

# University of the Free State: Redesigning the campus network for innovation, VoIP, multicast and IPv6

Andrew ALSTON<sup>1</sup>, Louis MARAIS<sup>2</sup>

<sup>1</sup>Alston Networks

*aa@alstonnetworks.net*

<sup>2</sup>University of Free State

*maraisl@ufs.ac.za*

## Abstract

Starting in July 2012 the University of the Free State embarked on a topological network redesign and implementation. This was in order to rectify performance issues on the network, and to facilitate the deployment of technologies such as IPv6 and Multicast. The redesign also had to cater for large scale innovation as well as network scalability. Once the planning had been done, the implementation of the new design was performed over a period of 3 weeks, with zero downtime to any member of staff or faculty. Today, to our knowledge, the university is now the largest destination of IPv6 sourced content on the continent with significant percentages of their traffic running over IPv6.

The paper explores the design process and the implementation process, as well as look at the benefits gained from the implementation now that it is complete. It explores the benefits behind some of the technologies that are now useable on the campus, with an emphasis on the multicast deployments done in conjunction with TENET.

## Keywords

IPV6,campus networks,multicast deployment,network design

## 1. Introduction

The University of the Free State has a large campus wide network, spanning approximately 160 buildings, and comprising of approximately 19000 wired network ports, 450+ CCTV Cameras, 500+ Switches and in excess of 400 wireless access points. At the start of the project described in this paper, this network was a single flat network utilizing a single /16 worth of IPv4 address space, in a single massive broadcast domain.

The single broadcast domain was causing significant issues, both in terms of stability and performance of the network. This paper looks at the process whereby the network was redesigned and subsequently redeployed to a more structured architecture. We will explore the motivation behind each of the choices made during this process, and look at the issues experienced during the course of the project.

## **2. The Original Network:**

As stated in the introduction, the network was originally flat with no segmentation. This meant that there was no routing protocol in use on the network, no VLAN segmentation and no true IP allocation planning. Because of the flat nature of the network, and the single large broadcast domain, issues were being seen with overloaded ARP tables, regular broadcast storms and excessive CPU utilization on the switching infrastructure. In addition, again due to the lack of segmentation, spanning tree and topology loop issues were being experienced.

Apart from the stability and performance issues, the flat topology also limited what technologies could be efficiently deployed on the network, and in order to properly deploy technologies including, but not limited to, IPv6, Multicast and QoS, it was imperative that the network be redesigned.

## **3. Project planning:**

As with any project of this size and scope, it was critical to first define a set of objectives, and to be clear on the intended outcomes of the project. This was important both for the design phase and for a post implementation evaluation of the success of the project. Once the objectives were clearly defined, as discussed later in this paper, a design was drawn up and analyzed for prerequisites. When looking at the implementation prerequisites, we analyzed the following things:

- a) Was the hardware currently in place capable of performing the desired functions
- b) Was the software utilized by routers/switches covered by sufficient licenses to enable the necessary features
- c) What were the time frames needed for implementation, and how would these affect the users of the network during the implementation period
- d) Did we have sufficient address space to cater for the new design

Once we had looked at each of these aspects, it was possible to rectify the areas where prerequisites were not met, and then plan the implementation process in more detail.

The project planning phase was also divided into phases, since certain pre-implementation project planning could only be done once other prerequisites highlighted by the first stage of planning had been met. This was particularly relevant with regards to the detailed IP address planning for the new design, which will be discussed further down in this paper.

Only once the planning had been completed did we proceed to the implementation phase, and it was due to the careful planning done that the project was implemented in such a successful manner with minimal disruption to the campus user base.

## **4. Design Objectives:**

The design of the network was created to meet specific criteria, and while some of these criteria were fairly generic, knowing exactly what the objectives of the design were helped to ensure an optimized work flow in the design phase.

We used the following list of objectives as our design criteria:

- a) The design had to scale as the network grew, both in terms of traffic and in terms of the number of edge devices connected to the network.
- b) The design had to cater for IPv6, and this was intricately linked to scalability. It was acknowledged by the design team that with the global depletion of IPv4 space, in order to create a truly scalable network, it was critical to ensure that the design did not rely on IPv4.
- c) The design had to cater for multicast. While multicast is still not commonly used on the commodity internet, there is heavy use of it in the academic sector for academic video content distribution. In addition, after conversations with commodity video content providers, it was clear that multicast on campus would be advantageous.
- d) The design had to cater for VoIP. This included ensuring that in the scenario where the campus chose to outsource or virtualize their PABX, the design would allow for this. It also meant that the ability to implement QoS where necessary had to be in place.
- e) The design had to be flexible enough to support innovation and research, while at the same time being structured enough to be stable and easy to maintain.
- f) The design had to eliminate the use of network address translation. This decision was taken due to the complexity involved in the concurrent use of network address translation and other technologies described in the rest of the design criteria. The elimination of NAT was also prompted by an analysis of security issues created by such translation.

## 5. The chosen design

Having looked at the objectives described above, the following decisions were taken:

- a.) The network would be built on service provider principles, rather than standard LAN network architecture principles. This was motivated by the fact that service provider networks are typically more robust and more flexible than networks designed using LAN architecture principles.
- b.) The network would be a three layer network, consisting of a core, a distribution and an edge layer. In addition, where necessary, there would be an additional L2 aggregation layer, this additional layer would decrease the number of required distribution points and hence the costs involved.
- c.) OSPF was chosen as the underlying IGP for loopback distribution purposes. OSPF was second prize for the design team, with IS-IS being the protocol of preference, however, due to limitations in deployed network hardware, OSPF was found to be the only option.
- d.) It was decided that OSPF would only carry loopback addresses and point to point addressing. All other routing would be carried via BGP. The choice to use BGP was to allow for scalability and flexibility. It also opened the option for more delayed analysis of traffic on the network when using Netflow/IPFIX/SFlow.

- e.) OSPFv3 was the protocol chosen for all IPv6 routing. Ideally, the design team would have preferred to deploy the IPv6 in the same way the IPv4 was to be deployed, however due to lack of support for IPv6 in BGP on the currently deployed hardware, this was not possible.
- f.) A decision was taken to ensure that no VLAN's spanned between distributions. This was to both keep broadcast domains small and manageable, and also to reduce the possibility of potential topology loops. The functionality for point-to-point links between segments separated by the distribution/core layer was also catered for in the future MPLS deployment.
- g.) While ideally MSDP would have been utilized on the network for multicast session distribution, the deployed hardware did not support this, hence all multicast in the design is purely routed utilizing PIM (Sparse-Mode). It should be noted that there are plans to implement MSDP once the core/distribution layer are upgraded.

## 6. Prerequisite planning phase

Once the design was completed, a full analysis of the prerequisites was done. (It should be noted that during this phase, the design was modified when it became clear that certain prerequisites could not be met, and this was particularly relevant with regards to the choice IGP and the methods used to route IPv6 traffic).

In order to implement the proposed design, the following things were discovered and rectified during this phase:

- a) The core and distribution switches required license upgrades to enable the necessary routing support. This caused a fairly lengthy delay in the project due to the vendor taking an excessively long time to deliver the required licenses.
- b) Further IPv4 address space had to be applied for in order to make implementation possible. It was also clear that the already existent legacy /16 would be vastly insufficient for any of the planned network expansion beyond this project, so such an application was also justified in this manner.

As stated above, the decision to apply for further IPv4 space was, among other things, prompted by the fact that implementation without additional IPv4 space would have been impractical. This was due to the fact that in order to segment the network, there were two choices. Either it would be necessary to insert more specific routes on each edge device for each newly segmented segment of the network as we proceeded, or a renumbering process had to take place. Because the network was flat post implementation, every device had the same /16 subnet mask. This meant that the moment we segmented any area of the network without renumbering, the two segments would be unable to communicate, as anything with the /16 subnet mask would assume that the new segment was still "directly connected", and routing would not take place. Since the cyclic insertion of more specific routes across thousands of edge devices was considered impractical in the extreme, the only real option left was to renumber to a completely separate block of IP space divorced from the already implemented legacy /16 space.

## 7. Detailed design phase:

Once the processes necessary to meet the prerequisites were in play, the detailed design planning could commence. This was to take the high level design already discussed above, and turn it into a concrete implementation plan. In order to achieve this, the following steps were performed:

- a) Thirteen network distribution points were identified
- b) A detailed VLAN planning exercise was performed
- c) A detailed IP planning exercise was performed

The VLAN planning exercise was done in such a manner as to devise a VLAN plan that could be duplicated across the distributions in a uniform fashion. To this end it was decided that each building on the campus would be segmented into 6 distinct VLAN's. This divided the networks in each building into the following segments:

- a) Wired ports
- b) Wireless access
- c) Access Control Systems
- d) Building Management Systems
- e) VoIP segments
- f) CCTV systems

Once the VLAN planning exercise had been concluded, we were able to determine the exact number of VLAN's per distribution and the required total address space per each segment for each distribution. This allowed us to allocate 6 supernet blocks to each distribution.

The supernet's were then further broken down into subnet's for each VLAN on each distribution. The IP planning was all documented using opensource software known as IP Plan. The software allowed us to easily track allocations as we went along.

Due to the number of VLAN's and as a result, the number of subnet's, the decision was taken to only route the supernet's beyond the distributions. The reachability of the subnet's was taken care of through the more specific connected routes on each distribution point.

## 8. The Implementation phase:

Since the network was flat pre-implementation, and everything was effectively in VLAN 1 which was spanned everywhere on the network, this was left untouched. Any changes to VLAN 1 on the network during the implementation phase would have resulted in downtime that we wanted to avoid.

We created a series of point-to-point VLAN's between the core and the various distributions, and trunked those along with VLAN 1 between the core and the distributions. Then, on each point-to-point, we assigned a /30 IPv4 subnet and a /126 IPv6 subnet. We then created the various building VLAN's on each distribution point, and assigned the correct IPv4 and IPv6 space to each VLAN. Again, we left all edge devices in VLAN 1 at this point, meaning the production traffic was still live and untouched.

Once the building VLAN's had been created on the distributions, we enabled the OSPF, OSPFv3 and BGP between the core and the distribution layer, and ensured and the supernet routing to each distribution layer was in place and functioning correctly. We also enabled PIM routing to the distributions at this point.

The DHCP scopes for each building VLAN were created during this phase as well. We also ensured that the DHCP relays were setup correctly on each distribution.

Since DHCPv6 has limited support, we chose to use RA(EUI-64) for IPv6 address assignment.

Once all of this had been verified, the building VLAN's created on the distribution were trunked down to the correct edge switches, again, still leaving VLAN 1 in place so as to not affect live traffic.

Once all of this had been completed, a series of commands was issued to each edge switch to tag the edge ports into the correct building VLAN's and out of VLAN 1, directly followed by a command to disable and then re-enable the ports that had just been retagged. The action of disabling and re-enabling the ports forced the edge devices to request new DHCP leases inside their new VLAN's. This gave them the correct addresses against the already defined DHCP scopes. As a result, the clients experienced 30 seconds or less of downtime during this process.

After ensuring that the edge devices were getting correct IPv4 and IPv6 addresses, and that the edge devices still had correctly functioning access, technicians were dispatched to each renumbered building to sort out devices (mainly printers) that did not utilize DHCP and required static addressing. As each device that required static numbering was renumbered, the device was moved into the correct VLAN.

At this point, the network was fully functional, segmented, and the IPv4 and IPv6 connectivity had been verified. The next step was to enable IPv6 on the proxy and DNS servers. This was done by simply adding the addresses to the servers, and then adding Quad-A DNS entries for each server, in addition to their A records.

## **9. The End Result:**

The redesign and redeployment of the network caused the CPU load on the switching infrastructure to drop from an average 90% utilization to an average 3% utilization. We attribute this largely to a far lower number of broadcast storms, and far smaller ARP tables on each switch. The smaller ARP tables prevented a huge amount of ARP churn on the switches caused by ARP table overflow.

Traffic throughput on the network increased by around 70%, and this was largely thanks to the elimination of hundreds of thousands of broadcast packets flowing network wide.

The most surprising change we noticed however was the massive and noticeable move of Internet traffic from IPv4 to IPv6. Approximately 60% of all Internet traffic switched to IPv6 within 90 minutes of the proxy servers being IPv6 enabled and the relevant Quad-A records being inserted. The 60% figure was far higher than originally anticipated, however since the network design was created in such a way that IPv4 and IPv6 are treated exactly equally, this swing did not cause any significant problems.

## **10. Project costing / timing.**

The project was completed over a five-week timespan. The first three weeks were devoted to planning, which was the most time intensive section of the project.

Since during the planning phase it was determined that we did not need to do any hardware replacement, the costs on this project were kept low. It was however necessary to purchase upgraded software licenses for the core and distribution equipment. These license's we necessary to enable Layer 3 routing (OSPF/OSPFv3/BGP).

The total cost of the project inclusive of the software licenses, consultant costs and other miscellaneous costs amounted to approximately \$50,000.00 (USD)

As a fun fact, approximately 1% of the project budget was spent on redbull and pizza to keep the engineers going over the long nights as we implemented!

## **11. Next Steps / Where to from here.**

On the hardware layer, a full core and distribution replacement is planned for Q1 2013. The new core will be comprised of two routers, running in a virtual chassis mode, with 2 x 100G circuits connecting them together. The new distribution layer will be comprised of 13 MPLS enabled L3 switches. Each distribution will have two 10G circuits back to the core, with one circuit into each physical core device. The two 10G circuits between each distribution and the core will be running LACP to form a 20gig trunk per distribution.

The purpose of the 20gig LACP trunking is to reduce the oversubscription between the core and distribution layer. This design is also flexible enough to allow us to expand the capacity to any distribution simply by extending the LACP trunks to include additional 10G circuits.

In addition to this, we will be implementing full BGP into the core, and then a limited table down to the distribution layer. This will allow for full NetFlow analysis of traffic flowing through the core. In addition, the implementation of full table BGP into the core and the border routers will allow for easier multi-homing and more flexibility.

In January, once the core and distribution replacement is completed, the network will also be MPLS enabled across the core and distribution layers. The purpose of implementing MPLS is to allow for EoMPLS between the distributions and to allow for technologies such as MPLS-TE for specific traffic pathing.

The multicast network, which we believe is critical too much of what we wish to accomplish on the campus network, is not nearly as robust as it could be. Because of lack of support in the current core/distribution hardware we cannot implement a full SSM or MSDP style rollout. This will be changed in January after the core/distribution replacement when we will do a full MSDP rollout.

## **12. Lessons Learnt.**

Restructuring a network for IPv6, multicast and flexibility does not require anywhere near the downtime or the finances most people expect. It can be done cheaply, effectively and with minimal downtime. However, to achieve this, it is critical to have a vision of where exactly the network needs to go. In addition to this, planning is critical, it is necessary to know exactly what is currently deployed and how the migration from the one network scenario to the other is going to be achieved.

## **Biographies**

**Andrew Alston** has worked in IT for more than twenty years, in UNIX system administration, security consulting and auditing, and academic networking. He spent 6 years as CTO of South Africa's national research and education network, TENET, where he led the development of high-speed connectivity and advanced networking services. Since leaving TENET in October 2011, Andrew has been doing extensive part time work for the UbuntuNet Alliance as their Network Development Manager, while also contracting to both the academic and commercial market on a private basis. Andrew is known to like fast cars, fine cigars, and a good argument.

**Louis Marais** has been involved in IT at the University of the Free State for 12 years and has worked with various technologies over the many years including Novell and Microsoft Directory services, Identity Management, Storage area Networks and VMWare. He is also responsible for the campus network and network security solutions at the UFS