July 2025
Geoff Huston

IEPG at IETF 123

The Internet Engineering Task Force (IETF) meets three times a year to work on Internet Standards and related operational practice documents. In July of 2025 the IETF met in Madrid (finally, and after a number of thwarted mis-starts!) with more than a thousand folk in attendance through the week.

The IEPG meeting is held each Sunday at the start of the IETF week. There is no particular theme for these sessions, although subjects of operational relevance are encouraged (www.iepg.org). These are my impressions of the presentations that were made at this IEPG meeting.

BGP Forensics

Detecting External Disruptions in Internet Services Provider Networks - The role of the BGP routing protocol is to maintain a current view of the state of the reachable Internet within the local BGP speaker. When disruptions in the network occur that affect the reachability of certain destinations then this change is propagated across the entire network by BGP. This can be reverse engineered such that changes in reachability and topology that are propagated as BGP updates can be analysed to infer the impact of the root cause events. This was the topic of this presentation. The work focuses on the occurrence of BGP withdrawals and topology updates that infer the drop of inter AS adjacencies, combined with some aspect of traffic flow data. This is not exactly novel work, and such analysis of routing and traffic data has been around for many years. The challenge has always been in converting this type of analysis into a marketable product. I'm not convionced that this work breaks any new ground in this respect.

Triggering QUIC

The QUIC transport protocol is one of the few truly innovative changes in the area of transport services for many years. It had its origins in work within Google to improve the speed of connection establishment (SPDY), and as well as providing a faster path to establishing as TLS end-to- end connection, within the connection QUIC provides a number of innovations including parallel streams that do not suffer from head-of line blocking, reliable datagram services and an protected control channel. The challenge is how a client can determine that the server supports QUIC connections. Yes, the client can send a QUIC connection establishment packet to the server's UDP port 443 and wait for a response. but such a send and wait approach is not necessarily fast, which defeats some of the original objective of QUIC. There are other approaches to triggering QUIC that do not have attendant delays. One approach uses a directive (alt-svc) embedded in the HTTPS header (RFC 7838) which indicates that the server supports QUIC. This means that the second time the client accesses this service, it can immediately use QUIC with the confident expectation that the server will positively respond to the Client's QUIC hello packet. The problem here is that with session persistence in HTTP/1.0 and HTTP/2 there is often no second connection, but just a reuse of the original connection. This embedded alt-svc approach has been used by the Chrome browser for the past five years. The Safari browser introduced

support for QUIC more recently, but Apple elected to use the more recently defined HTTPS DNS query (RFC 9460). This presentation reported on the use of these two QUIC trigger mechanisms by these two dominant browsers. The measurement reveals that a very small proportion of Chrome browsers support the use of QUIC (despite Chrome documentation from 2022 that suggests the opposite). Safari browsers exclusively use the HTTPS query, but only three quarters of these browsers follow up with a QUIC connection. The implication for content servers is that both QUIC trigger mechanisms, alt-svc and DNS HTTPS need to be supported to use QUIC for HTTP connections. We should be doing better!

BGP Path Attributes

BGP is a venerable protocol these days, with more than 30 years of operational experience connecting the Internet. You might think that we had all the uncertainties in the use of this protocol, but that's not the case! The topic for this presentation is the treatment of BGP Path Attributes. We've had Path Attributes in BGP since its inception. They are BGP's way of associating route properties with network destinations. This is encoded as an 8-bit field, so there are 256 possible attributes values. A registry of these Path Attributes have 40 such attributes standardise. Path Attributes also have a "transitive bit" setting. If this bit is clear then the attribute is non-transitive and the receiving BGP speaker discards the attribute without further propagation, whether it recognises the attribute or not. If set, the attribute is passed onward as part of conventional route object propagation in BGP. The "partial bit" is used of the used if the attribute is not recognised locally. Transitive Path Attributes are how BGP incrementally deploys new features in a partial adoption scenario. The current standard, BGP-4 (RFC4271), defines a small set of Path Attributes that all implementations must recognise, namely ORIGIN, AS_PATH, LOCAL_PREF, MULTI_EXIT_DISC (MED), NEXT_HOP, ATOMIC_AGGREGATE, AGGREGATOR. BGP implementations generally understand more than this set, including route reflection, Communities, Multi-Protocol, and 4-byte AS_PATH. There have been Path Attributes created for local circumstances, such as link state and VPN attributes. However, inconsistent care has been taken to prevent such local attributes leaking from their appropriate scope, which can cause incorrect forwarding. Much of the original motivation for RFC7606 error handling procedures was to deal with "optional transitive nonsense". The problem here is that the issues associated with bugs or forwarding issues caused by new path attributes have led to implementations creating Path Attribute filtering features. Some implementations discard routes with specific attributes, while others strip those attributes. and some locally ignore those attributes but propagate the routes along with the attributes. Filtering Path Attributes breaks incremental deployment of new features. However, it's a silent feature, and there's no visibility when it's used. It would be useful if BGP speakers could publish their policies within each BGP peering session, and there's a proposal in IDR to discuss doing exactly this. What should the default policy be? Should unknown path attributes be propagated by default if the transitive but is set, whether its locally recognised or not? Or should we filter it by discarding the attribute or blocking the routes?

EU Standards Initiatives

It has often been said of standards work is that the good thing is that there are so many to choose from! Many years ago, the IETF was not aimed at producing a bevy of standard specifications as an end in and of itself but was intended to produce specifications that aid implementers in producing interoperable implementations of an IETF protocol. In many cases such standard specifications are not about describing every potential configuration setting, but in making choices that enable interoperability. In its early there was a perception that the IETF was little more than an American standards body with international pretentions, and the group has worked diligently in the ensuing decades to amass an impressive level of international participation and recognition, both in the generation of RFCs and in the communities of users and vendors that rely on these standards to ensure the correct functioning of product that is compliant with the relevant RFCS.

In this context it seems somewhat odd to see the European Commission advocating the formation of a regional multistakeholder forum tasked to identify the "best" available standards and deployment techniques. It's not as if the IETF is not working on both standards and Best Current Practice documents already, and even in the European context there are many groups working on these topics in the Internet sector. Personally, I'm at a loss to understand why we need another regional coordination body in the space! More such bodies don't necessarily produce better outcomes, and they run a strong risk that a worse, or at least a contradictory and confusing outcome, will eventuate.

Next

To prevent this report becoming too unwieldy I'll structure this report on IETF 123 into separate posts. Next up is the DNS!

Disclaimer

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

Author

Geoff Huston AM, M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net