

Class A Subnet Experiment
Results and Recommendations

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Discussion/Purpose

This memo documents some experiences with the RFC 1797 [1] subnet A experiment (performed by the Net39 Test Group (see credits)) and provides a number of recommendations on future direction for both the Internet Registries and the Operations community.

Not all proposed experiments in RFC 1797 were done. Only the "case one" type delegations were made. Additional experimentation was done within the DNS service, by supporting a root nameserver and the primary for the domain from within the subnetted address space. In addition, testing was done on classless delegation [2].

Internet Services offered over the RFC 1797 experiment were:

- Finger
- HTTP
- Telnet
- FTP server/client
- Gopher
- kerberos
- lpr (and its ilk)
- X
- DNS

F.Root-Servers.Net, a root name server had an interface defined as part of the RFC 1797 experiment. Attached is a report fragment on it's performance: "My root server has processed 400,000,000 queries in the last 38 days, and well over half of them were to the temporary 39.13.229.241 address (note that I retained the old 192.5.5.241 address since I knew a lot of folks would not update their root.cache files and I didn't want to create a black hole.)" - Paul Vixie

Initial predictions [3] seemed to indicate that the safest path for an ISP that participates in such a routing system is to have -all- of the ISP clients be either:

- a) singly connected to one upstream ISP
- OR
- b) running a classless interior routing protocol

It is also noted that a network with default route may not notice it has potential routing problems until it starts using subnets of traditional A's internally.

Problems & Solutions

Operations

There were initial problems in at least one RIPE181 [4] implementation. It is clear that operators need to register in the Internet Routing Registry (IRR) all active aggregates and delegations for any given prefix. Additionally, there need to be methods for determining who is authoritative for announcing any given prefix.

It is expected that problems identified within the confines of this experiment are applicable to some RFC 1597 prefixes or any "natural" class "A" space.

Use of traceroute (LSRR) was critical for network troubleshooting during this experiment. In current cisco IOS, coding the following statement will disable LSRR and therefore inhibit cross-provider troubleshooting:

```
no ip source-route
```

We recommend that this statement -NOT- be placed in active ISP cisco configurations.

In general, there are serious weaknesses in the Inter-Provider cooperation model and resolution of these problems is outside the scope of this document. Perhaps the IEPG or any/all of the national or continental operations bodies [5] will take this as an action item for the continued health and viability of the Internet.

Routing

A classic cisco configuration that has the following statements

```
ip route 39.1.28.0 255.255.255.0
router bgp 64000
redistribute static
```

will, by default, promote any classful subnet route to a full classful route (supernet routes will be left alone). This behaviour can be changed in at least the following two ways:

- 1:

```
ip route 39.1.28.0 255.255.255.0
router bgp 64000
no auto-summary
redistribute static
```
- 2:

```
ip route 39.1.28.0 255.255.255.0
router bgp 64000
network 39.1.28.0 mask 255.255.255.0
redistribute static route-map static-bgp
....
access-list 98 deny 39.1.28.0 0.255.255.255
access-list 98 permit any
....
route-map static-bgp
match ip address 98
```

Users of cisco gear currently need to code the following two statements:

```
ip classless
ip subnet-zero
```

The implication of the first directive is that it eliminates the idea that if you know how to talk to a subnet of a network, you know how to talk to ALL of the network.

The second is needed since it is no longer clear where the all-ones or all-zeros networks are [6].

Other infrastructure gear exhibited similar or worse behaviour. Equipment that depends on use of a classful routing protocol, such as RIPv1 are prone to misconfiguration. Tested examples are current Ascend and Livingston gear, which continue to use RIPv1 as the default/only routing protocol. RIPv1 use will create an aggregate

announcement.

This pernicious use of this classful IGP was shown to impact otherwise capable systems. When attempting to communicate between an Ascend and a cisco the promotion problem identified above, was manifest. The problem turned out to be that a classful IGP (RIPv1) was being used between the Ascends and ciscos. The Ascend was told to announce 39.1.28/24, but since RIPv1 can't do this, the Ascend instead sent 39/8. We note that RIPv1, as with all classful IGPs should be considered historic.

This validates the predictions discussed in [3].

Cisco Specific Examples

There are actually three ways to solve the unintended aggregation problem, as described with current cisco IOS. Which of them applies will depend on what software version is in the router. Workarounds can be implemented for ancient (e.g., 8.X) version software.

- o Preferred solution: turn on "ip classless" in the routers and use a default route inside the AS. The "ip classless" command prevents the existence of a single "subnet" route from blocking access via the default route to other subnets of the same old-style network. Default only works with single-homed ISPs.
- o Workaround for 9.1 or later software where the "ip classless" command is not available: install a "default network route" like this:
"ip route 39.0.0.0 255.0.0.0 <next-hop>" along the axis the default route would normally take. It appears an ISP can utilize the "recursive route lookups" so the "next-hop" may not actually need to be a directly connected neighbour -- the internal router can e.g., point to a loopback interface on the border router. This can become "really uncomfortably messy" and it may be necessary to use a distribute-list to prevent the announcement of the shorter mask.
- o Workaround for 9.0 or older software: create a "default subnet route": "ip route 39.x.y.0 <next-hop>" combined with "ip default-network 39.x.y.0", otherwise as the 9.1 fix.

Both of the latter solutions rely on manual configuration, and in the long run these will be impossible to maintain. In some topologies the use of manual configuration can be a problem (e.g., if there is

more than one possible exit point from the AS to choose from).

Recommendations:

The RFC 1797 experiment appears to have been a success. We believe it safe to start carving up "Class A" space, if the spaces are delegated according to normal IR conventions [7] and recommend the IANA consider this for future address delegations.

Credits:

Thanks to all the RFC 1797 participants. Particular thanks to Paul Vixie, Geert Jan de Groot, and the Staff of the IETF33 Terminal room. Other thanks to ACES, MCI, Altnet, IIJ, UUNET-Canada, Nothwestnet, BBN-Planet, cisco systems, RIPE, RIPE NCC, ESnet, Xlink, SURFnet, STUPI, Connect-AU, INBEnet, SUNET, EUnet, InterPath, VIX.COM, MindSpring. Especial thanks to Suzanne Woolf for cleanup.

References:

- [1] IANA, "Class A Subnet Experiment", RFC 1797, USC/Information Sciences Institute, April 1995.
- [2] Eidnes, H., and G. J. de Groot, "Classless in-addr.arpa delegation", Work in Progress, SINTEF RUNIT, RIPE NCC, May 1995.
- [3] Huston, G., "Observations on the use of Components of the Class A Address Space within the Internet", Work in Progress, AARnet, May 1995.
- [4] Bates, T., et.al, "Representation of IP Routing Policies in a Routing Registry", RFC 1786, MCI, March 1995.
- [5] <http://info.ra.net/div7/ra/Ops.html>, November 1995.
- [6] Baker, F., Editor, "Requirements for IP Version 4 Routers", RFC 1812, cisco systems, June 1995.
- [7] Hubbard, K., Kusters, M., Conrad, D., and D. Karrenberg, "Internet Registry Guidelines", Work in Progress, InterNIC, APNIC, RIPE, November 1995.

Security Considerations

Security issues were not considered in this experiment.

Editor's Address:

Bill Manning
Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292-6695
USA

Phone: +1 310-822-1511 x387

Fax: +1 310-823-6714

EMail: bmanning@isi.edu