Address Scheme Planning for an ISP backbone Network

Philip Smith – Consulting Engineering, Office of the CTO

Version 0.1 (draft)

LIST OF FIGURES 2

INTRODUCTION 3

BACKGROUND 3

BUSINESS MODEL 3

ADDRESS PLAN 4

Network Plan – Starting off 5

Network Plan – end of first year 6

Network Plan – end of second year 7

PUTTING A PLAN TOGETHER 9

Loopback Interfaces 9

WAN Links 9

LANs 9

Customer Networks 10

SUMMARY 10

APPENDIX – USEFUL INFORMATION 11

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

CISCO SYSTEMS

List of Figures

Figure 1 – Network Plan at deployment **5**
Figure 2 – Network plan at end of year 1 **7**
Figure 3 – Network Plan at end of year 2 **8**

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Introduction

Designing an IP addressing plan for an ISP backbone is probably one of the more complex issues an ISP faces when entering the competitive Internet Marketplace. Internet growth in recent years has been explosive, and coverage is now well beyond that of academic and government organisations in the US. Due to the finite amount of IPv4 address space available, this growth has resulted in the community's desire to be more careful about how the resource is shared amongst its users.

ICANN (the Internet Corporate for Assigned Numbers and Names, formerly IANA) is responsible for the allocation of all IP address space. To this end it has delegated three Regional Internet Registries (ARIN, RIPE and APNIC) to carry out its responsibilities in the three regions of the world. The three Regional Internet Registries and the user community are working to ensure that each user of public address space is efficient with their utilisation, with the aim of being fair to all that require this scarce resource.

It is hoped that this document will encourage ISPs to consider how to design a scalable addressing plan. Conservation and efficient utilisation of address space is often seen as problematic and even undesirable by ISPs trying to minimise the number of prefixes carried around in their network.

The document is only intended as an example of the considerations which need to be made when designing the addressing plan for an ISP network. It does not advise on how to go through the process of applying for address space from the three regional address registries, or how to configure the BGP routing protocol.

Background

The purpose of this document is to describe how an ISP should go about planning the addressing architecture of their network infrastructure. It is based on a training session given to some students of the UNDP/APDIP ISP Workshop in Kuala Lumpur in December 1998.

It is assumed that the reader is familiar with terms such as CIDR, classless addressing, and the "/N" network mask scheme. These are used throughout the document. Note that antiquated terminology such as "Class C" is not used here, and should not be used in the vocabulary of the Internet.

Business Model

It is generally accepted that it is reasonable to think ahead for 2 years when designing an ISP network. Beyond that, given the rapid growth of the Internet, it is very difficult to predict what new technologies will be available, or which direction the business may be headed. Two examples:

- In 1994 available address space was more limited than it is today, prior to the arrival of CIDR. There was real concern about the rate of address space consumption, with dire

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

predictions that IPv4 space would run out in 1998. This concern spurred on the introduction of CIDR and the classless Internet. Since then, the registries have been making classless network assignments, and DHCP has made addressing and renumbering on the LAN very trivial. What was a problem in 1994 was no longer one in 1996.

- An ISP may go into business to service a particular market place. However, local operating conditions may change, or new "killer applications" may appear on the Internet, with the result that the ISPs original business objective may change completely. Furthermore, address space requirements usually change when this happens.

The three Internet registries look for 2 year estimate of address requirements, and make assignments on that basis. They do so because most ISP business plans operate around similar concepts.

It is important to sit back for a few minutes and look seriously at the avenues of growth for any ISP business. This isn't easy in a marketplace which is growing exponentially. However, it pays in terms of engineering time to try and get a realistic picture of where the business will lead. Will there be more PoPs, or larger PoPs, more backbone links, new services, new equipment with increased port density, and so on. An ISP operation usually has a business plan to guarantee the funding from sponsors. The existence of a business plan implies that some thought must have been given to the direction of growth of the network – the two are interlinked.

Address Plan

This section documents an simple example of an address plan for a growing ISP business. Obviously it is only intended as a guide, but understanding the principles will help with designing most other types of networks.

Three Network Plans are presented, the current plan, the plan after a year, and the plan after two years. The addressing schema should take each plan into account as it helps the ISP plan the network growth over the coming 24 months. Notice, however, that each section produces a summary of required address space. Once the three plans are developed, they should be documented in a format similar to this whitepaper. Another important tenet of an ISP's business – there can never be too much documentation. With good documentation such as this, it helps the engineering team working for the ISP to understand the growth plans, and helps them engineer the routing protocols to suit, and it helps the address registries understand the address space requirements, and ease the assignment process.

It is assumed all along that the ISP is following the essential features of configuring routing protocols – typical examples may be found in Cisco's ISP/IXP Workshop tutorials. In general networks are designed with an IGP (such as OSPF or ISIS) being used to carry point-to-point link addresses, loopback interface addresses, and LAN addresses. BGP is used to carry the ISP's customer networks (iBGP), and carry any externally learned networks (eBGP). (While BGP probably won't be used until the ISP multihomes (connects to more than one upstream

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

ISP), it still pays to follow the design principle in the event that BGP will be implemented in the future.)

Network Plan – Starting off

The first stage is to look at the network design at the start of the ISP's operation. Figure 1 gives an example network – it has four routers, three switches with some hosts connected to them, and some customer leased line connections. There is also a dialup router. Finally the network has a link to an upstream ISP.

Also on the figure are the sizes of the subnets which have been allocated to each portion of the network. In detail these are:

- WAN point-to-point links have been assigned a /30 – there are two hosts on a point-to-point link, so the maximum address space required is a /30. Assigning a larger subnet will result in the waste of address space. (If there is trouble calculating how many hosts can fit into a subnet, please refer to the table in the appendix.)
- LANs have been assigned only the address space they require.
- One loopback interface is assigned per router.

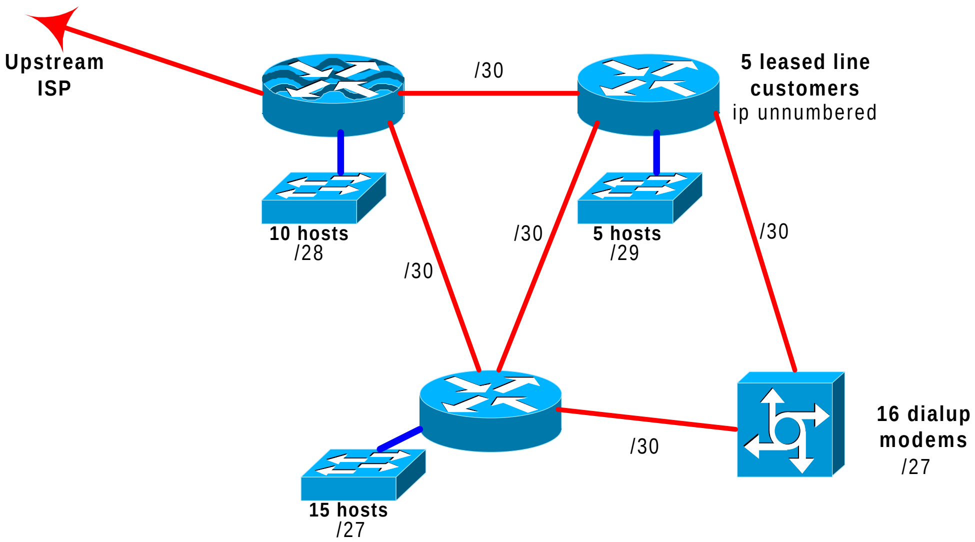

Figure 1 – Network Plan at deployment

The /30 address required for the ISP's upstream link is usually assigned by the upstream ISP, so is not required here. It is assumed that "*ip unnumbered*" is used to configure the point-to-point links going to a customer site. It is also assumed that one loopback interface is assigned to each of the four routers. It is also assumed that each ethernet switch has been assigned an IP address. These assumptions are common practice in most ISPs because:

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

- "ip unnumbered" simply means that the point-to-point link between the ISP and their customer routers is not assigned an IP address. Non-Cisco equipment will use other conventions, but are similarly capable. Using no address means that there is one network less in the ISPs IGP – IGP design always aims to have minimal networks in the routing table, for efficiency, and convergence speed.

- Loopback interface on a router divorces the routers administrative functions from any of its physical interfaces, thereby guaranteeing continued administrative access in the event of any physical link failure.

- All LAN equipment is assigned an IP address for administrative access. While all equipment will have a serial console, often an IP capable interface is extremely useful as a first line of entry for administration functions (many ISPs reserve console access as a last resort).

At this initial stage, we have five /30s, one /29, one /28 and two /27s, plus a /30 required for loopback addresses (loopback interfaces are assigned a /32). Following the sums through in detail:

- six /30s make a one /28 and one /29,
- one /28 and one /29, plus the single /29, make two /28s,
- the two /28s, plus the single /28, make one /27 and one /28,
- One /27 and one /28, together with the two /27s, make two /26s,
- The two /26s make a single /25.

In conclusion, this network infrastructure could be numbered out of a single /25 of address space.

Network Plan – end of first year

At the end of the first year, the ISP aims to show moderate growth in their network, and estimates that their network infrastructure will show additional equipment and sites. The network is intended to grow to that shown in Figure 2. The figure also documents the address plan intended for the network at the end of the first year.

Looking at the utilisation at this stage, there are eight /30s for point-to-point links, two /28s, two /27s and two /26s. There are also two new devices now requiring loopback interfaces (making a requirement for a /29). Following the sums through in detail:

- Eight /30s make one /27.
- Two /28s make one /27.
- Two /27s, plus the above two /27s, make a /25.
- Two /26s make a /25.
- The two /25s make a /24.
- The extra /29 for loopback interfaces should also be counted, making the total requirement a /23.

In conclusion, the address space requirement for this projected network at the end of the first year is a /23.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

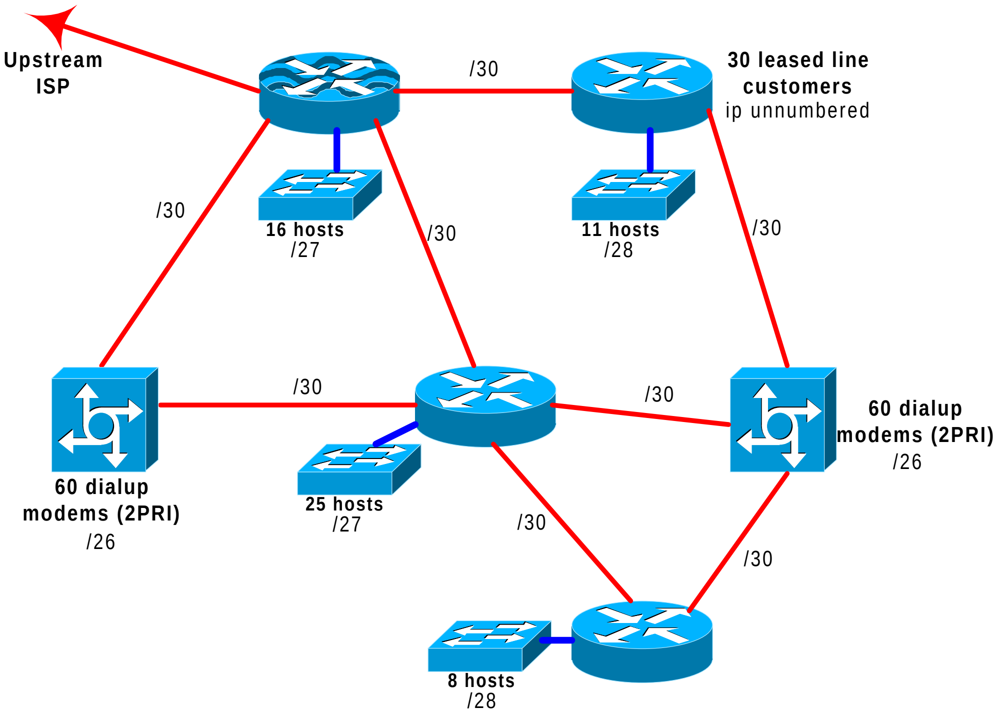

Figure 2 – Network plan at end of year 1

Network Plan – end of second year

The final stage to consider is the plan at the end of the second year of operation. Again this is a projection on what the business could be like. It is assumed that the ISP has started to do web hosting for its customers, as well as investing in a large dialup network (which the business plan calls for – say).

The ISP also plans to deploy "two of everything" by the second year of operation. This affords increased redundancy and resiliency in the network. (This isn't meant to be a tutorial in network design – these elements are simply included to show how the addressing plan might be done.)

Figure 3 shows how the network has grown, and shows the addressing which goes with it. Some values have been omitted for clarity. Some sites have dual routers now, and the links between them also require a /30 of address space each.

Looking at the utilisation now, there are fourteen /30s, two /28s, two /26s, two /24s, and twelve /32s for loopback interface addresses (making a requirement for a /28). Also, the switches are now dual homed to the routers, and the ISP is using HSRP (Hot Standby Routing Protocol) to create a virtual default gateway for his LAN devices. (Using HSRP consumes another /32 per LAN for the virtual gateway address so this should be remembered

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

in the LAN host count. BTW, LAN host count includes the IP address required for the router in every case.)

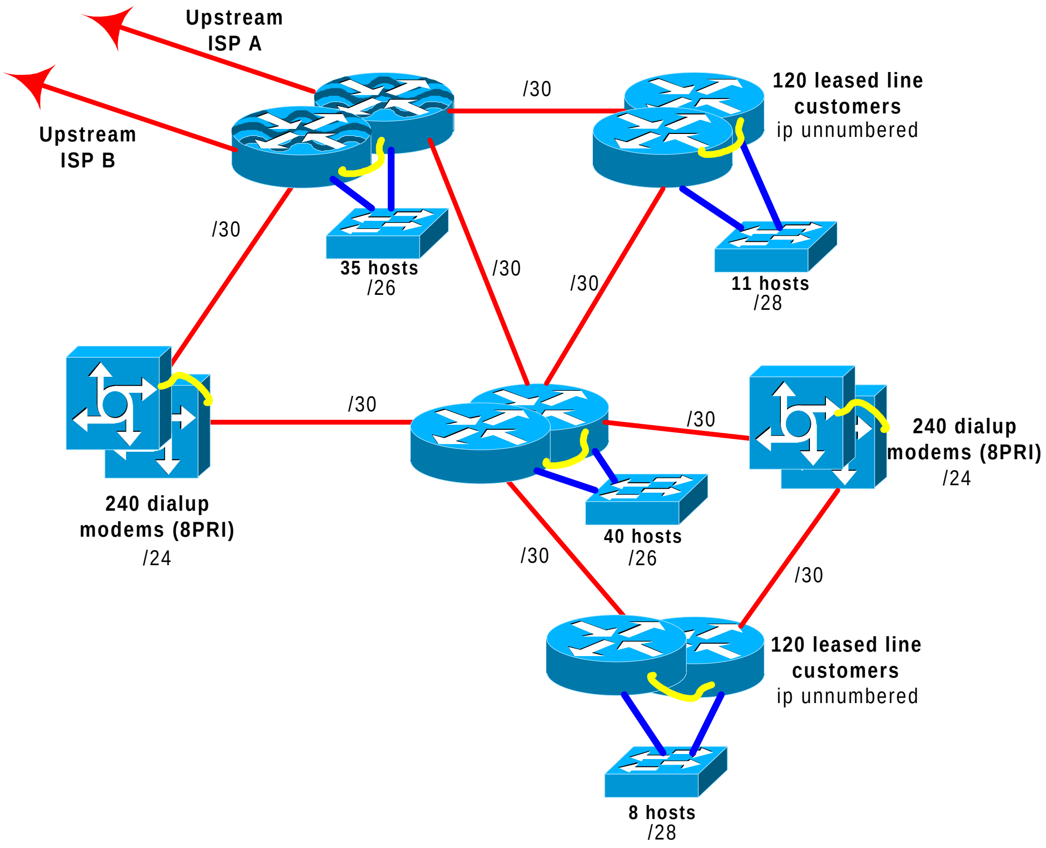

Figure 3 – Network Plan at end of year 2

Going through the sums in detail gives:

- Fourteen /30s make seven /29s, which make one /27, one /28, and one /29.
- Two /28s make a /27.
- Two /26s make a /25.
- Two /24s make a /23.
- The loopback /28 together with the /28 from above make a /27. With the /27 from above this makes a /26.
- The /26, with the remaining /29, round up to a /25.
- The two /25s make a /24.
- The /24 with the /23 above round up to a /22.

In conclusion, the address requirements after 2 years are a /22. Notice how the network subnets have been packed as efficiently as possible. There is quite a complex network, yet with judicious use of addressing, it has consumed only /22 of address space.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Furthermore, as the ISP now has two connections to the Internet, they would have applied for an AS number, and used that to multihome to the two upstream ISPs. This allows the layering of the network to use an IGP, iBGP, and eBGP as mentioned above.

Notice also the continued use of IP unnumbered. The ISP should make a reasonable estimate of how many links "*ip unnumbered*" can be used for. Sometimes, customer routers may not be able to support the concept.

Also, notice that IP addresses are dynamically assigned on the dialup access servers – while the ISP may be operating at a ratio of 20 users per modem, the requirement is only for one IP address to be available per modem line. And notice that while a PRI may have 32 channels, generally only 30 are available for use in an ISDN/PSTN service – quite convenient for addressing.

Putting a Plan Together

The above examples have established the address requirements for the sample ISP network which has been the subject of this paper. The final thing which needs to be done is take this and convert it into a reasonable infrastructure assignment process.

Loopback Interfaces

Always the first consideration on an ISP network is the loopback interface on the router. It is a useful general purpose feature used for many things, including iBGP peering, and source address for packets originating from the router (useful for authentication or filtering).

By the end of two years, the ISP estimated that they would have twelve routers in their network, so it is reasonable to assign a /28 to be used for that purpose from the very beginning. It is advisable, for technical reasons, to choose the highest part of the address assignment from the registry for the loopback addresses.

The loopback /32s will be carried around only in the IGP, and pose minimal overhead to the routing protocol.

WAN Links

The second consideration is the WAN links. Unless the network is easily separable into distinct regions, there is little to be gained by summarising the WAN /30 address space. Also, as WAN address space grows more slowly than other parts of the network, WAN links should be addressed out of another contiguous block from the allocation (say the chunk below the loopback network block).

LANs

Finally to the LANs and inter-router links. LANs probably gain hosts faster than other parts of the network, so it is prudent to allow greater headroom here than else where in the

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

network. Again, no benefit can be had from summarisation, unless the network can be split into a regional layout.

Unless there is a good reason not to, DHCP should be used to assign IP addresses. A DHCP server can take care of changing IP addresses and network masks on a LAN, as well as changes in the DNS resolver addresses.

Customer Networks

Customer network assignments have not been covered here as this document concentrates only on infrastructure. However, once an ISP is running BGP within their network, there is no benefit either to be had by aggregating or assigning networks on a regional basis. Customers move, they want their connections to move to different parts of the country, so there is little purpose in trying to allocate address space to them regionally. Besides, BGP is a very powerful routing protocol, and can easily handle huge numbers of customer assigned networks.

Naturally, all these customer networks would be aggregated on the edge of the ISP's network. The Regional Internet Registries will assign blocks of address space to their ISP membership, and the Internet community expects all ISPs to only announce these blocks to other ISPs. Announcing more specific networks from a block is generally frowned upon, although is acceptable when aiding situations such as multihoming.

Summary

This whitepaper has given an example of how to work out an addressing scheme for a developing ISP network. This is intended to help the growing ISP business work out how to apply addressing to their network, and how to allocate assigned address space to their infrastructure. Indeed, following these processes through should aid the application process for address space from the regional Internet registries.

Appendix – Useful Information

This section gives some useful information which the reader may find helpful.

1. Regional Internet Registries:
   - http://www.apnic.net
   - http://www.ripe.net
   - http://www.arin.net
2. Classful versus classless
   - Class A network is now known as a /8
   - Class B network is now known as a /16
   - Class C network is now known as a /24
   - Address space can be assigned by the registries from potentially anywhere in the IPv4 unicast address range. This unicast range starts at 1.0.0.0 and goes all the way to 223.255.255.255. The exceptions are 10/8, 172.16/12 and 192.168/20, as documented in RFC1918.
3. Subnetting
   - In the classful world, meant the subdivision of a class A, class B or class C network into smaller portions. Subnets usually had to be of an equal size to be routed properly.
   - In the classless world, this means the subdivision of assigned address space into smaller portions. Nothing more. For example, a /26 network could be split into two /27 subnets.
4. Hosts per subnet
   - The all zeros and all ones subnets could not be used in the old classful world – they can in the classless world, as that world has no such concept. So there is less waste of address space.
   - Subnetting a /24 gives the following network versus host count:

Network Mask	Subnets	Host count
/24	**1**	**254**
/25	**2**	**126**
/26	**4**	**62**
/27	**8**	**30**
/28	**16**	**14**
/29	**32**	**6**
/30	**64**	**2**
/31	**128**	**None**
/32	**255**	**1**

So, for example, if there are only 2 hosts on a network (for example a point-to-point serial link), a /30 network mask is required. Or, if there are 25 hosts on a network (typical LAN), a /27 network mask is required.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100