TSVWG   J. Touch
Internet Draft   USC/ISI
Intended status: Best Current Practice   July 28, 2011
Expires: January 2012

                 Recommendations for Transport Port Uses
                    draft-touch-tsvwg-port-use-01.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 August 28, 2011.

Copyright Notice

Abstract

   This document provides recommendations to application and service
   designers on how to use the transport protocol port number space to
   help in its preservation. **NOTE THAT THIS CURRENT VERSION IS
   LARGELY AN OUTLINE OF ISSUES**.

Table of Contents

1. Introduction

   (TBD)

2. History

   The term 'port' was first used in RFC33 to describe a simplex
   communication path from a process. 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 source and
   destination socket identifiers in packets. RFC38 explains this in
   detail, in which processes might have more than one of these paths,
   and that more than one 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 the incoming data could be demultiplexed to the

appropriate process. RFC38 further suggested that 32 bits would be used for these identifiers. RFC48 discusses the current notion of listening on a given port, but does not discuss how the issue of port 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 for common services.

RFC76 addresses this issue more constructively, proposing a "telephone book" by which an index would allow ports to be used by name, but still assumes that both source and destination ports are fixed by such a system. RFC333 suggests that the port pair, rather than an individual port, would be used on both sides of the connection for demultiplexing messages. This is the final view in RFC793 (and its predecessors, including IEN 112), and brings us to their current meaning. RFC739 introduces the notion of generic reserved ports, used for groups of protocols, such as "any private RJE server". Although the overall range of such ports 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 a list of such well-known ports, 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 assignments is presented. By RFC900, they 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 assignments individually (although the assumption was, and remains, that once assigned a port 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 space 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.

3. Current Port Use

The current IANA website (www.iana.org) indicates three ranges of
port assignments:

```
   Binary          Hex

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

   0-1023          0x03FF          Well-Known (a.k.a. 'system')

   1024-49151      0x0300-0xBFFF   Registered (a.k.a. 'user')

   49152-65535     0xC000-0xFFFF   Dynamic/Private
```

Well-known encompasses the range 0..1023. On some systems, use of
these ports requires privileged access, e.g., that the process run
as 'root', which is why these are referred to as 'system' ports. The
ports from 1024..49151 denotes non-privileged services, known as
'registered'; because these ports do not run with special
privileges, they are often referred to as 'user' ports. Dynamic or
Private ports are not registered through IANA.

Both Well-Known and Registered ports are registered through IANA, so
both are sometimes called "registered ports". As a result, the term
'registered' is ambiguous, referring either to the entire range 0-
49151 or to the user ports. Complicating matters further, 'system'
ports do not always require special (i.e., 'root') privilege.
Regardless, for clarity, throughout the remainder of this document
we will refer to the port ranges as 'system', 'user', and 'private'.

4. What is a Port?

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

   o Demultiplexing transport connections within an end host

   o Identifying a service

The first reason requires that each transport connection between a
given pair of IP addresses use a different pair of ports, but does
not require either coordination or registration of port use. It is
the second reason that drives the need for a common registry.

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

access that service at that host. Such information can be
distributed out of band, e.g., in the URL, such as:

    http:51509//www.example.com/

Ultimately, it's important to keep in mind that the correlation of a
service with a port number is an agreement between the two endpoints
of the connection only. The rest of the world might think that
you're sending DNS packets on port 53, but you can run a web server
on that port just fine, provided the server and client both decide
that port 53 is for HTTP web server traffic.

Which brings us to the concept of a service. A service is the
combination of ISO Layers 5-7 that represent an application protocol
capability. For example www (port 80) is a service that uses HTTP as
an application protocol, and provides a common web server. 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 registers ports so that endpoints on the Internet do not need
to pairwise, explicitly coordinate the meaning of their port
numbers. This is the primary reason for registering ports with IANA
- to have a common agreement between all endpoints on the Internet
as to the meaning of a port.

Ports are used for other purposes as well, however. The other
primary reason for registering ports with IANA is to simplify end
system configuration, so individual installations do not need to
coordinate their use of arbitrary ports. A similar reason is to
simplify firewall management, so that a single, fixed firewall
configuration can either permit or deny a service.

[5](). Conservation

(statistics of port allocations)

Ways to conserve, e.g., use service names (DNS SRV, TCP portnames,
etc.), use portmapper, Bonjour, or other services for demuxing

[6](). How to Use Registered Ports

[6.1](). Do You Need a Port?

How to carefully use experimental ports (include a large nonce)

Reasons NOT to register a port, e.g., not for a copy of an existing service,

Reasons why only server ports are registered (not client)

6.2. How Many Ports?

Reasons NOT to have multiple ports (performance, etc.)

Techniques to reduce port use:

o   When can you use a discovery service

o   When can you use multiplexing

o   When can you use handoff with in-band IDs

6.3. Picking a Port Number

Would you still want one if you can't pick the value?

Would you still want the UDP / TCP one if it didn't match the value for a previously assigned TCP / UDP one?

6.4. Support for Security

Why this is generally expected

Why this should/should not use a separate port (it's a performance issue, and performance would argue for multiple ports anyway, and ports are a limited resource)

TLS allows optional security

6.5. Support for Future Versions

Reasons NOT to include the version number in the name

6.6. Transport Protocols

UDP vs. TCP vs. others - and when the transport you lookup is not always the one you end up using

When/why you need multiples

When UDP is -disc

Caveats about using broadcast as discovery.

6.7. When to Register

What range to use before registering

When are you ready to register (basically when you have enough information to fill out the application)

Reasons NOT to squat

6.8. Other Considerations

Higher bar for system ports

Changing a name

Why aliases are bad and now deprecated

Providing enough information for IANA review, e.g., to avoid Internet congestion, fit in MTUs, deal with reordering, etc.

Provide enough information that's stand-alone; don't describe a protocol by a URL, e.g. - how to do a high-level description (what are we looking for?)

Why a heartbeat port MUST be on the same port as a service.

Relation of this doc to the IANA Port Experts review process (this is just a summary from the user/designer viewpoint, and is NOT binding to IANA or its Expert Review team)

7. Recommendations for Future Allocation

Abolish the distinction of system ports? BIG QUESTION HERE...

Why NOT to allocate ports or names for use as examples

8. Security Considerations

This document discusses ways to conserve port numbers, notably through encouraging demultiplexing within a single port. As such, there may be cases where two variants of a protocol - insecure and secure, are suggested to share the same port (e.g., HTTP and HTTPS, though currently those are assigned different ports).

This document reminds protocol designers that port numbers are not a substitute for security, and should not alone be used to avoid denial of service or firewall traffic, notably because their use is not regulated or authenticated.

9. IANA Considerations

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

10. Conclusions

<Add any conclusions>

11. References

11.1. Normative References

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

11.2. Informative References

(some of these may be normative if this is BCP...?)

[IEN112]

[RFC?] Dccp

[RFC?] HTTP

[RFC?] HTTPS

[RFC?] Sctp

[RFC?] SSL

[RFC?] TLS

[RFC33]

[RFC37]

[RFC38]

[RFC48]

[RFC61]

[RFC76]

[RFC333]

[RFC739]

[RFC758]

[RFC768] UDP

[RFC793] TCP

[RFC820]

[RFC900]

[RFC1122]

[RFC1340]

[RFC1700]

[RFC1918]

[RFC2780]

[RFC3232]

ISO-TP4

Simon Lam's pre-SSL (see his web page)

12. Acknowledgments

TBD

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
    URL: http://www.isi.edu/touch