Network Working Group   G. Huston
Request for Comments: 2036   Telstra Internet
Category: Informational   October 1996

              Observations on the use of Components of the Class A
                     Address Space within the Internet

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.

Abstract

   This document is a commentary on the recommendation that IANA
   commence allocation of the presently unallocated components of the
   Class A address space to registries, for deployment within the
   Internet as class-less address blocks.

   The document examines the implications for service providers and end
   clients within this environment. The document notes the major
   conclusion that widespread adoption of class-less routing protocols
   is required, within a relatively rapid timeframe for this
   recommendation to be effective.

Introduction

   The Address Lifetime Expectancy (ALE) Working Group of the IETF has
   recorded the allocation of Internet addresses from the unallocated
   address pool. ALE has noted that the existing practice of drawing
   addresses from the Class C space (192/3 address prefix) will result
   in near to medium term exhaustion of this section of the unallocated
   address pool. The largest remaining pool is in the Class A space,
   where some 25% of Internet addresses (the upper half of the Class A
   space) remain, to date, unallocated.

   This document is a commentary on the potential recommendation that
   the Internet Assigned Numbers Authority (IANA), through delegated
   registries, commence allocation of the presently unallocated
   components of the Class A address space to registries, for
   deployment within the Internet through the mechanism of allocation of
   class-less address prefixes.

   The deployment of class-less address prefixes from the Class A space
   within the Internet will require some changes to the routing
   structure within Internet component network domains. The motivation

Huston   Informational   [Page 1]

for, and nature of, such changes as they effect network domains and
network service providers are outlined in this document.

Current Practice with Address Allocations

To date the allocation of class-less network prefixed address blocks
has followed a conservative practice of using address allocations
which are compatible superblocks of Class C addresses, while the
allocation of addresses within the space of Class A and Class B
networks has continued to be aligned with the class-based prefix
structure.

Within this address allocation environment for non-transit network
domains there is accordingly the option to continue to use address
deployment strategies which involve fixed subnet address structures
within contiguous areas, and use Class-full interior routing
protocols. In the situation where variable length subnet masks or
disconnected subnets are deployed within the network domain's routing
structure, interior routing protocols which use subnet-based routing
of Class-full networks can still be successfully deployed and the end
network has the option of using an explicit or implicit sink subnet
default route. Where such non-transit network domains are connected
to the Internet infrastructure the boundary exchange between the
non-transit network and the network service provider (this term is
used as a synonym for a transit network domain, which provides a
traffic transit service to other non-transit and peer transit network
domains) is either a class-full advertisement of routes, or an
aggregated address advertisement where the aggregate is a superblock
of the deployed component class-full networks. At the boundary points
of the non-transit network it is a requirement that the non-transit
network's subnet default route (if used explicitly) not be directed
to the network service provider's domain, to avoid a routing loop at
the domain boundary point.

For network service providers the interior routing protocol can use
either aggregated routing or explicit class-full routing within this
environment. At the network service provider's boundary peering
points the strongly recommended practice is to advertise aggregated
routes to transit peers, which in turn may be further aggregated
across the Internet, within the parameters of permissible policies.

Implications of Address Allocation from the Class A space

Network Service Providers Must Use Class-less Routing

   For network service providers within the deployed Internet the
   implications from this recommendation to deploy prefixes from the
   Class A address space add more pressure to the requirement to
   uniformly deploy class-less routing protocols. While this is already
   a mandatory requirement for any domain which operates without a
   default route (ie. the provider carries full Internet routing and
   effectively calculates default), other providers currently can use
   an imported default route and operate within a class-full routing
   configuration. This mode of operation is sub-optimal, in so far as
   the task of aggregating routes falls on peer network service
   providers performing proxy aggregation of contiguous class-full
   address blocks.

   In deploying components of the Class A the use of proxy aggregation
   is no longer sufficient. Where a domain sees a default route and a
   subnet of a Class A route the routing structure, in a class-full
   configuration, may not necessarily follow the default route to reach
   other parts of the Class A network not covered by the advertised
   Class A subnet route.

   Accordingly for Network Service Providers operating within the
   Internet domain the deployment of components of the Class A space
   entails a requirement to deploy class-less routing protocols, even in
   the presence of a default route. It is noted that this absolute
   requirement is not the case at present.

Consideration of Non-Transit Network Configurations

   For disconnected network environments, where the network domain is
   operated with no links to any peer networking domain, such networks
   can continue to use class-full interior routing protocols with subnet
   support. Allocation of addresses using prefix blocks from the Class A
   space within such environments is possible without adding any
   additional routing or address deployment restrictions on the network
   domain.

Huston   Informational   [Page 3]

   For non-transit network domains which are connected to one or more
   peer network domains the situation does involve consideration of
   additional factors. The observation which is made in the context of
   this consideration is that there are at present relatively few non-
   transit networks operating a fully class-less interior routing
   protocol, as there has been no absolute requirement for this
   functionality when using single class-full network addresses, or when
   using block prefixed address allocations which are clusters of class-
   full network addresses.

   For non-transit network domains which support external peer
   connections to a network service provider, deployment of a component
   of the Class A space would be supportable using a fully class-less
   interior routing protocol.

   In this case there is an additional constraint placed on the external
   connection such that the non-transit domain either agrees that the
   network service will undertake proxy aggregation of the advertised
   class-less address components, or the network domain is configured to
   advertise to the provider an aggregate route. In both cases the
   aggregate route must be either the allocated address block, or a
   fully contained sub-block. Advertising aggregatable address blocks
   without proxy aggregation permission, or advertising multiple sub-
   blocks of the registry allocated address block is considered overly
   deleterious to the provider's internetworking environment due to
   considerations of consequent growth in routing table size.

   If the externally connected non-transit network domain uses class-
   full interior routing protocols then deployment of Class A address
   space prefixes implies that the domain must configure the Class A
   subnet default route along the same path as the default route to the
   network service provider (which is noted to be the exact opposite of
   the necessary routing configuration for those address prefixes which
   are either aligned to class-full address boundaries or are super
   blocks of such class-full address blocks). The network service
   provider may also receive leaked explicit subnet reachability
   information in such a routing configuration, potentially placing the
   responsibility for advertising the correct aggregate address block
   with the network service provider as a case of proxied aggregation.

   Within this configuration model, even when explicit subnet default
   routing is deployed, there is the risk of unintentional traffic
   leakage and routing loops. If the network service provider is
   undertaking proxy aggregation using the registry allocated address
   block then traffic originating within the non-transit domain which is
   (mis)directed to non-deployed components of the address block will
   loop at the interface between the network domain and the provider. If
   the network service provider is configured to explicitly route only

   those address components which are also explicitly routed within the
   non-transit domain, such (mis)directed traffic will be passed through
   the internetworking environment along the default route until a
   default-less routing point is encountered, where it can then be
   discarded. The outcome of this consideration is that the non-transit
   network domain should explicitly configure sink subnet routes for all
   non-deployed components of the allocated address block, and
   conservative operational practice would be to configure the proxy
   aggregation undertaken by the network service provider to aggregate
   according to the registry allocated address block.

   There is an additional constraint placed on the non-transit network
   domain using class-full interior routing protocols, such that the
   domain has no other exterior peer connections to other network
   domains which deploy class-full routing interior routing protocols.

   There is the further constraint placed on the of use of interior
   class-full routing protocols within a non-transit network domain. In
   the case where the non-transit network domain has multiple exterior
   connections to Network Service Providers (ie the network domain is
   multiply homed within a number of network providers) there is the
   possibility that each provider may wish to announce components of the
   same Class A parent. Accordingly the network domain must use a class-
   less interior routing protocol in the case where the network domain
   is multiply homed within network service providers.

   There are also additional constraints placed on the non-transit
   network domain where the network has exterior connections to other
   peer networks. Even in the case where the network domain uses a
   class-less interior routing protocol, there is the additional
   consideration that this requirement for use of a class-less routing
   domain is transitive to other connected network domains. An second
   network domain, externally connected to the class-less domain routing
   part of the Class A space, will interpret the boundary reachability
   advertisement as a complete Class A network advertisement, if using
   class-full routing. Even if both network domains are connected to the
   same network provider the provider's default routing advertisement
   default to the class-full domain will be overridden by the assumed
   class A advertisement through the domain-to-domain connection,
   leading to unintended traffic diversion. The diversion occurs in this
   case as the traffic directed to parts of the Class A network which
   are not deployed within the first domain will transit the first
   domain before entering the network service provider's domain.

   It is also possible to have configurations with unintended routing
   holes. An example of such a configuration is two stub clients of
   different network service providers, both using class-less interior
   routing (X and Y), both directly connected to a third network domain

   (Z), which uses class-full interior routing, which is configured as a
   transit between X and Y. X's advertisement of a component of a Class
   A to Z will be assumed by Z to be a complete Class A network, and as
   such will be advertised to Y, overriding Y's default route received
   from the network service provider. Y will pass all Class A addressed
   traffic to Z, who will in turn pass it to X. As X is configured as a
   non-transit stub network X must discard all non-locally addressed
   traffic.

   Thus reasonable operational practice would be to ensure that if a
   network domain deploys a component of the Class A address space, the
   network domain is configured to use class-less interior routing
   protocols, and the network has a single exterior connection to a
   class-less network provider domain, with the boundary configured as a
   class-less routing exchange. Multiply homed network domains do infer
   a common requirement of class-less routing exchanges and interior
   class-less routing protocols across all peer connected network
   domains.

   It is possible to propose that multi homed network domains should
   probably not get subnets of a class A for these reasons, although
   with an increasing diversity of network service providers instances
   of multi-homed network domains may become more prevalent, and the
   requirement to transition to an interior class-less routing structure
   as a consequence of moving to a multi-homed configuration may not be
   explicitly apparent to all network domains.

Potential Guidelines for Allocation of an Address Prefix from the Class
   A Address Space

   To summarise the possible guidelines for allocation from the Class A
   space, such addresses should only be assigned to network domains
   which:

     - have no exterior connection (in which case the domain can use
       either class-full or class-less interior routing protocols without
       further implication),

    or

     - are a component of a private internet domain which uses class-full
       routing exchanges and no other part of the same Class A is
       assigned into the domain (this is probably an unlikely scenario
       given a probable direction to use the Class A space as the major
       resource for the unallocated pool of addresses for allocation),

        or

       - have a single default exterior connection to a class-less routing
         domain, use class-full routing protocols and explicitly direct a
         subnet default route to the exterior connection,

        or

       - use class-less interior routing protocols and connect only to
         other network domains which also use class-less interior routing
         protocols.

    It is a reasonable objective to nominate a transition objective to
    the final configuration (uniform use of class-less routing domains
    within the Internet) which would enable deployment of components of
    the Class A space uniformly across the Internet.

Related Potential Activities

    Given the pressures on the remaining Class C address space in the
    unallocated address pool, it is noted that there would be widespread
    deployment of components of the remaining Class A space in class-less
    allocation guidelines. There is a consequent requirement for
    widespread deployment of class-less interior routing protocols in
    order to ensure continued correct operation of the routed Internet.
    This is a more significant transition than that deployed to date with
    the network service providers' deployment of Class-less Inter-Domain
    Routing (CIDR) protocols, in that there is a necessary transition to
    deploy Class-less Interior Routing Protocols (CIRP) within a large
    number of network domains which are currently configured with class-
    full routing.

    However this would appear to be a necessary task if we wish to
    continue to utilise a pool of globally unique Internet addresses to
    allocate to new systems and networks, but one requiring significant
    effort considering the space of the routing transition required to
    make this work.

    There are a number of directed activities which can assist in this
    transition:

     - The network registries commence initial class-less allocation from
       the unallocated Class A space to those entities who either:

       o operate a CIRP environment, and either have no external
          connectivity, or are singly homed to a network service provider
          using a CIDR environment, with no other exterior connections,

          or

          o operate a class-full routing protocol, and either have no
             external connectivity, or are singly homed to a network service
             provider using a CIDR environment, with no other exterior
             connections, and are willing to point the subnet default route
             towards the network service provider.

        - In deploying the Class A space there is a requirement within the
          vendors' product sets to allow explicit configuration of whether
          the router operates in a class-less or class-full mode, with
          correct behaviour of the default route in each case. Class-full
          mode of operation must also allow explicit configuration of
          subnet default behaviour as to whether to follow the default
          route, or to operate a subnet default sink.

        - There is a similar, but longer term, activity within the host
          configuration environment to support a mode of address
          configuration which uses a local network prefix and host address,
          possibly in addition to the current configuration mode of class-
          full network, subnet and host address

        - Internet Service Providers also must support full class-less
          configurations in both interior routing configurations and
          interdomain peering routing exchanges, and provide support to
          client network domains operating a class-less boundary routing
          exchange configuration and be able to undertake proxy aggregation
          as permitted.

Security Considerations

   Correct configuration of the routing environment of the Internet is
   essential to the secure operation of the Internet.

   The potential use of the Class A space raises no additional
   considerations in this area.

References

   [CIDR]
        Fuller, V., T. Li, J. Yu, and K. Varadhan, "Classless Inter-
        Domain Routing (CIDR): an Address Assignment and Aggregation
        Strategy", RFC 1519, BARRnet, cisco, MERIT, OARnet, September
        1993.

Author's Address

        Geoff Huston
        Telstra Internet
        Locked Bag 5744
        Canberra ACT 2601
        Australia

        phone: +61 6 208 1908
        email: gih@telstra.net