



# Clever Addressing Schemes

Peter J. Welcher

## Introduction

Last month we looked at measuring IPSec performance in the article at <http://www.netcraftsmen.net/welcher/papers/ipsec-perf-01.html>. The two prior articles were security related, delving into EAP and 802.1x. And before that, SDM and PDM, also security related. I guess that's a bit of a trend, my attempt to reflect what I'm seeing in the way of folks' interests and needs.

Since I'm writing this in early July, I thought a simpler article was in order (think "vacation"). I plan to discuss at a high level some addressing schemes I've used at a couple of consulting customer sites. In the title I've awarded myself the word "clever", fully bearing in mind that too much cleverness in networking can sometimes come back and bite you. So don't overdo the ideas in this article! I do hope as well that the subnet masks in this article might be of some interest for those working on their CCNA certification.

If there's an overall moral to this article, it's that **a little attention to detail and planning up front in your addressing scheme can come back and pay nice rewards for you when you end up doing access lists, either for security or for QoS.** In particular, don't reserve blocks of 10 addresses in a row, work with summarizable blocks and powers of two. Make those binary bit patterns work for you!

## Summarizable Blocks of Addresses

In a previous article, I talked about summarizable blocks of subnets for OSPF. The same sort of thing applies when you want a block of addresses that lends itself nicely to ACL wildcard masking. You want a bunch of addresses (or subnets) in a row. If there are N numbers in a row, you want N to be a power of two, and you want the first of the consecutive numbers to be a multiple of N. Thus 8-15 is summarizable (8 numbers in a row, starting with a multiple of 8). But 8-23 is not (8 is not a multiple of 16, so you'd have to describe this block of numbers as 2 blocks of 8).

Suppose we have decided to use addresses 64-95. That's 32 addresses (95 minus 64, plus 1) in a row. Starting with a multiple of 32 (64 = 2 x 32). To find a matching subnet style mask, we're working with 32 things at a time, it's just like subnets being multiples of 32, our subnet style mask involves 256 - 32 = 224. So if we use 172.18.10.64-95, the subnet style mask would be 255.255.255.224 (224 in the last octet, where the range of addresses lies). To convert to ACL wildcard masks, take the ones complement by subtracting each octet from 255, to get 0.0.0.31. In general, when you have N numbers in a row, and N is a power of two, then the wildcard mask is going to have N-1 in it. That is, 32 numbers in a row leads to a 31 in the wildcard mask.

You can even use addition to check this. The ACL entry would involve the starting address and the wildcard mask: 172.18.10.64 0.0.0.31. The range starts with 172.18.10.64 and ends with 172.18.10.64+31. That trick unfortunately doesn't always work, but does work when your wildcard mask octet is two to a power, minus one. Or in binary: all right-contiguous one bits. (Is everyone now equally confused?)

What makes all this work is the way binary bit patterns fit with powers of two. Consider the range 8-23 mentioned above.

8	01000
9	01001

10	01010
11	01011
12	01100
13	01101
14	01110
15	01111
16	10000
17	10001
18	10010
19	10011
20	10100
21	10101
22	10110
23	10111

Note that the leading bits change from 01 to 10 at 16 (which I've indicated with colored backgrounds). The 01 holds constant for the first 8 numbers, the 10 holds constant for the last 8. That's why we can't summarize this as a block of 16, but only as two blocks of 8. This range straddles the multiple of 16 in there. That's why 0-15 and 16-31 are fine for blocks of 16, but 8-23 is not.

To sum up the rule for recognizing a summarizable block of numbers:

- 1 a block of consecutive numbers
- 1 the number of numbers is a power of two (call it N)
- 1 the first number is a multiple of the power of two, N

## Case Study #1: Deploying IP PBX's

At one site we've been working with staff to provide security and Quality of Service (QoS) for an IP Telephony deployment. The deployment uses several Nortel Succession PBX's with a large Cisco infrastructure, mainly based on L2 and L3 use of Catalyst 6500 and 4000 switches.

It turns out that you can use a DHCP option to tell Nortel IP phones to put themselves into a VLAN, similar to the AUX VLAN commonly used with all-Cisco IP telephony. We configured the AUX VLAN into the Cisco switches and it all worked, at the cost of a double DHCP for the phones. This allowed us to give the IP phones distinctive addresses. The company in question used RFC 1918 private addressing. We had over time assigned summarizable blocks of subnets to buildings and sites, and had rather exhausted the /16 blocks we were using. Each wiring closet has an odd and an even subnet and VLAN, going to L3 switches that route (and summarize) into an OSPF core. We used a structured scheme in assigning VLAN's to wiring closets. So we assigned the telephone new VLAN's within that scheme, having left some room for expansion. The subnets for the IP phones matched the data subnets, but had a different /16 prefix. Let's say 172.16.0.0 and 172.17.0.0 were the data address blocks, with 172.16.0.0 being used for most of the users. Then say 172.18.0.0 was the block we used for IP phones. If a wiring closet had subnet 172.16.a.0 in it, then we added 172.18.a.0 for voice. This was to keep the addressing complexity down, so that if you know the data subnet, you know the voice subnet as well.

What this also bought us was the ability to use Access Lists (ACL's) to secure the IP phones. The Nortel IP phones uses a very limited range of ports, as does their Soft Phone software. Since the Soft Phone software runs on any PC and uses the PC's address, you have limited ability to secure it or lock down QoS for it. Having fewer ports helps but doesn't prevent some other PC application from using the relevant UDP ports for a QoS-centric Denial of Service. Having a known /16 prefix for the physical IP phone handsets let us provide a fairly good level of security isolation for them. We were able to

do this with a few statements, rather than having to list all relevant subnets, the way you would if you grabbed available subnets right out of the range of data subnets.

The main drawback to this whole approach is that we added some routes to the routing table. They summarize as nicely as the originals do, so they're only of local impact. And I like to be able to look at the local routing table and see which subnets are up or down, which is really what the presence of a route is telling you.

Here's a sample showing what this looks like. Note that I've simplified this somewhat. OTM is Nortel Optivity Telephony Manager. The ACL is applied outbound on the AUX VLAN, so destination addresses really don't need to be specified.

```
ip access-list extended Nortel-IP-Phone-Access
  remark Allow OTM server Access to phones
  permit ip 172.16.A.B 0.0.0.0 172.18.0.0 0.0.255.255
  remark Allow traffic between phones and to PBX components
  permit ip 172.18.0.0 0.0.255.255 172.18.0.0 0.0.255.255
  remark permit DNS and DHCP
  permit udp 172.16.C.D 0.0.0.0 172.18.0.0 0.0.255.255 eq bootpc
  permit icmp 172.16.C.D 0.0.0.0 172.18.0.0 0.0.255.255
  permit udp 172.16.E.F 0.0.0.0 172.18.0.0 0.0.255.255 eq domain
  deny tcp any any log
  deny udp any any log
  deny ip any any log
```

The last lines there log blocked traffic, which you may want to only do selectively (CPU impact, volume of log traffic). The actual access list used specified TCP or UDP ports and was more specific about allowed traffic. Painful to write and check but you should more or less only have to do so once.

The other idea we had was to try to restrict PBX application traffic tightly. This has a cost to it, of requiring administration to verify correct operation when new PBX software releases come out. But if you're a company of 10,000 people, your phones had better work reliably!

The Nortel IP PBX's use two VLAN's. The "ELAN" acts as virtual PBX backplane, and is used for traffic between PBX components. Nortel recommends securing it against all but unnecessary traffic, isolating it from broadcast traffic, etc. The only traffic we allowed to be routed to an ELAN was either traffic from another ELAN, or from the PBX administrative management station. The second VLAN used is the "TLAN", which is what outside devices talk to.

To make this more manageable, we reserved a block out of the telephony /16 for TLAN subnets, and another block for ELAN subnets (wisdom suggests you avoid saying "TLAN VLAN and ELAN VLAN" too often). Let's say 172.18.224.0 /20 and 172.18.240.0 /20 respectively. We used a /25 subnet mask. In retrospect, a /26 or even /27 mask could have been used. We assigned some 8 to 10 subnets out of each block, corresponding to the various PBX locations.

We were then able to write ACL's on phone VLAN's, allowing all 172.18.0.0 traffic on appropriate ports, then allowing 172.18.0.0 traffic to/from anywhere on the same ports. Our thought process was that should some security issue occur, we could cut off use of Soft Phones if we needed to better protect the handsets, physical IP phones. We could also use very similar ACL's to classify and mark traffic for QoS purposes.

On the router interfaces leading to ELAN subnets, we only allowed other ELAN or management station (OTM) traffic. The following is an edited version of the outbound ACL:

```
ip access-list extended Nortel-ELAN-VoIP-ACL
  remark Allow Nortel OTM server Access to ELAN
  permit ip 172.16.A.B 0.0.0.0 172.18.240.0 0.0.15.255
  remark allow access between ELAN to ELAN
  permit ip 172.18.240.0 0.0.15.255 172.18.240.0 0.0.15.255
  remark Allow replies to DNS queries
  permit udp host 172.16.C.D eq domain 172.18.240.0 0.0.15.255
  remark deny and log everything else coming onto this subnet
  deny tcp any any log
  deny udp any any log
  deny ip any any log
```

The destinations above could have been "any", since ELAN addresses are all that can be reached out that interface in any case. But this shows the use of the appropriate mask for the ELAN range of addresses. Note that our range of ELAN

subnets starts with 240 and ends with  $240 + 15 = 255$ .

We actually assigned addresses within blocks to the signaling and the H.323 gateway components of the IP PBX. Say the signaling server components are in the range 48-63, and the H.323 gateways are in the range 64-127. These are summarizable blocks: 48-63 is 16 numbers in a row, starting with a multiple of 16. And 64-127 is 64 numbers, starting with a multiple of 64. So these blocks are summarizable, and as I noted earlier, summarizable blocks make for simpler ACL's.

This allowed us to tighten up security by allowing only Nortel UNISTIM (etc.) to the signaling components, and H.323 and RTP to the gateway components. For the signaling, the ACL entries looked somewhat like:

```
permit udp 172.18.0.0 0.0.255.255 172.18.224.48 0.0.15.15 eq A
```

And for the H.323, the entries were similar:

```
permit udp 172.18.0.0 0.0.255.255 172.18.224.64 0.0.15.63 eq B
```

In these wildcard masks, the first 15 wildcards the 224, specifying that any value from 224 to  $224+15=239$  is ok in the third octet. And the second 15 wildcards the 48, allowing addresses 48 to  $48+15=63$ . The final 63 wildcards the 64, allowing addresses 64 to  $64+63=127$  to match the ACL rule.

If you think that's simple, I'll note in passing that our subnet masks were actually 255.255.255.128, so we added 128 to the last octets above, making those final masks 0.0.15.143 and 0.0.15.191. Adding in the 128 bit means we don't care if we're in the low or high half of the last octet (the subnets being either 0-127 and 128-255).

## Design Impact of L3 Switching

It's time for an intermission, followed by the second act. I've been looking at network design and noticing that I'm starting to see Layer3 switches in the wiring closets, as they get cheaper and cheaper. One can still use L2 to the wiring closet, with routing / Layer 3 switching at the distribution layer. With that approach, you have some Spanning Tree in your network. With L3 to the closet, you end up with perhaps more subnets and with more routes and routing. I happen to prefer the latter approach.

As I noted two articles back, Identity-Based networking is becoming attractive. The Cisco approach allows you to specify that a group of users, say Students, is VLAN 5. That's fine if you have one central L3 switch and all your closets understand and belong to VLAN 5. But if you do that, you end up with a big VLAN 5 and a large STP, which I strongly prefer not to do, for stability and ease of troubleshooting reasons.

One answer is to have lots of VLAN 5's at the edge, with the adjacent L3 switch making each of them a different subnet. So login ties to VLAN 5, which depending on location ties to an appropriate subnet.

If your campus has 8 groups of users (student, faculty/admin, network management, printers and appliances, etc.), then you get 8 subnets at each L3 switch, leading to a rapid consumption of subnets. So this approach can be made to work, but don't get too carried away.

If you like this approach, note that older gear, e.g. the Cisco 2940 switch, may only support a limited number of VLAN's. In the case of the 2940, "limited" equals 4.

## Case Study #2: Controlling College Students

I've been working with a university to design for a campus-wide rollout of Cisco switches and access points. To be "802.1x ready" we've done L3 switching where possible and configured several VLAN's. Worst case, we don't end up using some of them. The campus does have a few 2950's connected to the core 6500, and those use VLAN's that are L3 switched by the 6500.

For simplicity, we ended up with the addressing scheme `10.building.vlan.x /24`. Here "building" is usually a building number, except that some buildings have two or three L3 switches in them, and each such switch gets its own "building" number. Many of the L3 switches are 4500's with Supervisor 2+, shaving some cost off. (The price compares favorably to a stack of smaller switches, and there are fewer boxes to manage. I personally am mildly allergic to the Cisco stacking software.) We're using static routing, since there is no redundancy, so remote L3 switches have static default routes back

to the core 6500, and it has a static summary route to 10.building.0.0 255.255.0.0. That is, go from the core to the "building", and that L3 switch knows its locally-connected subnets. That's the first win from this addressing scheme.

This campus previously had an Enterasys Layer 2 deployment, with its share of issues. One nice feature was that the Student VLAN ran through an IDS and firewall before getting to the servers, the faculty VLAN, or to the Internet. I was asked to continue this sort of protection of faculty and administration from student computers, the main concern being not hackers but mitigating the impact of virus and worm storms that have plagued many college campuses recently. The policy is therefore that students can only talk to designated servers and to the Internet. Unlike the previous design however, student traffic is not forced to go through one central router. It now hits a L3 switch, and has a chance to be routed onto another subnet. So we're using VLAN ACL's to prevent such cross-VLAN traffic. Due to the nice addressing scheme, we can however write simple rules like the following:

```
deny ip 10.0.5.0 0.255.0.255 any
```

The student VLAN is 5. This rule says that any traffic from 10.something.5.something, i.e. any of the many VLAN 5 subnets, is not allowed (out to whatever other VLAN this is being applied to). This sure beats having to list each possible student subnet as a source for traffic! I kind of like it as well for using a rather interesting wildcard mask.

Anticipating networked devices that cannot do 802.1x authentication, we created a VLAN for such devices, e.g. printers. This is the first time I've really appreciated print servers. PC administrators like them for not having to install print drivers on many student or employee machines, and perhaps for print accounting. I've disliked them since consultants have to get a domain login, etc. -- one more barrier to quick printing. The advantage here is that we've set up an ACL so that printers can only talk to the central print server(s). That means that anyone borrowing their Ethernet connection is going to find it rather useless to them -- part of what you want when you're using 802.1x to lock network connections down. This really doesn't have much to do with addressing, but it's something I thought I'd share.

It did turn out this addressing scheme had one gotcha (to date). The site now has two Packeteers for traffic shaping, having previously seen student worms cause their sole Packeteer to melt down under too many flows. So we're using policy-based routing to force student traffic through the student Packeteer. If it melts, all the faculty and non-student traffic can flow through the surviving Packeteer. One the other side, we need static routes in the exterior device. You cannot put in a static route with a subnet mask like 255.0.255.0, so we ended up having to list all the student subnets there. We considered using OSPF, but decided static was simpler.

## Summary

I hope this has been interesting, and maybe given you some ideas on how to get creative with addressing. You can do this in almost any situation where you can group things. If you're thinking about servers, maybe divide up "the server subnet" into ranges. Perhaps one range for servers that students might access, and another for faculty only servers. (Noting that faculty may also access some of the servers used by students as well). Or do the same for corporate user groups. Then you could use entries in ACL's to more easily write policy rules such as "students are only allowed to go to these servers". It sure beats ACL's with rules listing relevant servers, one per line.

Admittedly, in the college case I'd prefer to see different subnets for such servers. My idea is that if a student hacks into a student server, they still aren't in the same subnet as an administrative server, so probably buys them little in terms of using dsniff or hacking into the admin server.

I haven't really written much about it in this article, but since ACL's are used for classification, some of the benefits noted above apply to QoS as well. For instance, perhaps at some point we'll write a rule giving faculty/admin web traffic higher QoS priority. It'll be a lot easier if I can just use one entry to describe all the faculty subnets!

Next month's topic may be determined by random selection from the long list of potential topics I've built. But you can influence the topic selection by sending me email with your suggestions for what you'd like an article on!

I've been meaning to do one or two articles on the Wireless LAN Solution Engine (WLSE) appliance and software for managing Wireless Access Points. I've got a WLSE and some WAPs all set up in my lab, screen captures, etc. See also <http://www.cisco.com/en/US/products/sw/cscowork/ps3915/index.html>. I've also been meaning to do an article with some lab work on the recent AutoQoS features for serial links, see also <http://www.cisco.com/warp/public/732/Tech/qos/>. Neat stuff!

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014, CCIP) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner dedicated to quality consulting and knowledge transfer. NetCraftsmen has eight CCIE's, with expertise including large network high-availability routing/switching and design, VoIP, QoS, MPLS, IPSec VPN, wireless LAN and bridging, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher> . New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to **[pjw <at> netcraftsmen <dot> net](mailto:pjw@netcraftsmen.net)**.

7/6/2004

Copyright (C) 2004 Peter J. Welcher