

Minutes of IPP Working Group Meeting

May 26-27, 1999

1. Meeting Attendees

Lee Farrell	Canon Information Systems
William Zhang	Canon Information Systems
Greg LeClair	Epson
Atsushi Uchino	Epson
Mike Moldovan	Genoa
Shivaun Albright*	Hewlett Packard
Ben Brezinski	Hewlett Packard
Ira McDonald*	High North Inc
Ron Bergman	Hitachi-Koki
Harry Lewis	IBM
Keith Moore*	IETF
Yuji Sasaki	Japan Computer Industry
Don Wright	Lexmark
Hugo Parra	Novell
Eric Random	Peerless
Chuck Adams	Tektronix
Tom Hastings	Xerox
Bob Herriot	Xerox
Carl-Uno Manros (Chairman)	Xerox
Pete Zehler	Xerox

* attended via teleconference

2. Day 1

Carl Uno-Manros opened the IPP meeting and provided the suggested agenda topics:

- Introduction of the IPP Test Specification draft
- Final discussion of any remaining IPP/1.1 Issues
- Scope for next version of the IPP Implementer's Guide
- Introduction of the new drafts for IPP Notifications
- Suggestions for any other IPP extensions
- Should we start up work on an IPP MIB specification?
- Discuss a revised IPP WG Charter for the IETF (requested by the IESG chair)

2.1 IPP Test Specification Draft

Mike Moldovan explained that Genoa has submitted a draft Test Specification for review and comment by the IPP WG members. (This document is available at: ftp://ftp.pwg.org/pub/pwg/ipp/new_TES/IPP1.0-Test-Spec/IPPSPEC.ZIP.) The document is very large—more than 1300 pages. Mike identified that the document is organized into two main sections: Mandatory Tests and Optional Tests.

He briefly referenced a few General Testing Issues that are identified in the document.

Q: Is there a “minimum test environment” identified?

A: No, the required environment would be specific to the individual implementation.

Q: Will this document need to become a Standard?

A: It is up to the group. It is doubtful that the IETF would support the Test Specification document as an Internet-Draft or RFC. However, Genoa would (at least) like the PWG to endorse the document.

Tom Hastings pointed out that there is a lot of redundancy in the current draft. Mike agrees that much of the text could be reduced by using references to repeated test steps.

Mike said that he would appreciate any feedback on the document and the test cases identified. He will attempt to provide an updated version of the specification before the next meeting in July.

2.2 Final Discussion of Remaining IPP/1.1 Issues

Tom identified that six Issues remain open.

2.2.1 ISSUE 17 – Client display of absolute time for job attributes

Alternative 1: An IPP Printer MUST support “time-at-created(integer)”, “time-at-processing(integer)”, and “time-at-completed(integer)” Job attributes.

Alternative 2: An IPP Printer MUST support either:

- a) the set of 3 “time-at-xxx(integer)” and “job-printer-up-time” Job attributes; OR
- b) the set of 3 “date-time-at-xxx(dateTime)” Job attributes and the “printer-current-time(dateTime)” Printer attribute

In favor of ensuring interoperability, the group agreed to require Alternative 1.

2.2.2 ISSUE 30 – Should “job-state-reasons” and “printer-state-reasons” be REQUIRED for an IPP/1.1 Printer?

At the last meeting, job-state-reasons was voted to be REQUIRED. A question was raised about whether any of the specific values should also be REQUIRED. Tom will reflect the resulting details in the next draft update.

2.2.3 ISSUE 31 – How should we indicate a (RIPped) job that is waiting for the Marker?

Proposed modifications for the text on three states (“idle”, “processing” and “stopped”) were considered. The group spent much time discussing the risks of trying to provide “too much detail” about printer states, concerned that one printer implementation or another won’t exactly fit the model.

Carl-Uno suggested that a “simple” approach should be adopted, with any clarifications for complicated scenarios placed in the Implementer’s Guide. Chuck Adams extended this view, with the additional suggestion that the results are compatible to the IPP/1.0 definitions.

Most people agreed that the definition of “processing” used in IPP/1.0 needs additional clarification because it presumes that only one job can be in the processing state at a time. However, it was generally agreed that such clarification detail should be moved to the Implementer’s Guide.

After more than an hour of discussion, the group agreed to defer the topic until the next day. At that point (and after a bit of further refinement) the group agreed on the following definitions:

- Idle: Indicates that new jobs can start processing without waiting.
- Processing: Indicates that jobs are processing; new jobs will wait before processing.
- Stopped: Indicates that no jobs can be processed and intervention is required.

2.2.4 ISSUE 35 – What error code should be returned on Print-URI or Send-URI if the document is not accessible?

Pete Zehler noted that this issue raises the fact that we should really have a “document” object in our model in order to address it properly.

It was suggested that an optional “document-access-error” Job Description and “document-access-error(text)” Print-URI/Send-URI attribute should be added. For protocol errors, standard values would be the “reference-uri” ‘uriScheme’ followed by a hyphen (-) and an error number (e.g. “http-xxx” and “ftp-xxx”).

The group felt that the error information should be “passed along” as an error number (or appropriate value) without trying to map its meaning, but identifying the protocol in which the error originated. Carl-Uno suggested that a URL should also be included in the provided information returned.

After more discussion it was decided that the proposal (and the related Issue) does not adequately address document errors within multi-document Jobs.

The next day, Tom made a new proposal that was accepted by the group:

- An optional “document-access-error” Job Description and “document-access-error(text)” Print-URI/Send-URI attribute should be added. For protocol errors, standard values would be the error code in parentheses, followed by the URL.

Additionally, the group accepted the following proposal:

- Add a “detailed-status-message (text)” operation attribute and Job Description attribute that contains non-localized detailed information.

2.2.5 ISSUE 36 – Don’t require 1.0 support and add REQUIRED “ipp-version-supported” attribute

The proposal suggested was to consider adding “option b)” below:

If the major version number matches, but the minor version number does not match any of the keyword values, the Printer MUST either:

- a) reject the request and return the “server-error-version-not-supported” status
- b) OR accept the request, process it normally, and return the normal status codes

A few people said that *if* the attribute is added, then option b) could be included. However, there were a few strong opinions that the attribute is unnecessary—there are already other (simpler?) methods available to determine the supported version level.

After another long discussion, the following was agreed:

If the IPP object does not support the major version, the object **MUST** respond with a status code of “server-error-version-not-supported.” If the IPP object does not support the minor version, it **SHOULD** accept and perform the request if supported, and return the nearest version number supported.

3. Day 2

3.1 Scope for Next Version of the IPP Implementer’s Guide

Carl-Uno suggested that the next draft of the Implementer’s Guide should include information on both IPP/1.0 and IPP/1.1. There were no objections.

3.2 Should We Start Work on an IPP MIB Specification?

The purpose of an IPP MIB would be to “manage the IPP protocol.” However, many of the attendees did not feel that there was any need for protocol management. It was generally agreed that we should respond to the IESG (Keith Moore) on this topic by saying that any management that might be necessary is already supplied by the Printer MIB or HTTP capabilities.

3.3 Discuss a Revised IPP WG Charter for the IETF (Requested by the IESG Chair)

It was noted that a new Charter should include the activities planned for the next 6-12 months. There was no attempt to draft a new IPP Charter at the meeting.

3.4 IPP Issues – cont’d

3.4.1 ISSUE 32 – Is Digest **REQUIRED for an IPP Client and an IPP Printer to support?**

3.4.1.1 Teleconference – the Prequel

Carl-Uno summarized the current status. A majority of the IPP members seem willing to have Digest Authentication mandated at both ends. However, some members have expressed concerns on this topic. They would prefer to have authentication as a **SHOULD**, rather than a **MUST**. At least one person believes this requirement should be driven by the market—not by the IETF.

Concern was expressed about the related overhead effort required for documentation and support, as well as code size—especially if it is expected that printers will be configured to ignore the security capability. A few individuals expressed a preference to allow their company to make the space tradeoff associated with supporting or not supporting security—and still be able to claim conformance with the IPP standard.

It was suggested that two levels of conformance could be defined: one for “IPP protocol”, and one for “IPP protocol with security support.”

After some discussion, the group took a poll of the current opinion:

Digest Authentication as a **MUST** requirement – 4 in favor, 6 against

3.4.1.2 Teleconference with Keith Moore (IETF Applications Area Director)

Keith Moore explained the rationale behind the IESG position that authentication is necessary. He said that problems have occurred in the past with protocols that allow passwords to be passed in clear text. He finds it hard to believe that there are any printers that “don’t need protection”—even the ones used at home.

Shivaun Albright explained that she would like to see security as an optional functional feature of the protocol, not a mandated part of it.

Keith expressed that he (and the IESG?) does not understand why anyone feels that there is a “significant cost or burden” to implementing this capability. His investigations seem to indicate that the implementation burden is small (the estimate of 4K was referenced.) He suggested that if anyone could provide a reasonable explanation as to why this is “a significant burden,” then the IESG might be more understanding.

Chuck Adams explained that it is not just the absolute size of the implementation that is critical—it is the flexibility associated with trade-off decisions about which protocol(s) to support.

Don Wright asked if two separate levels of conformance could be identified, one for IPP and one for IPP with security. Keith indicated that the IETF community would be reluctant to do this.

In general, Keith believes that the IESG will not support allowing a Standards Track document that does not mandate authentication. As an alternative, Keith suggested that in addition to the IPP/1.1 documents (with mandatory authentication), a separate Informational RFC could be published to define an “IPP Lite” protocol that is suited to lower-end devices.

Again, Keith expressed that he does not feel that the IESG will change their opinion on this issue. He indicated that at the very least, the IETF Security Area Director would need to be convinced.

Shivaun asked about the alternative in which it is specified that the Client MUST and the Server SHOULD support authentication. Keith did not think this would be accepted either—even after he began to understand that it might not be practical for a low-end printer to implement authentication. However, he did say that he is willing to support the group in their attempt to get this alternative accepted. He encouraged the group to submit the IPP/1.1 draft with the “SHOULD” statement, but advised that an explanation/justification should also be submitted to the IESG. Keith suggested that this explanation could be submitted in the form of a letter or e-mail.

Don, Shivaun, Carl-Uno and Chuck all volunteered to work together to create the explanation and/or justification message. They will post a draft for general review prior to submission to the IESG.

Don suggested that this explanatory information should be associated with the document by placing it in the “Security Considerations” section of the specification. This would avoid the possibility of losing the association between the specification and the explanation. As a fall-back alternative, it was suggested the information could be placed in the Implementer’s Guide.

3.5 IPP Notifications

Tom Hastings provided a brief background of previous discussions about IPP Notifications. Because IPP traffic goes from client to server, it is difficult to support sending information in the other direction—unless a Printer implements an HTTP client. Since it is desirable for the client to obtain error condition information and certain job status information, it has generally been agreed that Event Notification should be supported.

The group reviewed the previously published document on Notification Requirements. (This document is available at: ftp://ftp.pwg.org/pub/pwg/ipp/new_NOT/notify-req-00.pdf .) As the requirements were evaluated, several topics were identified for future consideration and/or possible modification to the Requirements document:

- Additional events:
 - * job state changes, like “paused”
 - * device problems for which the job is destined
 - * job (interpreter) issues
- “Out-of-band” registration—-independent of job submission
- Granularity and categorization of events—in context of anticipated event frequency
- Do we really need “mixed notification”? Or is machine-readable sufficient for the specification? What about supporting e-mail?
- Should IPP Notifications be able to support reliable accounting (i.e. guaranteed delivery—not just delivery over a reliable transport)?
- Policy: should printing be refused if notification can’t be guaranteed?
- Quality of service: requested by subscription, or fixed by printer?

Don Wright asked whether we really want to include Notifications at all. Given that we are being requested to update the Charter, he wonders if we should remove the topic altogether from our current Charter.

The group reviewed the latest document draft on IPP Event Notification. (This document is available at: ftp://ftp.pwg.org/pub/pwg/ipp/new_NOT/ipp-notification-990518.pdf.) Tom led a discussion about the document content and reviewed the Issues identified in it. The discussion led to many new issues that were recorded by Tom, and will be noted in his next update. However, no significant decisions or agreements were reached.

3.6 IPP Extensions

A new operation called “Set-Job-Attributes” has been suggested for consideration. It would allow the client to modify attributes of previously submitted jobs. For example, if a large job aborts in the middle, “page-ranges” or “copies” could be modified and the job could be restarted.

The group discussed possible issues that might need to be addressed with this type of operation. Some items that were raised include:

- Prohibiting job attributes from being set after a job completes
- Should setting (grouped) attributes be an atomic operation?
- How to handle dependency conflicts resulting from new values set? Would a new job “validation” step be necessary?

Additional operations that were submitted by Carl Kugler via e-mail were also reviewed and briefly discussed:

- Set-Printer-Attributes
- Cancel-Current-Job
- Shutdown
- Backspace-Current-Job

IPP meeting adjourned.