

Geoff Huston
April 2016

Declaring IPv4 “Historic”

At the IETF 95 meeting at the start of April I was in a meeting of the IPv4 Sunset Working Group, and heard Lee Howard present on a proposal that recommended that IP version 4, or to be specific, that the technical protocol specification documented in RFC 791, be declared “Historic”.

In the context of the Internet Standards Process the term “Historic” has a particular meaning. RFC 2026 defines “Historic” to mean that:

```
A specification that has been superseded by a more recent specification or
is for any other reason considered to be obsolete is assigned to the
"Historic" level.
```

The rationale for this proposed re-designation of IPv4 was that this protocol has indeed been superseded by a more recent specification, namely IP version 6. Furthermore, it was hoped that this action would add to the impetus to deploy IPv6. The take up of IPv6 is not overly uniform, and while some service providers have been enthusiastic proponents of IPv6, many more providers are still at best hesitant, and at worse just ignoring it and hoping that it will all just go away! As a clear message to the laggards, and potentially their customer base, it was argued that its time for the IETF to send another clear message that its time to move on to IPv6. And one way to achieve that is to declare that IPv4 is now an “Historic” technical specification.

If that was all there was to it then perhaps there is a case to be made that the IPv4 technical specification should be declared as “Historic”. But there are a few additional aspects to consider here. As also pointed out by RFC2026:

```
Not Recommended: A Technical Specification that is considered to be
inappropriate for general use is labeled "Not Recommended". This may be
because of its limited functionality, specialized nature, or historic
status.
```

This text seems to imply that a status of “Historic” implies also “Not Recommended,” which is perhaps sending the wrong signal to the existing user base that relies on IPv4.

This proposed re-designation would also throw the IPv4 specification out of the Internet Standards set:

```
Specifications that are not on the standards track are labeled with one of
three "off-track" maturity levels: "Experimental", "Informational", or
"Historic". The documents bearing these labels are not Internet Standards
in any sense.
```

So maybe this is a bigger step than just observing that IPv6 supersedes IPv4. As one commenter in the Working Group session pointed out, declaring IPv4 “Historic” would likely backfire and serve no better purpose other than exposing the IETF to ridicule. And certainly there is some merit in wondering why a standards body would take a protocol specification used by over 3 billion people, and by some estimated 10 billion devices each and every day and declare it to be “Historic”. In any other

context such adoption figures for a technology would conventionally be called “outstandingly successful”!

Perhaps we can put this into some broader context by looking at other “Historic” specifications. Unfortunately, the IETF does not have an obviously consistent story in its actions of declaring technical specifications to be “Historic”. For example, some very old and now completely unused services are described in RFCs that are not declared to be “Historic”. We can go a long way back in time to RFC 162, and find a specification of that the completely forgotten protocol “netbugger3” (I’m not sure if the spaces were dropped accidentally or not in the name of that protocol) is not historic. If there are any extant implementations of this curiously-named protocol, I would be keen to learn of them. And then there is Gopher, a specification that enjoyed a brief moment in the sun in the early 90’s before the juggernaut that is today’s web completely superseded it. Are either of these specifications “Historic”? Nope, not according to the IETF.

So if some pretty obviously defunct protocols are not “Historic”, then what is? A browse of the historic RFCs reveals a collection of TCP extension specifications, such as RFC1072 and RFC1106 as “Historic”. They were declared historic by RFC 6247, and the rationale given there was that they “have never seen widespread use”. This is fine, but that’s not applicable to IPv4 by any stretch of the imagination! Browsing the historic RFC list at <https://www.rfc-editor.org>, it’s evident that there are more than a few RFC documents that never even saw the light of day as “current” specifications as they were declared “Historic” at the outset. Again, quite obviously, this is not applicable to IPv4.

What about other Internet Standards? There does seem to be a reasonable case to be made that Internet Standards numbers 23, and 24 (Quote of the Day, and Finger respectively) have long passed out of common use. “Historic” seems to be entirely applicable for both of those now quite venerable standards, as their previous widespread use has waned to the point of almost complete invisibility.

So we appear to treat “Historic” status somewhat whimsically. We could be a little more consistent here and in that vein perhaps there is a case to be made to push Finger, Quote of the Day, Gopher, Netbugger3 and their like into “Historic”. The world has moved on and these protocols are stuck in an older world. But not IPv4. It’s not just that it is in use by an unprecedented number of people and devices today and each and every one of us rely on this particular protocol. It’s not just that.

It’s the observation that we probably haven’t finished with IPv4 yet.

While many folk, including myself would dearly like to see an all-IPv6 Internet today, I would like to think I’m pragmatic enough to understand that we are stuck with a dual stack Internet for many years to come. But that pragmatic observation has its consequences. So far we have managed to cram some 10 billion unique devices into today’s Internet. The silicon industry is not going to stop and wait for us to complete this IPv6 transition and in the meantime we can readily foresee imagine a near future which puts every new computer, smartphone, personal pad, every television, every new car, and a whole heap of other applications that may appear, on this dual stack Internet. We may well need to push the IPv4 Internet to encompass 20 billion or so devices on this strange and protracted dual stack journey to IPv6. One immediate change so far is the semantics of IPv4 addresses. An IPv4 address is increasingly an ephemeral short term element of a conversation stream identifier. It has lost any concept of being a stable endpoint identifier. The more devices we try and push into the network the more we are changing the way IPv4 behaves. As we try and make IPv4 stretch just that little bit further we may need to make some further subtle changes. They may impact the current specification for IPv4, or they may not. We don’t know just yet. But what we do know is that the story is by no means over for IPv4 right now.

Now as a standards body it may well sound like a good idea for the IETF to send a strong signal to the industry about the need to take this transition seriously by declaring IPv4 to be “Historic”. But if this is what the IETF does, then the work on IPv4 will probably continue. The risk is that it will continue

without the benefit and support of the acknowledged Internet Standards organization, namely the IETF. And we have some prior experience what happens then.

The last time the IETF turned its back on a technology specification was the development of the specification of Network Address Translation (NAT). The result was that implementers could not rely on a complete and coherent specification and were forced to make it up as they went along. Every NAT product had subtle behavioural differences with every other NAT product. The losers here were application developers and ultimately we, the users. Applications had to work across NATs and they had to negotiate functionality across a diverse set of undocumented and at times inconsistent behaviours. The resulting applications can be brittle and fail in unanticipated ways. Standards help us to understand how to interoperate, and how to rely upon predictable ways to interoperate. It takes a lot of the guesswork out of technology specification and can eliminate large amount of complexity in implementing technology. I would be horrified if we managed to repeat this mistake in this chapter of IPv4's life history.

IPv4 Historic?

No, the reports of its demise are severely exaggerated!

As a colleague observed, it was kind of a shame that the Working Group was unable to meet on April 1. This proposal would've made a whole lot more sense on that day!

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990's. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001 and chaired a number of IETF Working Groups. He has worked as an Internet researcher, as an ISP systems architect and a network operator at various times.

www.potaroo.net

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.