

HTTP
Internet-Draft
Obsoletes: [3205](#) (if approved)
Intended status: Best Current Practice
Expires: October 4, 2018

M. Nottingham
April 2, 2018

On the use of HTTP as a Substrate
draft-ietf-httpbis-bcp56bis-03

Abstract

HTTP is often used as a substrate for other application protocols. This document specifies best practices for these protocols' use of HTTP.

This document obsoletes [RFC 3205](#).

Note to Readers

Discussion of this draft takes place on the HTTP working group mailing list (ietf-http-wg@w3.org), which is archived at <https://lists.w3.org/Archives/Public/ietf-http-wg/> [1].

Working Group information can be found at <http://httpwg.github.io/> [2]; source code and issues list for this draft can be found at <https://github.com/httpwg/http-extensions/labels/bcp56bis> [3].

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 October 4, 2018.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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	3
1.1.	Notational Conventions	4
2.	Is HTTP Being Used?	4
3.	What's Important About HTTP	5
3.1.	Generic Semantics	5
3.2.	Links	6
3.3.	Rich Functionality	7
4.	Best Practices for Using HTTP	7
4.1.	Specifying the Use of HTTP	8
4.2.	Defining HTTP Resources	8
4.3.	Specifying Client Behaviours	9
4.4.	HTTP URLs	10
4.4.1.	Initial URL Discovery	11
4.4.2.	URL Schemes	11
4.4.3.	Transport Ports	12
4.5.	HTTP Methods	12
4.5.1.	GET	13
4.6.	HTTP Status Codes	14
4.7.	HTTP Header Fields	15
4.8.	Defining Message Payloads	16
4.9.	HTTP Caching	16
4.10.	Application State	17
4.11.	Client Authentication	17
4.12.	Co-Existing with Web Browsing	18
4.13.	Application Boundaries	19
5.	IANA Considerations	20
6.	Security Considerations	20
7.	References	20
7.1.	Normative References	20
7.2.	Informative References	22
7.3.	URIs	25

Nottingham

Expires October 4, 2018

[Page 2]

Appendix A . Changes from RFC 3205	25
Author's Address	25

1. Introduction

HTTP [[RFC7230](#)] is often used as a substrate for other application protocols. This is done for a variety of reasons, including:

- o familiarity by implementers, specifiers, administrators, developers and users,
- o availability of a variety of client, server and proxy implementations,
- o ease of use,
- o ubiquity of Web browsers,
- o reuse of existing mechanisms like authentication and encryption,
- o presence of HTTP servers and clients in target deployments, and
- o its ability to traverse firewalls.

In many cases, these protocols are ad hoc; they are intended for only deployment on the server side, and consumption by a limited set of clients. A body of practices and tools has arisen around defining such "HTTP APIs" that favours these conditions.

However, when such a protocol is standardised, it is typically deployed on multiple servers, implemented a number of times, and might be consumed by a broader variety of clients. Such diversity brings a different set of concerns, and tools and practices intended for a single-server deployment might not be suitable.

In particular, standards-defined HTTP APIs need to more carefully consider how extensibility and evolution will be handled, how different deployment requirements will be accommodated, and how clients will evolve with the API.

At the same time, the Internet community has a tradition of protocol reuse (e.g., Telnet [[RFC0854](#)] as a substrate for FTP [[RFC0959](#)] and SMTP [[RFC2821](#)]), but less experience using HTTP as a substrate. Because HTTP is extensible in many ways, a number of questions arise, such as:

- o Should an application using HTTP define a new URL scheme? Use new ports?

Nottingham

Expires October 4, 2018

[Page 3]

- o Should it use standard HTTP methods and status codes, or define new ones?
- o How can the maximum value be extracted from the use of HTTP?
- o How does it coexist with other uses of HTTP - especially Web browsing?
- o How can interoperability problems and "protocol dead ends" be avoided?

This document contains best current practices regarding the use of HTTP by applications other than Web browsing. [Section 2](#) defines what applications it applies to; [Section 3](#) surveys the properties of HTTP that are important to preserve, and [Section 4](#) conveys best practices for those applications that do use HTTP.

It is written primarily to guide IETF efforts to define application protocols using HTTP for deployment on the Internet, but might be applicable in other situations. Note that the requirements herein do not necessarily apply to the development of generic HTTP extensions.

[1.1.](#) Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

[2.](#) Is HTTP Being Used?

Different applications have different goals when using HTTP. In this document, we say an application is using HTTP when any of the following conditions are true:

- o The transport port in use is 80 or 443,
- o The URL scheme "http" or "https" is used,
- o The ALPN protocol ID [[RFC7301](#)] generically identifies HTTP (e.g., "http/1.1", "h2", "h2c"), or
- o The IANA registries defined for HTTP are updated or modified.

When an application is using HTTP, all of the requirements of the HTTP protocol suite (including but not limited to [[RFC7230](#)],

Nottingham

Expires October 4, 2018

[Page 4]

[[RFC7231](#)], [[RFC7232](#)], [[RFC7233](#)], [[RFC7234](#)], [[RFC7235](#)] and [[RFC7540](#)]) are in force.

An application might not be `_using HTTP_` according to this definition, but still relying upon the HTTP specifications in some manner. For example, an application might wish to avoid re-specifying parts of the message format, but change others; or, it might want to use a different set of methods.

Such applications are referred to as `_protocols based upon HTTP_` in this document. These have more freedom to modify protocol operation, but are also likely to lose at least a portion of the benefits outlined above, as most HTTP implementations won't be easily adaptable to these changes, and as the protocol diverges from HTTP, the benefit of mindshare will be lost.

Protocols that are based upon HTTP MUST NOT reuse HTTP's URL schemes, transport ports, ALPN protocol IDs or IANA registries; rather, they are encouraged to establish their own.

3. What's Important About HTTP

There are many ways that applications using HTTP are defined and deployed, and sometimes they are brought to the IETF for standardisation. In that process, what might be workable for deployment in a limited fashion isn't appropriate for standardisation and the corresponding broader deployment.

This section examines the facets of the protocol that are important to preserve in these situations.

3.1. Generic Semantics

When writing an application's specification, it's often tempting to specify exactly how HTTP is to be implemented, supported and used.

However, this can easily lead to an unintended profile of HTTP's behaviour. For example, it's common to see specifications with language like this:

A ``200 OK`` response means that the widget has successfully been updated.

This sort of specification is bad practice, because it is adding new semantics to HTTP's status codes and methods, respectively; a recipient - whether it's an origin server, client library, intermediary or cache - now has to know these extra semantics to understand the message.

Some applications even require specific behaviours, such as:

A ``POST`` request MUST result in a ``201 Created`` response.

This forms an expectation in the client that the response will always be "201 Created", when in fact there are a number of reasons why the status code might differ in a real deployment. If the client does not anticipate this, the application's deployment is brittle.

Much of the value of HTTP is in its `_generic semantics_` - that is, the protocol elements defined by HTTP are potentially applicable to every resource, not specific to a particular context. Application-specific semantics are expressed in the payload; mostly, in the body, but also in header fields.

This allows a HTTP message to be examined by generic HTTP software (e.g., HTTP servers, intermediaries, client implementations), and its handling to be correctly determined. It also allows people to leverage their knowledge of HTTP semantics without special-casing them for a particular application.

Therefore, applications that use HTTP MUST NOT re-define, refine or overlay the semantics of defined protocol elements. Instead, they should focus their specifications on protocol elements that are specific to that application; namely their HTTP resources.

See [Section 4.2](#) for details.

[3.2.](#) Links

Another common practice is assuming that the HTTP server's name space (or a portion thereof) is exclusively for the use of a single application. This effectively overlays special, application-specific semantics onto that space, precludes other applications from using it.

As explained in [[RFC7320](#)], such "squatting" on a part of the URL space by a standard usurps the server's authority over its own resources, can cause deployment issues, and is therefore bad practice in standards.

Instead of statically defining URL components like paths, it is RECOMMENDED that applications using HTTP define links in payloads, to allow flexibility in deployment.

Using runtime links in this fashion has a number of other benefits. For example, navigating with a link allows a request to be routed to a different server without the overhead of a redirection, thereby

Nottingham

Expires October 4, 2018

[Page 6]

supporting deployment across machines well. It becomes possible to "mix" different applications on the same server, and offers a natural path for extensibility, versioning and capability management.

3.3. Rich Functionality

The simplest possible use of HTTP is to POST data to a single URL, thereby effectively tunnelling through the protocol.

This "RPC" style of communication does get some benefit from using HTTP - namely, message framing and the availability of implementations - but fails to realise many others when used exclusively:

- o Caching for server scalability, latency and bandwidth reduction, and reliability;
- o Granularity of access control (through use of a rich space of URLs);
- o Partial content to selectively request part of a response;
- o Definition of an information space using URLs; and
- o The ability to interact with the application easily using a Web browser.

Using such a high-level protocol to tunnel simple semantics has downsides too; because of its more advanced capabilities, breadth of deployment and age, HTTP's complexity can cause interoperability problems that could be avoided by using a simpler substrate (e.g., WebSockets [\[RFC6455\]](#), if browser support is necessary, or TCP [\[RFC0793\]](#) if not), or making the application be _based upon HTTP_, instead of using it (as defined in [Section 2](#)).

Applications that use HTTP are encouraged to accommodate the various features that the protocol offers, so that their users receive the maximum benefit from it. This document does not require specific features to be used, since the appropriate design tradeoffs are highly specific to a given situation. However, following the practices in [Section 4](#) will help make them available.

4. Best Practices for Using HTTP

This section contains best practices regarding the use of HTTP by applications, including practices for specific HTTP protocol elements.

4.1. Specifying the Use of HTTP

When specifying the use of HTTP, an application SHOULD use [[RFC7230](#)] as the primary reference; it is not necessary to reference all of the specifications in the HTTP suite unless there are specific reasons to do so (e.g., a particular feature is called out).

Applications using HTTP MAY specify a minimum version to be supported (HTTP/1.1 is suggested), and MUST NOT specify a maximum version, to preserve the protocol's ability to evolve.

Likewise, applications need not specify what HTTP mechanisms - such as redirection, caching, authentication, proxy authentication, and so on - are to be supported. For example, an application can specify that it uses HTTP like this:

Foo Application uses HTTP [[RFC7230](#)]. Implementations MUST support HTTP/1.1, and MAY support later versions.

When specifying examples of protocol interactions, applications SHOULD document both the request and response messages, with full headers, preferably in HTTP/1.1 format. For example:

```
GET /thing HTTP/1.1
Host: example.com
Accept: application/things+json
User-Agent: Foo/1.0

HTTP/1.1 200 OK
Content-Type: application/things+json
Content-Length: 500
Server: Bar/2.2

[payload here]
```

4.2. Defining HTTP Resources

Applications that use HTTP should focus on defining the following application-specific protocol elements:

- o Media types [[RFC6838](#)], often based upon a format convention such as JSON [[RFC8259](#)],
- o HTTP header fields, as per [Section 4.7](#), and
- o The behaviour of resources, as identified by link relations [[RFC8288](#)].

By composing these protocol elements, an application can define a set of resources, identified by link relations, that implement specified behaviours, including:

- o Retrieval of their state using GET, in one or more formats identified by media type;
- o Resource creation or update using POST or PUT, with an appropriately identified request body format;
- o Data processing using POST and identified request and response body format(s); and
- o Resource deletion using DELETE.

For example, an application might specify:

Resources linked to with the "example-widget" link relation type are Widgets. The state of a Widget can be fetched in the "application/example-widget+json" format, and can be updated by PUT to the same link. Widget resources can be deleted.

The "Example-Count" response header field on Widget representations indicates how many Widgets are held by the sender.

The "application/example-widget+json" format is a JSON [[RFC8259](#)] format representing the state of a Widget. It contains links to related information in the link indicated by the Link header field value with the "example-other-info" link relation type.

[4.3.](#) Specifying Client Behaviours

HTTP does not mandate some behaviours that have nevertheless become very common; if these are not explicitly specified by applications using HTTP, there may be confusion and interoperability problems. This section recommends default handling for these mechanisms.

- o Redirect handling - applications using HTTP SHOULD specify that 3xx redirect status codes be followed automatically. See [\[RFC7231\]](#), [Section 6.4](#).
- o Redirect methods - applications using HTTP SHOULD specify that 301 and 302 redirect status codes rewrite the POST method to GET. See [\[RFC7231\]](#), [Section 6.4](#).
- o Cookies - Applications using HTTP MUST explicitly reference the Cookie specification [[RFC6265](#)] if they are required.

- o Certificates - Applications using HTTP MUST specify that TLS certificates are to be checked according to [\[RFC2818\]](#) when HTTPS is used.

In general, applications using HTTP ought to align their usage as closely as possible with Web browsers, to avoid interoperability issues when they are used. See [Section 4.12](#).

If an application using HTTP has browser compatibility as a goal, client interaction ought to be defined in terms of [\[FETCH\]](#), since that is the abstraction that browsers use for HTTP; it enforces many of these best practices.

Applications using HTTP MUST NOT require HTTP features that are usually negotiated to be supported. For example, requiring that clients support responses with a certain content-encoding ([\[RFC7231\]](#), [Section 3.1.2.2](#)) instead of negotiating for it ([\[RFC7231\]](#), [Section 5.3.4](#)) means that otherwise conformant clients cannot interoperate with the application. Applications MAY encourage the implementation of such features, though.

[4.4.](#) HTTP URLs

In HTTP, URLs are opaque identifiers under the control of the server. As outlined in [\[RFC7320\]](#), standards cannot usurp this space, since it might conflict with existing resources, and constrain implementation and deployment.

In other words, applications that use HTTP shouldn't associate application semantics with specific URL paths on arbitrary servers. Doing so inappropriately conflates the identity of the resource (its URL) with the capabilities that resource supports, bringing about many of the same interoperability problems that [\[RFC4367\]](#) warns of.

For example, specifying that a "GET to the URL /foo retrieves a bar document" is bad practice. Likewise, specifying "The widget API is at the path /bar" violates [\[RFC7320\]](#).

Instead, applications that use HTTP are encouraged to ensure that URLs are discovered at runtime, allowing HTTP-based services to describe their own capabilities. One way to do this is to use typed links [\[RFC8288\]](#) to convey the URIs that are in use, as well as the semantics of the resources that they identify. See [Section 4.2](#) for details.

Nottingham

Expires October 4, 2018

[Page 10]

4.4.1. Initial URL Discovery

Generally, a client will begin interacting with a given application server by requesting an initial document that contains information about that particular deployment, potentially including links to other relevant resources.

Applications that use HTTP are encouraged to allow an arbitrary URL to be used as that entry point. For example, rather than specifying "the initial document is at `/foo/v1`", they should allow a deployment to use any URL as the entry point for the application.

In cases where doing so is impractical (e.g., it is not possible to convey a whole URL, but only a hostname) standard applications that use HTTP can request a well-known URL [[RFC5785](#)] as an entry point.

4.4.2. URL Schemes

Applications that use HTTP will typically employ the "http" and/or "https" URL schemes. "https" is preferred to provide authentication, integrity and confidentiality, as well as mitigate pervasive monitoring attacks [[RFC7258](#)].

However, application-specific schemes can be defined as well.

When defining an URL scheme for an application using HTTP, there are a number of tradeoffs and caveats to keep in mind:

- o Unmodified Web browsers will not support the new scheme. While it is possible to register new URL schemes with Web browsers (e.g. `registerProtocolHandler()` in [[HTML5](#)], as well as several proprietary approaches), support for these mechanisms is not shared by all browsers, and their capabilities vary.
- o Existing non-browser clients, intermediaries, servers and associated software will not recognise the new scheme. For example, a client library might fail to dispatch the request; a cache might refuse to store the response, and a proxy might fail to forward the request.
- o Because URLs occur in and are generated in HTTP artefacts commonly, often without human intervention (e.g., in the "Location" response header), it can be difficult to assure that the new scheme is used consistently.
- o The resources identified by the new scheme will still be available using "http" and/or "https" URLs. Those URLs can "leak" into use, which can present security and operability issues. For example,

Nottingham

Expires October 4, 2018

[Page 11]

using a new scheme to assure that requests don't get sent to a "normal" Web site is likely to fail.

- o Features that rely upon the URL's origin [[RFC6454](#)], such as the Web's same-origin policy, will be impacted by a change of scheme.
- o HTTP-specific features such as cookies [[RFC6265](#)], authentication [[RFC7235](#)], caching [[RFC7234](#)], and CORS [[FETCH](#)] might or might not work correctly, depending on how they are defined and implemented. Generally, they are designed and implemented with an assumption that the URL will always be "http" or "https".
- o Web features that require a secure context [[SECCTXT](#)] will likely treat a new scheme as insecure.

See [[RFC7595](#)] for more information about minting new URL schemes.

[4.4.3.](#) Transport Ports

Applications that use HTTP can use the applicable default port (80 for HTTP, 443 for HTTPS), or they can be deployed upon other ports. This decision can be made at deployment time, or might be encouraged by the application's specification (e.g., by registering a port for that application).

In either case, non-default ports will need to be reflected in the authority of all URLs for that resource; the only mechanism for changing a default port is changing the scheme (see [Section 4.4.2](#)).

Using a port other than the default has privacy implications (i.e., the protocol can now be distinguished from other traffic), as well as operability concerns (as some networks might block or otherwise interfere with it). Privacy implications should be documented in Security Considerations.

See [[RFC7605](#)] for further guidance.

[4.5.](#) HTTP Methods

Applications that use HTTP MUST confine themselves to using registered HTTP methods such as GET, POST, PUT, DELETE, and PATCH.

New HTTP methods are rare; they are required to be registered with IETF Review (see [[RFC7232](#)]), and are also required to be `_generic_`. That means that they need to be potentially applicable to all resources, not just those of one application.

Nottingham

Expires October 4, 2018

[Page 12]

While historically some applications (e.g., [[RFC4791](#)]) have defined non-generic methods, [[RFC7231](#)] now forbids this.

When authors believe that a new method is required, they are encouraged to engage with the HTTP community early, and document their proposal as a separate HTTP extension, rather than as part of an application's specification.

[4.5.1](#). GET

GET is one of the most common and useful HTTP methods; its retrieval semantics allow caching, side-effect free linking and forms the basis of many of the benefits of using HTTP.

A common use of GET is to perform queries, often using the query component of the URL; this is this a familiar pattern from Web browsing, and the results can be cached, improving efficiency of an often expensive process.

In some cases, however, GET might be unwieldy for expressing queries, because of the limited syntax of the URL; in particular, if binary data forms part of the query terms, it needs to be encoded to conform to URL syntax.

While this is not an issue for short queries, it can become one for larger query terms, or ones which need to sustain a high rate of requests. Additionally, some HTTP implementations limit the size of URLs they support - although modern HTTP software has much more generous limits than previously (typically, considerably more than 8000 octets, as required by [[RFC7230](#)] [Section 3.1.1](#)).

In these cases, an application using HTTP might consider using POST to express queries in the request body; doing so avoids encoding overhead and URL length limits in implementations. However, in doing so it should be noted that the benefits of GET such as caching and linking to query results are lost. Therefore, applications using HTTP that feel a need to allow POST queries ought consider allowing both methods.

Applications that use HTTP SHOULD NOT define GET requests to have side effects, since implementations can and do retry HTTP GET requests that fail.

Finally, note that while HTTP allows GET requests to have a body syntactically, this is done only to allow parsers to be generic; as per [[RFC7231](#)], [Section 4.3.1](#), a body on a GET has no meaning, and will be either ignored or rejected by generic HTTP software. As a

Nottingham

Expires October 4, 2018

[Page 13]

result, applications that use HTTP SHOULD NOT define GET to have any side effects upon their resources.

4.6. HTTP Status Codes

Applications that use HTTP MUST only use registered HTTP status codes.

As with methods, new HTTP status codes are rare, and required (by [\[RFC7231\]](#)) to be registered with IETF review. Similarly, HTTP status codes are generic; they are required (by [\[RFC7231\]](#)) to be potentially applicable to all resources, not just to those of one application.

When authors believe that a new status code is required, they are encouraged to engage with the HTTP community early, and document their proposal as a separate HTTP extension, rather than as part of an application's specification.

The primary function of status codes is to convey HTTP semantics for the benefit of generic HTTP software, not application-specific semantics. Therefore, applications MUST NOT specify additional semantics or refine existing semantics for status codes.

In particular, specifying that a particular status code has a specific meaning in the context of an application is harmful, as these are not generic semantics, since the consumer needs to be in the context of the application to understand them.

Furthermore, applications using HTTP MUST NOT re-specify the semantics of HTTP status codes, even if it is only by copying their definition. They MUST NOT require specific reason phrases to be used; the reason phrase has no function in HTTP, and is not guaranteed to be preserved by implementations. The reason phrase is not carried in the [\[RFC7540\]](#) message format.

Typically, applications using HTTP will convey application-specific information in the message body and/or HTTP header fields, not the status code. [\[RFC7807\]](#) provides one way for applications using HTTP to do this.

Specifications sometimes also create a "laundry list" of potential status codes, in an effort to be helpful. The problem with doing so is that such a list is never complete; for example, if a network proxy is interposed, the client might encounter a "407 Proxy Authentication Required" response; or, if the server is rate limiting the client, it might receive a "429 Too Many Requests" response.

Nottingham

Expires October 4, 2018

[Page 14]

Since the list of HTTP status codes can be added to, it's safer to specify behaviours in terms of general response classes (e.g., "successful response" for 2xx; "client error" for 4xx and "server error" for 5xx), pointing out that clients SHOULD be able to handle all applicable protocol elements gracefully (i.e., falling back to the generic "n00" semantics of a given status code; e.g., "499" can be safely handled as "400" by clients that don't recognise it).

4.7. HTTP Header Fields

Applications that use HTTP MAY define new HTTP header fields, following the advice in [\[RFC7231\]](#), [Section 8.3.1](#).

Typically, using HTTP header fields is appropriate in a few different situations:

- o Their content is useful to intermediaries (who often wish to avoid parsing the body), and/or
- o Their content is useful to generic HTTP software (e.g., clients, servers), and/or
- o It is not possible to include their content in the message body (usually because a format does not allow it).

New header fields MUST be registered, as per [\[RFC7231\]](#) and [\[RFC3864\]](#).

It is RECOMMENDED that header field names be short (even when HTTP/2 header compression is in effect, there is an overhead) but appropriately specific. In particular, if a header field is specific to an application, an identifier for that application SHOULD form a prefix to the header field name, separated by a "-".

For example, if the "example" application needs to create three headers, they might be called "example-foo", "example-bar" and "example-baz". Note that the primary motivation here is to avoid consuming more generic header names, not to reserve a portion of the namespace for the application; see [\[RFC6648\]](#) for related considerations.

The semantics of existing HTTP header fields MUST NOT be re-defined without updating their registration or defining an extension to them (if allowed). For example, an application using HTTP cannot specify that the "Location" header has a special meaning in a certain context.

Nottingham

Expires October 4, 2018

[Page 15]

See [Section 4.9](#) for the interaction between headers and HTTP caching; in particular, request headers that are used to "select" a response have impact there, and need to be carefully considered.

See [Section 4.10](#) for considerations regarding header fields that carry application state (e.g., `Cookie`).

4.8. Defining Message Payloads

There are many potential formats for payloads; for example, JSON [[RFC8259](#)], XML [[XML](#)], and CBOR [[RFC7049](#)]. Best practices for their use are out of scope for this document.

Applications SHOULD register distinct media types for each format they define; this makes it possible to identify them unambiguously and negotiate for their use. See [[RFC6838](#)] for more information.

4.9. HTTP Caching

HTTP caching [[RFC7234](#)] is one of the primary benefits of using HTTP for applications; it provides scalability, reduces latency and improves reliability. Furthermore, HTTP caches are readily available in browsers and other clients, networks as forward and reverse proxies, Content Delivery Networks and as part of server software.

Assigning even a short freshness lifetime ([\[RFC7234\], Section 4.2](#)) - e.g., 5 seconds - allows a response to be reused to satisfy multiple clients, and/or a single client making the same request repeatedly. In general, if it is safe to reuse something, consider assigning a freshness lifetime; cache implementations take active measures to expire content intelligently when they are out of space, so "it will fill up the cache" is not a valid concern.

Understand that stale responses (e.g., one with "Cache-Control: max-age=0") can be reused when the cache is disconnected from the origin server; this can be useful for handling network issues. See [\[RFC7234\], Section 4.2.4](#), and also [[RFC5861](#)] for additional controls over stale content.

Stale responses can be refreshed by assigning a validator, saving both transfer bandwidth and latency for large responses; see [[RFC7232](#)].

In some situations, responses without explicit cache directives (e.g., Cache-Control or Expires) will be stored and served using a heuristic freshness lifetime; see [\[RFC7234\], Section 4.2.2](#). As the heuristic is not under control of the application, it is generally preferable to set an explicit freshness lifetime.

Nottingham

Expires October 4, 2018

[Page 16]

If caching of a response is not desired, the appropriate response directive is "Cache-Control: no-store". This only need be sent in situations where the response might be cached; see [\[RFC7234\]](#), [Section 3](#).

If an application defines a request header field that might be used by a server to change the response's headers or body, authors should point out that this has implications for caching; in general, such resources need to either make their responses uncacheable (e.g., with the "no-store" cache-control directive defined in [\[RFC7234\]](#), [Section 5.2.2.3](#)) or consistently send the Vary response header ([\[RFC7231\]](#), [Section 7.1.4](#)).

When an application has a need to express a lifetime that's separate from the freshness lifetime, this should be expressed separately, either in the response's body or in a separate header field. When this happens, the relationship between HTTP caching and that lifetime need to be carefully considered, since the response will be used as long as it is considered fresh.

Like other functions, HTTP caching is generic; it does not have knowledge of the application in use. Therefore, caching extensions need to be backwards-compatible, as per [\[RFC7234\]](#), [Section 5.2.3](#).

[4.10](#). Application State

Applications that use HTTP MAY use stateful cookies [\[RFC6265\]](#) to identify a client and/or store client-specific data to contextualise requests.

When used, it is important to carefully specify the scoping and use of cookies; if the application exposes sensitive data or capabilities (e.g., by acting as an ambient authority), exploits are possible. Mitigations include using a request-specific token to assure the intent of the client.

Applications MUST NOT make assumptions about the relationship between separate requests on a single transport connection; doing so breaks many of the assumptions of HTTP as a stateless protocol, and will cause problems in interoperability, security, operability and evolution.

[4.11](#). Client Authentication

Applications that use HTTP MAY use HTTP authentication [\[RFC7235\]](#) to identify clients. The Basic authentication scheme [\[RFC7617\]](#) MUST NOT be used unless the underlying transport is authenticated, integrity-protected and confidential (e.g., as provided the "HTTPS" URL scheme,

or another using TLS). The Digest scheme [[RFC7616](#)] MUST NOT be used unless the underlying transport is similarly secure, or the chosen hash algorithm is not "MD5".

When used, it is important to carefully specify the scoping and use of authentication; if the application exposes sensitive data or capabilities (e.g., by acting as an ambient authority), exploits are possible. Mitigations include using a request-specific token to assure the intent of the client.

4.12. Co-Existing with Web Browsing

Even if there is not an intent for an application that uses HTTP to be used with a Web browser, its resources will remain available to browsers and other HTTP clients.

This means that all such applications need to consider how browsers will interact with them, particularly regarding security.

For example, if an application's state can be changed using a POST request, a Web browser can easily be coaxed into making that request by a HTML form on an arbitrary Web site.

Or, if a resource reflects data from the request into a response, that can be used to perform a Cross-Site Scripting attack on Web browsers directed to it.

This is only a small sample of the kinds of issues that applications using HTTP must consider. Generally, the best approach is to consider the application *as* a Web application, and to follow best practices for their secure development.

A complete enumeration of such practices is out of scope for this document, but some considerations include:

- o Using Strict Transport Security [[RFC6797](#)] to assure that HTTPS is used
- o Using Content-Security-Policy [[W3C.WD-CSP3-20160913](#)] to constrain the capabilities of content, thereby mitigating Cross-Site Scripting attacks (which are possible if client-provided data is exposed in any part of a response in the application)
- o Using X-Frame-Options [[RFC7034](#)] to prevent content from being included in a HTML frame from another origin, thereby enabling "clickjacking"

- o Using Referrer-Policy [[W3C.CR-referrer-policy-20170126](#)] to prevent sensitive data in URLs from being leaked in the Referer request header
- o Using the 'HttpOnly' flag on Cookies to assure that cookies are not exposed to browser scripting languages [[RFC6265](#)]

Depending on how they are intended to be deployed, specifications for applications using HTTP might require the use of these mechanisms in specific ways, or might merely point them out in Security Considerations.

If an application using HTTP has browser compatibility as a goal, client interaction ought to be defined in terms of [[FETCH](#)], since that is the abstraction that browsers use for HTTP; it enforces many of these best practices.

[4.13.](#) Application Boundaries

Because the origin [[RFC6454](#)] is how many HTTP capabilities are scoped, applications also need to consider how deployments might interact with other applications (including Web browsing) on the same origin.

For example, if Cookies [[RFC6265](#)] are used to carry application state, they will be sent with all requests to the origin by default, unless scoped by path, and the application might receive cookies from other applications on the origin. This can lead to security issues, as well as collision in cookie names.

One solution to these issues is to require a dedicated hostname for the application, so that it has a unique origin. However, it is often desirable to allow multiple applications to be deployed on a single hostname; doing so provides the most deployment flexibility and enables them to be "mixed" together (See [[RFC7320](#)] for details). Therefore, applications using HTTP should strive to allow multiple applications on an origin.

To enable this, when specifying the use of Cookies, HTTP authentication realms [[RFC7235](#)], or other origin-wide HTTP mechanisms, applications using HTTP SHOULD NOT mandate the use of a particular identifier, but instead let deployments configure them. Consideration SHOULD be given to scoping them to part of the origin, using their specified mechanisms for doing so.

Modern Web browsers constrain the ability of content from one origin to access resources from another, to avoid leaking private information. As a result, applications that wish to expose cross-

Nottingham

Expires October 4, 2018

[Page 19]

origin data to browsers will need to implement the CORS protocol; see [\[FETCH\]](#).

5. IANA Considerations

This document has no requirements for IANA.

6. Security Considerations

[Section 4.10](#) discusses the impact of using stateful mechanisms in the protocol as ambient authority, and suggests a mitigation.

[Section 4.4.2](#) requires support for 'https' URLs, and discourages the use of 'http' URLs, to provide authentication, integrity and confidentiality, as well as mitigate pervasive monitoring attacks.

[Section 4.12](#) highlights the implications of Web browsers' capabilities on applications that use HTTP.

[Section 4.13](#) discusses the issues that arise when applications are deployed on the same origin as Web sites (and other applications).

Applications that use HTTP in a manner that involves modification of implementations - for example, requiring support for a new URL scheme, or a non-standard method - risk having those implementations "fork" from their parent HTTP implementations, with the possible result that they do not benefit from patches and other security improvements incorporated upstream.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", [RFC 2818](#), DOI 10.17487/RFC2818, May 2000, <<https://www.rfc-editor.org/info/rfc2818>>.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), DOI 10.17487/RFC3864, September 2004, <<https://www.rfc-editor.org/info/rfc3864>>.

Nottingham

Expires October 4, 2018

[Page 20]

- [RFC6454] Barth, A., "The Web Origin Concept", [RFC 6454](#), DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/info/rfc6454>>.
- [RFC6648] Saint-Andre, P., Crocker, D., and M. Nottingham, "Deprecating the "X-" Prefix and Similar Constructs in Application Protocols", [BCP 178](#), [RFC 6648](#), DOI 10.17487/RFC6648, June 2012, <<https://www.rfc-editor.org/info/rfc6648>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", [BCP 13](#), [RFC 6838](#), DOI 10.17487/RFC6838, January 2013, <<https://www.rfc-editor.org/info/rfc6838>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", [RFC 7230](#), DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", [RFC 7231](#), DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.
- [RFC7232] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", [RFC 7232](#), DOI 10.17487/RFC7232, June 2014, <<https://www.rfc-editor.org/info/rfc7232>>.
- [RFC7233] Fielding, R., Ed., Lafon, Y., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", [RFC 7233](#), DOI 10.17487/RFC7233, June 2014, <<https://www.rfc-editor.org/info/rfc7233>>.
- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", [RFC 7234](#), DOI 10.17487/RFC7234, June 2014, <<https://www.rfc-editor.org/info/rfc7234>>.
- [RFC7235] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", [RFC 7235](#), DOI 10.17487/RFC7235, June 2014, <<https://www.rfc-editor.org/info/rfc7235>>.

- [RFC7301] 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>>.
- [RFC7320] Nottingham, M., "URI Design and Ownership", [BCP 190](#), [RFC 7320](#), DOI 10.17487/RFC7320, July 2014, <<https://www.rfc-editor.org/info/rfc7320>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", [RFC 7540](#), DOI 10.17487/RFC7540, May 2015, <<https://www.rfc-editor.org/info/rfc7540>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8288] Nottingham, M., "Web Linking", [RFC 8288](#), DOI 10.17487/RFC8288, October 2017, <<https://www.rfc-editor.org/info/rfc8288>>.

7.2. Informative References

- [FETCH] WHATWG, "Fetch - Living Standard", n.d., <<https://fetch.spec.whatwg.org>>.
- [HTML5] WHATWG, "HTML - Living Standard", n.d., <<https://html.spec.whatwg.org>>.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.
- [RFC0854] Postel, J. and J. Reynolds, "Telnet Protocol Specification", STD 8, [RFC 854](#), DOI 10.17487/RFC0854, May 1983, <<https://www.rfc-editor.org/info/rfc854>>.
- [RFC0959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, [RFC 959](#), DOI 10.17487/RFC0959, October 1985, <<https://www.rfc-editor.org/info/rfc959>>.
- [RFC2821] Klensin, J., Ed., "Simple Mail Transfer Protocol", [RFC 2821](#), DOI 10.17487/RFC2821, April 2001, <<https://www.rfc-editor.org/info/rfc2821>>.

Nottingham

Expires October 4, 2018

[Page 22]

- [RFC3205] Moore, K., "On the use of HTTP as a Substrate", [BCP 56](#), [RFC 3205](#), DOI 10.17487/RFC3205, February 2002, <<https://www.rfc-editor.org/info/rfc3205>>.
- [RFC4367] Rosenberg, J., Ed. and IAB, "What's in a Name: False Assumptions about DNS Names", [RFC 4367](#), DOI 10.17487/RFC4367, February 2006, <<https://www.rfc-editor.org/info/rfc4367>>.
- [RFC4791] Daboo, C., Desruisseaux, B., and L. Dusseault, "Calendaring Extensions to WebDAV (CalDAV)", [RFC 4791](#), DOI 10.17487/RFC4791, March 2007, <<https://www.rfc-editor.org/info/rfc4791>>.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", [RFC 5785](#), DOI 10.17487/RFC5785, April 2010, <<https://www.rfc-editor.org/info/rfc5785>>.
- [RFC5861] Nottingham, M., "HTTP Cache-Control Extensions for Stale Content", [RFC 5861](#), DOI 10.17487/RFC5861, May 2010, <<https://www.rfc-editor.org/info/rfc5861>>.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), DOI 10.17487/RFC6265, April 2011, <<https://www.rfc-editor.org/info/rfc6265>>.
- [RFC6455] Fette, I. and A. Melnikov, "The WebSocket Protocol", [RFC 6455](#), DOI 10.17487/RFC6455, December 2011, <<https://www.rfc-editor.org/info/rfc6455>>.
- [RFC6797] Hodges, J., Jackson, C., and A. Barth, "HTTP Strict Transport Security (HSTS)", [RFC 6797](#), DOI 10.17487/RFC6797, November 2012, <<https://www.rfc-editor.org/info/rfc6797>>.
- [RFC7034] Ross, D. and T. Gondrom, "HTTP Header Field X-Frame-Options", [RFC 7034](#), DOI 10.17487/RFC7034, October 2013, <<https://www.rfc-editor.org/info/rfc7034>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", [RFC 7049](#), DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.

- [RFC7595] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", [BCP 35](#), [RFC 7595](#), DOI 10.17487/RFC7595, June 2015, <<https://www.rfc-editor.org/info/rfc7595>>.
- [RFC7605] Touch, J., "Recommendations on Using Assigned Transport Port Numbers", [BCP 165](#), [RFC 7605](#), DOI 10.17487/RFC7605, August 2015, <<https://www.rfc-editor.org/info/rfc7605>>.
- [RFC7616] Shekh-Yusef, R., Ed., Ahrens, D., and S. Bremer, "HTTP Digest Access Authentication", [RFC 7616](#), DOI 10.17487/RFC7616, September 2015, <<https://www.rfc-editor.org/info/rfc7616>>.
- [RFC7617] Reschke, J., "The 'Basic' HTTP Authentication Scheme", [RFC 7617](#), DOI 10.17487/RFC7617, September 2015, <<https://www.rfc-editor.org/info/rfc7617>>.
- [RFC7807] Nottingham, M. and E. Wilde, "Problem Details for HTTP APIs", [RFC 7807](#), DOI 10.17487/RFC7807, March 2016, <<https://www.rfc-editor.org/info/rfc7807>>.
- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, [RFC 8259](#), DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [SECCTXT] West, M., "Secure Contexts", World Wide Web Consortium CR CR-secure-contexts-20160915, September 2016, <<https://www.w3.org/TR/2016/CR-secure-contexts-20160915>>.
- [W3C.CR-referrer-policy-20170126]
Eisinger, J. and E. Stark, "Referrer Policy", World Wide Web Consortium CR CR-referrer-policy-20170126, January 2017, <<https://www.w3.org/TR/2017/CR-referrer-policy-20170126>>.
- [W3C.WD-CSP3-20160913]
West, M., "Content Security Policy Level 3", World Wide Web Consortium WD WD-CSP3-20160913, September 2016, <<https://www.w3.org/TR/2016/WD-CSP3-20160913>>.
- [XML] Bray, T., Paoli, J., Sperberg-McQueen, M., Maler, E., and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Fifth Edition)", World Wide Web Consortium Recommendation REC-xml-20081126, November 2008, <<http://www.w3.org/TR/2008/REC-xml-20081126>>.

Nottingham

Expires October 4, 2018

[Page 24]

7.3. URIs

- [1] <https://lists.w3.org/Archives/Public/ietf-http-wg/>
- [2] <http://httpwg.github.io/>
- [3] <https://github.com/httpwg/http-extensions/labels/bcp56bis>

Appendix A. Changes from [RFC 3205](#)

[RFC3205] captured the Best Current Practice in the early 2000's, based on the concerns facing protocol designers at the time. Use of HTTP has changed considerably since then, and as a result this document is substantially different. As a result, the changes are too numerous to list individually.

Author's Address

Mark Nottingham

Email: mnot@mnot.net

URI: <https://www.mnot.net/>

