TSVWG   J. Touch
Internet Draft   USC/ISI
Intended status: Best Current Practice   January 23, 2015
Expires: July 2015

                Recommendations for Transport Port Number Uses
                    draft-ietf-tsvwg-port-use-07.txt

Status of this Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF), its areas, and its working groups. Note that
   other groups may also distribute working documents as Internet-
   Drafts.

   Internet-Drafts are draft documents valid for a maximum of six
   months and may be updated, replaced, or obsoleted by other documents
   at any time. It is inappropriate to use Internet-Drafts as
   reference material or to cite them other than as "work in progress."

   The list of current Internet-Drafts can be accessed at
   http://www.ietf.org/ietf/1id-abstracts.txt

   The list of Internet-Draft Shadow Directories can be accessed at
   http://www.ietf.org/shadow.html

   This Internet-Draft will expire on July 23, 2015.

Abstract

   This document provides recommendations to application and service
   designers on how to use the transport protocol port number space. It
   complements (but does not update) RFC6335, which focuses on IANA
   process.

Table of Contents

1. Introduction

   This document provides information and advice to application and
   service designers on the use of transport port numbers. It provides
   a detailed historical background of the evolution of transport port
   numbers and their multiple meanings. It also provides specific
   recommendations to designers on how to use assigned port numbers.
   Note that this document provides information to potential port
   number applicants that complements the IANA process described in
   BCP165 [RFC6335], but it does not update that document.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

In this document, the characters ">>" preceding an indented line(s) indicates a compliance requirement statement using the key words listed above. This convention aids reviewers in quickly identifying or finding the explicit compliance requirements of this RFC.

3. History

The term 'port' was first used in [RFC33] to indicate a simplex communication path from an individual process and originally applied to only the Network Control Program (NCP) connection-oriented protocol. At a meeting described in [RFC37], an idea was presented to decouple connections between processes and links that they use as paths, and thus to include numeric source and destination socket identifiers in packets. [RFC38] provides further detail, describing how processes might have more than one of these paths and that more than one path may be active at a time. As a result, there was the need to add a process identifier to the header of each message so that incoming messages could be demultiplexed to the appropriate process. [RFC38] further suggested that 32 bit numbers would be used for these identifiers. [RFC48] discusses the current notion of listening on a specific port number, but does not discuss the issue of port number determination. [RFC61] notes that the challenge of knowing the appropriate port numbers is "left to the processes" in general, but introduces the concept of a "well-known" port number for common services.

[RFC76] proposed a "telephone book" by which an index would allow port numbers to be used by name, but still assumed that both source and destination port numbers are fixed by such a system. [RFC333] proposed that a port number pair, rather than an individual port number, would be used on both sides of the connection for demultiplexing messages. This is the final view in [RFC793] (and its predecessors, including [IEN112]), and brings us to their current meaning. [RFC739] introduced the notion of generic reserved port numbers for groups of protocols, such as "any private RJE server" [RFC739]. Although the overall range of such port numbers was (and

remains) 16 bits, only the first 256 (high 8 bits cleared) in the range were considered assigned.

[RFC758] is the first to describe port numbers as being used for TCP (previous RFCs all refer to only NCP). It includes a list of such well-known port numbers, as well as describing ranges used for different purposes:

```
   Binary    Octal

   ------------------------------------------------------------

   0-63      0-77       Network Wide Standard Function

   64-127    100-177    Hosts Specific Functions

   128-223   200-337    Reserved for Future Use

   224-255   340-377    Any Experimental Function
```

In [RFC820] those range meanings disappeared, and a single list of number assignments is presented. This is also the first time that port numbers are described as applying to a connectionless transport (UDP) rather than only connection-oriented transports.

By [RFC900] the ranges appeared as decimal numbers rather than the octal ranges used previously. [RFC1340] increased this range from 0..255 to 0..1023, and began to list TCP and UDP port number assignments individually (although the assumption was that once assigned a port number applies to all transport protocols, including TCP, UDP, recently SCTP and DCCP, as well as ISO-TP4 for a brief period in the early 1990s). [RFC1340] also established the Registered range of 1024-59151, though it notes that it is not controlled by the IANA at that point. The list provided by [RFC1700] in 1994 remained the standard until it was declared replaced by an on-line version, as of [RFC3232] in 2002.

4. Current Port Number Use

RFC6335 indicates three ranges of port number assignments:

```
Binary         Hex

-----------------------------------------------------------

0-1023         0x0000-0x03FF   System (also Well-Known)

1024-49151     0x0400-0xBFFF   User (also Registered)

49152-65535    0xC000-0xFFFF   Dynamic (also Private)
```

System (also Well-Known) encompasses the range 0..1023. On some
systems, use of these port numbers requires privileged access, e.g.,
that the process run as 'root' (i.e., as a privileged user), which
is why these are referred to as System port numbers. The port
numbers from 1024..49151 denotes non-privileged services, known as
User (also Registered), because these port numbers do not run with
special privileges. Dynamic (also Private) port numbers are not
assigned.

Both System and User port numbers are assigned through IANA, so both
are sometimes called 'registered port numbers'. As a result, the
term 'registered' is ambiguous, referring either to the entire range
0-49151 or to the User port numbers. Complicating matters further,
System port numbers do not always require special (i.e., 'root')
privilege. For clarity, the remainder of this document refers to the
port number ranges as System, User, and Dynamic, to be consistent
with IANA process [RFC6335].

5. **What is a Port Number?**

A port number is a 16-bit number used for two distinct purposes:

   o Demultiplexing transport endpoint associations within an end
      host

   o Identifying a service

The first purpose requires that each transport endpoint association
(e.g., TCP connection or UDP pairwise association) using a given
transport between a given pair of IP addresses use a different pair
of port numbers, but does not require either coordination or
registration of port number use. It is the second purpose that
drives the need for a common registry.

Consider a user wanting to run a web server. That service could run
on any port number, provided that all clients knew what port number

to use to access that service at that host. Such information can be
distributed out-of-band, e.g., in the URI:

   http://www.example.com:51509/

Ultimately, the correlation of a service with a port number is an
agreement between just the two endpoints of the association. A web
server can run on port number 53, which might appear as DNS traffic
to others but will connect to browsers that know to use port number
53 rather than 80.

As a concept, a service is the combination of ISO Layers 5-7 that
represents an application protocol capability. For example www (port
number 80) is a service that uses HTTP as an application protocol
and provides access to a web server [RFC7230]. However, it is
possible to use HTTP for other purposes, such as command and
control. This is why some current service names (HTTP, e.g.) are a
bit overloaded - they describe not only the application protocol,
but a particular service.

IANA assigns port numbers so that Internet endpoints do not need
pairwise, explicit coordination of the meaning of their port
numbers. This is the primary reason for requesting assigned port
numbers with IANA - to have a common agreement between all endpoints
on the Internet as to the default meaning of a port number.

Port numbers are sometimes used by intermediate devices on a network
path, either to monitor available services, to monitor traffic
(e.g., to indicate the data contents), or to intercept traffic (to
block, proxy, relay, aggregate, or otherwise process it). In each
case, the intermediate device interprets traffic based on the port
number. It is important to recognize that any interpretation of port
numbers - except at the endpoints - may be incorrect, because port
numbers are meaningful only at the endpoints. Further, port numbers
may not be visible to these intermediate devices, such as when the
transport protocol is encrypted (as in network- or link-layer
tunnels), or when a packet is fragmented (in which case only the
first fragment has the port number information). Such port number
invisibility may interfere with these in-network port number-based
capabilities.

Port numbers can also be useful for other purposes. Assigned port
numbers can simplify end system configuration, so that individual
installations do not need to coordinate their use of arbitrary port
numbers. Such assignments can also simplify firewall management, so
that a single, fixed firewall configuration can either permit or
deny a service.

It is useful to differentiate a port number from a service name. The
former is a numeric value that is used directly in transport
protocol headers as a demultiplexing and service identifier. The
latter is primarily a user convenience, where the default map
between the two is considered static and resolved using a cached
index. This document focuses on the former because it is the
fundamental network resource. Dynamic maps between the two, i.e.,
using DNS SRV records, are discussed further in Section 7.1.

6. Conservation

Assigned port numbers are a limited resource that is globally shared
by the entire Internet community. As of 2014, approximately 5850 TCP
and 5570 UDP port numbers have been assigned out of a total range of
49151. As a result of past conservation, current port use is small
and the current rate of assignment avoids the need for transition to
larger number spaces. This conservation also helps avoid the need
for IANA to rely on port number reclamation, which is practically
impossible even though procedurally permitted [RFC6335].

IANA aims to assign only one port number per service, including
variants [RFC6335], but there are other benefits to using fewer port
numbers for a given service. Use of multiple port numbers can make
applications more fragile, especially when firewalls block a subset
of those port numbers or use ports numbers to route or prioritize
traffic differently. As a result:

>> Each port requested MUST be justified as independently necessary.

6.1. Guiding Principles

This document provides recommendations for users that also help
conserve port number space. Again, this document does not update
BCP165 [RFC6335], which describes the IANA procedures for managing
transport port numbers and services. Port number conservation is
based on a number of basic principles:

   o A single assigned port number can support different functions
      over separate endpoint associations, determined using in-band
      information. An FTP data connection can transfer binary or
      text files, the latter translating line-terminators, as
      indicated in-band over the control port number [RFC959].

   o A single assigned port number can indicate the Dynamic port
      number(s) on which different capabilities are supported, as
      with passive-mode FTP [RFC959].

   o Several existing services can indicate the Dynamic port
      number(s) on which other services are supported, such as with
      mDNS and portmapper [RFC1833] [RFC6762] [RFC6763].

   o Copies of an existing service can be differentiated by using
      different IP addresses, either on different hosts or as
      different real or virtual interfaces (or even operating
      systems) on the same host.

   o Copies of some existing services can be differentiated using
      in-band information (e.g., URIs in HTTP Host field and TLS
      Server Name Indication extension) [RFC7230] [RFC6066].

   o Services requiring varying performance properties can already
      be supported using separate endpoint associations (connections
      or other associations), each configured to support the desired
      properties.

   Port numbers are intended to differentiate services, not variations
   of performance, replicas, pairwise endpoint associations, or payload
   types. Port numbers are also a small space compared to other
   Internet number spaces; it is never appropriate to consume port
   numbers to conserve larger spaces such as IP addresses.

6.2. Firewall and NAT Considerations

   Assigned port numbers are useful for configuring firewalls and other
   port-based systems for access control. Ultimately, these port
   numbers indicate services only to the endpoints, and any
   intermediate device that assigns meaning to a value can be
   incorrect. End systems might agree to run web services (HTTP) over
   port number 53 (typically used for DNS) rather than port number 80,
   at which point a firewall that blocks port number 80 but permits
   port number 53 would not have the desired effect. However, assigned
   port numbers often are important in helping configure firewalls.

   Using Dynamic port numbers, or explicitly-indicated port numbers
   indicated in-band over another service (such as with FTP) often
   complicates firewall and NAT interactions [RFC959]. FTP over
   firewalls often requires direct support for deep-packet inspection
   (to snoop for the Dynamic port number for the NAT to correctly map)
   or passive-mode FTP (in which both connections are opened from the
   client side).

7. How to Use Assigned Port Numbers

Port numbers are assigned by IANA by a set of documented procedures [RFC6335]. The following section describes the steps users can take to help assist with the use of assigned port numbers, and with preparing an application for a port number assignment.

7.1. Is a port number assignment necessary?

First, it is useful to consider whether a port number assignment is required. In many cases, a new number assignment may not be needed, for example:

   o Is this really a new service, or can an existing service
      suffice?

   o Is this an experimental service [RFC3692]? If so, consider
      using the current experimental ports [RFC2780].

   o Is this service independently useful? Some systems are
      composed from collections of different service capabilities,
      but not all component functions are useful as independent
      services. Port numbers are typically shared among the smallest
      independently-useful set of functions. Different service uses
      or properties can be supported in separate pairwise endpoint
      associations after an initial negotiation, e.g., to support
      software decomposition.

   o Can this service use a Dynamic port number that is coordinated
      out-of-band, e.g.:

       o By explicit configuration of both endpoints.

       o By internal mechanisms within the same host (e.g., a
         configuration file, indicated within a URI, or using
         interprocess communication).

       o Using information exchanged on a related service: FTP, SIP,
         etc. [RFC959] [RFC3261].

       o Using an existing port discovery service: portmapper, mDNS,
         etc. [RFC1833] [RFC6762] [RFC6763].

There are a few good examples of reasons that more directly suggest that not only is a port number not necessary, but it is directly counter-indicated:

o Port numbers are not intended to differentiate performance
   variations within the same service, e.g., high-speed vs.
   ordinary speed. Performance variations can be supported within
   a single port number in context of separate pairwise endpoint
   associations.

o Additional port numbers are not intended to replicate an
   existing service. For example, if a device is configured to
   use a typical web browser then it the port number used for
   that service is a copy of the http service that is already
   assigned to port number 80 and does not warrant a new
   assignment. However, an automated system that happens to use
   HTTP framing - but is not primarily accessed by a browser -
   might be a new service. A good way to tell is "can an
   unmodified client of the existing service interact with the
   proposed service"? If so, that service would be a copy of an
   existing service and would not merit a new assignment.

o Port numbers not intended for intra-machine communication.
   Such communication can already be supported by internal
   mechanisms (interprocess communication, shared memory, shared
   files, etc.). When Internet communication within a host is
   desired, the server can bind to a Dynamic port that is
   indicated to the client using these internal mechanisms.

o Separate port numbers are not intended for insecure versions
   of existing (or new) secure services. A service that already
   requires security would be made more vulnerable by having the
   same capability accessible without security.

   Note that the converse is different, i.e., it can be useful to
   create a new, secure service that replicates an existing
   insecure service on a new port number assignment. This can be
   necessary when the existing service is not backward-compatible
   with security enhancements, such as the use of TLS [RFC5246].

o Port numbers are not intended for indicating different service
   versions. Version differentiation should be handled in-band,
   e.g., using a version number at the beginning of an
   association (e.g., connection or other transaction). This may
   not be possible with legacy assignments, but all new
   assignments should incorporate support for version indication.

Some users may not need assigned port numbers at all, e.g., SIP
allows voice calls to use Dynamic ports [RFC3261]. Some systems can
register services in the DNS, using SRV entries. These services can
be discovered by a variety of means, including mDNS, or via direct

query [RFC6762] [RFC6763]. In such cases, users can more easily
request a SRV name, which are assigned first-come, first-served from
a much larger namespace.

IANA assigns port numbers, but this assignment is typically used
only for servers, i.e., the host that listens for incoming
connections or other associations. Clients, i.e., hosts that
initiate connections or other associations, typically refer to those
assigned port numbers but do not need port number assignments for
their endpoint.

Finally, an assigned port number is not a guarantee of exclusive
use. Traffic for any service might appear on any port number, due to
misconfiguration or deliberate misuse. Application and service
designers are encouraged to validate traffic based on its content.

7.2. How Many Port Numbers?

As noted earlier, systems might require a single port number
assignment, but rarely require multiple port numbers. There are a
variety of known ways to reduce port number use. Although some may
be cumbersome or inefficient, they are always preferable to
consuming additional port numbers.

Such techniques include:

o Use of a discovery service, either a shared service (mDNS), or
   a discovery service for a given system [RFC6762] [RFC6763].

o Multiplex packet types using in-band information, either on a
   per-message or per-connection basis. Such demultiplexing can
   even hand-off different messages and connections among
   different processes, such as is done with FTP [RFC959].

There are some cases where it is still important to have assigned
port numbers, largely to traverse either NATs or firewalls. Although
automatic configuration protocols have been proposed and developed
(e.g., STUN [RFC5389], TURN [RFC5766], and ICE [RFC5245]),
application and service designers cannot yet rely on their presence.

In the past, some services were assigned multiple port numbers or
sometimes fairly large port ranges (e.g., X11). This occurred for a
variety of reasons: port number conservation was not as widely
appreciated, assignments were not as ardently reviewed, etc. This no
longer reflects current practice and such assignments are not
considered to constitute a precedent for future assignments.

7.3. Picking a Port Number

   Given a demonstrated need for a port number assignment, the next
   question is how to pick the desired port number. An application for
   a port number assignment does not need to include a desired port
   number; in that case, IANA will select from those currently
   available.

   Users should consider whether the requested port number is
   important. For example, would an assignment be acceptable if IANA
   picked the port number value? Would a TCP (or other transport
   protocol) port number assignment be useful by itself? If so, a TCP
   (UDP) port number can be assigned whose port number is already (or
   can be subsequently) assigned to a different transport protocol.

   The most critical issue in picking a number is selecting the desired
   range, i.e., System vs. User port numbers. The distinction was
   intended to indicate a difference in privilege; originally, System
   port numbers required privileged ('root') access, while User port
   numbers did not. That distinction has since blurred because some
   current systems do not limit access control to System port numbers
   and because some System services have been replicated on User
   numbers (e.g., IRC). Even so, System port number assignments have
   continued at an average rate of 3-4 per year over the past 7 years
   (2007-2013), indicating that the desire to keep this distinction
   continues.

   As a result, the difference between System and User port numbers
   needs to be treated with caution. Developers are advised to treat
   services as if they are always run without privilege. As a result:

   >> Developers SHOULD NOT apply for System port numbers because the
   increased privilege they are intended to provide is not always
   enforced.

   Even when developers seek a System port number, it may be very
   difficult to obtain. System port number assignment requires IETF
   Review or IESG Approval and justification that both User and Dynamic
   port number ranges are insufficient [RFC6335].

   >> System implementers SHOULD enforce the need for privilege for
   processes to listen on System port numbers.

   At some future date, it might be useful to deprecate the distinction
   between System and User port numbers altogether. Services typically
   require elevated ('root') privileges to bind to a System port
   number, but many such services go to great lengths to immediately

drop those privileges just after connection or other association establishment to reduce the impact of an attack using their capabilities. Such services might be more securely operated on User port numbers than on System port numbers. Further, if System port numbers were no longer assigned, as of 2014 it would cost only 180 of the 1024 System values (17%), or 180 of the overall 49152 assigned (System and User) values (<0.04%).

7.4. Support for Security

Just as a service is a way to obtain information or processing from a host over a network, a service can also be the opening through which to attack that host. This vulnerability can be mitigated a number of ways:

>> New services SHOULD support security, either directly or via a secure transport such as TLS [RFC5246].

>> Insecure versions of new or existing secure services SHOULD be avoided because of the new vulnerability they create.

>> When simultaneously requesting both a secure and an insecure port, strong justification MUST be provided for the insecure port. Precedent (citing other protocols that use an insecure port) is not strong justification by itself. A strong case for utility of the insecure service is REQUIRED for approval of the insecure port.

>> Security SHOULD NOT rely on port number distinctions alone; every service, whether secure or not, is likely to be attacked.

There is debate as to how to secure legacy insecure services [RFC6335]. Some argue that secure variants should share the existing port number assignment, such that security is enabled on a per-connection or other association basis [RFC2817]. Others argue that security should be supported on a new port number assignment and be enabled by default. Either approach is currently permitted, although use of a single port number is consistent with port number conservation. A separate port number might be important for security coordination (e.g., firewall management), but this might further argue for deprecation of the insecure variant.

Optional security can penalize performance, requiring additional round-trip exchanges before a connection or other association can be established. As discussed earlier, port numbers are a critical resource and it is inappropriate to consume assignments to increase performance. As a result, the need for separate ports for both secure and insecure variants is not justified merely for performance

- either for the connection or association establishment performance
or differences in data performance between secure and insecure
variants.

Note however that a new service might not be eligible for IANA
assignment of both an insecure and a secure variant of the same
service, and similarly applications requesting assignment for both
an insecure port number for a secure service might not be
appropriate. In both cases, security of the service is compromised
by adding the insecure port number assignment.

7.5. Support for Future Versions

Current IANA assignments are expected to support the multiple
versions on the same assigned port number [RFC6335]. Versions are
typically indicated in-band, either at the beginning of a connection
or other association, or in each protocol message.

>> Version support SHOULD be included in new services.

>> Version numbers SHOULD NOT be included in either the service name
or service description.

Again, the port number space is far too limited to be used as an
indicator of protocol version or message type. Although this has
happened in the past (e.g., for NFS), it should be avoided in new
requests.

7.6. Transport Protocols

IANA assigns port numbers specific to one or more transport
protocols, typically UDP and TCP, but also SCTP, DCCP, and any other
standard transport protocol [RFC768] [RFC793] [RFC4340] [RFC4960].
Originally, IANA port number assignments were concurrent for both
UDP and TCP; other transports were not indicated. However, to
conserve space and to reflect increasing use of other transports,
assignments are now specific only to the transport being used.

In general, a service should request assignments for multiple
transports using the same service name and description on the same
port number only when they all reflect essentially the same service.
Good examples of such use are DNS and NFS, where the difference
between the UDP and TCP services are specific to supporting each
transport. E.g., the UDP variant of a service might add sequence
numbers and the TCP variant of the same service might add in-band
message delimiters. This document does not describe the appropriate
selection of a transport protocol for a service.

>> Service names and descriptions for multiple transport port number
assignments SHOULD match only when they describe the same service,
excepting only enhancements for each supported transport.

When the services differ, their service names and descriptions
should reflect that difference. E.g., if TCP is used for the basic
control protocol and UDP for an alarm protocol, then the services
might be "name-ctl" and "name-alarm". A common example is when TCP
is used for a service and UDP is used to determine whether that
service is active (e.g., via a unicast, broadcast, or multicast test
message) [RFC1122]. The following convention has been used by IANA
for several years to distinguish discovery services, such as are
used to identify endpoints capable of a given service:

>> Names of discovery services SHOULD use an identifiable suffix;
the suggestion is "-disc".

Some services are used for discovery, either in conjunction with a
TCP service or as a stand-alone capability. Such services will be
more reliable when using multicast rather than broadcast (over IPv4)
because IP routers do not forward "all nodes" (all 1's, i.e.,
255.255.255.255 for IPv4) broadcasts and have not been required to
support subnet-directed broadcasts since 1999 [RFC1812] [RFC2644].
This issue is relevant only for IPv4 because IPv6 does not support
broadcast.

>> UDP over IPv4 multi-host services SHOULD use multicast rather
than broadcast.

Designers should be very careful in creating services over
transports that do not support congestion control or error recovery,
notably UDP. There are several issues that should be considered in
such cases, as summarized in Table 1 in [RFC5405]. In addition, the
following recommendations apply to service design:

>> Services that use multipoint communication SHOULD be scalable,
and SHOULD NOT rely solely on the efficiency of multicast
transmission for scalability.

>> Services SHOULD NOT use UDP as a performance enhancement over
TCP, i.e., to circumnavigate TCP's congestion control.

7.7. When to Request an Assignment

Assignments are typically requested when a user has enough
information to reasonably answer the questions in the IANA
application. IANA applications typically take up to a few weeks to

process, with some complex cases taking up to a month. The process typically involves a few exchanges between the IANA Ports Expert Review team and the applicant.

An application needs to include a description of the service, as well as to address key questions designed to help IANA determine whether the assignment is justified. The application should be complete and not refer solely to the Internet Draft, RFC, a website, or any other external documentation.

Services that are independently developed can be requested at any time, but are typically best requested in the last stages of design and initial experimentation, before any deployment has occurred that cannot easily be updated.

>> Users MUST NOT deploy implementations that use assigned port numbers prior their assignment by IANA.

>> Users MUST NOT deploy implementations that default to using the experimental System port numbers (1021 and 1022 [RFC4727]) outside a controlled environment where they can be updated with a subsequent assigned port [RFC3692].

Deployments that use port numbers before deployment complicate IANA management of the port number space. Keep in mind that this recommendation protects existing assignees, users of current services, and applicants for new assignments; it helps ensure that a desired number and service name are available when assigned. The list of currently unassigned numbers is just that - *currently* unassigned. It does not reflect pending applications. Waiting for an official IANA assignment reduces the chance that an assignment request will conflict with another deployed service.

Applications made through Internet Draft / RFC publication (in any stream) typically use a placeholder ("PORTNUM") in the text, and implementations use an experimental port number until a final assignment has been made [RFC6335]. That assignment is initially indicated in the IANA Considerations section of the document, which is tracked by the RFC Editor. When a document has been approved for publication and proceeds to IESG Approval, that request is forwarded to IANA for handling. IANA will make the new assignment accordingly. At that time, IANA may also request that the applicant fill out the application form on their website, e.g., when the RFC does not directly address the information expected as per [RFC6335]. "Early" assignments can be made when justified, e.g., for early interoperability testing, according to existing process [RFC7120] [RFC6335].

>> Users writing specifications SHOULD use symbolic names for port numbers and service names until an IANA assignment has been completed. Implementations SHOULD use experimental port numbers during this time, but those numbers MUST NOT be cited in documentation except as interim.

7.8. Squatting

"Squatting" describes the use of a number from the assigned range in deployed software without IANA assignment. It is hazardous because IANA cannot track such usage and thus cannot avoid making legitimate assignments that conflict with such unauthorized usage.

Such "squatted" port numbers remain unassigned, and IANA retains the right to assign them when requested by applicants. Application and service designers are reminded that is never appropriate to use port numbers that have not been directly assigned [RFC6335]. In particular, any unassigned code from the assigned ranges will be assigned by IANA, and any conflict will be easily resolved as the protocol designer's fault once that happens (because they would not be the assignee). This may reflect in the public's judgment on the quality of their expertise and cooperation with the Internet community.

Regardless, there are numerous services that have squatted on such numbers that are in widespread use. Designers who are using such port numbers are encouraged to apply for an assignment. Note that even widespread de-facto use may not justify a later IANA assignment of that value, especially if either the value has already been assigned to a legitimate applicant or if the service would not qualify for an assignment of its own accord.

7.9. Other Considerations

As noted earlier, System port numbers should be used sparingly, and it is better to avoid them altogether. This avoids the potentially incorrect assumption that the service on such port numbers run in a privileged mode.

Port numbers are not intended to be changed; this includes the corresponding service name. Once deployed, it can be very difficult to recall every implementation, so the assignment should be retained. However, in cases where the current assignee of a name or number has reasonable knowledge of the impact on such uses, and is willing to accept that impact, the name or number of an assignment can be changed [RFC6335]

Aliases, or multiple service names for the same port number, are no longer considered appropriate [RFC6335].

8. Security Considerations

This document discusses ways to conserve port numbers, notably through encouraging demultiplexing within a single port number. As such, there may be cases where two variants of a protocol - insecure and secure (such as using optional TLS) or different versions - are suggested to share the same port number.

This document reminds application and service designers that port numbers do not protect against denial of service overload or guarantee that traffic should be trusted. Using assigned numbers for port filtering isn't a substitute for authentication, encryption, and integrity protection. The port number alone should not be used to avoid denial of service or firewall traffic because their use is not regulated or validated.

The use of assigned port numbers is the antithesis of privacy because they are intended to explicitly indicate the desired application or service. Strictly, port numbers are meaningful only at the endpoints, so any interpretation elsewhere in the network can be arbitrarily incorrect. However, those numbers can also expose information about available services on a given host. This information can be used by intermediate devices to monitor and intercept traffic as well as to potentially identify key endpoint software properties ("fingerprinting"), which can be used to direct other attacks.

9. IANA Considerations

The entirety of this document focuses on suggestions that help ensure the conservation of port numbers and provide useful hints for issuing informative requests thereof.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2780] Bradner, S., and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.

[RFC3692] Narten, T., "Assigning Experimental and Testing Numbers
          Considered Useful", BCP 82, RFC 3962, Jan. 2004.

[RFC4727] Fenner, B., "Experimental Values in IPv4, IPv6, ICMPv4,
          ICMPv6, UDP, and TCP Headers", RFC 4727, November 2006.

[RFC5405] Eggert, L., and G. Fairhurst, "Unicast UDP Usage
          Guidelines for Application Designers", BCP 145, RFC 5405,
          Nov. 2008.

[RFC6335] Cotton, M., L. Eggert, J. Touch, M. Westerlund, and S.
          Cheshire, "Internet Assigned Numbers Authority (IANA)
          Procedures for the Management of the Service Name and
          Transport Protocol Port Number Registry", BCP 165, RFC
          6335, August 2011.

10.2. Informative References

[IEN112] Postel, J., "Transmission Control Protocol", IEN 112,
          August 1979.

[RFC33]   Crocker, S., "New Host-Host Protocol", RFC 33 February
          1970.

[RFC37]   Crocker, S., "Network Meeting Epilogue", RFC 37, March
          1970.

[RFC38]   Wolfe, S., "Comments on Network Protocol from NWG/RFC
          #36", RFC 38, March 1970.

[RFC48]   Postel, J., and S. Crocker, "Possible protocol plateau",
          RFC 48, April 1970.

[RFC61]   Walden, D., "Note on Interprocess Communication in a
          Resource Sharing Computer Network", RFC 61, July 1970.

[RFC76]   Bouknight, J., J. Madden, and G. Grossman, "Connection by
          name: User oriented protocol", RFC 76, October 1970.

[RFC333] Bressler, R., D. Murphy, and D. Walden. "Proposed
          experiment with a Message Switching Protocol", RFC 333,
          May 1972.

[RFC739] Postel, J., "Assigned numbers", RFC 739, November 1977.

[RFC758] Postel, J., "Assigned numbers", RFC 758, August 1979.

   [RFC768]   Postel, J., "User Datagram Protocol", RFC 768, August
              1980.

   [RFC793]   Postel, J., "Transmission Control Protocol" RFC 793,
              September 1981

   [RFC820]   Postel, J., "Assigned numbers", RFC 820, August 1982.

   [RFC900]   Reynolds, J., and J. Postel, "Assigned numbers", RFC 900,
              June 1984.

   [RFC959]   Postel, J., and J. Reynolds, "FILE TRANSFER PROTOCOL
              (FTP)", RFC 959, October 1985.

   [RFC1122] Braden, B. (Ed.), "Requirements for Internet Hosts --
              Communication Layers", RFC 1122, October 1989.

   [RFC1340] Reynolds, J., and J. Postel, "Assigned numbers", RFC 1340,
              July 1992.

   [RFC1700] Reynolds, J., and J. Postel, "Assigned numbers", RFC 1700,
              October 1994.

   [RFC1812] Baker, F. (Ed.), "Requirements for IP Version 4 Routers",
              RFC 1812, June 1995.

   [RFC1833] Srinivasan, R., "Binding Protocols for ONC RPC Version 2",
              RFC 1833, August 1995.

   [RFC2644] Senie, D., "Changing the Default for Directed Broadcasts
              in Routers", RFC 2644, August 1999.

   [RFC2817] Khare, R., and S. Lawrence, "Upgrading to TLS Within
              HTTP/1.1", RFC 2817, May 2000.

   [RFC3232] Reynolds, J. (Ed.), "Assigned Numbers: RFC 1700 is
              Replaced by an On-line Database", RFC 3232, January 2002.

   [RFC3261] Rosenberg, J., H. Schulzrinne, G. Camarillo, A. Johnston,
              J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP:
              Session Initiation Protocol", RFC 3261, June 2002.

   [RFC4340] Kohler, E., M. Handley, and S. Floyd, "Datagram Congestion
              Control Protocol (DCCP)", RFC 4340, March 2006.

   [RFC4960] Stewart, R. (Ed.), "Stream Control Transmission Protocol",
              RFC 4960, September 2007.

   [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment
             (ICE): A Protocol for Network Address Translator (NAT)
             Traversal for Offer/Answer Protocols", RFC 5245, April
             2010.

   [RFC5246] Dierks, T., and E. Rescorla, "The Transport Layer Security
             (TLS) Protocol Version 1.2", RFC 5246, August 2008.

   [RFC5389] Rosenberg, J., R. Mahy, P. Matthews, and D. Wing, "Session
             Traversal Utilities for NAT", RFC 5389, October 2008.

   [RFC5766] Mahy, R., P. Matthews, and J. Rosenberg, "Traversal Using
             Relays around NAT (TURN): Relay Extensions to Session
             Traversal Utilities for NAT (STUN)", RFC 5766, April 2010.

   [RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS)
             Extensions: Extension Definitions", RFC 6066, January
             2011.

   [RFC6762] Cheshire, S., and M. Krochmal, "Multicast DNS", RFC 6762,
             February 2013.

   [RFC6763] Cheshire, S., and M. Krochmal, "DNS-Based Service
             Discovery", RFC 6763, February 2013.

   [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code
             Points", BCP 100, RFC 7120, January 2014.

   [RFC7230] Fielding, R., (Ed.), and J. Reshke, (Ed.), "Hypertext
             Transfer Protocol (HTTP/1.1): Message Syntax and Routing",
             RFC 7230, June 2014.

11. Acknowledgments

   This work benefitted from the feedback from David Black, Lars
   Eggert, Gorry Fairhurst, and Eliot Lear, as well as discussions of
   the IETF TSVWG WG.

   This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

    Joe Touch
    USC/ISI
    4676 Admiralty Way
    Marina del Rey, CA 90292-6695
    U.S.A.

    Phone: +1 (310) 448-9151
    EMail: touch@isi.edu