           A Set of Possible Requirements for a Future Routing Architecture

Abstract

   The requirements for routing architectures described in this document
   were produced by two sub-groups under the IRTF Routing Research Group
   (RRG) in 2001, with some editorial updates up to 2006. The two sub-
   groups worked independently, and the resulting requirements represent
   two separate views of the problem and of what is required to fix the
   problem. This document may usefully serve as part of the recommended
   reading for anyone who works on routing architecture designs for the
   Internet in the future.

   The document is published with the support of the IRTF RRG as a
   record of the work completed at that time, but with the understanding
   that it does not necessarily represent either the latest technical
   understanding or the technical consensus of the research group at the
   date of publication.

Status of This Memo

Copyright Notice

Table of Contents

1. Background

     In 2001, the IRTF Routing Research Group (IRTF RRG) chairs, Abha
     Ahuja and Sean Doran, decided to establish a sub-group to look at
     requirements for inter-domain routing (IDR). A group of well-known
     routing experts was assembled to develop requirements for a new
     routing architecture. Their mandate was to approach the problem
     starting from a blank slate. This group was free to take any
     approach, including a revolutionary approach, in developing
     requirements for solving the problems they saw in inter-domain
     routing.

     Simultaneously, an independent effort was started in Sweden with a
     similar goal. A team, calling itself Babylon, with participation
     from vendors, service providers, and academia assembled to understand
     the history of inter-domain routing, to research the problems seen by
     the service providers, and to develop a proposal of requirements for
     a follow-on to the current routing architecture. This group's remit
     required an evolutionary approach starting from current routing
     architecture and practice. In other words, the group limited itself
     to developing an evolutionary strategy. The Babylon group was later
     folded into the IRTF RRG as Sub-Group B to distinguish it from the
     original RRG Sub-Group A.

     One of the questions that arose while the groups were working in
     isolation was whether there would be many similarities between their
     sets of requirements. That is, would the requirements that grew from
     a blank sheet of paper resemble those that started with the
     evolutionary approach? As can be seen from reading the two sets of
     requirements, there were many areas of fundamental agreement but some
     areas of disagreement.

     There were suggestions within the RRG that the two teams should work
     together to create a single set of requirements. Since these
     requirements are only guidelines to future work, however, some felt
     that doing so would risk losing content without gaining any
     particular advantage. It is not as if any group, for example, the
     IRTF RRG or the IETF Routing Area, was expected to use these
     requirements as written and to create an architecture that met these
     requirements. Rather, the requirements were in practice strong

recommendations for a way to proceed in creating a new routing
architecture. In the end, the decision was made to include the
results of both efforts, side by side, in one document.

This document contains the two requirement sets produced by the
teams. The text has received only editorial modifications; the
requirements themselves have been left unaltered. Whenever the
editors felt that conditions had changed in the few years since the
text was written, an editors' note has been added to the text.

In reading this document, it is important to keep in mind that all of
these requirements are suggestions, which are laid out to assist
those interested in developing new routing architectures. It is also
important to remember that, while the people working on these
suggestions have done their best to make intelligent suggestions,
there are no guarantees. So a reader of this document should not
treat what it says as absolute, nor treat every suggestion as
necessary. No architecture is expected to fulfill every
"requirement". Hopefully, though, future architectures will consider
what is offered in this document.

The IRTF RRG supported publication of this document as a historical
record of the work completed on the understanding that it does not
necessarily represent either the latest technical understanding or
the technical consensus of the research group at the time of
publication. The document has had substantial review by members of
the two teams, other members of the IRTF RRG, and additional experts
over the years.

Finally, this document does not make any claims that it is possible
to have a practical solution that meets all the listed requirements.

2. Results from Group A

   This section presents the results of the work done by Sub-Group A of
   the IRTF RRG during 2001-2002. The work originally appeared under
   the title: "Requirements For a Next Generation Routing and Addressing
   Architecture" and was edited by Frank Kastenholz.

2.1. Group A - Requirements for a Next Generation Routing and
      Addressing Architecture

   The requirements presented in this section are not presented in any
   particular order.

2.1.1. Architecture

   The new routing and addressing protocols, data structures, and
   algorithms need to be developed from a clear, well thought-out, and
   documented architecture.

   The new routing and addressing system must have an architectural
   specification that describes all of the routing and addressing
   elements, their interactions, what functions the system performs, and
   how it goes about performing them. The architectural specification
   does not go into issues such as protocol and data structure design.

   The architecture should be agnostic with regard to specific
   algorithms and protocols.

   Doing architecture before doing detailed protocol design is good
   engineering practice. This allows the architecture to be reviewed
   and commented upon, with changes made as necessary, when it is still
   easy to do so. Also, by producing an architecture, the eventual
   users of the protocols (the operations community) will have a better
   understanding of how the designers of the protocols meant them to be
   used.

2.1.2. Separable Components

   The architecture must place different functions into separate
   components.

   Separating functions, capabilities, and so forth into individual
   components and making each component "stand alone" is generally
   considered by system architects to be "A Good Thing". It allows
   individual elements of the system to be designed and tuned to do
   their jobs "very well". It also allows for piecemeal replacement and
   upgrading of elements as new technologies and algorithms become
   available.

   The architecture must have the ability to replace or upgrade existing
   components and to add new ones, without disrupting the remaining
   parts of the system. Operators must be able to roll out these
   changes and additions incrementally (i.e., no "flag days"). These
   abilities are needed to allow the architecture to evolve as the
   Internet changes.

   The architecture specification shall define each of these components,
   their jobs, and their interactions.

Some thoughts to consider along these lines are:

o Making topology and addressing separate subsystems. This may
   allow highly optimized topology management and discovery without
   constraining the addressing structure or physical topology in
   unacceptable ways.

o Separate "fault detection and healing" from basic topology. From
   Mike O'Dell:

      Historically the same machinery is used for both. While
      attractive for many reasons, the availability of exogenous
      topology information (i.e., the intended topology) should, it
      seems, make some tasks easier than the general case of starting
      with zero knowledge. It certainly helps with recovery in the
      case of constraint satisfaction. In fact, the intended
      topology is a powerful way to state certain kinds of policy.
      [ODell01]

o Making policy definition and application a separate subsystem,
   layered over the others.

The architecture should also separate topology, routing, and
addressing from the application that uses those components. This
implies that applications such as policy definition, forwarding, and
circuit and tunnel management are separate subsystems layered on top
of the basic topology, routing, and addressing systems.

2.1.3. Scalable

Scaling is the primary problem facing the routing and addressing
architecture today. This problem must be solved and it must be
solved for the long term.

The architecture must support a large and complex network. Ideally,
it will serve our needs for the next 20 years. Unfortunately:

1. we do not know how big the Internet will grow over that time, and

2. the architecture developed from these requirements may change the
    fundamental structure of the Internet and therefore its growth
    patterns. This change makes it difficult to predict future
    growth patterns of the Internet.

As a result, we can't quantify the requirement in any meaningful way.
Using today's architectural elements as a mechanism for describing
things, we believe that the network could grow to:

   1. tens of thousands of ASs

          Editors' Note: As of 2005, this level had already been
          reached.

   2. tens to hundreds of millions of prefixes, during the lifetime of
       this architecture.

   These sizes are given as a "flavor" for how we expect the Internet to
   grow. We fully believe that any new architecture may eliminate some
   current architectural elements and introduce new ones.

   A new routing and addressing architecture designed for a specific
   network size would be inappropriate. First, the cost of routing
   calculations is based only in part on the number of ASs or prefixes
   in the network. The number and locations of the links in the network
   are also significant factors. Second, past predictions of Internet
   growth and topology patterns have proven to be wildly inaccurate, so
   developing an architecture to a specific size goal would at best be
   shortsighted.

      Editors' Note: At the time of these meetings, the BGP statistics
      kept at sites such as www.routeviews.org either did not exist or
      had been running for only a few months. After 5 years of
      recording public Internet data trends in AS growth, routing table
      growth can be observed (past) with some short-term prediction. As
      each year of data collection continues, the ability to observe and
      predict trends improves. This architecture work pointed out the
      need for such statistics to improve future routing designs.

   Therefore, we will not make the scaling requirement based on a
   specific network size. Instead, the new routing and addressing
   architecture should have the ability to constrain the increase in
   load (CPU, memory space and bandwidth, and network bandwidth) on ANY
   SINGLE ROUTER to be less than these specific functions:

   1. The computational power and memory sizes required to execute the
       routing protocol software and to contain the tables must grow
       more slowly than hardware capabilities described by Moore's Law,
       doubling every 18 months. Other observations indicate that
       memory sizes double every 2 years or so.

   2. Network bandwidth and latency are some key constraints on how
       fast routing protocol updates can be disseminated (and therefore
       how fast the routing system can adapt to changes). Raw network
       bandwidth seems to quadruple every 3 years or so. However, it
       seems that there are some serious physics problems in going
       faster than 40 Gbit/s (OC768); we should not expect raw network

link speed to grow much beyond OC768. On the other hand, for
economic reasons, large swathes of the core of the Internet will
still operate at lower speeds, possibly as slow as DS3.

> Editors' Note: Technology is running ahead of imagination and
> higher speeds are already common.

Furthermore, in some sections of the Internet, even lower speed
links are found. Corporate access links are often T1, or slower.
Low-speed radio links exist. Intra-domain links may be T1 or
fractional-T1 (or slower).

Therefore, the architecture must not make assumptions about the
bandwidth available.

3. The speeds of high-speed RAMs (Static RAMs (SRAMs), used for
    caches and the like) are growing, though slowly. Because of
    their use in caches and other very specific applications, these
    RAMs tend to be small, a few megabits, and the size of these RAMs
    is not increasing very rapidly.

    On the other hand, the speed of "large" memories (Dynamic RAMs
    (DRAMs)) is increasing even slower than that for the high-speed
    RAMs. This is because the development of these RAMs is driven by
    the PC market, where size is very important, and low speed can be
    made up for by better caches.

    Memory access rates should not be expected to increase
    significantly.

    > Editors' Note: Various techniques have significantly increased
    > memory bandwidth. 800 MHz is now possible, compared with less
    > than 100 MHz in the year 2000. This does not, however,
    > contradict the next paragraph, but rather just extends the
    > timescales somewhat.

The growth in resources available to any one router will eventually
slow down. It may even stop. Even so, the network will continue to
grow. The routing and addressing architecture must continue to scale
in even this extreme condition. We cannot continue to add more
computing power to routers forever. Other strategies must be
available. Some possible strategies are hierarchy, abstraction, and
aggregation of topology information.

2.1.4. Lots of Interconnectivity

   The new routing and addressing architecture must be able to cope with
   a high degree of interconnectivity in the Internet. That is, there
   are large numbers of alternate paths and routes among the various
   elements. Mechanisms are required to prevent this interconnectivity
   (and continued growth in interconnectivity) from causing tables,
   compute time, and routing protocol traffic to grow without bound.
   The "cost" to the routing system of an increase in complexity must be
   limited in scope; sections of the network that do not see, or do not
   care about, the complexity ought not pay the cost of that complexity.

   Over the past several years, the Internet has seen an increase in
   interconnectivity. Individual end sites (companies, customers,
   etc.), ISPs, exchange points, and so on, all are connecting to more
   "other things". Companies multi-home to multiple ISPs, ISPs peer
   with more ISPs, and so on. These connections are made for many
   reasons, such as getting more bandwidth, increased reliability and
   availability, policy, and so on. However, this increased
   interconnectivity has a price. It leads to more scaling problems as
   it increases the number of AS paths in the networks.

   Any new architecture must assume that the Internet will become a
   denser mesh. It must not assume, nor can it dictate, certain
   patterns or limits on how various elements of the network
   interconnect.

   Another facet of this requirement is that there may be multiple
   valid, loop-free paths available to a destination. See Section 2.1.7
   for a further discussion.

   We wryly note that one of the original design goals of IP was to
   support a large, heavily interconnected network, which would be
   highly survivable (such as in the face of a nuclear war).

2.1.5. Random Structure

   The routing and addressing architecture must not place any
   constraints on or make assumptions about the topology or
   connectedness of the elements comprising the Internet. The routing
   and addressing architecture must not presume any particular network
   structure. The network does not have a "nice" structure. In the
   past, we used to believe that there was this nice "backbone/tier-1/
   tier-2/end-site" sort of hierarchy. This is not so. Therefore, any
   new architecture must not presume any such structure.

Some have proposed that a geographic addressing scheme be used,
requiring exchange points to be situated within each geographic
"region". There are many reasons why we believe this to be a bad
approach, but those arguments are irrelevant. The main issue is that
the routing architecture should not presume a specific network
structure.

2.1.6. Multi-Homing

The architecture must provide multi-homing for all elements of the
Internet. That is, multi-homing of hosts, subnetworks, end-sites,
"low-level" ISPs, and backbones (i.e., lots of redundant
interconnections) must be supported. Among the reasons to multi-home
are reliability, load sharing, and performance tuning.

The term "multi-homing" may be interpreted in its broadest sense --
one "place" has multiple connections or links to another "place".

The architecture must not limit the number of alternate paths to a
multi-homed site.

When multi-homing is used, it must be possible to use one, some (more
than one but less than all), or all of the available paths to the
multi-homed site. The multi-homed site must have the ability to
declare which path(s) are used and under what conditions (for
example, one path may be declared "primary" and the other "backup"
and to be used only when the primary fails).

A current problem in the Internet is that multi-homing leads to undue
increases in the size of the BGP routing tables. The new
architecture must support multi-homing without undue routing table
growth.

2.1.7. Multi-Path

As a corollary to multi-homing, the architecture must allow for
multiple paths from a source to a destination to be active at the
same time. These paths need not have the same attributes. Policies
are to be used to disseminate the attributes and to classify traffic
for the different paths.

There must be a rich "language" for specifying the rules for
classifying the traffic and assigning classes of traffic to different
paths (or prohibiting it from certain paths). The rules should allow
traffic to be classified based upon, at least, the following:

   o IPv6 FlowIDs,

   o Diffserv Code Point (DSCP) values,

   o source and/or destination prefixes, or

   o random selections at some probability.

   A mechanism is needed that allows operators to plan and manage the
   traffic load on the various paths. To start, this mechanism can be
   semi-automatic or even manual. Eventually, it ought to become fully
   automatic.

   When multi-path forwarding is used, options must be available to
   preserve packet ordering where appropriate (such as for individual
   TCP connections).

   Please refer to Section 2.2.7 for a discussion of dynamic load-
   balancing and management over multiple paths.

2.1.8. Convergence

   The speed of convergence (also called the "stabilization time") is
   the time it takes for a router's routing processes to reach a new,
   stable, "solution" (i.e., forwarding information base) after a change
   someplace in the network. In effect, what happens is that the output
   of the routing calculations stabilizes -- the Nth iteration of the
   software produces the same results as the N-1th iteration.

   The speed of convergence is generally considered to be a function of
   the number of subnetworks in the network and the amount of
   connections between those networks. As either number grows, the time
   it takes to converge increases.

   In addition, a change can "ripple" back and forth through the system.
   One change can go through the system, causing some other router to
   change its advertised connectivity, causing a new change to ripple
   through. These oscillations can take a while to work their way out
   of the network. It is also possible that these ripples never die
   out. In this situation, the routing and addressing system is
   unstable; it never converges.

   Finally, it is more than likely that the routers comprising the
   Internet never converge simply because the Internet is so large and
   complex. Assume it takes S seconds for the routers to stabilize on a
   solution for any one change to the network. Also, assume that
   changes occur, on average, every C seconds. Because of the size and
   complexity of the Internet, C is now less than S. Therefore, if a

change, C1, occurs at time T, the routing system would stabilize at
time T+S, but a new change, C2, will occur at time T+C, which is
before T+S. The system will start processing the new change before
it's done with the old.

This is not to say that all routers are constantly processing
changes. The effects of changes are like ripples in a pond. They
spread outward from where they occur. Some routers will be
processing just C1, others C2, others both C1 and C2, and others
neither.

We have two separate scopes over which we can set requirements with
respect to convergence:

1. Single Change: In this requirement, a single change of any type
    (link addition or deletion, router failure or restart, etc.) is
    introduced into a stabilized system. No additional changes are
    introduced. The system must re-stabilize within some measure of
    bounded time. This requirement is a fairly abstract one as it
    would be impossible to test in a real network. Definition of the
    time constraints remains an open research issue.

2. System-Wide: Defining a single target for maximum convergence
    time for the real Internet is absurd. As we mentioned earlier,
    the Internet is large enough and diverse enough so that it is
    quite likely that new changes are introduced somewhere before the
    system fully digests old ones.

So, the first requirement here is that there must be mechanisms to
limit the scope of any one change's visibility and effects. The
number of routers that have to perform calculations in response to a
change is kept small, as is the settling time.

The second requirement is based on the following assumptions:

- the scope of a change's visibility and impact can be limited.
   That is, routers within that scope know of the change and
   recalculate their tables based on the change. Routers outside of
   the scope don't see it at all.

- Within any scope, S, network changes are constantly occurring and
   the average inter-change interval is Tc seconds.

- There are Rs routers within scope S.

- A subset of the destinations known to the routers in S, Ds, are
   impacted by a given change.

- We can state that for Z% of the changes, within Y% of Tc seconds
  after a change, C, X% of the Rs routers have their routes to Ds
  settled to a useful answer (useful meaning that packets can get to
  Ds, though perhaps not by the optimal path -- this allows some
  "hunting" for the optimal solution).

  X, Y, and Z are yet to be defined. Their definition remains a
  research issue.

This requirement implies that the scopes can be kept relatively small
in order to minimize Rs and maximize Tc.

The growth rate of the convergence time must not be related to the
growth rate of the Internet as a whole. This implies that the
convergence time either:

1. not be a function of basic network elements (such as prefixes and
    links/paths), and/or

2. that the Internet be continuously divisible into chunks that
    limit the scope and effect of a change, thereby limiting the
    number of routers, prefixes, links, and so on, involved in the
    new calculations.

2.1.9. Routing System Security

The security of the Internet's routing system is paramount. If the
routing system is compromised or attacked, the entire Internet can
fail. This is unacceptable. Any new architecture must be secure.

Architectures by themselves are not secure. It is the implementation
of an architecture, its protocols, algorithms, and data structures
that are secure. These requirements apply primarily to the
implementation. The architecture must provide the elements that the
implementation needs to meet these security requirements. Also, the
architecture must not prevent these security requirements from being
met.

Security means different things to different people. In order for
this requirement to be useful, we must define what we mean by
security. We do this by identifying the attackers and threats we
wish to protect against. They are:

Masquerading
    The system, including its protocols, must be secure against
    intruders adopting the identity of other known, trusted
    elements of the routing system and then using that position of
    trust for carrying out other attacks. Protocols must use
    cryptographically strong authentication.

Denial-of-Service (DoS) Attacks
    The architecture and protocols should be secure against DoS
    attacks directed at the routers.

    The new architecture and protocols should provide as much
    information as it can to allow administrators to track down
    sources of DoS and Distributed DOS (DDoS) attacks.

No Bad Data
    Any new architecture and protocols must provide protection
    against the introduction of bad, erroneous, or misleading data
    by attackers. Of particular importance, an attacker must not
    be able to redirect traffic flows, with the intent of:

    o directing legitimate traffic away from a target, causing a
       denial-of-service attack by preventing legitimate data from
       reaching its destination,

    o directing additional traffic (going to other destinations
       that are "innocent bystanders") to a target, causing the
       target to be overloaded, or

    o directing traffic addressed to the target to a place where
       the attacker can copy, snoop, alter, or otherwise affect the
       traffic.

Topology Hiding
    Any new architecture and protocols must provide mechanisms to
    allow network owners to hide the details of their internal
    topologies, while maintaining the desired levels of service
    connectivity and reachability.

Privacy
    By "privacy" we mean privacy of the routing protocol exchanges
    between routers.

    When the routers are on point-to-point links, with routers at
    each end, there may not be any need to encrypt the routing
    protocol traffic as the possibility of a third party

intercepting the traffic is limited, though not impossible. We
do believe, however, that it is important to have the ability
to protect routing protocol traffic in two cases:

1. When the routers are on a shared network, it is possible
    that there are hosts on the network that have been
    compromised. These hosts could surreptitiously monitor the
    protocol traffic.

2. When two routers are exchanging information "at a distance"
    (over intervening routers and, possibly, across
    administrative domain boundaries). In this case, the
    security of the intervening routers, links, and so on,
    cannot be assured. Thus, the ability to encrypt this
    traffic is important.

Therefore, we believe that the option to encrypt routing
protocol traffic is required.

Data Consistency
    A router should be able to detect and recover from any data
    that is received from other routers that is inconsistent. That
    is, it must not be possible for data from multiple routers,
    none of which is malicious, to "break" another router.

Where security mechanisms are provided, they must use methods that
are considered to be cryptographically secure (e.g., using
cryptographically strong encryption and signatures -- no cleartext
passwords!).

Use of security features should not be optional (except as required
above). This may be "social engineering" on our part, but we believe
it to be necessary. If a security feature is optional, the
implementation of the feature must default to the "secure" setting.

2.1.10. End Host Security

The architecture must not prevent individual host-to-host
communications sessions from being secured (i.e., it cannot interfere
with things like IPsec).

2.1.11. Rich Policy

Before setting out policy requirements, we need to define the term.
Like "security", "policy" means different things to different people.
For our purposes, "policy" is the set of administrative influences
that alter the path determination and next-hop selection procedures
of the routing software.

The main motivators for influencing path and next-hop selection seem
to be transit rules, business decisions, and load management.

The new architecture must support rich policy mechanisms.
Furthermore, the policy definition and dissemination mechanisms
should be separated from the network topology and connectivity
dissemination mechanisms. Policy provides input to and controls the
generation of the forwarding table and the abstraction, filtering,
aggregation, and dissemination of topology information.

Note that if the architecture is properly divided into subsystems,
then at a later time, new policy subsystems that include new features
and capabilities could be developed and installed as needed.

We divide the general area of policy into two sub-categories: routing
information and traffic control. Routing Information Policies
control what routing information is disseminated or accepted, how it
is disseminated, and how routers determine paths and next-hops from
the received information. Traffic Control Policies determine how
traffic is classified and assigned to routes.

2.1.11.1. Routing Information Policies

There must be mechanisms to allow network administrators, operators,
and designers to control receipt and dissemination of routing
information. These controls include, but are not limited to:

- Selecting to which other routers routing information will be
  transmitted.

- Specifying the "granularity" and type of transmitted information.
  The length of IPv4 prefixes is an example of granularity.

- Selection and filtering of topology and service information that
  is transmitted. This gives different "views" of internal
  structure and topology to different peers.

- Selecting the level of security and authenticity for transmitted
  information.

- Being able to cause the level of detail that is visible for some
  portion of the network to reduce the farther you get from that
  part of the network.

- Selecting from whom routing information will be accepted. This
  control should be "provisional" in the sense of "accept routes
  from "foo" only if there are no others available".

     - Accepting or rejecting routing information based on the path the
        information traveled (using the current system as an example, this
        would be filtering routes based on an AS appearing anywhere in the
        AS path). This control should be "use only if there are no other
        paths available".

     - Selecting the desired level of granularity for received routing
        information (this would include, but is not limited to, things
        similar in nature to the prefix-length filters widely used in the
        current routing and addressing system).

     - Selecting the level of security and authenticity of received
        information in order for that information to be accepted.

     - Determining the treatment of received routing information based on
        attributes supplied with the information.

     - Applying attributes to routing information that is to be
        transmitted and then determining treatment of information (e.g.,
        sending it "here" but not "there") based on those tags.

     - Selection and filtering of topology and service information that
        is received.

2.1.11.2. Traffic Control Policies

   The architecture should provide mechanisms that allow network
   operators to manage and control the flow of traffic. The traffic
   controls should include, but are not limited to:

   - The ability to detect and eliminate congestion points in the
      network (by redirecting traffic around those points).

   - The ability to develop multiple paths through the network with
      different attributes and then assign traffic to those paths based
      on some discriminators within the packets (discriminators include,
      but are not limited to, IP addresses or prefixes, IPv6 flow ID,
      DSCP values, and MPLS labels).

   - The ability to find and use multiple, equivalent paths through the
      network (i.e., they would have the "same" attributes) and allocate
      traffic across the paths.

   - The ability to accept or refuse traffic based on some traffic
      classification (providing, in effect, transit policies).

Traffic classification must at least include the source and
destination IP addresses (prefixes) and the DSCP value. Other
fields may be supported, such as:

* Protocol and port-based functions,

* DSCP/QoS (Quality of Service) tuple (such as ports),

* Per-host operations (i.e., /32 s for IPv4 and /128 s for IPv6),
   and

* Traffic matrices (e.g., traffic from prefix X and to prefix Y).

2.1.12. Incremental Deployment

   The reality of the Internet is that there can be no Internet-wide
   cutover from one architecture and protocol to another. This means
   that any new architecture and protocol must be incrementally
   deployable; ISPs must be able to set up small sections of the new
   architecture, check it out, and then slowly grow the sections.
   Eventually, these sections will "touch" and "squeeze out" the old
   architecture.

   The protocols that implement the architecture must be able to
   interoperate at "production levels" with currently existing routing
   protocols. Furthermore, the protocol specifications must define how
   the interoperability is done.

   We also believe that sections of the Internet will never convert over
   to the new architecture. Thus, it is important that the new
   architecture and its protocols be able to interoperate with "old
   architecture" regions of the network indefinitely.

   The architecture's addressing system must not force existing address
   allocations to be redone: no renumbering!

2.1.13. Mobility

   There are two kinds of mobility: host mobility and network mobility.
   Host mobility is when an individual host moves from where it was to
   where it is. Network mobility is when an entire network (or
   subnetwork) moves.

   The architecture must support network-level mobility. Please refer
   to Section 2.2.9 for a discussion of host mobility.

Editors' Note: Since the time of this work, the Network Mobility
(NEMO) extensions to Mobile-IP [RFC3963] to accommodate mobile
networks have been developed.

2.1.14. Address Portability

One of the big "hot items" in the current Internet political climate
is portability of IP addresses (both v4 and v6). The short
explanation is that people do not like to renumber when changing
connection point or provider and do not trust automated renumbering
tools.

The architecture must provide complete address portability.

2.1.15. Multi-Protocol

The Internet is expected to be "multi-protocol" for at least the next
several years. IPv4 and IPv6 will co-exist in many different ways
during a transition period. The architecture must be able to handle
both IPv4 and IPv6 addresses. Furthermore, protocols that supplant
IPv4 and IPv6 may be developed and deployed during the lifetime of
the architecture. The architecture must be flexible and extensible
enough to handle new protocols as they arise.

Furthermore, the architecture must not assume any given relationships
between a topological element's IPv4 address and its IPv6 address.
The architecture must not assume that all topological elements have
IPv4 addresses/prefixes, nor can it assume that they have IPv6
addresses/prefixes.

The architecture should allow different paths to the same destination
to be used for different protocols, even if all paths can carry all
protocols.

In addition to the addressing technology, the architecture need not
be restricted to only packet-based multiplexing/demultiplexing
technology (such as IP); support for other multiplexing/
demultiplexing technologies may be added.

2.1.16. Abstraction

The architecture must provide mechanisms for network designers and
operators to:

o Group elements together for administrative control purposes,

o Hide the internal structure and topology of those groupings for
   administrative and security reasons,

   o Limit the amount of topology information that is exported from the
      groupings in order to control the load placed on external routers,

   o Define rules for traffic transiting or terminating in the
      grouping.

   The architecture must allow the current Autonomous System structure
   to be mapped into any new abstraction schemes.

   Mapping mechanisms, algorithms, and techniques must be specified.

2.1.17. Simplicity

   The architecture must be simple enough so that someone who is
   extremely knowledgeable in routing and who is skilled at creating
   straightforward and simple explanations can explain all the important
   concepts in less than an hour.

   This criterion has been chosen since developing an objective measure
   of complexity for an architecture can be very difficult and is out of
   scope for this document.

   The requirement is that the routing architecture be kept as simple as
   possible. This requires careful evaluation of possible features and
   functions with a merciless weeding out of those that "might be nice"
   but are not necessary.

   By keeping the architecture simple, the protocols and software used
   to implement the architecture are simpler. This simplicity in turn
   leads to:

   1. Faster implementation of the protocols. If there are fewer bells
       and whistles, then there are fewer things that need to be
       implemented.

   2. More reliable implementations. With fewer components, there is
       less code, reducing bug counts, and fewer interactions between
       components that could lead to unforeseen and incorrect behavior.

2.1.18. Robustness

   The architecture, and the protocols implementing it, should be
   robust. Robustness comes in many different flavors. Some
   considerations with regard to robustness include (but are not limited
   to):

   o Continued correct operation in the face of:

      * Defective (even malicious) trusted routers.

      * Network failures. Whenever possible, valid alternate paths are
         to be found and used.

   o Failures must be localized. That is, the architecture must limit
      the "spread" of any adverse effects of a misconfiguration or
      failure. Badness must not spread.

   Of course, the general robustness principle of being liberal in
   what's accepted and conservative in what's sent must also be applied.

      Original Editor's Note: Some of the contributors to this section
      have argued that robustness is an aspect of security. I have
      exercised editor's discretion by making it a separate section.
      The reason for this is that to too many people "security" means
      "protection from break-ins" and "authenticating and encrypting
      data". This requirement goes beyond those views.

2.1.19. Media Independence

   While it is an article of faith that IP operates over a wide variety
   of media (such as Ethernet, X.25, ATM, and so on), IP routing must
   take an agnostic view toward any "routing" or "topology" services
   that are offered by the medium over which IP is operating. That is,
   the new architecture must not be designed to integrate with any
   media-specific topology management or routing scheme.

   The routing architecture must assume, and must work over, the
   simplest possible media.

   The routing and addressing architecture can certainly make use of
   lower-layer information and services, when and where available, and
   to the extent that IP routing wishes.

2.1.20. Stand-Alone

   The routing architecture and protocols must not rely on other
   components of the Internet (such as DNS) for their correct operation.
   Routing is the fundamental process by which data "finds its way
   around the Internet" and most, if not all, of those other components
   rely on routing to properly forward their data. Thus, routing cannot
   rely on any Internet systems, services, or capabilities that in turn
   rely on routing. If it did, a dependency loop would result.

2.1.21. Safety of Configuration

   The architecture, protocols, and standard implementation defaults
   must be such that a router installed "out of the box" with no
   configuration, etc., by the operators will not cause "bad things" to
   happen to the rest of the routing system (e.g., no dial-up customers
   advertising routes to 18/8!).

2.1.22. Renumbering

   The routing system must allow topological entities to be renumbered.

2.1.23. Multi-Prefix

   The architecture must allow topological entities to have multiple
   prefixes (or the equivalent under the new architecture).

2.1.24. Cooperative Anarchy

   As RFC 1726[RFC1726] states:

      A major contributor to the Internet's success is the fact that
      there is no single, centralized, point of control or promulgator
      of policy for the entire network. This allows individual
      constituents of the network to tailor their own networks,
      environments, and policies to suit their own needs. The
      individual constituents must cooperate only to the degree
      necessary to ensure that they interoperate.

   This decentralization, called "cooperative anarchy", is still a key
   feature of the Internet today. The new routing architecture must
   retain this feature. There can be no centralized point of control or
   promulgator of policy for the entire Internet.

2.1.25. Network-Layer Protocols and Forwarding Model

   For the purposes of backward compatibility, any new routing and
   addressing architecture and protocols must work with IPv4 and IPv6
   using the traditional "hop-by-hop" forwarding and packet-based
   multiplex/demultiplex models. However, the architecture need not be
   restricted to these models. Additional forwarding and multiplex/
   demultiplex models may be added.

2.1.26. Routing Algorithm

   The architecture should not require a particular routing algorithm
   family. That is to say, the architecture should be agnostic about
   link-state, distance-vector, or path-vector routing algorithms.

2.1.27. Positive Benefit

   Finally, the architecture must show benefits in terms of increased
   stability, decreased operational costs, and increased functionality
   and lifetime, over the current schemes. This benefit must remain
   even after the inevitable costs of developing and debugging the new
   protocols, enduring the inevitable instabilities as things get shaken
   out, and so on.

2.1.28. Administrative Entities and the IGP/EGP Split

   We explicitly recognize that the Internet consists of resources under
   control of multiple administrative entities. Each entity must be
   able to manage its own portion of the Internet as it sees fit.
   Moreover, the constraints that can be imposed on routing and
   addressing on the portion of the Internet under the control of one
   administration may not be feasibly extended to cover multiple
   administrations. Therefore, we recognize a natural and inevitable
   split between routing and addressing that is under a single
   administrative control and routing and addressing that involves
   multiple administrative entities. Moreover, while there may be
   multiple administrative authorities, the administrative authority
   boundaries may be complex and overlapping, rather than being a strict
   hierarchy.

   Furthermore, there may be multiple levels of administration, each
   with its own level of policy and control. For example, a large
   network might have "continental-level" administrations covering its
   European and Asian operations, respectively. There would also be
   that network's "inter-continental" administration covering the
   Europe-to-Asia links. Finally, there would be the "Internet" level
   in the administrative structure (analogous to the "exterior" concept
   in the current routing architecture).

   Thus, we believe that the administrative structure of the Internet
   must be extensible to many levels (more than the two provided by the
   current IGP/EGP split). The interior/exterior property is not
   absolute. The interior/exterior property of any point in the network
   is relative; a point on the network is interior with respect to some
   points on the network and exterior with respect to others.

   Administrative entities may not trust each other; some may be almost
   actively hostile toward each other. The architecture must
   accommodate these models. Furthermore, the architecture must not
   require any particular level of trust among administrative entities.

2.2. Non-Requirements

   The following are not required or are non-goals. This should not be
   taken to mean that these issues must not be addressed by a new
   architecture. Rather, addressing these issues or not is purely an
   optional matter for the architects.

2.2.1. Forwarding Table Optimization

   We believe that it is not necessary for the architecture to minimize
   the size of the forwarding tables (FIBs). Current memory sizes,
   speeds, and prices, along with processor and Application-specific
   Integrated Circuit (ASIC) capabilities allow forwarding tables to be
   very large, O(E6), and allow fast (100 M lookups/second) tables to be
   built with little difficulty.

2.2.2. Traffic Engineering

   "Traffic engineering" is one of those terms that has become terribly
   overloaded. If one asks N people what traffic engineering is, one
   would get something like N! disjoint answers. Therefore, we elect
   not to require "traffic engineering", per se. Instead, we have
   endeavored to determine what the ultimate intent is when operators
   "traffic engineer" their networks and then make those capabilities an
   inherent part of the system.

2.2.3. Multicast

   The new architecture is not designed explicitly to be an inter-domain
   multicast routing architecture. However, given the notable lack of a
   viable, robust, and widely deployed inter-domain multicast routing
   architecture, the architecture should not hinder the development and
   deployment of inter-domain multicast routing without an adverse
   effect on meeting the other requirements.

   We do note however that one respected network sage [Clark91] has said
   (roughly):

      When you see a bunch of engineers standing around congratulating
      themselves for solving some particularly ugly problem in
      networking, go up to them, whisper "multicast", jump back, and
      watch the fun begin...

2.2.4. Quality of Service (QoS)

   The architecture concerns itself primarily with disseminating network
   topology information so that routers may select paths to destinations
   and build appropriate forwarding tables. Quality of Service (QoS) is
   not a part of this function and we make no requirements with respect
   to QoS.

   However, QoS is an area of great and evolving interest. It is
   reasonable to expect that in the not too distant future,
   sophisticated QoS facilities will be deployed in the Internet. Any
   new architecture and protocols should be developed with an eye toward
   these future evolutions. Extensibility mechanisms, allowing future
   QoS routing and signaling protocols to "piggy-back" on top of the
   basic routing system are desired.

   We do require the ability to assign attributes to entities and then
   do path generation and selection based on those attributes. Some may
   call this QoS.

2.2.5. IP Prefix Aggregation

   There is no specific requirement that CIDR-style (Classless Inter-
   Domain Routing) IP prefix aggregation be done by the new
   architecture. Address allocation policies, societal pressure, and
   the random growth and structure of the Internet have all conspired to
   make prefix aggregation extraordinarily difficult, if not impossible.
   This means that large numbers of prefixes will be sloshing about in
   the routing system and that forwarding tables will grow quite big.
   This is a cost that we believe must be borne.

   Nothing in this non-requirement should be interpreted as saying that
   prefix aggregation is explicitly prohibited. CIDR-style IP prefix
   aggregation might be used as a mechanism to meet other requirements,
   such as scaling.

2.2.6. Perfect Safety

   Making the system impossible to misconfigure is, we believe, not
   required. The checking, constraints, and controls necessary to
   achieve this could, we believe, prevent operators from performing
   necessary tasks in the face of unforeseen circumstances.

   However, safety is always a "good thing", and any results from
   research in this area should certainly be taken into consideration
   and, where practical, incorporated into the new routing architecture.

2.2.7. Dynamic Load Balancing

   History has shown that using the routing system to perform highly
   dynamic load balancing among multiple more-or-less-equal paths
   usually ends up causing all kinds of instability, etc., in the
   network. Thus, we do not require such a capability.

   However, this is an area that is ripe for additional research, and
   some believe that the capability will be necessary in the future.
   Thus, the architecture and protocols should be "malleable" enough to
   allow development and deployment of dynamic load-balancing
   capabilities, should we ever figure out how to do it.

2.2.8. Renumbering of Hosts and Routers

   We believe that the routing system is not required to "do
   renumbering" of hosts and routers. That's an IP issue.

   Of course, the routing and addressing architecture must be able to
   deal with renumbering when it happens.

2.2.9. Host Mobility

   In the Internet architecture, host mobility is handled on a per-host
   basis by a dedicated, Mobile-IP protocol [RFC3344]. Traffic destined
   for a mobile-host is explicitly forwarded by dedicated relay agents.
   Mobile-IP [RFC3344] adequately solves the host-mobility problem and
   we do not see a need for any additional requirements in this area.
   Of course, the new architecture must not impede or conflict with
   Mobile-IP.

2.2.10. Backward Compatibility

   For the purposes of development of the architecture, we assume that
   there is a "clean slate". Unless specified in Section 2.1, there are
   no explicit requirements that elements, concepts, or mechanisms of
   the current routing architecture be carried forward into the new one.

3. Requirements from Group B

   The following is the result of the work done by Sub-Group B of the
   IRTF RRG in 2001-2002. It was originally released under the title:
   "Future Domain Routing Requirements" and was edited by Avri Doria and
   Elwyn Davies.

3.1. Group B - Future Domain Routing Requirements

   It is generally accepted that there are major shortcomings in the
   inter-domain routing of the Internet today and that these may result
   in meltdown within an unspecified period of time. Remedying these
   shortcomings will require extensive research to tie down the exact
   failure modes that lead to these shortcomings and identify the best
   techniques to remedy the situation.

      Reviewer's Note: Even in 2001, there was a wide difference of
      opinion across the community regarding the shortcomings of inter-
      domain routing. In the years between writing and publication,
      further analysis, changes in operational practice, alterations to
      the demands made on inter-domain routing, modifications made to
      BGP, and a recognition of the difficulty of finding a replacement
      may have altered the views of some members of the community.

   Changes in the nature and quality of the services that users want
   from the Internet are difficult to provide within the current
   framework, as they impose requirements never foreseen by the original
   architects of the Internet routing system.

   The kind of radical changes that have to be accommodated are
   epitomized by the advent of IPv6 and the application of IP mechanisms
   to private commercial networks that offer specific service guarantees
   beyond the best-effort services of the public Internet. Major
   changes to the inter-domain routing system are inevitable to provide
   an efficient underpinning for the radically changed and increasingly
   commercially-based networks that rely on the IP protocol suite.

3.2. Underlying Principles

   Although inter-domain routing is seen as the major source of
   problems, the interactions with intra-domain routing, and the
   constraints that confining changes to the inter-domain arena would
   impose, mean that we should consider the whole area of routing as an
   integrated system. This is done for two reasons:

   - Requirements should not presuppose the solution. A continued
      commitment to the current definitions and split between inter-
      domain and intra-domain routing would constitute such a
      presupposition. Therefore, this part of the document uses the
      name Future Domain Routing (FDR).

   - It is necessary to understand the degree to which inter-domain and
      intra-domain routing are related within today's routing
      architecture.

We are aware that using the term "domain routing" is already fraught
with danger because of possible misinterpretation due to prior usage.
The meaning of "domain routing" will be developed implicitly
throughout the document, but a little advance explicit definition of
the word "domain" is required, as well as some explanation on the
scope of "routing".

This document uses "domain" in a very broad sense, to mean any
collection of systems or domains that come under a common authority
that determines the attributes defining, and the policies
controlling, that collection. The use of "domain" in this manner is
very similar to the concept of region that was put forth by John
Wroclawski in his Metanet model [Wroclawski95]. The idea includes
the notion that certain attributes will characterize the behavior of
the systems within a domain and that there will be borders between
domains. The idea of domain presented here does not presuppose that
two domains will have the same behavior. Nor does it presuppose
anything about the hierarchical nature of domains. Finally, it does
not place restrictions on the nature of the attributes that might be
used to determine membership in a domain. Since today's routing
domains are an example of the concept of domains in this document,
there has been no attempt to create a new term.

Current practice in routing-system design stresses the need to
separate the concerns of the control plane and the forwarding plane
in a router. This document will follow this practice, but we still
use the term "routing" as a global portmanteau to cover all aspects
of the system. Specifically, however, "routing" will be used to mean
the process of discovering, interpreting, and distributing
information about the logical and topological structure of the
network.

3.2.1. Inter-Domain and Intra-Domain

   Throughout this section, the terms "intra-domain" and "inter-domain"
   will be used. These should be understood as relative terms. In all
   cases of domains, there will be a set of network systems that are
   within that domain; routing between these systems will be termed
   "intra-domain". In some cases there will be routing between domains,
   which will be termed "inter-domain". It is possible that the routing
   exchange between two network systems can be viewed as intra-domain
   from one perspective and as inter-domain from another perspective.

3.2.2. Influences on a Changing Network

   The development of the Internet is likely to be driven by a number of
   changes that will affect the organization and the usage of the
   network, including:

   - Ongoing evolution of the commercial relationships between
      (connectivity) service providers, leading to changes in the way in
      which peering between providers is organized and the way in which
      transit traffic is routed.

   - Requirements for traffic engineering within and between domains
      including coping with multiple paths between domains.

   - Addition of a second IP addressing technique, in the form of IPv6.

   - The use of VPNs and private address space with IPv4 and IPv6.

   - Evolution of the end-to-end principle to deal with the expanded
      role of the Internet, as discussed in [Blumenthal01]: this paper
      discusses the possibility that the range of new requirements,
      especially the social and techno-political ones that are being
      placed on the future, may compromise the Internet's original
      design principles. This might cause the Internet to lose some of
      its key features, in particular, its ability to support new and
      unanticipated applications. This discussion is linked to the rise
      of new stakeholders in the Internet, especially ISPs; new
      government interests; the changing motivations of the ever growing
      user base; and the tension between the demand for trustworthy
      overall operation and the inability to trust the behavior of
      individual users.

   - Incorporation of alternative forwarding techniques such as the
      explicit routing (pipes) supplied by the MPLS [RFC3031] and GMPLS
      [RFC3471] environments.

   - Integration of additional constraints into route determination
      from interactions with other layers (e.g., Shared Risk Link Groups
      [InferenceSRLG]). This includes the concern that redundant routes
      should not fate-share, e.g., because they physically run in the
      same trench.

   - Support for alternative and multiple routing techniques that are
      better suited to delivering types of content organized in ways
      other than into IP-addressed packets.

   Philosophically, the Internet has the mission of transferring
   information from one place to another. Conceptually, this
   information is rarely organized into conveniently sized, IP-addressed
   packets, and the FDR needs to consider how the information (content)
   to be carried is identified, named, and addressed. Routing
   techniques can then be adapted to handle the expected types of
   content.

3.2.3. High-Level Goals

   This section attempts to answer two questions:

   - What are we trying to achieve in a new architecture?

   - Why should the Internet community care?

   There is a third question that needs to be answered as well, but that
   has seldom been explicitly discussed:

   - How will we know when we have succeeded?

3.2.3.1. Providing a Routing System Matched to Domain Organization

   Many of today's routing problems are caused by a routing system that
   is not well matched to the organization and policies that it is
   trying to support. Our goal is to develop a routing architecture
   where even a domain organization that is not envisioned today can be
   served by a routing architecture that matches its requirements. We
   will know when this goal is achieved when the desired policies,
   rules, and organization can be mapped into the routing system in a
   natural, consistent, and easily understood way.

3.2.3.2. Supporting a Range of Different Communication Services

   Today's routing protocols only support a single data forwarding
   service that is typically used to deliver a best-effort service in
   the public Internet. On the other hand, Diffserv for example, can
   construct a number of different bit transport services within the
   network. Using some of the per-domain behaviors (PDB)s that have
   been discussed in the IETF, it is possible to construct services such
   as Virtual Wire [DiffservVW] and Assured Rate [DiffservAR].

   Providers today offer rudimentary promises about traffic handling in
   the network, for example, delay and long-term packet loss guarantees.
   As time goes on, this becomes even more relevant. Communicating the
   service characteristics of paths in routing protocols will be
   necessary in the near future, and it will be necessary to be able to
   route packets according to their service requirements.

   Thus, a goal of this architecture is to allow adequate information
   about path service characteristics to be passed between domains and
   consequently, to allow the delivery of bit transport services other
   than the best-effort datagram connectivity service that is the
   current common denominator.

3.2.3.3. Scalable Well Beyond Current Predictable Needs

   Any proposed FDR system should scale beyond the size and performance
   we can foresee for the next ten years. The previous IDR proposal as
   implemented by BGP, has, with some massaging, held up for over ten
   years. In that time the Internet has grown far beyond the
   predictions that were implied by the original requirements.

   Unfortunately, we will only know if we have succeeded in this goal if
   the FDR system survives beyond its design lifetime without serious
   massaging. Failure will be much easier to spot!

3.2.3.4. Alternative Forwarding Mechanisms

   With the advent of circuit-based technologies (e.g., MPLS [RFC3031]
   and GMPLS [RFC3471]) managed by IP routers there are forwarding
   mechanisms other than the datagram service that need to be supported
   by the routing architecture.

   An explicit goal of this architecture is to add support for
   forwarding mechanisms other then the current hop-by-hop datagram
   forwarding service driven by globally unique IP addresses.

3.2.3.5. Separation of Topology Map from Connectivity Service

   It is envisioned that an organization can support multiple services
   within a single network. These services can, for example, be of
   different quality, of different connectivity type, or of different
   protocols (e.g., IPv4 and IPv6). For all these services, there may
   be common domain topology, even though the policies controlling the
   routing of information might differ from service to service. Thus, a
   goal with this architecture is to support separation between creation
   of a domain (or organization) topology map and service creation.

3.2.3.6. Separation between Routing and Forwarding

   The architecture of a router is composed of two main separable parts:
   control and forwarding. These components, while inter-dependent,
   perform functions that are largely independent of each other.
   Control (routing, signaling, and management) is typically done in
   software while forwarding typically is done with specialized ASICs or
   network processors.

   The nature of an IP-based network today is that control and data
   protocols share the same network and forwarding regime. This may not
   always be the case in future networks, and we should be careful to
   avoid building in this sharing as an assumption in the FDR.

A goal of this architecture is to support full separation of control
and forwarding, and to consider what additional concerns might be
properly considered separately (e.g., adjacency management).

3.2.3.7. Different Routing Paradigms in Different Areas of the Same
          Network

A number of routing paradigms have been used or researched, in
addition to the conventional shortest-path-by-hop-count paradigm that
is the current mainstay of the Internet. In particular, differences
in underlying transport networks may mean that other kinds of routing
are more relevant, and the perceived need for traffic engineering
will certainly alter the routing chosen in various domains.

Explicitly, one of these routing paradigms should be the current
routing paradigm, so that the new paradigms will inter-operate in a
backward-compatible way with today's system. This will facilitate a
migration strategy that avoids flag days.

3.2.3.8. Protection against Denial-of-Service and Other Security
          Attacks

Currently, existence of a route to a destination effectively implies
that anybody who can get a packet onto the network is entitled to use
that route. While there are limitations to this generalization, this
is a clear invitation to denial-of-service attacks. A goal of the
FDR system should be to allow traffic to be specifically linked to
whole or partial routes so that a destination or link resources can
be protected from unauthorized use.

   Editors' Note: When sections like this one and the previous ones
   on quality differentiation were written, the idea of separating
   traffic for security or quality was considered an unqualified
   advantage. Today, however, in the midst of active discussions on
   Network Neutrality, it is clear that such issues have a crucial
   policy component that also needs to be understood. These, and
   other similar issues, are open to further research.

3.2.3.9. Provable Convergence with Verifiable Policy Interaction

It has been shown both analytically, by Griffin, et al. (see
[Griffin99]), and practically (see [RFC3345]) that BGP will not
converge stably or is only meta-stable (i.e., will not re-converge in
the face of a single failure) when certain types of policy constraint
are applied to categories of network topology. The addition of
policy to the basic distance-vector algorithm invalidates the proofs
of convergence that could be applied to a policy-free implementation.

It has also been argued that global convergence may no longer be a
necessary goal and that local convergence may be all that is
required.

A goal of the FDR should be to achieve provable convergence of the
protocols used that may involve constraining the topologies and
domains subject to convergence. This will also require vetting the
policies imposed to ensure that they are compatible across domain
boundaries and result in a consistent policy set.

   Editors' Note: This requirement is very optimistic in that it
   implies that it is possible to get operators to cooperate even it
   is seen by them to be against their business practices. Though
   perhaps Utopian, this is a good goal.

3.2.3.10. Robustness Despite Errors and Failures

From time to time in the history of the Internet, there have been
occurrences where misconfigured routers have destroyed global
connectivity.

A goal of the FDR is to be more robust to configuration errors and
failures. This should probably involve ensuring that the effects of
misconfiguration and failure can be confined to some suitable
locality of the failure or misconfiguration.

3.2.3.11. Simplicity in Management

The policy work ([rap-charter02], [snmpconf-charter02], and
[policy-charter02]) that has been done at IETF provides an
architecture that standardizes and simplifies management of QoS.
This kind of simplicity is needed in a Future Domain Routing
architecture and its protocols.

A goal of this architecture is to make configuration and management
of inter-domain routing as simple as possible.

   Editors' Note: Snmpconf and rap are the hopes of the past. Today,
   configuration and policy hope is focused on netconf
   [netconf-charter].

3.2.3.12. The Legacy of RFC 1126

RFC 1126 outlined a set of requirements that were used to guide the
development of BGP. While the network has changed in the years since
1989, many of the same requirements remain. A future domain routing
solution has to support, as its base requirement, the level of
function that is available today. A detailed discussion of RFC 1126

and its requirements can be found in [RFC5773]. Those requirements, while specifically spelled out in that document, are subsumed by the requirements in this document.

3.3. High-Level User Requirements

This section considers the requirements imposed by the target audience of the FDR both in terms of organizations that might own networks that would use FDR, and the human users who will have to interact with the FDR.

3.3.1. Organizational Users

The organizations that own networks connected to the Internet have become much more diverse since RFC 1126 [RFC1126] was published. In particular, major parts of the network are now owned by commercial service provider organizations in the business of making profits from carrying data traffic.

3.3.1.1. Commercial Service Providers

The routing system must take into account the commercial service provider's need for secrecy and security, as well as allowing them to organize their business as flexibly as possible.

Service providers will often wish to conceal the details of the network from other connected networks. So far as is possible, the routing system should not require the service providers to expose more details of the topology and capability of their networks than is strictly necessary.

Many service providers will offer contracts to their customers in the form of Service Level Agreements (SLAs). The routing system must allow the providers to support these SLAs through traffic engineering and load balancing as well as multi-homing, providing the degree of resilience and robustness that is needed.

Service providers can be categorized as:

- Global Service Providers (GSPs) whose networks have a global reach. GSPs may, and usually will, wish to constrain traffic between their customers to run entirely on their networks. GSPs will interchange traffic at multiple peering points with other GSPs, and they will need extensive policy-based controls to control the interchange of traffic. Peering may be through the use of dedicated private lines between the partners or, increasingly, through Internet Exchange Points.

- National, or regional, Service Providers (NSPs) that are similar
  to GSPs but typically cover one country. NSPs may operate as a
  federation that provides similar reach to a GSP and may wish to be
  able to steer traffic preferentially to other federation members
  to achieve global reach.

- Local Internet Service Providers (ISPs) operate regionally. They
  will typically purchase transit capacity from NSPs or GSPs to
  provide global connectivity, but they may also peer with
  neighboring, and sometimes distant, ISPs.

The routing system should be sufficiently flexible to accommodate the
continually changing business relationships of the providers and the
various levels of trustworthiness that they apply to customers and
partners.

Service providers will need to be involved in accounting for Internet
usage and monitoring the traffic. They may be involved in government
action to tax the usage of the Internet, enforce social mores and
intellectual property rules, or apply surveillance to the traffic to
detect or prevent crime.

3.3.1.2. Enterprises

The leaves of the network domain graph are in many cases networks
supporting a single enterprise. Such networks cover an enormous
range of complexity. Some multi-national companies own networks that
rival the complexity and reach of a GSP, whereas many fall into the
Small Office-Home Office (SOHO) category. The routing system should
allow simple and robust configuration and operation for the SOHO
category, while effectively supporting the larger enterprise.

Enterprises are particularly likely to lack the capability to
configure and manage a complex routing system, and every effort
should be made to provide simple configuration and operation for such
networks.

Enterprises will also need to be able to change their service
provider with ease. While this is predominantly a naming and
addressing issue, the routing system must be able to support seamless
changeover; for example, if the changeover requires a change of
address prefix, the routing system must be able to cope with a period
when both sets of addresses are in use.

Enterprises will wish to be able to multi-home to one or more
providers as one possible means of enhancing the resilience of their
network.

Enterprises will also frequently need to control the trust that they
place both in workers and external connections through firewalls and
similar mid-boxes placed at their external connections.

3.3.1.3. Domestic Networks

Increasingly domestic, i.e., non-business home, networks are likely
to be 'always on' and will resemble SOHO enterprises networks with no
special requirements on the routing system.

The routing system must also continue to support dial-up users.

3.3.1.4. Internet Exchange Points

Peering of service providers, academic networks, and larger
enterprises is happening increasingly at specific Internet Exchange
Points where many networks are linked together in a relatively small
physical area. The resources of the exchange may be owned by a
trusted third party or owned jointly by the connecting networks. The
routing systems should support such exchange points without requiring
the exchange point to either operate as a superior entity with every
connected network logically inferior to it or by requiring the
exchange point to be a member of one (or all) connected networks.
The connecting networks have to delegate a certain amount of trust to
the exchange point operator.

3.3.1.5. Content Providers

Content providers are at one level a special class of enterprise, but
the desire to deliver content efficiently means that a content
provider may provide multiple replicated origin servers or caches
across a network. These may also be provided by a separate content
delivery service. The routing system should facilitate delivering
content from the most efficient location.

3.3.2. Individual Users

This section covers the most important human users of the FDR and
their expected interactions with the system.

3.3.2.1. All End Users

The routing system must continue to deliver the current global
connectivity service (i.e., any unique address to any other unique
address, subject to policy constraints) that has always been the
basic aim of the Internet.

End user applications should be able to request, or have requested on
their behalf by agents and policy mechanisms, end-to-end
communication services with QoS characteristics different from the
best-effort service that is the foundation of today's Internet. It
should be possible to request both a single service channel and a
bundle of service channels delivered as a single entity.

3.3.2.2. Network Planners

   The routing system should allow network planners to plan and
   implement a network that can be proved to be stable and will meet
   their traffic engineering requirements.

3.3.2.3. Network Operators

   The routing system should, so far as is possible, be simple to
   configure, operate and troubleshoot, behave in a predictable and
   stable fashion, and deliver appropriate statistics and events to
   allow the network to be managed and upgraded in an efficient and
   timely fashion.

3.3.2.4. Mobile End Users

   The routing system must support mobile end users. It is clear that
   mobility is becoming a predominant mode for network access.

3.4. Mandated Constraints

   While many of the requirements to which the protocol must respond are
   technical, some aren't. These mandated constraints are those that
   are determined by conditions of the world around us. Understanding
   these requirements requires an analysis of the world in which these
   systems will be deployed. The constraints include those that are
   determined by:

   - environmental factors,

   - geography,

   - political boundaries and considerations, and

   - technological factors such as the prevalence of different levels
      of technology in the developed world compared to those in the
      developing or undeveloped world.

3.4.1. The Federated Environment

   The graph of the Internet network, with routers and other control
   boxes as the nodes and communication links as the edges, is today
   partitioned administratively into a large number of disjoint domains.

   A common administration may have responsibility for one or more
   domains that may or may not be adjacent in the graph.

   Commercial and policy constraints affecting the routing system will
   typically be exercised at the boundaries of these domains where
   traffic is exchanged between the domains.

   The perceived need for commercial confidentiality will seek to
   minimize the control information transferred across these boundaries,
   leading to requirements for aggregated information, abstracted maps
   of connectivity exported from domains, and mistrust of supplied
   information.

   The perceived desire for anonymity may require the use of zero-
   knowledge security protocols to allow users to access resources
   without exposing their identity.

   The requirements should provide the ability for groups of peering
   domains to be treated as a complex domain. These complex domains
   could have a common administrative policy.

3.4.2. Working with Different Sorts of Networks

   The diverse Layer 2 networks, over which the Layer 3 routing system
   is implemented, have typically been operated totally independently
   from the Layer 3 network and often with their own routing mechanisms.
   Consideration needs to be given to the desirable degree and nature of
   interchange of information between the layers. In particular, the
   need for guaranteed robustness through diverse routing layers implies
   knowledge of the underlying networks.

   Mobile access networks may also impose extra requirements on Layer 3
   routing.
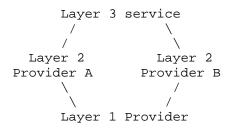
3.4.3. Delivering Resilient Service

   The routing system operates at Layer 3 in the network. To achieve
   robustness and resilience at this layer requires that, where multiple
   diverse routes are employed as part of delivering the resilience, the
   routing system at Layer 3 needs to be assured that the Layer 2 and
   lower routes are really diverse. The "diamond problem" is the

simplest form of this problem -- a Layer 3 provider attempting to
provide diversity buys Layer 2 services from two separate providers
who in turn buy Layer 1 services from the same provider:

```
                        Layer 3 service
                        /             \
                       /               \
                  Layer 2           Layer 2
                  Provider A        Provider B
                       \               /
                        \             /
                        Layer 1 Provider
```

Now, when the backhoe cuts the trench, the Layer 3 provider has no
resilience unless he had taken special steps to verify that the
trench wasn't common. The routing system should facilitate avoidance
of this kind of trap.

Some work is going on to understand the sort of problems that stem
from this requirement, such as the work on Shared Risk Link Groups
[InferenceSRLG]. Unfortunately, the full generality of the problem
requires diversity be maintained over time between an arbitrarily
large set of mutually distrustful providers. For some cases, it may
be sufficient for diversity to be checked at provisioning or route
instantiation time, but this remains a hard problem requiring
research work.

3.4.4.   When Will the New Solution Be Required?

There is a full range of opinion on this subject. An informal survey
indicates that the range varies from 2 to 6 years. And while there
are those, possibly outliers, who think there is no need for a new
routing architecture as well as those who think a new architecture
was needed years ago, the median seems to lie at around 4 years. As
in all projections of the future, this is not provable at this time.

   Editors' Note: The paragraph above was written in 2002, yet could
   be written without change in 2006. As with many technical
   predictions and schedules, the horizon has remained fixed through
   this interval.

3.5.   Assumptions

In projecting the requirements for the Future Domain Routing, a
number of assumptions have been made. The requirements set out
should be consistent with these assumptions, but there are doubtless
a number of other assumptions that are not explicitly articulated
here:

1.   The number of hosts today is somewhere in the area of 100
     million. With dial-in, NATs, and the universal deployment of
     IPv6, this is likely to become up to 500 million users (see
     [CIDR]). In a number of years, with wireless accesses and
     different appliances attaching to the Internet, we are likely to
     see a couple of billion (10^9) "users" on the Internet. The
     number of globally addressable hosts is very much dependent on
     how common NATs will be in the future.

2.   NATs, firewalls, and other middle-boxes exist, and we cannot
     assume that they will cease being a presence in the networks.

3.   The number of operators in the Internet will probably not grow
     very much, as there is a likelihood that operators will tend to
     merge. However, as Internet-connectivity expands to new
     countries, new operators will emerge and then merge again.

4.   At the beginning of 2002, there are around 12000 registered ASs.
     With current use of ASs (e.g., multi-homing) the number of ASs
     could be expected to grow to 25000 in about 10 years [Broido02].
     This is down from a previously reported growth rate of 51% per
     year [RFC3221]. Future growth rates are difficult to predict.

        Editors' Note: In the routing report table of August 2006,
        the total number of ASs present in the Internet Routing Table
        was 23000. In 4 years, this is substantial progress on the
        prediction of 25000 ASs. Also, there are significantly more
        ASs registered than are visibly active, i.e., in excess of
        42000 in mid-2006. It is possible, however, that many are
        being used internally.

5.   In contrast to the number of operators, the number of domains is
     likely to grow significantly. Today, each operator has
     different domains within an AS, but this also shows in SLAs and
     policies internal to the operator. Making this globally visible
     would create a number of domains; 10-100 times the number of
     ASs, i.e., between 100,000 and 1,000,000.

6.   With more and more capacity at the edge of the network, the IP
     network will expand. Today, there are operators with several
     thousands of routers, but this is likely to be increased. Some
     domains will probably contain tens of thousands of routers.

7.   The speed of connections in the (fixed) access will technically
     be (almost) unconstrained. However, the cost for the links will
     not be negligible so that the apparent speed will be effectively
     bounded. Within a number of years, some will have multi-gigabit
     speed in the access.

8.   At the same time, the bandwidth of wireless access still has a
     strict upper-bound. Within the foreseeable future each user
     will have only a tiny amount of resources available compared to
     fixed accesses (10 kbps to 2 Mbps for a Universal Mobile
     Telecommunications System (UMTS) with only a few achieving the
     higher figure as the bandwidth is shared between the active
     users in a cell and only small cells can actually reach this
     speed, but 11 Mbps or more for wireless LAN connections). There
     may also be requirements for effective use of bandwidth as low
     as 2.4 Kbps or lower, in some applications.

9.   Assumptions 7 and 8 taken together suggest a minimum span of
     bandwidth between 2.4 kbps to 10 Gbps.

10. The speed in the backbone has grown rapidly, and there is no
     evidence that the growth will stop in the coming years.
     Terabit-speed is likely to be the minimum backbone speed in a
     couple of years. The range of bandwidths that need to be
     represented will require consideration on how to represent the
     values in the protocols.

11. There have been discussions as to whether Moore's Law will
     continue to hold for processor speed. If Moore's Law does not
     hold, then communication circuits might play a more important
     role in the future. Also, optical routing is based on circuit
     technology, which is the main reason for taking "circuits" into
     account when designing an FDR.

12. However, the datagram model still remains the fundamental model
     for the Internet.

13. The number of peering points in the network is likely to grow,
     as multi-homing becomes important. Also, traffic will become
     more locally distributed, which will drive the demand for local
     peering.

        Editors' Note: On the other hand, peer-to-peer networking may
        shift the balance in demand for local peering.

14. The FDR will achieve the same degree of ubiquity as the current
     Internet and IP routing.

3.6. Functional Requirements

   This section includes a detailed discussion of new requirements for a
   Future Domain Routing architecture. The nth requirement carries the
   label "R(n)". As discussed in Section 3.2.3.12, a new architecture

must build upon the requirements of the past routing framework and
must not reduce the functionality of the network. A discussion and
analysis of the RFC 1126 requirements can be found in [RFC5773].

3.6.1. Topology

3.6.1.1. Routers Should Be Able to Learn and to Exploit the Domain
          Topology

   R(1) Routers must be able to acquire and hold sufficient information
         on the underlying topology of the domain to allow the
         establishment of routes on that topology.

   R(2) Routers must have the ability to control the establishment of
         routes on the underlying topology.

   R(3) Routers must be able, where appropriate, to control Sub-IP
         mechanisms to support the establishment of routes.

   The OSI Inter-Domain Routing Protocol (IDRP) [ISO10747] allowed a
   collection of topologically related domains to be replaced by an
   aggregate domain object, in a similar way to the Nimrod [Chiappa02]
   domain hierarchies. This allowed a route to be more compactly
   represented by a single collection instead of a sequence of
   individual domains.

   R(4) Routers must, where appropriate, be able to construct
         abstractions of the topology that represent an aggregation of
         the topological features of some area of the topology.

3.6.1.2. The Same Topology Information Should Support Different Path
          Selection Ideas

   The same topology information needs to provide the more flexible
   spectrum of path selection methods that we might expect to find in a
   future Internet, including distributed techniques such as hop-by-hop,
   shortest path, local optimization constraint-based, class of service,
   source address routing, and destination address routing, as well as
   the centralized, global optimization constraint-based "traffic
   engineering" type. Allowing different path selection techniques will
   produce a much more predictable and comprehensible result than the
   "clever tricks" that are currently needed to achieve the same
   results. Traffic engineering functions need to be combined.

   R(5) Routers must be capable of supporting a small number of
         different path selection algorithms.

3.6.1.3. Separation of the Routing Information Topology from the Data
          Transport Topology

   R(6)   The controlling network may be logically separate from the
          controlled network.

   The two functional "planes" may physically reside in the same nodes
   and share the same links, but this is not the only possibility, and
   other options may sometimes be necessary. An example is a pure
   circuit switch (that cannot see individual IP packets) combined with
   an external controller. Another example may be multiple links
   between two routers, where all the links are used for data forwarding
   but only one is used for carrying the routing session.

3.6.2. Distribution

3.6.2.1. Distribution Mechanisms

   R(7)   Relevant changes in the state of the network, including
          modifications to the topology and changes in the values of
          dynamic capabilities, must be distributed to every entity in
          the network that needs them, in a reliable and trusted way, at
          the earliest appropriate time after the changes have occurred.

   R(8)   Information must not be distributed outside areas where it is
          needed, or believed to be needed, for the operation of the
          routing system.

   R(9)   Information must be distributed in such a way that it minimizes
          the load on the network, consistent with the required response
          time of the network to changes.

3.6.2.2. Path Advertisement

   R(10) The router must be able to acquire and store additional static
          and dynamic information that relates to the capabilities of
          the topology and its component nodes and links and that can
          subsequently be used by path selection methods.

   The inter-domain routing system must be able to advertise more kinds
   of information than just connectivity and domain paths.

   R(11) The routing system must support service specifications, e.g.,
          the Service Level Specifications (SLSs) developed by the
          Differentiated Services working group [RFC3260].

Careful attention should be paid to ensuring that the distribution of
additional information with path advertisements remains scalable as
domains and the Internet get larger, more numerous, and more
diversified.

R(12)   The distribution mechanism used for distributing network state
        information must be scalable with respect to the expected size
        of domains and the volume and rate of change of dynamic state
        that can be expected.

The combination of R(9) and R(12) may result in a compromise between
the responsiveness of the network to change and the overhead of
distributing change notifications. Attempts to respond to very rapid
changes may damage the stability of the routing system.

Possible examples of additional capability information that might be
carried include:

- QoS information

   To allow an ISP to sell predictable end-to-end QoS service to any
   destination, the routing system should have information about the
   end-to-end QoS. This means that:

R(13)   The routing system must be able to support different paths for
        different services.

R(14)   The routing system must be able to forward traffic on the path
        appropriate for the service selected for the traffic, either
        according to an explicit marking in each packet (e.g., MPLS
        labels, Diffserv Per-Hop Behaviors (PHBs) or DSCP values) or
        implicitly (e.g., the physical or logical port on which the
        traffic arrives).

R(15)   The routing system should also be able to carry information
        about the expected (or actually, promised) characteristics of
        the entire path and the price for the service.

   (If such information is exchanged at all between network operators
   today, it is through bilateral management interfaces, and not
   through the routing protocols.)

   This would allow for the operator to optimize the choice of path
   based on a price/performance trade-off.

   In addition to providing dynamic QoS information, the system
   should be able to use static class-of-service information.

- Security information

  Security characteristics of other domains referred to by
  advertisements can allow the routing entity to make routing
  decisions based on political concerns. The information itself is
  assumed to be secure so that it can be trusted.

- Usage and cost information

  Usage and cost information can be used for billing and traffic
  engineering. In order to support cost-based routing policies for
  customers (i.e., peer ISPs), information such as "traffic on this
  link or path costs XXX per Gigabyte" needs to be advertised, so
  that the customer can choose a more or a less expensive route.

- Monitored performance

  Performance information such as delay and drop frequency can be
  carried. (This may only be suitable inside a domain because of
  trust considerations.) This should support at least the kind of
  delay-bound contractual terms that are currently being offered by
  service providers. Note that these values refer to the outcome of
  carrying bits on the path, whereas the QoS information refers to
  the proposed behavior that results in this outcome.

- Multicast information

  R(16) The routing system must provide information needed to create
         multicast distribution trees. This information must be
         provided for one-to-many distribution trees and should be
         provided for many-to-many distribution trees.

  The actual construction of distribution trees is not necessarily
  done by the routing system.

3.6.2.3. Stability of Routing Information

  R(17) The new network architecture must be stable without needing
         global convergence, i.e., convergence is a local property.

The degree to which this is possible and the definition of "local"
remain research topics. Restricting the requirement for convergence
to localities will have an effect on all of the other requirements in
this section.

   R(18)   The distribution and the rate of distribution of changes must
           not affect the stability of the routing information. For
           example, commencing redistribution of a change before the
           previous one has settled must not cause instability.

3.6.2.3.1. Avoiding Routing Oscillations

   R(19)   The routing system must minimize oscillations in route
           advertisements.

3.6.2.3.2. Providing Loop-Free Routing and Forwarding

   In line with the separation of routing and forwarding concerns:

   R(20)   The distribution of routing information must be, so far as is
           possible, loop-free.

   R(21)   The forwarding information created from this routing
           information must seek to minimize persistent loops in the
           data-forwarding paths.

   It is accepted that transient loops may occur during convergence of
   the protocol and that there are trade-offs between loop avoidance and
   global scalability.

3.6.2.3.3. Detection, Notification, and Repair of Failures

   R(22)   The routing system must provide means for detecting failures
           of node equipment or communication links.

   R(23)   The routing system should be able to coordinate failure
           indications from Layer 3 mechanisms, from nodal mechanisms
           built into the routing system, and from lower-layer mechanisms
           that propagate up to Layer 3 in order to determine the root
           cause of the failure. This will allow the routing system to
           react correctly to the failure by activating appropriate
           mitigation and repair mechanisms if required, while ensuring
           that it does not react if lower-layer repair mechanisms are
           able to repair or mitigate the fault.

   Most Layer 3 routing protocols have utilized keepalives or "hello"
   protocols as a means of detecting failures at Layer 3. The keepalive
   mechanisms are often complemented by analog mechanisms (e.g., laser-
   light detection) and hardware mechanisms (e.g., hardware/software
   watchdogs) that are built into routing nodes and communication links.
   Great care must be taken to make the best possible use of the various
   failure repair methods available while ensuring that only one repair
   mechanism at a time is allowed to repair any given fault.

Interactions between, for example, fast reroute mechanisms at Layer 3
and Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/
SDH) repair at Layer 1 are highly undesirable and are likely to cause
problems in the network.

R(24)   Where a network topology and routing system contains multiple
        fault repair mechanisms, the responses of these systems to a
        detected failure should be coordinated so that the fault is
        repaired by the most appropriate means, and no extra repairs
        are initiated.

R(25)   Where specialized packet exchange mechanisms (e.g., Layer 3
        keepalive or "hello" protocol mechanisms) are used to detect
        failures, the routing system must allow the configuration of
        the rate of transmission of these keepalives. This must
        include the capability to turn them off altogether for links
        that are deliberately broken when no real user or control
        traffic is present (e.g., ISDN links).

This will allow the operator to compromise between the speed of
failure detection and the proportion of link bandwidth dedicated to
failure detection.

3.6.3. Addressing

3.6.3.1. Support Mix of IPv4, IPv6, and Other Types of Addresses

R(26)   The routing system must support a mix of different kinds of
        addresses.

This mix will include at least IPv4 and IPv6 addresses, and
preferably various types of non-IP addresses, too. For instance,
networks like SDH/SONET and Wavelength Division Multiplexing (WDM)
may prefer to use non-IP addresses. It may also be necessary to
support multiple sets of "private" (e.g., RFC 1918) addresses when
dealing with multiple customer VPNs.

R(27)   The routing system should support the use of a single topology
        representation to generate routing and forwarding tables for
        multiple address families on the same network.

This capability would minimize the protocol overhead when exchanging
routes.

3.6.3.2. Support for Domain Renumbering/Readdressing

   R(28) If a domain is subject to address reassignment that would
          cause forwarding interruption, then the routing system should
          support readdressing (e.g., when a new prefix is given to an
          old network, and the change is known in advance) by
          maintaining routing during the changeover period [RFC2071]
          [RFC2072].

3.6.3.3. Multicast and Anycast

   R(29) The routing system must support multicast addressing, both
          within a domain and across multiple domains.

   R(30) The routing system should support anycast addressing within a
          domain. The routing system may support anycast addressing
          across domains.

   An open question is whether it is possible or useful to support
   anycast addressing between cooperating domains.

3.6.3.4. Address Scoping

   R(31) The routing system must support scoping of unicast addresses,
          and it should support scoping of multicast and anycast address
          types.

   The unicast address scoping that is being designed for IPv6 does not
   seem to cause any special problems for routing. IPv6 inter-domain
   routing handles only IPv6 global addresses, while intra-domain
   routing also needs to be aware of the scope of private addresses.

      Editors' Note: the original reference was to site-local addresses,
      but these have been deprecated by the IETF. Link-local addresses
      are never routed at all.

   More study may be needed to identify the requirements and solutions
   for scoping in a more general sense and for scoping of multicast and
   anycast addresses.

3.6.3.5. Mobility Support

   R(32) The routing system must support system mobility. The term
          "system" includes anything from an end system to an entire
          domain.

We observe that the existing solutions based on renumbering and/or
tunneling are designed to work with the current routing, so they do
not add any new requirements to future routing. But the requirement
is general, and future solutions may not be restricted to the ones we
have today.

3.6.4. Statistics Support

R(33) Both the routing and forwarding parts of the routing system
       must maintain statistical information about the performance of
       their functions.

3.6.5. Management Requirements

While the tools of management are outside the scope of routing, the
mechanisms to support the routing architecture and protocols are
within scope.

R(34) Mechanisms to support Operational, Administrative, and
       Management control of the routing architecture and protocols
       must be designed into the original fabric of the architecture.

3.6.5.1. Simple Policy Management

The basic aims of this specification are:

- to require less manual configuration than today, and

- to satisfy the requirements for both easy handling and maximum
   control. That is:

   - All the information should be available,

   - but should not be visible except for when necessary.

   - Policies themselves should be advertised and not only the
      result of policy, and

   - policy-conflict resolution must be provided.

R(35) The routing system must provide management of the system by
       means of policies. For example, policies that can be
       expressed in terms of the business and services implemented on
       the network and reflect the operation of the network in terms
       of the services affected.

Editors' Note: This requirement is optimistic in that it
implies that it is possible to get operators to
cooperate even if it is seen by them to be against their
business practices.

R(36) The distribution of policies must be amenable to scoping to
protect proprietary policies that are not relevant beyond the
local set of domains.

3.6.5.2. Startup and Maintenance of Routers

A major problem in today's networks is the need to perform initial
configuration on routers from a local interface before a remote
management system can take over. It is not clear that this imposes
any requirements on the routing architecture beyond what is needed
for a ZeroConf host.

Similarly, maintenance and upgrade of routers can cause major
disruptions to the network routing because the routing system and
management of routers is not organized to minimize such disruption.
Some improvements have been made, such as graceful restart mechanisms
in protocols, but more needs to be done.

R(37) The routing system and routers should provide mechanisms that
minimize the disruption to the network caused by maintenance
and upgrades of software and hardware. This requirement
recognizes that some of the capabilities needed are outside
the scope of the routing architecture (e.g., minimum impact
software upgrade).

3.6.6. Provability

R(38) The routing system and its component protocols must be
demonstrated to be locally convergent under the permitted
range of parameter settings and policy options that the
operator(s) can select.

There are various methods for demonstration and proof that include,
but are not limited to: mathematical proof, heuristic, and pattern
recognition. No requirement is made on the method used for
demonstrating local convergence properties.

R(39) Routing protocols employed by the routing system and the
overall routing system should be resistant to bad routing
policy decisions made by operators.

Tools are needed to check compatibility of routing policies. While
these tools are not part of the routing architecture, the mechanisms
to support such tools are.

Routing policies are compatible if their interaction does not cause
instability. A domain or group of domains in a system is defined as
being convergent, either locally or globally, if and only if, after
an exchange of routing information, routing tables reach a stable
state that does not change until the routing policies or the topology
changes again.

To achieve the above-mentioned goals:

R(40)   The routing system must provide a mechanism to publish and
        communicate policies so that operational coordination and
        fault isolation are possible.

Tools are required that verify the stability characteristics of the
routing system in specified parts of the Internet. The tools should
be efficient (fast) and have a broad scope of operation (check large
portions of Internet). While these tools are not part of the
architecture, developing them is in the interest of the architecture
and should be defined as a Routing Research Group activity while
research on the architecture is in progress.

Tools analyzing routing policies can be applied statically or
(preferably) dynamically. A dynamic solution requires tools that can
be used for run time checking for oscillations that arise from policy
conflicts. Research is needed to find an efficient solution to the
dynamic checking of oscillations.

3.6.7. Traffic Engineering

The ability to do traffic engineering and to get the feedback from
the network to enable traffic engineering should be included in the
future domain architecture. Though traffic engineering has many
definitions, it is, at base, another alternative or extension for the
path selection mechanisms of the routing system. No fundamental
changes to the requirements are needed, but the iterative processes
involved in traffic engineering may require some additional
capabilities and state in the network.

Traffic engineering typically involves a combination of off-line
network planning and administrative control functions in which the
expected and measured traffic flows are examined, resulting in
changes to static configurations and policies in the routing system.

During operations, these configurations control the actual flow of
traffic and affect the dynamic path selection mechanisms; the results
are measured and fed back into further rounds of network planning.

3.6.7.1. Support for, and Provision of, Traffic Engineering Tools

At present, there is an almost total lack of effective traffic
engineering tools, whether in real time for network control or off-
line for network planning. The routing system should encourage the
provision of such tools.

R(41)   The routing system must generate statistical and accounting
        information in such a way that traffic engineering and network
        planning tools can be used in both real-time and off-line
        planning and management.

3.6.7.2. Support of Multiple Parallel Paths

R(42)   The routing system must support the controlled distribution
        over multiple links or paths of traffic toward the same
        destination. This applies to domains with two or more
        connections to the same neighbor domain, and to domains with
        connections to more than one neighbor domain. The paths need
        not have the same metric.

R(43)   The routing system must support forwarding over multiple
        parallel paths when available. This support should extend to
        cases where the offered traffic is known to exceed the
        available capacity of a single link, and to the cases where
        load is to be shared over paths for cost or resiliency
        reasons.

R(44)   Where traffic is forwarded over multiple parallel paths, the
        routing system must, so far as is possible, avoid the
        reordering of packets in individual micro-flows.

R(45)   The routing system must have mechanisms to allow the traffic
        to be reallocated back onto a single path when multiple paths
        are not needed.

3.6.7.3. Peering Support

R(46)   The routing system must support peer-level connectivity as
        well as hierarchical connections between domains.

The network is becoming increasingly complex, with private peering
arrangements set up between providers at every level of the hierarchy
of service providers and even by certain large enterprises, in the
form of dedicated extranets.

R(47) The routing system must facilitate traffic engineering of peer
       routes so that traffic can be readily constrained to travel as
       the network operators desire, allowing optimal use of the
       available connectivity.

3.6.8. Support for Middleboxes

One of our assumptions is that NATs and other middle-boxes such as
firewalls, web proxies, and address family translators (e.g., IPv4 to
IPv6) are here to stay.

R(48) The routing system should work in conjunction with middle-
       boxes, e.g., NAT, to aid in bi-directional connectivity
       without compromising the additional opacity and privacy that
       the middle-boxes offer.

This problem is closely analogous to the abstraction problem, which
is already under discussion for the interchange of routing
information between domains.

3.7. Performance Requirements

Over the past several years, the performance of the routing system
has frequently been discussed. The requirements that derive from
those discussions are listed below. The specific values for these
performance requirements are left for further discussion.

R(49) The routing system must support domains of at least N systems.
       A system is taken to mean either an individual router or a
       domain.

R(50) Local convergence should occur within T units of time.

R(51) The routing system must be measurably reliable. The measure
       of reliability remains a research question.

R(52) The routing system must be locally stable to a measured
       degree. The degree of measurability remains a research issue.

R(53) The routing system must be globally stable to a measured
       degree. The degree of measurability remains a research issue.

R(54) The routing system should scale to an indefinitely large
       number of domains.

There has been very little data or statistical evidence for many of
the performance claims made in the past. In recent years, several
efforts have been initiated to gather data and do the analyses
required to make scientific assessments of performance issues and
requirements. In order to complete this section of the requirements
analysis, the data and analyses from these studies needs to be
gathered and collated into this document. This work has been started
but has yet to be completed.

    Editors' Note: This work was never completed due to corporate
    reorganizations.

3.8. Backward Compatibility (Cutover) and Maintainability

This area poses a dilemma. On one hand, it is an absolute
requirement that:

R(55) The introduction of the routing system must not require any
       flag days.

R(56) The network currently in place must continue to run at least
       as well as it does now while the new network is being
       installed around it.

However, at the same time, it is also an requirement that:

R(57) The new architecture must not be limited by the restrictions
       that plague today's network.

It has to be admitted that R(57) is not a well-defined requirement,
because we have not fully articulated what the restrictions might be.
Some of these restrictions can be derived by reading the discussions
for the positive requirements above. It would be a useful exercise
to explicitly list all the restrictions and irritations with which we
wish to do away. Further, it would be useful to determine if these
restrictions can currently be removed at a reasonable cost or whether
we are actually condemned to live with them.

Those restrictions cannot be allowed to become permanent baggage on
the new architecture. If they do, the effort to create a new system
will come to naught. It may, however, be necessary to live with some
of them temporarily for practical reasons while providing an
architecture that will eventually allow them to be removed. The last
three requirements have significance not only for the transition

strategy but also for the architecture itself. They imply that it
must be possible for an internet such as today's BGP-controlled
network, or one of its ASs, to exist as a domain within the new FDR.

3.9. Security Requirements

As previously discussed, one of the major changes that has overtaken
the Internet since its inception is the erosion of trust between end
users making use of the net, between those users and the suppliers of
services, and between the multiplicity of providers. Hence,
security, in all its aspects, will be much more important in the FDR.

It must be possible to secure the routing communication.

R(58)   The communicating entities must be able to identify who sent
        and who received the information (authentication).

R(59)   The communicating entities must be able to verify that the
        information has not been changed on the way (integrity).

Security is more important in inter-domain routing where the operator
has no control over the other domains, than in intra-domain routing
where all the links and the nodes are under the administration of the
operator and can be expected to share a trust relationship. This
property of intra-domain trust, however, should not be taken for
granted:

R(60)   Routing communications must be secured by default, but an
        operator must have the option to relax this requirement within
        a domain where analysis indicates that other means (such as
        physical security) provide an acceptable alternative.

R(61)   The routing communication mechanism must be robust against
        denial-of-service attacks.

R(62)   It should be possible to verify that the originator of the
        information was authorized to generate the information.

Further considerations that may impose further requirements include:

- whether no one else but the intended recipient is able to access
   (privacy) or understand (confidentiality) the information,

- whether it is possible to verify that all the information has been
   received and that the two parties agree on what was sent
   (validation and non-repudiation),

- whether there is a need to separate security of routing from
  security of forwarding, and

- whether traffic flow security is needed (i.e., whether there is
  value in concealing who can connect to whom, and what volumes of
  data are exchanged).

Securing the BGP session, as done today, only secures the exchange of
messages from the peering domain, not the content of the information.
In other words, we can confirm that the information we got is what
our neighbor really sent us, but we do not know whether or not this
information (that originated in some remote domain) is true.

A decision has to be made on whether to rely on chains of trust (we
trust our peers who trust their peers who..), or whether we also need
authentication and integrity of the information end-to-end. This
information includes both routes and addresses. There has been
interest in having digital signatures on originated routes as well as
countersignatures by address authorities to confirm that the
originator has authority to advertise the prefix. Even understanding
who can confirm the authority is non-trivial, as it might be the
provider who delegated the prefix (with a whole chain of authority
back to ICANN) or it may be an address registry. Where a prefix
delegated by a provider is being advertised through another provider
as in multi-homing, both may have to be involved to confirm that the
prefix may be advertised through the provider who doesn't have any
interest in the prefix!

R(63) The routing system must cooperate with the security policies
       of middle-boxes whenever possible.

This is likely to involve further requirements for abstraction of
information. For example, a firewall that is seeking to minimize
interchange of information that could lead to a security breach. The
effect of such changes on the end-to-end principle should be
carefully considered as discussed in [Blumenthal01].

R(64) The routing system must be capable of complying with local
       legal requirements for interception of communication.

3.10. Debatable Issues

This section covers issues that need to be considered and resolved in
deciding on a Future Domain Routing architecture. While they can't
be described as requirements, they do affect the types of solution
that are acceptable. The discussions included below are very open-
ended.

3.10.1. Network Modeling

   The mathematical model that underlies today's routing system uses a
   graph representation of the network. Hosts, routers, and other
   processing boxes are represented by nodes and communications links by
   arcs. This is a topological model in that routing does not need to
   directly model the physical length of the links or the position of
   the nodes; the model can be transformed to provide a convenient
   picture of the network by adjusting the lengths of the arcs and the
   layout of the nodes. The connectivity is preserved and routing is
   unaffected by this transformation.

   The routing algorithms in traditional routing protocols utilize a
   small number of results from graph theory. It is only recently that
   additional results have been employed to support constraint-based
   routing for traffic engineering.

   The naturalness of this network model and the "fit" of the graph
   theoretical methods may have tended to blind us to alternative
   representations and inhibited us from seeking alternative strands of
   theoretical thinking that might provide improved results.

   We should not allow this habitual behavior to stop us from looking
   for alternative representations and algorithms; topological
   revolutions are possible and allowed, at least in theory.

3.10.2. System Modeling

   The assumption that object modeling of a system is an essential first
   step to creating a new system is still novel in this context.
   Frequently, the object modeling effort becomes an end in itself and
   does not lead to system creation. But there is a balance, and a lot
   that can be discovered in an ongoing effort to model a system such as
   the Future Domain Routing system. It is recommended that this
   process be included in the requirements. It should not, however, be
   a gating event to all other work.

   Some of the most important realizations will occur during the process
   of determining the following:

   - Object classification

   - Relationships and containment

   - Roles and Rules

3.10.3. One, Two, or Many Protocols

   There has been a lot of discussion of whether the FDR protocol
   solution should consist of one (probably new) protocol, two (intra-
   and inter-domain) protocols, or many protocols. While it might be
   best to have one protocol that handles all situations, this seems
   improbable. On the other hand, maintaining the "strict" division
   evident in the network today between the IGP and EGP may be too
   restrictive an approach. Given this, and the fact that there are
   already many routing protocols in use, the only possible answer seems
   to be that the architecture should support many protocols. It
   remains an open issue, one for the solution, to determine if a new
   protocol needs to be designed in order to support the highest goals
   of this architecture. The expectation is that a new protocol will be
   needed.

3.10.4. Class of Protocol

   If a new protocol is required to support the FDR architecture, the
   question remains open as to what kind of protocol this ought to be.
   It is our expectation that a map distribution protocol will be
   required to augment the current path-vector protocol and shortest
   path first protocols.

3.10.5. Map Abstraction

   Assuming that a map distribution protocol, as defined in [RFC1992] is
   required, what are the requirements on this protocol? If every
   detail is advertised throughout the Internet, there will be a lot of
   information. Scalable solutions require abstraction.

   - If we summarize too much, some information will be lost on the
      way.

   - If we summarize too little, then more information than required is
      available, contributing to scaling limitations.

   - One can allow more summarization, if there also is a mechanism to
      query for more details within policy limits.

   - The basic requirement is not that the information shall be
      advertised, but rather that the information shall be available to
      those who need it. We should not presuppose a solution where
      advertising is the only possible mechanism.

3.10.6. Clear Identification for All Entities

   As in all other fields, the words used to refer to concepts and to
   describe operations about routing are important. Rather than
   describe concepts using terms that are inaccurate or rarely used in
   the real world of networking, it is necessary to make an effort to
   use the correct words. Many networking terms are used casually, and
   the result is a partial or incorrect understanding of the underlying
   concept. Entities such as nodes, interfaces, subnetworks, tunnels,
   and the grouping concepts such as ASs, domains, areas, and regions,
   need to be clearly identified and defined to avoid confusion.

   There is also a need to separate identifiers (what or who) from
   locators (where) from routes (how to reach).

      Editors' Note: Work was undertaken in the shim6 working group of
      the IETF on this sort of separation. This work needs to be taken
      into account in any new routing architecture.

3.10.7. Robustness and Redundancy

   The routing association between two domains should survive even if
   some individual connection between two routers goes down.

   The "session" should operate between logical "routing entities" on
   each domain side, and not necessarily be bound to individual routers
   or addresses. Such a logical entity can be physically distributed
   over multiple network elements. Or, it can reside in a single
   router, which would default to the current situation.

3.10.8. Hierarchy

   A more flexible hierarchy with more levels and recursive groupings in
   both upward and downward directions allows more structured routing.
   The consequence is that no single level will get too big for routers
   to handle.

   On the other hand, it appears that the real-world Internet is
   becoming less hierarchical, so that it will be increasingly difficult
   to use hierarchy to control scaling.

   Note that groupings can look different depending on which aspect we
   use to define them. A Diffserv area, an MPLS domain, a trusted
   domain, a QoS area, a multicast domain, etc., do not always coincide;
   nor are they strict hierarchical subsets of each other. The basic
   distinction at each level is "this grouping versus everything
   outside".

3.10.9. Control Theory

   Is it possible to apply a control theory framework to analyze the
   stability of the control system of the whole network domain, with
   regard to, e.g., convergence speed and the frequency response, and
   then use the results from that analysis to set the timers and other
   protocol parameters?

   Control theory could also play a part in QoS routing, by modifying
   current link-state protocols with link costs dependent on load and
   feedback. Control theory is often used to increase the stability of
   dynamic systems.

   It might be possible to construct a new, totally dynamic routing
   protocol solely on a control theoretic basis, as opposed to the
   current protocols that are based in graph theory and static in
   nature.

3.10.10. Byzantium

   Is solving the Byzantine Generals problem a requirement? This is the
   problem of reaching a consensus among distributed units if some of
   them give misleading answers. The current intra-domain routing
   system is, at one level, totally intolerant of misleading
   information. However, the effect of different sorts of misleading or
   incorrect information has vastly varying results, from total collapse
   to purely local disconnection of a single domain. This sort of
   behavior is not very desirable.

   There are, possibly, other network robustness issues that must be
   researched and resolved.

3.10.11. VPN Support

   Today, BGP is also used for VPNs, for example, as described in RFC
   4364 [RFC4364].

   Internet routing and VPN routing have different purposes and most
   often exchange different information between different devices. Most
   Internet routers do not need to know VPN-specific information. The
   concepts should be clearly separated.

   But when it comes to the mechanisms, VPN routing can share the same
   protocol as ordinary Internet routing; it can use a separate instance
   of the same protocol or it can use a different protocol. All
   variants are possible and have their own merits. These requirements
   are silent on this issue.

3.10.12. End-to-End Reliability

   The existing Internet architecture neither requires nor provides end-
   to-end reliability of control information dissemination. There is,
   however, already a requirement for end-to-end reliability of control
   information distribution, i.e., the ends of the VPN established need
   to have an acknowledgment of the success in setting up the VPN.
   While it is not necessarily the function of a routing architecture to
   provide end-to-end reliability for this kind of purpose, we must be
   clear that end-to-end reliability becomes a requirement if the
   network has to support such reliable control signaling. There may be
   other requirements that derive from requiring the FDR to support
   reliable control signaling.

3.10.13. End-to-End Transparency

   The introduction of private addressing schemes, Network Address
   Translators, and firewalls has significantly reduced the end-to-end
   transparency of the network. In many cases, the network is also no
   longer symmetric, so that communication between two addresses is
   possible if the communication session originates from one end but not
   from the other. This impedes the deployment of new peer-to-peer
   services and some "push" services where the server in a client-
   server arrangement originates the communication session. Whether a
   new routing system either can or should seek to restore this
   transparency is an open issue.

   A related issue is the extent to which end-user applications should
   seek to control the routing of communications to the rest of the
   network.

4. Security Considerations

   We address security issues in the individual requirements. We do
   require that the architecture and protocols developed against this
   set of requirements be "secure". Discussion of specific security
   issues can be found in the following sections:

   o Group A: Routing System Security - Section 2.1.9

   o Group A: End Host Security - Section 2.1.10

   o Group A: Routing Information Policies - Section 2.1.11.1

   o Group A: Abstraction - Section 2.1.16

   o Group A: Robustness - Section 2.1.18

   o Group B: Protection against Denial-of-Service and Other Security
      Attacks - Section 3.2.3.8

   o Group B: Commercial Service Providers - Section 3.3.1.1

   o Group B: The Federated Environment - Section 3.4.1

   o Group B: Path Advertisement - Section 3.6.2.2

   o Group B: Security Requirements - Section 3.9

5. IANA Considerations

   This document is a set of requirements from which a new routing and
   addressing architecture may be developed. From that architecture, a
   new protocol, or set of protocols, may be developed.

   While this note poses no new tasks for IANA, the architecture and
   protocols developed from this document probably will have issues to
   be dealt with by IANA.

6. Acknowledgments

   This document is the combined effort of two groups in the IRTF.
   Group A, which was formed by the IRTF Routing Research chairs, and
   Group B, which was self-formed and later was folded into the IRTF
   Routing Research Group. Each group has it own set of
   acknowledgments.

   Group A Acknowledgments

      This originated in the IRTF Routing Research Group's sub-group on
      Inter-domain routing requirements. The members of the group were:

         Abha Ahuja   Danny McPherson
         J. Noel Chiappa   David Meyer
         Sean Doran   Mike O'Dell
         JJ Garcia-Luna-Aceves   Andrew Partan
         Susan Hares   Radia Perlman
         Geoff Huston   Yakov Rehkter
         Frank Kastenholz   John Scudder
         Dave Katz   Curtis Villamizar
         Tony Li   Dave Ward

      We also appreciate the comments and review received from Ran
      Atkinson, Howard Berkowitz, Randy Bush, Avri Doria, Jeffery Haas,
      Dmitri Krioukov, Russ White, and Alex Zinin. Special thanks to
      Yakov Rehkter for contributing text and to Noel Chiappa.

7. Informative References

   [Blumenthal01]
             Blumenthal, M. and D. Clark, "Rethinking the design of the
             Internet: The end to end arguments vs. the brave new
             world", May 2001,
             <http://dspace.mit.edu/handle/1721.1/1519>.

   [Broido02]
             Broido, A., Nemeth, E., Claffy, K., and C. Elves,
             "Internet Expansion, Refinement and Churn", February 2002.

   [CIDR]   Telcordia Technologies, "CIDR Report",
             <http://www.cidr-report.org/>.

   [Chiappa02]
             Chiappa, N., "A New IP Routing and Addressing
             Architecture", July 1991,
             <http://ana-3.lcs.mit.edu/~jnc/nimrod/overview.txt>.

   [Clark91] Clark, D., "Quote reportedly from IETF Plenary
             discussion", 1991.

   [DiffservAR]
             Seddigh, N., Nandy, B., and J. Heinanen, "An Assured Rate
             Per-Domain Behaviour for Differentiated Services", Work
             in Progress, July 2001.

   [DiffservVW]
             Jacobson, V., Nichols, K., and K. Poduri, "The 'Virtual
             Wire' Per-Domain Behavior", Work in Progress, July 2000.

   [Griffin99]
             Griffin, T. and G. Wilfong, "An Analysis of BGP
             Convergence Properties", SIGCOMM 1999.

   [ISO10747]
             ISO/IEC, "Protocol for Exchange of Inter-Domain Routeing
             Information among Intermediate Systems to Support
             Forwarding of ISO 8473 PDUs", International Standard
             10747 ISO/IEC JTC 1, Switzerland, 1993.

   [InferenceSRLG]
             Papadimitriou, D., Poppe, F., J. Jones, J., S.
             Venkatachalam, S., S. Dharanikota, S., Jain, R., Hartani,
             R., and D. Griffith, "Inference of Shared Risk Link
             Groups", Work in Progress, November 2001.

   [ODell01] O'Dell, M., "Private Communication", 2001.

   [RFC1126] Little, M., "Goals and functional requirements for inter-
              autonomous system routing", RFC 1126, October 1989.

   [RFC1726] Partridge, C. and F. Kastenholz, "Technical Criteria for
              Choosing IP The Next Generation (IPng)", RFC 1726,
              Dec 1994.

   [RFC1992] Castineyra, I., Chiappa, N., and M. Steenstrup, "The
              Nimrod Routing Architecture", RFC 1992, August 1996.

   [RFC2071] Ferguson, P. and H. Berkowitz, "Network Renumbering
              Overview: Why would I want it and what is it anyway?",
              RFC 2071, January 1997.

   [RFC2072] Berkowitz, H., "Router Renumbering Guide", RFC 2072,
              January 1997.

   [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol
              Label Switching Architecture", RFC 3031, January 2001.

   [RFC3221] Huston, G., "Commentary on Inter-Domain Routing in the
              Internet", RFC 3221, December 2001.

   [RFC3260] Grossman, D., "New Terminology and Clarifications for
              Diffserv", RFC 3260, April 2002.

   [RFC3344] Perkins, C., "IP Mobility Support.", RFC 3344,
              August 2002.

   [RFC3345] McPherson, D., Gill, V., Walton, D., and A. Retana,
              "Border Gateway Protocol (BGP) Persistent Route
              Oscillation Condition", RFC 3345, August 2002.

   [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching
              (GMPLS) Signaling Functional Description", RFC 3471,
              January 2003.

   [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P.
              Thubert, "Network Mobility (NEMO) Basic Support Protocol",
              RFC 3963, January 2005.

   [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private
              Networks (VPNs)", RFC 4364, February 2006.

   [RFC5773] Davies, E. and A. Doria, "Analysis of Inter-Domain Routing
              Requirements and History", RFC 5773, February 2010.

   [Wroclawski95]
            Wroclowski, J., "The Metanet White Paper - Workshop on
            Research Directions for the Next Generation Internet",
            1995.

   [netconf-charter]
            Internet Engineering Task Force, "IETF Network
            Configuration working group", 2005,
            <http://www.ietf.org/html.charters/netconf-charter.html>.

   [policy-charter02]
            Internet Engineering Task Force, "IETF Policy working
            group", 2002, <http://www.ietf.org/html.charters/OLD/
            policy-charter.html>.

   [rap-charter02]
            Internet Engineering Task Force, "IETF Resource Allocation
            Protocol working group", 2002,
            <http://www.ietf.org/html.charters/OLD/rap-charter.html>.

   [snmpconf-charter02]
            Internet Engineering Task Force, "IETF Configuration
            management with SNMP working group", 2002, <http://
            www.ietf.org/html.charters/OLD/snmpconf-charter.html>.

Authors' Addresses

    Avri Doria
    LTU
    Lulea 971 87
    Sweden

    Phone: +46 73 277 1788
    EMail: avri@ltu.se

    Elwyn B. Davies
    Folly Consulting
    Soham, Cambs
    UK

    Phone: +44 7889 488 335
    EMail: elwynd@dial.pipex.com

    Frank Kastenholz
    BBN Technologies
    10 Moulton St.
    Cambridge, MA 02183
    USA

    Phone: +1 617 873 8047
    EMail: frank@bbn.com