. Tews, C. (2007). Security through VLAN Segmentation: Isolating and Securing Critical Assets Without Loss of Usability. Retrieved from https://cdn.selinc.com/assets/Literature/Publications/Technical%20Papers/6277_SecurityVLAN_GL_20070329_Web.pdf

4 Amin, Mohammad Rohul. (2007) Analysis on the Virtual Local Area Network to Reduce Collision Domain and Manage Broadcast Domain. Asian Journal of Information Technology, 6(10) 1003-1009. Retrieved from http://docsdrive.com/pdfs/medwelljournals/ajit/2007/1003-1009.pdf

5 Klincewicz, J.G., Schmitt, J.A. & Wong, R.T. Telecommunication Systems (2002) 20: 81. https://doi-org.jproxy.lib.ecu.edu/10.1023/A:1015441400785

6 Odi,k A. C. Nwogbaga, N. E. Chukwuka O. N. (2013) Thr Proposed Roles of VLAN and Inter0VLAN Routing in Effective Distribution of Network Services in Ebonyi State University. International Journal of Science and Research. Retrieved from http://www.ijsr.net/archive/v4i7/SUB157109.pdf

7 Ali, M. N. Hossain, M. E. Parvez, M. M. (2015) Design and implementation of a Secure Campus Network. International Journal of Emerging technology and Advanced Engineering, 5 (7) 370-374. Retrieved from http://www.ijetae.com/files/Volume5Issue7/IJETAE_0715_63.pdf

8 – Rossi, Jeremy. (Aug. 15, 1009) VRF is the new Black: How I learned to Stop Worrying and Love the Complexity. Retrieved from https://jeremyrossi.com/blog/2009/08/15/vrf-is-the-new-black-how-i-learned-to-stop-worrying-and-love-the-complexity/

9 – Harrison, Reuven. (June 16, 2014) Network Segmentation Key to Good Network Hygiene. Network Computing.com. Retrieved from https://www.networkcomputing.com/networking/network-segmentation-key-good-network-hygiene/2096831069

10 – Glover, Gary. (March 11, 2015) How Does Network Segmentation Affect PCI Scope. Securitymetrics.com. Retrieved from http://blog.securitymetrics.com/2015/03/network-segmentation-pci-scope.html