

---

Stream: Internet Architecture Board (IAB)  
RFC: [9170](#)  
Category: Informational  
Published: December 2021  
ISSN: 2070-1721  
Authors: M. Thomson T. Pauly

# RFC 9170

## 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. This document explores how regular use of new protocol features can ensure that it remains possible to deploy changes to a protocol. Examples are given where lack of use caused changes to be more difficult or costly.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. It represents the consensus of the Internet Architecture Board (IAB). Documents approved for publication by the IAB are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9170>.

### 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.

## 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
  - 2.3. Multi-party Interactions and Middleboxes
3. Active Use
  - 3.1. Dependency Is Better
  - 3.2. Version Negotiation
  - 3.3. Falsifying Active Use
  - 3.4. Examples of Active Use
  - 3.5. Restoring Active Use
4. Complementary Techniques
  - 4.1. Fewer Extension Points
  - 4.2. Invariants
  - 4.3. Limiting Participation
  - 4.4. Effective Feedback
5. Security Considerations
6. IANA Considerations
7. Informative References
- Appendix A. Examples
  - A.1. DNS
  - A.2. HTTP
  - A.3. IP
  - A.4. SNMP
  - A.5. TCP
  - A.6. TLS
- IAB Members at the Time of Approval
- Acknowledgments

[Authors' Addresses](#)

## 1. Introduction

A successful protocol [SUCCESS] needs to change in ways that allow it to continue to fulfill the changing 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.

An extension point is a mechanism that allows a protocol to be changed or enhanced. This document examines the specific conditions that determine whether protocol maintainers have the ability to design and deploy new or modified protocols via their specified extension points. [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 describes how successful protocols make frequent use of new extensions and code points. [Section 4](#) outlines several additional 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 with limited numbers of participants. Though similar issues are present in many protocols that operate at scale, the trade-offs involved with applying some of the suggested techniques can be more complex when there are many participants, such as at the network layer or in routing systems.

## 2. Imperfect Implementations Limit Protocol Evolution

It can be extremely difficult to deploy a change to a protocol if implementations with which the new deployment needs to interoperate do not operate predictably. Variation in how new code points or extensions are handled can be the result of bugs in implementation or specifications. Unpredictability can manifest as errors, crashes, timeouts, abrupt termination of sessions, or disappearances of endpoints.

The risk of interoperability problems can in turn make it infeasible to deploy certain protocol changes. If deploying a new code point or extension makes an implementation less reliable than others, even if only in rare cases, it is far less likely that implementations will adopt the change.

Deploying a change to a protocol could require implementations to 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.

This document only addresses cases where extensions are not deliberately blocked. Some deployments or implementations apply policies that explicitly prohibit the use of unknown capabilities. This is especially true of functions that seek to make security guarantees, like firewalls.

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 extensions. 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.

There are a number of protocols for which this has proven to be insufficient in practice. These protocols have imperfect implementations of these mechanisms. 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 early focus on a reduced feature set, 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.

[Appendix A](#) includes examples of disuse in a number of widely deployed Internet protocols.

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 of a protocol or the order in which these values appear 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.

## 2.3. Multi-party Interactions and Middleboxes

One of the key challenges in deploying new features is ensuring compatibility with all actors that could be involved in the protocol. Even the most superficially simple protocols can often involve more actors than is immediately apparent.

The design of extension points needs to consider what actions middleboxes might take in response to a protocol change as well as the effect those actions could have on the operation of the protocol.

Deployments of protocol extensions also need to consider the impact of the changes on entities beyond protocol participants and middleboxes. Protocol changes can affect the behavior of applications or systems that don't directly interact with the protocol, such as when a protocol change modifies the formatting of data delivered to an application.

## 3. Active Use

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. Implementation 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.

Only by using the extension capabilities of a protocol is the availability of that capability assured. "Using" here includes specifying, implementing, and deploying capabilities that rely on the extension capability. Protocols that fail to use a mechanism, or a protocol that only rarely uses a mechanism, could lead to that mechanism being unreliable.

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.

Protocols that routinely add new extensions and code points rarely have trouble adding additional ones especially when the handling of new versions or extensions are well defined. The definition of mechanisms alone is insufficient; it is the assured implementation and active use of those mechanisms that determines their availability.

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.

### **3.1. Dependency Is Better**

The easiest 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 they do 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.

### 3.2. 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 code point being usable. For instance, the IP protocol number is known to be unreliable and therefore not suitable [[NEW-PROTOCOLS](#)].

### 3.3. 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 code points, 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 noncritical 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, highest mutually supported version (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.

### 3.4. Examples of Active Use

Header fields in email [[SMTP](#)], HTTP [[HTTP](#)], and SIP [[SIP](#)] all derive from the same basic design, which amounts to a list of 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 generally ignore headers they do not understand or need. The widespread deployment of SIP back-to-back user agents (B2BUAs), which generally do not ignore unknown fields, means that new SIP header fields do not reliably reach peers. This does not necessarily cause interoperability issues in SIP but rather causes features to remain unavailable until the B2BUA is updated. All three protocols are still able to deploy new features reliably, but SIP features are deployed more slowly due to the larger number of active participants that need to support new features.

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 show that a good design is not required for 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](#)].

### 3.5. Restoring Active Use

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

Extension Mechanisms for DNS ([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 became inaccessible to most clients.

## 4. Complementary Techniques

The protections to protocol evolution that come from [active use](#) ([Section 3](#)) can be improved through the use of other defensive techniques. The techniques listed here might not prevent ossification on their own, but they can make active use more effective.

### 4.1. Fewer Extension Points

A successful protocol will include many potential types of extensions. Designing multiple types of extension mechanisms, 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 3.4](#).

### 4.2. Invariants

Documenting aspects of the protocol that cannot or will not change as extensions or new versions are added can be a useful exercise. [Section 2.2](#) of [[RFC5704](#)] defines invariants as:

Invariants are core properties that are consistent across the network and do not change over extremely long time-scales.

Understanding what aspects of a protocol are invariant can help guide the process of identifying those parts of the protocol that might change. [[QUIC-INVARIANTS](#)] and [Section 9.3](#) of [[TLS13](#)] are both examples of documented invariants.

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. A protocol that declares protocol invariants relies on implementations understanding and respecting those invariants. If active use is not possible for all non-invariant parts of the protocol, greasing ([Section 3.3](#)) might be used to improve the chance that invariants are respected.

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

### 4.3. Limiting Participation

Reducing the number of entities that can participate in a protocol or limiting the extent of participation can reduce the number of entities that might affect extensibility. Using TLS or other cryptographic tools can therefore reduce the number of entities that can influence whether new features are usable.

[\[PATH-SIGNALS\]](#) also 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.

### 4.4. Effective Feedback

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

The visibility of errors is critical to the success of techniques like grease (see [Section 3.3](#)). 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, as seen in [\[DMARC\]](#) and [\[SMTP-TLS-REPORTING\]](#).

## 5. 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, i.e., 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 the replacement of weak hash algorithms was made more difficult through misuse [HASH].

## 6. IANA Considerations

This document has no IANA actions.

## 7. 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/info/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/info/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/info/rfc7838>>.
- [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/info/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/info/rfc7489>>.
- [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/info/rfc6891>>.
- [EXT-TCP]** Honda, M., Nishida, Y., Raiciu, C., Greenhalgh, A., Handley, M., and H. Tokuda, "Is it still possible to extend TCP?", IMC '11: Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference, DOI 10.1145/2068816.2068834, November 2011, <<https://doi.org/10.1145/2068816.2068834>>.
- [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/info/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/info/rfc8701>>.
- [HASH]** Bellovin, S. and E. Rescorla, "Deploying a New Hash Algorithm", Proceedings of NDSS, 2006, <<https://www.cs.columbia.edu/~smb/papers/new-hash.pdf>>.
- [HTTP]** Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", Work in Progress, Internet-Draft, draft-ietf-httpbis-semantics-19, September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-semantics-19>>.
- [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/info/rfc8941>>.
- [HTTP11]** Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1", Work in Progress, Internet-Draft, draft-ietf-httpbis-messaging-19, September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-messaging-19>>.
- [INTOLERANCE]** Kario, H., "Re: [TLS] Thoughts on Version Intolerance", July 2016, <<https://mailarchive.ietf.org/arch/msg/tls/bOJ2JQc3HjAHFFWCiNTIb0JuMZc>>.
- [MPTCP]** Ford, A., Raiciu, C., Handley, M., Bonaventure, O., and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 8684, DOI 10.17487/RFC8684, March 2020, <<https://www.rfc-editor.org/info/rfc8684>>.
- [MPTCP-HOW-HARD]** Raiciu, C., Paasch, C., Barre, S., Ford, A., Honda, M., Duchene, F., Bonaventure, O., and M. Handley, "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP", April 2012, <<https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/raiciu>>.
- [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]** Hardie, T., Ed., "Transport Protocol Path Signals", RFC 8558, DOI 10.17487/RFC8558, April 2019, <<https://www.rfc-editor.org/info/rfc8558>>.
- [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/info/rfc9000>>.
- [QUIC-INVARIANTS]** Thomson, M., "Version-Independent Properties of QUIC", RFC 8999, DOI 10.17487/RFC8999, May 2021, <<https://www.rfc-editor.org/info/rfc8999>>.
- [RAv4]** Katz, D., "IP Router Alert Option", RFC 2113, DOI 10.17487/RFC2113, February 1997, <<https://www.rfc-editor.org/info/rfc2113>>.
- [RAv6]** Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, DOI 10.17487/RFC2711, October 1999, <<https://www.rfc-editor.org/info/rfc2711>>.
- [RFC0791]** Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC1112]** Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.
- [RFC2464]** Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.
- [RFC5704]** Bryant, S., Ed., Morrow, M., Ed., and IAB, "Uncoordinated Protocol Development Considered Harmful", RFC 5704, DOI 10.17487/RFC5704, November 2009, <<https://www.rfc-editor.org/info/rfc5704>>.
- [RRTYPE]** Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", RFC 3597, DOI 10.17487/RFC3597, September 2003, <<https://www.rfc-editor.org/info/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/info/rfc3261>>.
- [SMTP]** Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <<https://www.rfc-editor.org/info/rfc5321>>.
- [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/info/rfc8460>>.
- [SNI]** Langley, A., "[TLS] Accepting that other SNI name types will never work.", March 2016, <[https://mailarchive.ietf.org/arch/msg/tls/1t79gzNItZd71DwwoaqcQQ\\_4Yxc](https://mailarchive.ietf.org/arch/msg/tls/1t79gzNItZd71DwwoaqcQQ_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/info/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/info/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/info/rfc5218>>.
- [TCP]** Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/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/info/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/info/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/info/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/info/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/info/rfc8170>>.

## Appendix A. Examples

This appendix contains a brief study of problems in a range of Internet protocols at different layers of the stack.

### A.1. DNS

Ossified DNS code bases and systems resulted in new Resource Record Codes (RRCodes) being unusable. A new code point would take years of coordination between implementations and deployments before it could be relied upon. Consequently, use of the TXT record was overloaded in order to avoid the effort and delays involved in allocating new code points; this approach was used in the Sender Policy Framework [SPF] and other protocols.

It was not until after the standard mechanism for dealing with new RRCodes [RRTYPE] was considered widely deployed that new RRCodes could be safely created and used.

### A.2. 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 [\[HTTP\]](#)).

### A.3. IP

The version field in IP 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.

Protocol identifiers or code points that are reserved for future use can be especially problematic. Reserving values 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 224/3 address space in IPv4 that [[RFC0791](#)] reserved without assigning any semantics. [[RFC1112](#)] partially reclaimed that reserved address space for use in multicast (224/4), but the remaining address space (240/4) has not been successfully reclaimed for any purpose.

For protocols that can use negotiation to attribute semantics to values, it is possible that unused code points 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.

IP Router Alerts [[RAv4](#)][[RAv6](#)] use IP options or extension headers to indicate that data is intended for consumption by the next-hop router rather than the addressed destination. In part, the deployment of router alerts was unsuccessful due to the realities of processing IP packets at line rates, combined with bad assumptions in the protocol design about these performance constraints. However, this was not exclusively down to design problems or bugs, as the capability was also deliberately blocked at some routers.

### A.4. SNMP

As a counter example, the first version of the Simple Network Management Protocol (SNMP) [[SNMPv1](#)] states 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.

## A.5. TCP

Extension points in TCP [TCP] have been rendered difficult to use largely due to middlebox interactions; see [EXT-TCP].

For instance, multipath TCP ([MPTCP]) can only be deployed opportunistically; see [MPTCP-HOW-HARD]. Since MPTCP is a protocol enhancement that doesn't impair the connection if it is blocked, network path intolerance of the extension only results in the multipath functionality becoming unavailable.

In comparison, the deployment of TCP Fast Open ([TFO]) critically depends on extension capability being widely available. Though very few network paths were intolerant of the extension in absolute terms, TCP Fast Open could not be deployed as a result.

## A.6. 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 proportion 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] 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 extensions that is used successfully in other parts of the TLS protocol. The original design of SNI anticipated the ability to include multiple names of different types.

SNI was originally 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].

## IAB Members at the Time of Approval

Internet Architecture Board members at the time this document was approved for publication were:

Jari Arkko  
Deborah Brungard  
Ben Campbell

Lars Eggert  
Wes Hardaker  
Cullen Jennings  
Mirja Kühlewind  
Zhenbin Li  
Jared Mauch  
Tommy Pauly  
David Schinazi  
Russ White  
Jiankang Yao

## Acknowledgments

Toerless Eckert, Wes Hardaker, Mirja Kühlewind, Eliot Lear, Mark Nottingham, and Brian Trammell made significant contributions to this document.

## Authors' Addresses

### **Martin Thomson**

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

### **Tommy Pauly**

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