

Issues

1. What is our relationship to Internet Fax? What is their final charter? What is our final charter? Are there any other groups that we need to liaison with?
2. Randy suggested looking at draft-ietf-asid-string-filter-v2-00.txt. Did any one review this and have comments about IPP use of LDAP?
3. On 11/14/96, Roger posted httpnew.ps. Has this been reviewed and incorporated?
4. Have we resolved the issue of "Extending HTTP vs. Use Existing HTTP Methods

Larry Masinter:

What would help is to have a specific proposal of how to extend

># HTTP to have the following DPA/LDPA operations:

># SubmitJob

># CancelJob

># ListJobAttributes

># ModifyJob

># ResubmitJob

>

>This is similar to the discussion we're having in the "Distributed
>Authoring and Versioning" group, which is trying to build a profile
>for doing document management at web sites (e.g., such as FrontPage,
>MKS Source Integrity, Documentum, etc. are doing). In that group, as
>here, the issue is: New methods, or POST with new content?

> You might want to look at the home page for that working group:

>

> <http://www.ics.uci.edu/~ejw/versioning/>

>

>SubmitJob could easily be a POST which returns a job attributes and a
>Location: header that is the Job URL. CancelJob could be done with a
>DELETE of the Job URL. ListJobAttributes might be GET of the Job URL,
>ModifyJob could be a PUT of the Job URL with a new set of Job
>Attributes (or Post). I don't know what ResubmitJob might be, but
>perhaps it's a POST too.

Henning commnets:

The MMUSIC group is currently considering extending HTTP for conference control (see my RTSP draft and a yet-to-be-published conference initiation draft). We'll discuss this with the HTTP WG chair, but our impression is that extending HTTP does not necessarily imply changing the base draft. The advantage is conciseness and the immediate ability to extend web servers to do the job. Your method is then a "first class citizen", with appropriate access control, query about existing methods, etc..

5. Are there any data transfer issues associated with running IPP over HTTP/1.0? Can the contents of a POST be arbitrarily large (several documents of many Mb each)?
6. On 11/15/96, Roger posted ipp-http-stuff.doc has this been reviewed or incorporated?
7. Have we resolved all of the Attribute Syntax issues?

8. Do we need to add a fan in picture back into the document?
9. Some directory entry attributes are set by the Printer. Some are set by the Administrator? Do we know which are which? Do we care?

Notes from Kieth: The original set of Printer object attributes were hardware-oriented i.e. a Printer object that is really a physical printer could specify these attributes. Per our phone call, there are now attributes (i.e. "speed", "quality" or "cost" = "high", "medium", "low") that an administrator - not a physical printer - must specify i.e. It seems that an administrator must determine how to assign "high", "medium" or "low" to "speed", "quality" and "cost" while a printer can only return its actual speed and resolution upon which the administrator determines the appropriate attribute value. "Location" can be specified and stored in some - but not all - printers so in some cases the "location" attribute must also be specified by an administrator. I believe the "speed", "quality", "cost" and "location" attributes help the user more easily locate Printer objects. The point is there are now attributes for a Printer object directory entry some of which are defined by the hardware and others by the administrator. A Printer object that is really a printer would only add/modify the hardware attributes for its directory entry and not add/update the administrator attributes for its entry.

10. Where are we on the proposal to establish, through IANA, of a well known port (380 recommended) for printing via IPP over HTTP?

Dons Comments: Dedicating a port to printing on the surface sounds good. But, not being an expert on firewalls, I ask is that a problem? My understanding is that some firewalls keep things under control by not allowing certain ports to be used (i.e. nothing can be used unless explicitly permitted.) Is this going to be a problem in an inter-enterprise environment? Could someone with a better understanding of firewalls comment on this?

Harry's comments:

I would like to propose we attempt to standardize a well known port for printing via http. The notion would be that whatever we define as printing via http will work over http well known port 80 (along with everything else that flows via http), but, if directed to well known port (say 380 - assuming this port is not already defined), then only PRINTING, not any other form of http, would be expected.

This proposal closely follows the motivation described in the internet draft <draft-mellquist-web-sys-00.txt> submitted June 1996, which recommends well known port 280 for SNMP over HTTP. Read (3) of this draft for a more thorough dissertation on the merits of this approach.

As for text to add to the IPP draft (tonight), I suggest a new head level created between 2 (Distributed Printing) and 3 (IPP Objects) called 2 (Printing via HTTP). Probably, a lot could be lifted from Roger's document and put into this section, I'm not going to try and flesh this out here. Also, in this section, however, we should put some words like...

"It will be suggested (in section 5) that Clients identify Printer objects using an HTTP type URL. One element of this proposal will be to further recommend the establishment, through IANA, of a well known port (380 recommended) for printing via HTTP. The purpose of this well known port would be to distinguish printing from non-printing content. While any acceptable HTTP content could be inter-mixed over HTTP well known port 80,

only HTTP printing would be acceptable on port 380.

The remainder of this draft will define the IPP content for HTTP printing, including IPP objects, operations, naming and attributes."

Carl-Uno comments:

I brought up this subject with Larry Masinter (who is the IETF HTTP Chair) when I met him last week. He claims that the HTTP does not really care about the port number and that hence it would be no point in using a different number for the printing protocol. Can somebody else, who is supplying Web servers, verify this?

Jay comments:

For the sake of global sanity, IANA also registers ports above 1023, as I recall. So, it's probably best we eventually coordinate through them if we are dealing on the Internet.

11. We still have issues about Subjective values versus Objective values, especially in the Directory Entry. What is the set of values for each attribute?
12. On 11/19/96, Roger posted some sample flows. Is this what is in the current I-D? Have we reviewed the encoding of content within an HTTP operation?

Roger comments:

Two things that I noted in doing this:

My http proposal puts the document in a separate "container" rather than shipping it as an attribute as proposed in the current attribute write-up. I think this is preferable as it allows more structure of the data on the wire. This will make it easier to provide real document objects later on and add additional document attributes.

As part of the above, I carry the document format in the document-content header. This is more consistent with http, which puts mime-like content identifiers in all of its headers.

I noted that we have no attribute which describes the length of the queue. One of the main things that I look at when I personally submit print jobs is how busy the printer is. If a printer is very busy, I find one that is not so backed up. What about an attribute called <jobs-in-queue>?

Bob comments:

After thinking about the protocol issue some more, it seems like there are two ways to organize the IPP Job Object in HTTP Protocol, the current proposal or my proposal described below.

1) current proposal: The way proposed in the current document where the job and all document data is a part of a HTTP single entity-body.

2) my proposal: A job is an HTTP document whose top level content type is multipart/mixed (an existing type) or perhaps multipart/ipp (a new type and probably a bad idea). In either case, the first entity has a content type of application/ippjob and contains the job attributes.

Each remaining entity contains a document and they are all standard MIME types, such as application/PostScript or text/plain. Such a document could in the future become multipart/ippdocument with a document attribute entity and a document content entity if a document has attribute overrides. The document data entity could also have a content-type of multipart/alternative (an existing type) if a client wanted to have several representations of a document, e.g. HTTP and PDF and a printer supported such.

The advantage of my proposal is that the document parts are conventional MIME documents which any software could read. Only the application/ippjob and application/ippdocument require special software to process.

Roger comment:

Bob, I am not familiar enough with the real operation of HTTP protocols to know whether your idea is a good one or not. I had thought that I could only send one entity at a time, which would require leaving the connection open and turning the line around several times to complete transfer of the message when sending the document along with the print request. How would this affect performance? Obviously this is not a concern when pulling the document later.

What are the advantages of allowing "any" software to read the document data? Since I only apply the headers when sending the document to be printed I'm not sure what I gain? In my proposal, the document content itself still has a standard Mime type

13. What are we using for Device Id?

Jay's comments:

This issue/concept falls right into the same kind of discussion we recently conducted in the PWG regarding printer model identification, particularly with regard to supporting print job submission and job formatting (ie, PPD-like file support,etc).

Our experience tells us that users want to be able to view "real world" identification information, such as Make, Model and Type (although I'm not quite sure what "Type" would imply depending on the context).

The previous related discussion pointed out that an OID must be provided to unambiguously and succinctly define the specific printer model; within the Printer MIB context, this value is defined as the printerDeviceID object in the HR MIB. During the discussion someone questioned whether this was a workable solution, since the application would have to have knowledge of the OID mappings to determine the model type. Those of us who develop such application software explained that the OID approach was indeed a reasonable solution, citing the need for a clean, simple syntax that has both hierarchical potential and distributed administration within the name space. All of this is achieved via the standard SNMP SMI concept of the "Enterprises" tree.

If we use a P1284 Device ID for the IPP, will these same naming features exist? Can someone brief the PWG on the name space administration of the P1284 Device ID? For example, how would an app developer know which vendors have which products registered for which IDs?

14. What is the status of overlap with WEBDAV?
 15. What is the current write up on printer-name vs. URL and job-identifier vs URL?
 16. What is the current security story?
 17. Have we resolved all of the number-up, embellishments issues?
 18. Are we cleared up on the Job Retention time issues (cancel, etc.)?
 19. What is the status of the shortest-job-first vs smallest-job-fisrt issue?
 20. Do we agree with the proposal to have no attributes that have no values?
 21. Was there a proposal to change the name of IPP due to other IPPs being found?
 22. Questions from reviewers:
 - I don't see any references to fan-fold paper anywhere.
 - It appears there is currently nosupport fo accounting. One can certainly argue that this is a function of security which you say will be addressed later.
 - How do you plan to support color calibration and other features like what is provided through PPD files?
 - I suspect that PostScript printing control (like spooling, status queries, etc.) are not supported as part of the PostScript Content-Type, right? It probably should be spelled out.
- HTTP methods are case-sensitive (HTTP/1.1, Section 5.1.1); thus, your example needs updating
- Wouldn't it make more sense to define a content-type: applicatioipp (or similar), so that the request would be a valid HTTP request? (The print request is then part of the entity).
 - Have you considered the alternative of defining a new method, say, PRINT, with entity headers appropriate for the new task?
23. On 11/25/96 Roger posted an attribute summary. Has this been reviewed?
 24. We need to clarify Job Templates. Who wants to write up a new summary? With examples?
 25. Has anyone reviewed draft-mellquist-web-sys-01.txt? Tom suggests that there is some overlap with Managment and IPP/2.0?

