

# 1 HTTP 1.1 Conformance Issues for IPP

## 4 **Overview**

6 There are a number of implied conformance requirements for origin servers within the current HTTP 1.1  
7 specification. As anticipated by the IPP working group, a significant portion of the current HTTP 1.1  
8 specification is related to the proper operation of HTTP gateways and/or proxy servers. Very little of the  
9 document utilizes the word MUST when applied to origin servers. Throughout this document I will use  
10 the term "current specification" to mean RFC 2068.

11  
12 The current specification lists two types of compliance: conditional and unconditional. To quote from the  
13 text of RFC 2068:

14  
15 An implementation is not compliant if it fails to satisfy one or more  
16 of the MUST requirements for the protocols it implements. An  
17 implementation that satisfies all the MUST and all the SHOULD  
18 requirements for its protocols is said to be "unconditionally  
19 compliant"; one that satisfies all the MUST requirements but not all  
20 the SHOULD requirements for its protocols is said to be  
21 "conditionally compliant."  
22  
23

24 The ratio of SHOULDs to MUSTs in the current specification is quite high so I would suggest the  
25 following criteria be used for our consideration of HTTP 1.1 compliance. If we envision the environment  
26 in which IPP clients and servers interoperate as being "closed", then we can assume that conditional  
27 compliance would be enough to satisfy our requirements. If the environment in which IPP servers and  
28 clients interoperate is "open", then we should lean towards unconditional compliance.

29  
30 The terms "closed" and "open" in the above paragraph refer to the operational environment in which IPP  
31 servers and clients are interoperating. In a "closed" environment, the HTTP "traffic" that is occurring  
32 between IPP clients and servers is dedicated to IPP. In an "open" environment, either the IPP client  
33 and/or IPP server is operating in a dual-role as both a generic HTTP client or server, and has knowledge  
34 of IPP/HTTP protocol (possibly using application/ipp tags).  
35

36 The rationale for the above proposal is that in a "closed" environment, we have a priori knowledge of  
37 exactly the types of HTTP methods, status codes, and to some extent, MIME tags that we can expect over  
38 a particular HTTP connection; and this set of methods, codes, and tags are a subset of HTTP, a subset that  
39 would fit quite easily within the realm of "conditional" HTTP compliance. In an "open" HTTP  
40 environment, we have less knowledge of the types of traffic (methods, URI types, status codes, etc.) that  
41 might have to be handled, and the traditional network protocol guidelines should be met: "Be conservative  
42 with what you send, but robust in what you can receive". Therefore, unconditional compliance would give  
43 us the best chance at interoperating within "open" environments.  
44

45 One other question that should be considered if we use HTTP as a mapping for IPP is whether or not IPP  
46 is "loosely" coupled to HTTP, or "tightly" coupled. "Loosely" coupled in this context means that we are  
47 only using HTTP as a pure transport (which is the method my earlier draft employed), and that the "real"  
48 IPP packets are encoded within "application/ipp" entities. In the loosely coupled model, we would define  
49 our own separate protocol operations, semantics, and status codes. Using HTTP as a pure transport in the

50 loosely coupled model would allow rapid implementation of IPP over other transports (directly to TCP,  
51 SPX, etc.).

52

53 In a tightly coupled specification, we would leverage as much of the existing HTTP protocol as possible,  
54 extending the semantics of certain HTTP methods to be IPP-specific when applied to IPP URIs, as well as  
55 reusing all of the existing status codes and MIME-like packaging characteristics employed by the current  
56 HTTP 1.1 specification. In the tightly coupled model, implementations of "dual-role" servers (HTTP/IPP)  
57 would be easier to deploy and HTTP/IPP clients might be easier to implement. We are also not  
58 "reinventing the wheel" if a particular need of our model is sufficiently addressed by HTTP 1.1. In the  
59 tightly coupled model, it might be more difficult to port IPP onto some other application-level transport if  
60 the need ever arose.

61

62 One idea that might have been overlooked is that we should attempt to create conformance requirements  
63 for both IPP servers, as well as clients. The group is currently looking at creating a mapping of the IPP  
64 model document directly to HTTP 1.1. It is this goal that this document is trying to address with regards  
65 to scope and capability. The 'scope' of the effort being how much resources are needed (manpower, code  
66 size, RAM, CPU cycles, etc.) to implement IPP over HTTP. The 'capability' aspect of our requirements  
67 asks the question whether or not HTTP meets the transport-requirements or operational requirements  
68 implied by the IPP model document.

69

70 It is assumed that we currently have three alternatives for the choosing a protocol mapping for IPP 1.0:

71

72 1. Use HTTP 1.1 origin server capability for IPP servers; use some subset of HTTP 1.1 methods for IPP  
73 clients. In this scenario, we could not use the full capabilities of MIME since HTTP does not support  
74 MIME, per the MIME standard.

75

76 2. Use an HTTP-like protocol for both IPP servers and clients. We don't say that we're HTTP compliant,  
77 but we're so close that if we wanted to construct gateways between IPP and HTTP, the work would be  
78 more or less trivial. In this scenario, since we are not HTTP compliant, we could choose to use MIME  
79 structures that would not otherwise be supported by a conforming HTTP implementation.

80

81 3. We "do our own thing", with no resemblance to anything like HTTP. Our own custom encapsulation,  
82 headers, status codes, and protocol operation semantics.

83

84 At the IETF Plenary in Memphis, and in subsequent teleconferences in the protocol subgroup, it has been  
85 suggested that we avoid specifying multiple protocol mappings for IPP, at least for the first standards  
86 effort. Rather, our area directors suggested picking only one mapping and going with that. For this reason,  
87 this document assumes that the method for job delivery via RFC 1867 to support existing WEB browsers  
88 will not be supported. Instead, existing browsers will have to utilize IPP support built into the underlying  
89 operating system environment. New browser technology will either use underlying OS support for IPP, or  
90 incorporate IPP client capability directly into the browser.

91

92 Rather than echo the current HTTP 1.1 specification with regards to MUST and SHOULD requirements, I  
93 would advise the WG to review RFC 2068 for the words MUST and SHOULD, with special attention to  
94 where and how I have proposed conditional and unconditional implementations of HTTP for IPP clients  
95 and servers.

96

## 97 **HTTP Proxies**

98 One aspect of implementing IPP using HTTP that must be decided by the working group would be to take  
99 advantage of HTTP proxies. Overall inter- or intranet bandwidth could be reduced by allowing  
100 intermediate HTTP proxies to cache responses to IPP attribute requests. It is recognized that some  
101 intelligent use of response lifetimes would have to be utilized, but this capability is fully supported by

102 HTTP 1.1, and can be automatically (dynamically) set on a server-by-server basis. The disadvantage is  
103 that IPP servers would have to (possibly) implement more support for different types of headers and/or  
104 methods that would be utilized by proxies in between dedicated IPP servers and clients. Servers would also  
105 have to manage response lifetimes and possibly include support for content-encodings or other headers  
106 that proxies might inject into an HTTP transaction.

107  
108 The remainder of this document is a briefly paraphrased version of my internet-draft regarding key issues  
109 with implementing IPP tightly coupled to HTTP.

## 111 ***HTTP 1.1 Methods***

### 113 **GET**

114 HTTP 1.1 GET methods could be used to obtain attributes for different types of IPP objects. If a GET  
115 method is applied to an IPP printer object, then the attributes for the printer object are returned. If the  
116 GET method is applied to a job object, then the attributes for the job are returned as the response.

### 117 **POST**

118 The POST method would be used as a way to create IPP objects (printers, jobs, etc.). Initially, the POST  
119 operation would be used to map the "CreateJob" model operation. The POST method could also be used to  
120 map the "SendJob" model operation. Using HTTP 1.1 persistent connections, multiple POST operations  
121 could be used to efficiently deliver job data.

### 122 **HEAD**

123 HTTP 1.1 servers are required to implement the HEAD method. This method is often used to inquire as to  
124 find out meta-information about resources prior to actually performing a GET operation on the resource.  
125 This provides a way for clients to learn as much as possible about an object before actually retrieving the  
126 contents of the resource with a subsequent GET operation. The HEAD method, like the GET method,  
127 could be used to retrieve meta-information (attributes, etc.) about IPP objects before actually accessing it.

### 128 **PUT (OPTIONAL)**

129 Using the PUT method with an "Allow:" header can provide IPP clients to create IPP objects and specify  
130 what types of methods are allowed on the object.

131  
132  
133 According to section 5.1.1 of the current specification, only the GET and HEAD methods **MUST** be  
134 supported by servers. For a particular type of URI, only certain methods may be supported. In the event of  
135 an unsupported method being received for a particular URI, the server is required to return a status code  
136 of 405 (Method not allowed), and include an "Allow:" header field listing exactly which methods are  
137 supported for the URI in question.

138  
139 Lightweight embedded IPP server implementations could be unconditionally compliant by only supporting  
140 the GET, HEAD, and POST methods, and returning a status code of 501 when requests are received  
141 specifying any other unsupported method. The 501 response should include a "Public:" response header  
142 indicating which methods the server does support.

## 144 ***HTTP 1.1 General Headers***

145

146 Many of the headers specified by RFC 2068 do not have to be supported by general purpose HTTP servers.  
147 The following text clarifies what I think IPP should utilize and the HTTP 1.1 conformance issues for each.  
148  
149

### 150 Cache-Control:

151 IPP servers and clients will have to operate within a caching proxy environment. In order to ensure a pure  
152 client and server environment between IPP clients and servers, we will have to decide whether or not to  
153 take advantage of caching, or prohibit caching of any and all IPP traffic. If we choose to prohibit caching  
154 of IPP traffic (to keep things simple), then the "Cache-Control:" general header must be included in all  
155 IPP requests and responses. The value for the "Cache-Control:" header directive would be "no-cache".  
156 Also, to make sure there are no HTTP 1.0 caching proxies between HTTP 1.1 clients and servers, we must  
157 also include the "Pragma:" general header, also specifying "no-cache" as the value.  
158

159 It is conceivable that there would be some value in caching of attribute requests to IPP printer objects,  
160 since in a large environment, these requests might be very frequent. The IPP working group may want to  
161 consider the value in caching certain IPP object attribute requests. If caching of certain IPP response data  
162 is allowed, then we should also consider the use of the "no-transform" value for the "Cache-Control"  
163 directive.  
164

### 165 Connection:

166 The "Connection:" general header could be used by IPP servers or IPP clients to instruct either a remote  
167 client or server that the HTTP connection be closed. In HTTP 1.1, persistent connections are the rule, not  
168 the exception. If we decide to allow a "SendJob" operation to occur in multiple POST or PUT operations  
169 to an IPP server, then persistent connections would be very valuable for enhancing performance of job  
170 submission. In this type of multiple POST/PUT operation, the last POST/PUT operation required to  
171 deliver the job data could include the "Connection:" header with the value "close" to instruct the server  
172 that the connection will be closed after reception of this request.  
173

### 174 Content-Coding:

175 The "Content-Coding:" entity header field specifies how the entity body of a particular message is to be  
176 decoded. For HTTP, this is typically a compression encoding so the field would be "gzip" or  
177 "compressed". If the IPP working group wants to define a base set of content-codings, then the values for  
178 these codings would be specified in IPP messages via the "Content-Coding:" header.  
179

### 180 Content-Language:

181 Like the "Content-Coding:" header, this header also specifies information related to the enclosed entity.  
182 The "Content-Coding:" header describes the natural language in which the entity body has been encoded.  
183 IPP servers should only return entities in languages that have been "agreed" upon by a particular client in  
184 a previously received "Accept-Language:" request header.  
185

### 186 Content-Length:

187 The "Content-Length:" header specifies the size of a message body. IPP clients and servers will use the  
188 same algorithm as general-purpose HTTP 1.1 servers for determining the length of IPP messages. From  
189 RFC 2068, the description of the Content-Length header:  
190  
191 Applications SHOULD use this field to indicate the size of the  
192 message-body to be transferred, regardless of the media type of the

193 entity. It must be possible for the recipient to reliably determine  
194 the end of HTTP/1.1 requests containing an entity-body, e.g., because  
195 the request has a valid Content-Length field, uses Transfer-Encoding:  
196 chunked or a multipart body.  
197

## 198 Content-MD5

199 The Content-MD5 header can be used by IPP clients and servers to provide a more robust authentication  
200 method than just basic HTTP authentication. The current HTTP 1.1 specification states that use of the  
201 MD5 digest authentication is sufficient to protect against accidental modification of the message, but NOT  
202 sufficient to protect against malicious attempts to modify the message. The IPP working group should  
203 consider as few security mechanisms to provide a higher degree of interoperability between clients and  
204 servers. It would seem as if two levels of secure access to IPP objects would suffice: a simple method for  
205 moderate to insecure sites where security is not an issue, and a very robust method that is sufficient to  
206 meet the needs of sites requiring very high levels of security, including commercial transactions. The  
207 simple security could be provided by basic HTTP authentication, and a yet-to-be-decided method (maybe  
208 SSL or secure MIME) could be utilized in high security environments. With these two scenarios and  
209 methods, IPP clients and servers would not generate messages with the Content-MD5 header.  
210

## 211 Content-Type:

212 The Content-Type header would be used by IPP clients and servers to specify IPP-specific entities. The  
213 Content-Type value would be "application/ipp". IPP clients and servers would also supply a Content-Type  
214 modifier "charset", as part of the applicaion/ipp Content-Type. The "charset" modifier would specify the  
215 character set used within the application/ipp entity body.  
216

## 217 Date:

218 The "Date:" header field is currently specified in the HTTP 1.1 document as a MUST header by all  
219 compliant implementations. The date format used as the value of this header must be in RFC 1123 format.  
220 There is a recent internet draft that has been published that attempts to describe how some embedded,  
221 lightweight HTTP server implementations can still be "compliant" even if they don't contain any realtime  
222 clock or time capabilities.  
223

## 224 Pragma:

225 The "Pragma:" directive would only be used by IPP implementations for backwards compatibility with  
226 HTTP 1.0 caching proxies. The Pragma header would specify the value "no-cache", which is understood  
227 by HTTP 1.0 proxies to have the same semantics as the HTTP 1.1 "Cache-Control" directive with the  
228 value "no-cache".  
229

## 230 Transfer-Encoding:

231 For HTTP 1.1, the only "Transfer-Encoding" specified is the "chunked" encoding. Since an HTTP  
232 connection is "8-bit clean", the traditional rationale for transfer-encodings (like used in MIME) are  
233 unneeded. But when IPP implementations are attempting to send messages for which the total length of  
234 the message cannot be determined, then the message should be transferred as either "chunked" or via a  
235 multipart message with message boundaries. The current HTTP specification requires that all HTTP 1.1  
236 applications MUST be able to receive and decode the chunked transfer encoding.

## 237 **HTTP 1.1 Request Headers**

238

### 239 **Accept:**

240 The Accept: header is used to specify certain media types that a client is willing accept as a result of a  
241 request to a server. IPP clients should always specify (at a minimum) application/ipp, text/html, and  
242 text/plain.  
243

### 244 **Accept-Charset:**

245 This header indicates to servers what character sets a client is willing to accept in a response. According  
246 to the HTTP 1.1 specification, all clients should be able to support ISO-8859-1.  
247

### 248 **Accept-Encoding:**

249 Similar to "Accept:", the Accept-Encoding header is sent from client to server to inform the server what  
250 types of encoding of responses that the client can handle.  
251

### 252 **Accept-Language:**

253 IPP clients would send Accept-Language headers in IPP requests to notify IPP servers what type of  
254 localization is acceptable to the client.  
255

### 256 **Authorization:**

257 IPP servers may protect certain types of IPP objects via HTTP basic authentication. If an IPP client has  
258 knowledge that a requested resource requires basic authentication, then an appropriate "Authorization:"  
259 request header should be included in all IPP requests to the IPP object (URI) in question. The client can  
260 also dynamically learn of the authentication requirements for a particular object if the client attempts to  
261 access the object without an authentication header. IPP servers that receive un-authenticated requests for  
262 IPP objects that require basic authentication would return a status code of 401, which indicates to clients  
263 that authentication is required for accessing the requested object.  
264

265 It is assumed that, for the lifetime of a particular IPP object (URI), that the user's credentials (once  
266 successfully validated) will be valid. Therefore, on the first successful authenticated response to a request,  
267 IPP clients can cache the user's credentials and reuse these credentials on subsequent requests to the server  
268 for this object. Each subsequent request for the IPP object (URI) would include an "Authorization:" header  
269 specifying the cached credentials.  
270

### 271 **From:**

272 The "From:" header contains the internet e-mail address for the human individual that is responsible for  
273 the request being generated. The IPP working group has talked about using the "From:" header as a  
274 means for some type of authentication or access protection. The current HTTP 1.1 specification states that  
275 the "From:" header "SHOULD NOT be used as an insecure method of access protection". The  
276 specification goes on to say that "the interpretation of this field is that the request is being performed on  
277 behalf of the user specified by the "From:" header, who accepts responsibility for the operation being  
278 performed.". The following paragraph from RFC 2068 is especially relevant:  
279

280

281

Note: The client SHOULD not send the From header field without the

282 user's approval, as it may conflict with the user's privacy  
283 interests or their site's security policy. It is strongly  
284 recommended that the user be able to disable, enable, and modify  
285 the value of this field at any time prior to a request.  
286

#### 287 **Host:**

288 The Host: field typically comes on a separate line after the HTTP method specification. This field MUST  
289 be set by HTTP 1.1 clients with the network location of the specified URI in the method. All internet-  
290 based HTTP 1.1 servers MUST respond with a 400 status code to any HTTP 1.1 request message which  
291 lacks a "Host:" header. This header is used by newer WEB server sites for so-called "virtual host" access.  
292 IPP could utilize this field in some very interesting ways with regards to multiple logical printers serviced  
293 by a single IPP/HTTP server.  
294

#### 295 **Proxy-Authorization:**

296 When there is an HTTP 1.1 caching proxy operating in between an IPP client and server, it is possible  
297 that certain resources identified by a site administrator might require basic authentication. If an IPP client  
298 receives a 407 response to a valid IPP request, the client should format an authorization request back to  
299 the requested resource (URI) using the "Proxy-Authorization:" request header. Section 11 of the current  
300 HTTP 1.1 specification discusses HTTP authorization in detail.

### 301 ***HTTP 1.1 Response Headers***

#### 302 **Accept-Ranges:**

303 IPP servers could make use of the "Accept-Ranges:" response header for other purposes than just byte  
304 ranges. The "Accept-Ranges:" header includes a parameter that specifies what type of range the server is  
305 capable of handling; "bytes" is just one possible value for this field. Other possible values could include  
306 "pages", "cost", and other types of range values that would be applicable to printer or print job resources.  
307

#### 308 **Allow:**

309 The "Allow:" entity header field can be returned by IPP servers to notify IPP clients which HTTP methods  
310 are allowed to be executed on a particular URI (or IPP object). In the future, we may want to define  
311 conformance levels with respect to IPP, wherein some IPP servers implement all possible methods on IPP  
312 objects, and other lighter weight IPP servers are restricted in the domain of methods supported on IPP  
313 objects. The "Allow:" header permits interoperability between clients and servers of different capabilities.  
314 The client can adapt its behavior to the capabilities it learns from a particular server.  
315

#### 316 **Content-Location:**

317  
318 IPP servers can return a Content-Location header that specifies the URI of a job object created with the  
319 "CreateJob" operation. IPP clients can also use the Content-Location header to specify the target IPP  
320 object (URI) to which a particular IPP operation is to apply.  
321

#### 322 **Expires:**

323 If the working group decides that IPP responses can be cached by intermediate HTTP caching proxies,  
324 then appropriate use of the "Expires:" header should specify how long proxies (and possibly) clients can  
325 consider the response "valid". It is understandable that in the case of an embedded IPP/HTTP server that

326 does not have access to a time source, that the “Expires:” header would not be generated. In this case, the  
327 embedded server should disable caching of responses using “Cache-Control” headers.  
328

### 329 **Location:**

330 The “Location:” header would be used by IPP servers to dynamically redirect IPP clients to other URIs  
331 that can be contacted for completing the client’s request. The Location header could be used as a  
332 replacement for the multiple-URL facility discussed in the early IPP-over-HTTP internet draft. IPP  
333 implementations would follow the direction set forth by the current HTTP 1.1 specification:  
334

335 “for 201 (“Created”) responses, the “Location” is that of the new resource created by the request. For 3xx  
336 responses, the location SHOULD indicate the server’s preferred URL for automatic redirection to the  
337 resource.”  
338

339 The term “resource” used in the above paragraph would normally be a URI referencing an IPP job object.  
340

### 341 **Proxy-Authenticate:**

342 It is possible that, in the presence of caching HTTP 1.1 proxies, that IPP client implementations may have  
343 to deal with “Proxy-Authenticate” responses. The “Proxy-Authenticate” response header would be  
344 returned as part of a 407 (Proxy Authentication Required) response. (see also Proxy-Authorization request  
345 header).  
346

### 347 **Public:**

348 The “Public:” response header would be used by IPP servers to inform IPP clients what types of HTTP  
349 methods are supported by the server. The “Public” response header would typically be used by very  
350 lightweight HTTP/IPP server implementations that implement a minimal IPP capability.  
351

### 352 **Retry-After:**

353 The “Retry-After” response header would be used in tandem with the 503 (Service Unavailable) response  
354 code to indicate how long the resource (or service) is to remain unavailable. This could be used by IPP  
355 servers to indicate how long a printing service might be unavailable to IPP clients.  
356

### 357 **WWW-Authenticate:**

358 The WWW-Authenticate response header is used to initiate basic HTTP authentication. If an IPP client  
359 receives a 401 (Unauthorized) response to an IPP request, then the response MAY contain a “WWW-  
360 Authenticate” header with an appropriate challenge. The next request for this resource formulated by the  
361 IPP client should contain an “Authorization” header specifying appropriate credentials.

## 362 ***HTTP 1.1 Status Codes***

363  
364 The status codes recommended by Keith Carter, and in a subsequent document from Bob Herriot seem  
365 sufficient for a "closed" implementation of IPP clients and servers.  
366

## 367 ***Other issues***

368

369 **Administrative Framework**

370 The care and feeding of IPP client and server implementations should be taken into account during the  
371 design of the protocol. This framework can get very complicated, especially if proxy and security issues  
372 are taken into account. Before reaching final consensus on a protocol definition for IPP, the complexity of  
373 configuring clients and servers should be weighed appropriately.  
374

375 **Print-By-Reference**

376 There is a recent document, co-authored by Keith Moore, entitled “Definition of the URL MIME  
377 External-Body Access-Type” (RFC 2017), that discusses an easy way to support the “Print-by-Reference”  
378 capability that has been discussed in the IPP working group. To avoid replicating the text of the RFC, an  
379 example of the use of this method would look something like this:

380

381 `Content-Type: message/external-body; access-type="URL";`  
382 `URL="http://www.yahoo.com/daily-stock-quotes"`

383

384 This method seems to fit very well with the requirement for “Print-By-Reference”. One aspect of this  
385 method is that, if any secure access is to be applied to the retrieval of the external body, that any and all  
386 security mechanisms would have to be specified (encoded) somehow within the URL string.

387

388