Notes from an Internet Chair: Off oading Network Services to the Cloud for a Low Budget Conference Wireless Solution

Richard Gass
CoNEXT'12 Internet Chair
richardgass@gmail.com

Damien Saucez
INRIA
damien.saucez@inria.fr

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

The ACM 8th international conference on emerging Networking EXperiements and Technologies (CoNEXT) was organized in a lovely hotel in the south of France. Although it was in an excellent location in the city center of Nice with views to the sea, it suffered from poor Internet connectivity. In this paper we describe what happened to the network at CoNEXT and explain why Internet connectivity is usually a problem at small hotel venues. Next we highlight the usual issues with the network equipment that leads to the general network dissatisfaction of conference attendees. Finally we describe how we alleviated the problem by offloading network services and all network traffic into the cloud while supporting over 100 simultaneous connected devices on a single ADSL link with a device that is rated to only support around 15-20. Our experience shows that with simple offloading of certain network services, small conference venues with limited budget no longer have to be plagued by the usual factors that lead to an unsatisfactory Internet connectivity experience.

Categories and Subject Descriptors: C.2.1[Network Architecture and Design]: Wireless communication, Network topology

General Terms: Design, Performance

Keywords: Network service offloading, cloud computing

1. INTRODUCTION

Organizing a small conference is an extremely difficult task which is often criticized by many conference attendees for some reason or another. Many details of dissatisfaction cannot be addressed due to limited budget, physical constraints of the venue, and lack of time and previous experience at the venue to understand the issues that will undoubtedly arise. However, organizers of such events manage to make the experience as pleasing as humanly possible with all of these constraints. It is generally difficult to have an event where 100% of the attendees are completely satisfied without any complaints.

Conferences are the perfect venue to hear the latest cutting edge work, broaden our research interests, and meet people also doing similar research in the topic area. Larger conferences are mostly hosted in conference event centers that are specifically designed to handle large numbers of people both in terms of logistics as well has network infrastructure. However, many smaller scale conferences are hosted in smaller venues (i.e., hotels) due to the number of participants, location, cost, lack of a conference center, etc. While hotels offer an ideal place for people to gather and socialize, their networks are usually not specifically designed to handle large numbers of simultaneous connections from the hundreds of connected devices of the conference attendees.

With regard to network connectivity, installing a network to support a conference environment requires careful network planning, equipment provisioning, and general experience supporting a large wireless user base. In special purpose environments, there are dedicated teams with sufficient budget allocated to support these systems to guarantee a quality experience for their users. Once these environments are setup, they consistently perform well and can support the network connectivity requirements for a large wireless user base. However, in the case of smaller conference environments, these luxuries do not exist. Many times, venues are chosen based solely on other factors that do not overlap with network requirements which lead to a general dissatisfaction in this regard.

Many hotels have wireless networks available that cover most public areas of their property and guest rooms. Hotels with specific conference facilities have provisioned extra resources to handle larger numbers of users. However, if the venue does not have the necessary resources, organizers resort to other methods to attempt to make sure their conference attendees have the necessary connectivity in place before the event to handle their demands. This usually involves renting a temporary leased line and installing some extra access points around the heavily populated areas of the conference. Unless you are able to temporarily setup specialized equipment (e.g., special purpose access points, routers, etc) which typically carries a high price tag, the connectivity support can be lacking.

The main function of the Internet chair at these types of events is meant to handle small fires, take feedback from the user base regarding issues, and work with the hotel to tweak settings and resolve problems. The general setup is typically less than ideal and making major changes to the infrastructure is prohibited. This lack of control over the environment makes for an interesting challenge to squeeze any performance out of the system that is possible. Once the network begins to fail and every performance enhancement possible within the constraints of the allowable changes to

the hotel network has been explored, it is usually time to throw in the towel and succumb to being the butt of many jokes during the conference dinners.

The remainder of this paper follows the journey of the Internet chair and dedicated assistant and details the events of what happened to the network at CoNEXT'12, issues we faced, and how we ultimately found a usable solution that can essentially be implemented fairly easily to increase the overall wireless network performance with virtually no additional cost.

2. I HAVE NO INTERNET

Nowadays, an Internet connection during conferences is as crucial as providing water and organizers frequently rely on the venue network to provide Internet connectivity to their attendees. However, hotel networks are seldom designed to support hundreds of users and devices (especially computer scientists specializing in network research) in the same room. For CoNEXT'12, a 4Mbit/s symmetric link (maximum allowed with the allocated budget and venue location) was leased for the duration of the event.

The first day of the conference was dedicated to three workshops and attendance was significantly lower than what would be on the following days. This was the first trail run at this venue to understand the true capabilities of the hotel wireless. At the beginning of the day, everything appeared to run smoothly. Connectivity was sufficient to retrieve mails or news and there were no issues. Unfortunately, the quality rapidly degraded as more users started utilizing the resource. The attendees were noticing disconnections from the access points (APs), or while being connected, were incapable of accessing the web.

The first diagnostic revealed that the DNS resolver running on the gateway access point was not resolving names. The DNS resolver was automatically configured on clients by means of DHCP and unfortunately, it was not possible to configure the AP's DHCP server to advertise another DNS resolver other than itself (e.g., a Google public DNS server). In addition, we identified a second, more severe issue. We found that the AP was performing network address translation (NAT) and that the NAT table on the device was not able to support all the flows of the attendees. After calling the Internet service provider of our leased line, we were told that the AP they sent us was simply a home router only designed to support roughly 15-20 people.

In addition to the conference network, the hotel provided Internet connectivity with three access points. One access point was operated by an external company and was impossible to configure. The two other access points were operated directly by the hotel and it was only possible to configure their wireless settings (with heavy persuasion by our conference treasurer). Unfortunately, none of the networks were able to handle more than a few tens of users. Moreover, using a wireless sniffer, we noticed that all access points were using the same wireless channel, increasing the risk of interference. Therefore, in addition to the limitations of our ADSL gateway, we were facing a poorly configured wireless network with high levels of interference.

Throughout the day, we continued to tweak the settings in a futile attempt to make things work but had no success. Every idea that we proposed ended up not working due to the limited functionality of the equipment we had to work with. After exhausting all the possible options we were finally cornered with the problem of network services not being provided reliably due to the high load caused by the number of users. Our challenge now was to figure out what we could do with a spare laptop, an old Cisco AP, and a few cables to provide some kind of usable network to the conference attendees. The solution that we finally used was to offload all network services to the cloud.

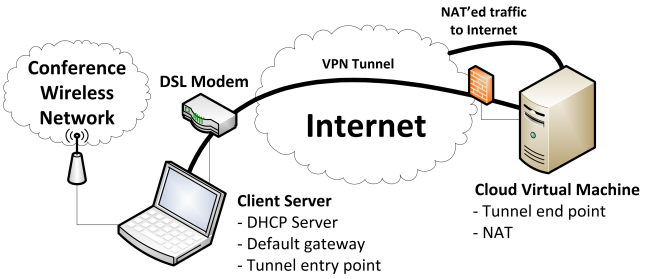

Figure 1: CoNEXT'12 network setup to offload all traffic to the cloud

3. OFFLOADING TO THE CLOUD

The main problem boiled down to the NAT on the gateway not being able to support all the flows generated by more than 15-20 users due to the lack of sufficient memory on the device. The only option was to reduce the number of flows moving through that gateway. One simple solution would be to require every attendee to use a virtual private network (VPN) such that each attendee consumes one flow per device. This option was of course not acceptable since it would still limit us to a small number of VPN connections and having an alternative that is transparent for the user is obviously preferable. To that aim, the solution we employed was to aggregate every flow using a tunnel such that the gateway limitation would be avoided. A tunnel had one end in the conference network and the other outside the conference network in the cloud (after the ADSL gateway). This meant that every packet would transparently be sent through the tunnel into the cloud where it will be finally injected into the Internet where we could have a virtual machine handle the NAT'ing. From a technical standpoint, the tunnel endpoint could be anywhere in the Internet but due to time constraints, we were not able to obtain the necessary authorization from our respective universities or institutions to carry transit traffic from unauthenticated users and thus decided to install the tunnel endpoint in the Amazon cloud [1].

There are several tunneling solutions that exist to accomplish this goal. Some of them are directly over IP (e.g., IPIP [5], GRE [3]...) and others over UDP or TCP (e.g., LISP [4], OpenVPN [2]...). The ADSL gateway only offered GRE and IPIP for tunnelling packets. However, Amazon did not support these protocols at that time and only supported UDP, TCP, or ICMP in their security groups. Therefore, neither the ADSL gateway nor the access point could be used for the conference network tunnel entry point. As depicted in Figure 1, a spare laptop computer (a commodity GNU/Linux laptop) was used in order for it to play the role of tunnel entry point. OpenVPN (open source VPN implementation) was selected to create the point-to-point tunnel which we established on the laptop and the virtual machine

running in the cloud. As we didn't know the capacity of the laptop to operate encrypted or compressed tunnels for all attendees traffic and time being of the essence, we deactivated all unnecessary OpenVPN options.

The laptop was configured as a DHCP server and was instructed to advertise the default gateway to be the interface of the laptop. Additionally, to ensure the scalability of the naming resolution, DHCP advertised the public Google DNS resolvers addresses (i.e., 8.8.8.8 and 8.8.4.4) for name resolution. Finally, IP forwarding on the laptop was activated and its next-hop default route was configured to be the OpenVPN tunnel entry point in order to forward all packets for the conference network.

After resolving some minor issues like blocking ICMP redirect packets to the clients because they were changing the route and bypassing our tunnel, everything starting working flawlessly. At some point we asked all conference attendees to connect to our conference access point to see if something could be broken. It was a pleasant surprise that the setup was able to support over 100 simultaneous connections through the access point without any issues.

4. OBSERVATIONS

Offloading the NAT to the cloud performed surprisingly well. We did manage to capture some data while the solution was in place in order to report on some of the findings. We found that in order to support around 100 devices, a very modest hardware footprint was required. On the conference side, we had a laptop with 4GB of memory running DHCP and the OpenVPN tunnel entry point. The cloud side was an Amazon Micro Instance which is rated to have 2 EC2 compute units and 613 MiB of RAM.

Both CPU and memory utilization was negligible on the conference side laptop. On the cloud side, there was a slightly higher CPU utilization but was still well below 10% for the entire duration of the data collection. The same was true for the amount of memory required where the average was well below 200 MBytes of total RAM used by NAT. This is extremely promising information for conference organizers trying to deploy a cloud based wireless network solution. The amount of resources required to support a large number of users is surprisingly low. As far as Amazon charges are concerned, they currently only charge for bits that leave the cloud and for the 3 days of running this solution, the total cost of transit traffic was less than 15 US dollars.

5. SOME GOTCHAS AND COMMENTS

Setting up a solution like this is fairly straightforward. However, there are a few lessons that we learned and would like to share for future implementers of this type of solution.

The first deals with the NAT and where traffic appears to be originating. Since Amazon's cloud service was utilized, all traffic was routed to their only European data center located in Ireland. This caused some confusion when other services attempted to geolocate its users via IP addresses. We heard several accounts where people using Google's email service needed to re-authenticate since the service previously thought they were in France and in the same day in Ireland.

Secondly, in the evening of the first day of forwarding all traffic into the cloud, we received an email from Amazon. It stated that we have reached the limit on the volume of email coming from the virtual machine that was acting as our OpenVPN endpoint and we would need to request these limit increased or removed. This is a precaution in Amazon that is in place to prevent the amazon public IP address range from being blacklisted by spam filters. Since the conference attendees where being routed through the cloud, many mail requests were originating from this address which caused the alerts within Amazon. This was mostly because we did not have a reverse DNS address setup which is typically an easy way of identifying a spamming system. A simple email to amazon explaining the situation was all that was needed to resolve this problem but it could be avoided in the future by following some simple precautions.

One interesting comment given was waking a device up from a sleep state and reconnecting to the network seemed much faster than normal. We attribute this to having our DHCP server allocating the addresses as opposed to the AP. Mac users were particularly surprised when they opened their laptops and were immediately connected to the AP, assigned an address, and were able to start browsing the web. We were not able to capture any logs to quantify this result but more than a few users pointed this out.

A few users commented and were impressed that they could even stream video which was not possible on the hotel wireless connection. Overall, we believe the general user satisfaction improved once the solution was in place.

6. CONCLUSION

Offloading network services to the cloud offers a potential solution and more options for future conference organizers that have a limited budget and want reliable connectivity to the Internet. We have seen that this type of offloading works extremely well with minimal resources both at the conference site and in the cloud. If the network were designed with this solution in mind, one could design a configuration to support a significantly higher number of users with much lower cost and resources that would be required to have a dedicated solution. We hope that the community finds this information useful and can take advantage of this network offloading solution if they ever find themselves in a similar predicament.

Acknowledgments

We would like send a special thanks to Guillaume Urvoy-Keller, Chadi Barakat, Christophe Diot, and Renata Teixeira for their support and comments as well as the CoNEXT conference attendees for their feedback and tolerance.

7. REFERENCES

[1] Amazon elastic compute cloud. `http://aws.amazon.com/ec2`.

[2] OpenVPN. `http://openvpn.net`.

[3] Hanks, S., Li, T., Farinacci, D., and Traina, P. Generic Routing Encapsulation (GRE). RFC 1701 (Informational), Oct. 1994.

[4] Saucez, D., Iannone, L., Bonaventure, O., and Farinacci, D. Designing a deployable future internet: the locator/identifier separation protocol (lisp) case. *IEEE Internet Computing 99*, PrePrints (2012).

[5] Simpson, W. IP in IP Tunneling. RFC 1853 (Informational), Oct. 1995.