

Workgroup: Network Working Group

Internet-Draft: draft-ietf-use-it-or-lose-it-01

Published: 14 July 2021

Intended Status: Informational

Expires: 15 January 2022

Authors: M. Thomson    T. Pauly

          Mozilla        Apple

## **Long-term Viability of Protocol Extension Mechanisms**

### **Abstract**

The ability to change protocols depends on exercising the extension and version negotiation mechanisms that support change. Protocols that don't use these mechanisms can find it difficult and costly to deploy changes.

### **Discussion Venues**

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the EDM Program mailing list (edm@iab.org), which is archived at <https://mailarchive.ietf.org/arch/browse/edm/>.

Source for this draft and an issue tracker can be found at <https://github.com/intarchboard/use-it-or-lose-it>.

### **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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 15 January 2022.

### **Copyright Notice**

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

- [1. Introduction](#)
- [2. Imperfect Implementations Limit Protocol Evolution](#)
  - [2.1. Good Protocol Design is Not Itself Sufficient](#)
  - [2.2. Disuse Can Hide Problems](#)
    - [2.2.1. TLS](#)
    - [2.2.2. DNS](#)
    - [2.2.3. SNMP](#)
    - [2.2.4. HTTP](#)
    - [2.2.5. IPv4](#)
  - [2.3. Multi-Party Interactions and Middleboxes](#)
- [3. Retaining Viable Protocol Evolution Mechanisms](#)
  - [3.1. Examples of Active Use](#)
  - [3.2. Dependency is Better](#)
  - [3.3. Restoring Active Use](#)
- [4. Active Use](#)
  - [4.1. Version Negotiation](#)
  - [4.2. Falsifying Active Use](#)
- [5. Complementary Techniques](#)
  - [5.1. Cryptography](#)
  - [5.2. Fewer Extension Points](#)
    - [5.2.1. Invariants](#)
  - [5.3. Effective Feedback](#)
- [6. Security Considerations](#)
- [7. IANA Considerations](#)
- [8. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

## 1. Introduction

A successful protocol [[SUCCESS](#)] needs to change in ways that allow it to continue to fulfill the needs of its users. New use cases, conditions and constraints on the deployment of a protocol can render a protocol that does not change obsolete.

Usage patterns and requirements for a protocol shift over time. In response, implementations might adjust usage patterns within the

constraints of the protocol, the protocol could be extended, or a replacement protocol might be developed. Experience with Internet-scale protocol deployment shows that each option comes with different costs. [\[TRANSITIONS\]](#) examines the problem of protocol evolution more broadly.

This document examines the specific conditions that determine whether protocol maintainers have the ability to design and deploy new or modified protocols. [Section 2](#) highlights some historical examples of difficulties in transitions to new protocol features. [Section 3](#) argues that ossified protocols are more difficult to update and successful protocols make frequent use of new extensions and code-points. [Section 4](#) and [Section 5](#) outline several strategies that might aid in ensuring that protocol changes remain possible over time.

The experience that informs this document is predominantly at "higher" layers of the network stack, in protocols that operate at very large scale and Internet-scale applications. It is possible that these conclusions are less applicable to protocol deployments that have less scale and diversity, or operate under different constraints.

## **2. Imperfect Implementations Limit Protocol Evolution**

It can be extremely difficult to deploy a change to a protocol if there are bugs in implementations with which the new deployment needs to interoperate. Bugs in how new codepoints or extensions are handled often mean that endpoints will react poorly to the use of extension mechanisms. This can manifest as abrupt termination of sessions, errors, crashes, or disappearances of endpoints and timeouts.

Interoperability with other implementations is usually highly valued, so deploying mechanisms that trigger adverse reactions can be untenable. Where interoperability is a competitive advantage, this is true even if the negative reactions happen infrequently or only under relatively rare conditions.

Deploying a change to a protocol could require implementations fix a substantial proportion of the bugs that the change exposes. This can involve a difficult process that includes identifying the cause of these errors, finding the responsible implementation(s), coordinating a bug fix and release plan, contacting users and/or the operator of affected services, and waiting for the fix to be deployed.

Given the effort involved in fixing problems, the existence of these sorts of bugs can outright prevent the deployment of some types of

protocol changes, especially for protocols involving multiple parties or that are considered critical infrastructure (e.g., IP, BGP, DNS, or TLS). It could even be necessary to come up with a new protocol design that uses a different method to achieve the same result.

The set of interoperable features in a protocol is often the subset of its features that have some value to those implementing and deploying the protocol. It is not always the case that future extensibility is in that set.

## **2.1. Good Protocol Design is Not Itself Sufficient**

It is often argued that the careful design of a protocol extension point or version negotiation capability is critical to the freedom that it ultimately offers.

RFC 6709 [[EXTENSIBILITY](#)] contains a great deal of well-considered advice on designing for extension. It includes the following advice:

This means that, to be useful, a protocol version-negotiation mechanism should be simple enough that it can reasonably be assumed that all the implementers of the first protocol version at least managed to implement the version-negotiation mechanism correctly.

This has proven to be insufficient in practice. Many protocols have evidence of imperfect implementation of critical mechanisms of this sort. Mechanisms that aren't used are the ones that fail most often. The same paragraph from RFC 6709 acknowledges the existence of this problem, but does not offer any remedy:

The nature of protocol version-negotiation mechanisms is that, by definition, they don't get widespread real-world testing until *after* the base protocol has been deployed for a while, and its deficiencies have become evident.

Indeed, basic interoperability is considered critical early in the deployment of a protocol. A desire to deploy can result in an engineering practice that values simplicity, which could result in deferring implementation of version negotiation and extension mechanisms. This leads to these mechanisms being particularly affected by this problem.

## **2.2. Disuse Can Hide Problems**

There are many examples of extension points in protocols that have been either completely unused, or their use was so infrequent that they could no longer be relied upon to function correctly.

### 2.2.1. TLS

Transport Layer Security (TLS) [[TLS12](#)] provides examples of where a design that is objectively sound fails when incorrectly implemented. TLS provides examples of failures in protocol version negotiation and extensibility.

Version negotiation in TLS 1.2 and earlier uses the "Highest mutually supported version (HMSV)" scheme exactly as it is described in [[EXTENSIBILITY](#)]. However, clients are unable to advertise a new version without causing a non-trivial proportions of sessions to fail due to bugs in server and middlebox implementations.

Intolerance to new TLS versions is so severe [[INTOLERANCE](#)] that TLS 1.3 [[TLS13](#)] has abandoned HMSV version negotiation for a new mechanism.

The server name indication (SNI) [[TLS-EXT](#)] in TLS is another excellent example of the failure of a well-designed extensibility point. SNI uses the same technique for extension that is used with considerable success in other parts of the TLS protocol. The original design of SNI includes the ability to include multiple names of different types.

What is telling in this case is that SNI was defined with just one type of name: a domain name. No other type has ever been standardized, though several have been proposed. Despite an otherwise exemplary design, SNI is so inconsistently implemented that any hope for using the extension point it defines has been abandoned [[SNI](#)].

Even where extension points have multiple valid values, if the set of permitted values does not change over time, there is still a risk that new values are not tolerated by existing implementations. If the set of values for a particular field remains fixed over a long period, some implementations might not correctly handle a new value when it is introduced. For example, implementations of TLS broke when new values of the `signature_algorithms` extension were introduced.

### 2.2.2. DNS

Ossified DNS code bases and systems resulted in fears that new Resource Record Codes (RRCodes) would take years of software propagation before new RRCodes could be used. The result for a long time was heavily overloaded use of the TXT record, such as in the Sender Policy Framework [[SPF](#)]. It wasn't until after the standard mechanism for dealing with new RRCodes [[RRTYPE](#)] was considered widely deployed that new RRCodes can be safely created and used.

### **2.2.3. SNMP**

As a counter example, the first version of the Simple Network Management Protocol (SNMP) [[SNMPv1](#)] defines that unparseable or unauthenticated messages are simply discarded without response:

It then verifies the version number of the SNMP message. If there is a mismatch, it discards the datagram and performs no further actions.

When SNMP versions 2, 2c and 3 came along, older agents did exactly what the protocol specifies. Deployment of new versions was likely successful because the handling of newer versions was both clear and simple.

### **2.2.4. HTTP**

HTTP has a number of very effective extension points in addition to the aforementioned header fields. It also has some examples of extension points that are so rarely used that it is possible that they are not at all usable.

Extension points in HTTP that might be unwise to use include the extension point on each chunk in the chunked transfer coding [Section 7.1](#) of [[HTTP11](#)], the ability to use transfer codings other than the chunked coding, and the range unit in a range request [Section 14](#) of [[HTTP](#)].

### **2.2.5. IPv4**

Codepoints that are reserved for future use can be especially problematic. Reserving codepoints without attributing semantics to their use can result in diverse or conflicting semantics being attributed without any hope of interoperability. An example of this is the "class E" address space in IPv4 [[RFC0988](#)], which was reserved without assigning any semantics.

For protocols that can use negotiation to attribute semantics to codepoints, it is possible that unused codepoints can be reclaimed for active use, though this requires that the negotiation include all protocol participants. For something as fundamental as addressing, negotiation is difficult or even impossible, as all nodes on the network path plus potential alternative paths would need to be involved.

## **2.3. Multi-Party Interactions and Middleboxes**

Even the most superficially simple protocols can often involve more actors than is immediately apparent. A two-party protocol has two ends, but even at the endpoints of an interaction, protocol elements

can be passed on to other entities in ways that can affect protocol operation.

One of the key challenges in deploying new features is ensuring compatibility with all actors that could be involved in the protocol.

Protocols deployed without active measures against intermediation will tend to become intermediated over time, as network operators deploy middleboxes to perform some function on traffic [[PATH-SIGNALS](#)]. In particular, one of the consequences of an unencrypted protocol is that any element on path can interact with the protocol. For example, HTTP was specifically designed with intermediation in mind, transparent proxies [[HTTP](#)] are not only possible but sometimes advantageous, despite some significant downsides. Consequently, transparent proxies for cleartext HTTP are commonplace. The DNS protocol was designed with intermediation in mind through its use of caching recursive resolvers [[DNS](#)]. What was less anticipated was the forced spoofing of DNS records by many middle-boxes such as those that inject authentication or pay-wall mechanisms as an authentication and authorization check, which are now prevalent in hotels, coffee shops and business networks.

Middleboxes are also protocol participants, to the degree that they are able to observe and act in ways that affect the protocol. The degree to which a middlebox participates varies from the basic functions that a router performs to full participation. For example, a SIP back-to-back user agent (B2BUA) [[B2BUA](#)] can be very deeply involved in the SIP protocol.

This phenomenon appears at all layers of the protocol stack, even when protocols are not designed with middlebox participation in mind. TCP's [[TCP](#)] extension points have been rendered difficult to use, largely due to middlebox interactions, as experience with Multipath TCP [[MPTCP](#)] and Fast Open [[TFO](#)] has shown. IP's version field was rendered useless when encapsulated over Ethernet, requiring a new ethertype with IPv6 [[RFC2464](#)], due in part to layer 2 devices making version-independent assumptions about the structure of the IPv4 header. The announcements of new optional transitive attributes in BGP caused significant routing instability [[RIPE-99](#)].

By increasing the number of different actors involved in any single protocol exchange, the number of potential implementation bugs that a deployment needs to contend with also increases. In particular, incompatible changes to a protocol that might be negotiated between endpoints in ignorance of the presence of a middlebox can result in a middlebox interfering in negative and unexpected ways.

Unfortunately, middleboxes can considerably increase the difficulty of deploying new versions or other changes to a protocol.

### 3. Retaining Viable Protocol Evolution Mechanisms

The design of a protocol for extensibility and eventual replacement [[EXTENSIBILITY](#)] does not guarantee the ability to exercise those options. The set of features that enable future evolution need to be interoperable in the first implementations and deployments of the protocol. Implementations of mechanisms that support evolution is necessary to ensure that they remain available for new uses, and history has shown this occurs almost exclusively through active mechanism use.

The conditions for retaining the ability to evolve a design is most clearly evident in the protocols that are known to have viable version negotiation or extension points. The definition of mechanisms alone is insufficient; it's the assured implementation through active use of those mechanisms that determines the existence of freedom. Protocols that routinely add new extensions and code points rarely have trouble adding additional ones, especially when the handling of new versions or extension is well defined.

#### 3.1. Examples of Active Use

For example, header fields in email [[SMTP](#)], HTTP [[HTTP](#)] and SIP [[SIP](#)] all derive from the same basic design, which amounts to a list name/value pairs. There is no evidence of significant barriers to deploying header fields with new names and semantics in email and HTTP as clients and servers can ignore headers they do not understand or need. The widespread deployment of SIP B2BUAs means that new SIP header fields do not reliably reach peers, however, which doesn't necessarily cause interoperability issues but rather causes feature deployment issues due to the lack of option passing [Section 2.3](#).

As another example, the attribute-value pairs (AVPs) in Diameter [[DIAMETER](#)] are fundamental to the design of the protocol. Any use of Diameter requires exercising the ability to add new AVPs. This is routinely done without fear that the new feature might not be successfully deployed.

These examples show extension points that are heavily used are also being relatively unaffected by deployment issues preventing addition of new values for new use cases.

These examples also confirm the case that good design does not guarantee success. On the contrary, success is often despite shortcomings in the design. For instance, the shortcomings of HTTP



header fields are significant enough that there are ongoing efforts to improve the syntax [[HTTP-HEADERS](#)].

Only by using a protocol's extension capabilities does it ensure the availability of that capability. Protocols that fail to use a mechanism, or a protocol that only rarely uses a mechanism, may suffer an inability to rely on that mechanism.

### **3.2. Dependency is Better**

The best way to guarantee that a protocol mechanism is used is to make the handling of it critical to an endpoint participating in that protocol. This means that implementations must rely on both the existence of extension mechanisms and their continued, repeated expansion over time.

For example, the message format in SMTP relies on header fields for most of its functions, including the most basic delivery functions. A deployment of SMTP cannot avoid including an implementation of header field handling. In addition to this, the regularity with which new header fields are defined and used ensures that deployments frequently encounter header fields that it does not yet (and may never) understand. An SMTP implementation therefore needs to be able to both process header fields that it understands and ignore those that it does not.

In this way, implementing the extensibility mechanism is not merely mandated by the specification, it is crucial to the functioning of a protocol deployment. Should an implementation fail to correctly implement the mechanism, that failure would quickly become apparent.

Caution is advised to avoid assuming that building a dependency on an extension mechanism is sufficient to ensure availability of that mechanism in the long term. If the set of possible uses is narrowly constrained and deployments do not change over time, implementations might not see new variations or assume a narrower interpretation of what is possible. Those implementations might still exhibit errors when presented with new variations.

### **3.3. Restoring Active Use**

With enough effort, active use can be used to restore capabilities.

EDNS [[EDNS](#)] was defined to provide extensibility in DNS. Intolerance of the extension in DNS servers resulted in a fallback method being widely deployed (see [Section 6.2.2](#) of [[EDNS](#)]), This fallback resulted in EDNS being disabled for affected servers. Over time, greater support for EDNS and increased reliance on it for different

features motivated a flag day [[DNSFLAGDAY](#)] where the workaround was removed.

The EDNS example shows that effort can be used to restore capabilities. This is in part because EDNS was actively used with most resolvers and servers. It was therefore possible to force a change to ensure that extension capabilities would always be available. However, this required an enormous coordination effort. A small number of incompatible servers and the names they serve also become inaccessible to most clients.

#### **4. Active Use**

As discussed in [Section 3](#), the most effective defense against ossification of protocol extension points is active use.

Implementations are most likely to be tolerant of new values if they depend on being able to frequently use new values. Failing that, implementations that routinely see new values are more likely to correctly handle new values. More frequent changes will improve the likelihood that incorrect handling or intolerance is discovered and rectified. The longer an intolerant implementation is deployed, the more difficult it is to correct.

What constitutes "active use" can depend greatly on the environment in which a protocol is deployed. The frequency of changes necessary to safeguard some mechanisms might be slow enough to attract ossification in another protocol deployment, while being excessive in others.

##### **4.1. Version Negotiation**

As noted in [Section 2.1](#), protocols that provide version negotiation mechanisms might not be able to test that feature until a new version is deployed. One relatively successful design approach has been to use the protocol selection mechanisms built into a lower-layer protocol to select the protocol. This could allow a version negotiation mechanism to benefit from active use of the extension point by other protocols.

For instance, all published versions of IP contain a version number as the four high bits of the first header byte. However, version selection using this field proved to be unsuccessful. Ultimately, successful deployment of IPv6 over Ethernet [[RFC2464](#)] required a different EtherType from IPv4. This change took advantage of the already-diverse usage of EtherType.

Other examples of this style of design include Application-Layer Protocol Negotiation ([\[ALPN\]](#)) and HTTP content negotiation ([Section 12](#) of [\[HTTP\]](#)).

This technique relies on the codepoint being usable. For instance, the IP protocol number is known to be unreliable and therefore not suitable [\[NEW-PROTOCOLS\]](#).

#### 4.2. Falsifying Active Use

"Grease" was originally defined for TLS [\[GREASE\]](#), but has been adopted by other protocols, such as QUIC [\[QUIC\]](#). Grease identifies lack of use as an issue (protocol mechanisms "rusting" shut) and proposes reserving values for extensions that have no semantic value attached.

The design in [\[GREASE\]](#) is aimed at the style of negotiation most used in TLS, where one endpoint offers a set of options and the other chooses the one that it most prefers from those that it supports. An endpoint that uses grease randomly offers options - usually just one - from a set of reserved values. These values are guaranteed to never be assigned real meaning, so its peer will never have cause to genuinely select one of these values.

More generally, greasing is used to refer to any attempt to exercise extension points without changing endpoint behavior, other than to encourage participants to tolerate new or varying values of protocol elements.

The principle that grease operates on is that an implementation that is regularly exposed to unknown values is less likely to be intolerant of new values when they appear. This depends largely on the assumption that the difficulty of implementing the extension mechanism correctly is as easy or easier than implementing code to identify and filter out reserved values. Reserving random or unevenly distributed values for this purpose is thought to further discourage special treatment.

Without reserved greasing codepoints, an implementation can use code points from spaces used for private or experimental use if such a range exists. In addition to the risk of triggering participation in an unwanted experiment, this can be less effective. Incorrect implementations might still be able to identify these code points and ignore them.

In addition to advertising bogus capabilities, an endpoint might also selectively disable non-critical protocol elements to test the ability of peers to handle the absence of certain capabilities.

This style of defensive design is limited because it is only superficial. As greasing only mimics active use of an extension point, it only exercises a small part of the mechanisms that support extensibility. More critically, it does not easily translate to all forms of extension points. For instance, HMSV negotiation cannot be

greased in this fashion. Other techniques might be necessary for protocols that don't rely on the particular style of exchange that is predominant in TLS.

Grease is deployed with the intent of quickly revealing errors in implementing the mechanisms it safeguards. Though it has been effective at revealing problems in some cases with TLS, the efficacy of greasing isn't proven more generally. Where implementations are able to tolerate a non-zero error rate in their operation, greasing offers a potential option for safeguarding future extensibility. However, this relies on there being a sufficient proportion of participants that are willing to invest the effort and tolerate the risk of interoperability failures.

## **5. Complementary Techniques**

The protections to protocol evolution that come from [active use](#) ([Section 4](#)) can be improved through the use of other defensive techniques. The techniques listed here might not prevent ossification on their own, but can make active use more effective.

### **5.1. Cryptography**

Cryptography can be used to reduce the number of middlebox entities that can participate in a protocol or limit the extent of participation. Using TLS or other cryptographic tools can therefore reduce the number of entities that can influence whether new features are usable.

[\[PATH-SIGNALS\]](#) recommends the use of encryption and integrity protection to limit participation. For example, encryption is used by the QUIC protocol [\[QUIC\]](#) to limit the information that is available to middleboxes and integrity protection prevents modification.

### **5.2. Fewer Extension Points**

A successful protocol will include many potential types of extension. Designing multiple types of extension mechanism, each suited to a specific purpose, might leave some extension points less heavily used than others.

Disuse of a specialized extension point might render it unusable. In contrast, having a smaller number of extension points with wide applicability could improve the use of those extension points. Use of a shared extension point for any purpose can protect rarer or more specialized uses.

Both extensions and core protocol elements use the same extension points in protocols like HTTP [[HTTP](#)] and DIAMETER [[DIAMETER](#)]; see [Section 3.1](#).

#### 5.2.1. Invariants

Documenting aspects of the protocol that cannot or will not change as extensions or new versions are added can be a useful exercise. Understanding what aspects of a protocol are invariant can help guide the process of identifying those parts of the protocol that might change.

As a means of protecting extensibility, a declaration of protocol invariants is useful only to the extent that protocol participants are willing to allow new uses for the protocol. Like with greasing, protocol participants could still purposefully block the deployment of new features. A protocol that declares protocol invariants relies on implementations understanding and respecting those invariants.

Protocol invariants need to be clearly and concisely documented. Including examples of aspects of the protocol that are not invariant, such as the appendix of [[QUIC-INVARIANTS](#)], can be used to clarify intent.

#### 5.3. Effective Feedback

While not a direct means of protecting extensibility mechanisms, feedback systems can be important to discovering problems.

Visibility of errors is critical to the success of techniques like grease (see [Section 4.2](#)). The grease design is most effective if a deployment has a means of detecting and reporting errors. Ignoring errors could allow problems to become entrenched.

Feedback on errors is more important during the development and early deployment of a change. It might also be helpful to disable automatic error recovery methods during development.

Automated feedback systems are important for automated systems, or where error recovery is also automated. For instance, connection failures with HTTP alternative services [[ALT-SVC](#)] are not permitted to affect the outcome of transactions. An automated feedback system for capturing failures in alternative services is therefore necessary for failures to be detected.

How errors are gathered and reported will depend greatly on the nature of the protocol deployment and the entity that receives the report. For instance, end users, developers, and network operations each have different requirements for how error reports are created, managed, and acted upon.

Automated delivery of error reports can be critical for rectifying deployment errors as early as possible, such as seen in [DMARC] and [SMTP-TLS-Reporting].

## 6. Security Considerations

Many of the problems identified in this document are not the result of deliberate actions by an adversary, but more the result of mistakes, decisions made without sufficient context, or simple neglect. Problems therefore not the result of opposition by an adversary. In response, the recommended measures generally assume that other protocol participants will not take deliberate action to prevent protocol evolution.

The use of cryptographic techniques to exclude potential participants is the only strong measure that the document recommends. However, authorized protocol peers are most often responsible for the identified problems, which can mean that cryptography is insufficient to exclude them.

The ability to design, implement, and deploy new protocol mechanisms can be critical to security. In particular, it is important to be able to replace cryptographic algorithms over time [AGILITY]. For example, preparing for replacement of weak hash algorithms was made more difficult through misuse [HASH].

## 7. IANA Considerations

This document makes no request of IANA.

## 8. Informative References

[AGILITY] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<https://www.rfc-editor.org/rfc/rfc7696>>.

[ALPN] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, DOI 10.17487/RFC7301, July 2014, <<https://www.rfc-editor.org/rfc/rfc7301>>.

[ALT-SVC] Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", RFC 7838, DOI 10.17487/RFC7838, April 2016, <<https://www.rfc-editor.org/rfc/rfc7838>>.

[B2BUA] Kaplan, H. and V. Pascual, "A Taxonomy of Session Initiation Protocol (SIP) Back-to-Back User Agents", RFC

7092, DOI 10.17487/RFC7092, December 2013, <<https://www.rfc-editor.org/rfc/rfc7092>>.

**[DIAMETER]** Fajardo, V., Ed., Arkko, J., Loughney, J., and G. Zorn, Ed., "Diameter Base Protocol", RFC 6733, DOI 10.17487/RFC6733, October 2012, <<https://www.rfc-editor.org/rfc/rfc6733>>.

**[DMARC]** Kucherawy, M., Ed. and E. Zwicky, Ed., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)", RFC 7489, DOI 10.17487/RFC7489, March 2015, <<https://www.rfc-editor.org/rfc/rfc7489>>.

**[DNS]** Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/rfc/rfc1034>>.

**[DNSFLAGDAY]** "DNS Flag Day 2019", May 2019, <<https://dnsflagday.net/2019/>>.

**[EDNS]** Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <<https://www.rfc-editor.org/rfc/rfc6891>>.

**[EXTENSIBILITY]** Carpenter, B., Aboba, B., Ed., and S. Cheshire, "Design Considerations for Protocol Extensions", RFC 6709, DOI 10.17487/RFC6709, September 2012, <<https://www.rfc-editor.org/rfc/rfc6709>>.

**[GREASE]** Benjamin, D., "Applying Generate Random Extensions And Sustain Extensibility (GREASE) to TLS Extensibility", RFC 8701, DOI 10.17487/RFC8701, January 2020, <<https://www.rfc-editor.org/rfc/rfc8701>>.

**[HASH]** Bellovin, S. and E. Rescorla, "Deploying a New Hash Algorithm", Proceedings of NDSS '06 , 2006, <<https://www.cs.columbia.edu/~smb/papers/new-hash.pdf>>.

**[HTTP]** Fielding, R. T., Nottingham, M., and J. Reschke, "HTTP Semantics", Work in Progress, Internet-Draft, draft-ietf-httpbis-semantics-16, 27 May 2021, <<https://>

[datatracker.ietf.org/doc/html/draft-ietf-httpbis-semantics-16](https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-semantics-16)>.

**[HTTP-HEADERS]** Nottingham, M. and P-H. Kamp, "Structured Field Values for HTTP", RFC 8941, DOI 10.17487/RFC8941, February 2021, <<https://www.rfc-editor.org/rfc/rfc8941>>.

**[HTTP11]** Fielding, R. T., Nottingham, M., and J. Reschke, "HTTP/1.1", Work in Progress, Internet-Draft, draft-ietf-httpbis-messaging-16, 27 May 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-messaging-16>>.

**[INTOLERANCE]** Kario, H., "Re: [TLS] Thoughts on Version Intolerance", 20 July 2016, <<https://mailarchive.ietf.org/arch/msg/tls/b0J2JQc3HjAHFFWCiNTIb0JuMZc>>.

**[MPTCP]** Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <<https://www.rfc-editor.org/rfc/rfc6824>>.

**[NEW-PROTOCOLS]** Barik, R., Welzl, M., Fairhurst, G., Elmokashfi, A., Dreibholz, T., and S. Gjessing, "On the usability of transport protocols other than TCP: A home gateway and internet path traversal study", Computer Networks Vol. 173, pp. 107211, DOI 10.1016/j.comnet.2020.107211, May 2020, <<https://doi.org/10.1016/j.comnet.2020.107211>>.

**[PATH-SIGNALS]** Wendt, C. and M. Barnes, "Personal Assertion Token (PaSSporT) Extension for Signature-based Handling of Asserted information using toKENS (SHAKEN)", RFC 8588, DOI 10.17487/RFC8588, May 2019, <<https://www.rfc-editor.org/rfc/rfc8588>>.

**[QUIC]** Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI



10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/rfc/rfc9000>>.

**[QUIC-INVARIANTS]** Thomson, M., "Version-Independent Properties of QUIC", RFC 8999, DOI 10.17487/RFC8999, May 2021, <<https://www.rfc-editor.org/rfc/rfc8999>>.

**[RFC0988]** Deering, S., "Host extensions for IP multicasting", RFC 988, DOI 10.17487/RFC0988, July 1986, <<https://www.rfc-editor.org/rfc/rfc988>>.

**[RFC2464]** Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/rfc/rfc2464>>.

**[RIPE-99]** Romijn, E., "RIPE NCC and Duke University BGP Experiment", 27 August 2010, <<https://labs.ripe.net/Members/erik/ripe-ncc-and-duke-university-bgp-experiment/>>.

**[RRTYPE]** Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", RFC 3597, DOI 10.17487/RFC3597, September 2003, <<https://www.rfc-editor.org/rfc/rfc3597>>.

**[SIP]**

Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<https://www.rfc-editor.org/rfc/rfc3261>>.

**[SMTP]** Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <<https://www.rfc-editor.org/rfc/rfc5322>>.

**[SMTP-TLS-Reporting]** Margolis, D., Brotman, A., Ramakrishnan, B., Jones, J., and M. Risher, "SMTP TLS Reporting", RFC 8460, DOI 10.17487/RFC8460, September 2018, <<https://www.rfc-editor.org/rfc/rfc8460>>.

**[SNI]** Langley, A., "Accepting that other SNI name types will never work", 3 March 2016, <[https://mailarchive.ietf.org/arch/msg/tls/1t79gzNIItZd71DwwoaqcQQ\\_4Yxc](https://mailarchive.ietf.org/arch/msg/tls/1t79gzNIItZd71DwwoaqcQQ_4Yxc)>.

**[SNMPv1]** Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol (SNMP)", RFC 1157,

DOI 10.17487/RFC1157, May 1990, <<https://www.rfc-editor.org/rfc/rfc1157>>.

**[SPF]** Kitterman, S., "Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1", RFC 7208, DOI 10.17487/RFC7208, April 2014, <<https://www.rfc-editor.org/rfc/rfc7208>>.

**[SUCCESS]** Thaler, D. and B. Aboba, "What Makes for a Successful Protocol?", RFC 5218, DOI 10.17487/RFC5218, July 2008, <<https://www.rfc-editor.org/rfc/rfc5218>>.

**[TCP]** Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/rfc/rfc793>>.

**[TFO]** Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/rfc/rfc7413>>.

**[TLS-EXT]** Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, DOI 10.17487/RFC6066, January 2011, <<https://www.rfc-editor.org/rfc/rfc6066>>.

**[TLS12]** Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/rfc/rfc5246>>.

**[TLS13]** Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.

**[TRANSITIONS]** Thaler, D., Ed., "Planning for Protocol Adoption and Subsequent Transitions", RFC 8170, DOI 10.17487/RFC8170, May 2017, <<https://www.rfc-editor.org/rfc/rfc8170>>.

## Acknowledgments

Wes Hardaker, Mirja Kuehlewind, Mark Nottingham, and Brian Trammell made significant contributions to this document.

## Authors' Addresses

Martin Thomson  
Mozilla

Email: [mt@lowentropy.net](mailto:mt@lowentropy.net)

Tommy Pauly  
Apple

Email: [tpauly@apple.com](mailto:tpauly@apple.com)