

Minutes of IPP Working Group Meeting

April 4-5, 2000

1. Meeting Attendees

| | |
|-------------------------|---------------------------|
| Koichi Masegi | Canon |
| Shigeru Ueda | Canon |
| Lee Farrell | Canon Information Systems |
| Bill Wagner | DPI/Net Silicon |
| Fumio Nagasaka | Epson |
| Masanori Tanizaki | Epson |
| Atsushi Uchino | Epson |
| Michael Wu | Heidelberg Digital |
| Ron Bergman | Hitachi-Koki |
| Yoshinori Nakai | JCI |
| Yuji Sasaki | JCI |
| Stuart Rowley | Kyocera |
| Don Wright | Lexmark |
| Kazuya Torikai | NEC |
| Hironori Goto | Panasonic |
| Minoru Ozaki | Panasonic |
| Paul Moore | Peerless |
| Howard Sidorsky | Peerless |
| Yuusuke Wada | Peerless |
| Nick Webb | Peerless |
| Satoshi Fujitani | Ricoh |
| Norimi Kawashima | Ricoh |
| Michael Reffke | SEH |
| Werner Schweer | SEH |
| Matthias Wolff | SEH |
| Eric Olbricht | Sharp |
| Craig Whittle | Sharp |
| Satoshi Ikawa | Techno Scope |
| Lin Chin Wei | Techno Scope |
| Bob Herriot | Xerox |
| Carl-Uno Manros (Chair) | Xerox |

2. Day 1

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

- Proposed Charter for IPPEXT
- Discussion of new IPP Operations
- IPP Event Notification
- IPP “Open Source”
- Document review

2.1 Registering Print Languages as MIME Types.

At the October meeting, it was agreed that Tom Hastings would pursue the registration of different vendor document print languages as MIME types—and that he would send an e-mail on the progress of this effort. To date, he has not made any progress on this task.

Carl-Uno asked if anyone else would take on this activity in Tom's place, but no one volunteered. Carl-Uno said that he would raise this topic on the e-mail list.

It was suggested that the Implementer's Guide could include a section that explains the process of registering a print language as a MIME type. Then the responsibility of registration could be passed back to the individual companies.

2.2 IETF Plenary

Carl-Uno announced that Keith Moore, the IETF Applications Area Director, has resigned his position. He has been replaced by Ned Freed—an individual that has been associated with the IETF for many years. Ned has acknowledged the fact that the IPP documents have been “in the queue” for standards progression for quite some time. He has suggested that it is ok for Carl-Uno to “periodically remind him” about the documents' status.

Carl-Uno reported that the IESG has not yet decided whether the IPP activity should continue its activity as a new Working Group (“IPPEXT”) or not. If so, a new Charter would need to be established and approved. The other alternative is to simply continue as the existing WG by updating the current Charter document. Carl-Uno noted that either approach is acceptable—but he is concerned that *some* decision needs to be made soon. In the meantime, the group will continue to make progress regardless of the IESG decision.

Carl-Uno reported that the Authentication, Authorization and Accounting (AAA) Working Group is not going to resolve “secure print by reference” as he had hoped. As a result, the IPP WG will leave this as an unresolved issue.

2.3 Requirements for IPP Notifications

Carl-Uno led a review of the IPP Notification Requirements document:

ftp://ftp.pwg.org/pub/pwg/ipp/new_NOT/ipp-notification-requirements-990811.pdf

ISSUE Ok if there isn't a way for an End-User to submit an empty Per-Printer Subscription, in case such a Subscription slot is a scarce commodity, and then enable the Per-Printer Subscription when the data arrives and disable later without deleting the subscription?

The group agreed that this Issue should be removed from the document.

ISSUE Ok if spec doesn't have means for a Notification Recipient acknowledging receipt of a notification to the Notification Source?

The group agreed that this Issue should also be removed from the document.

2.4 IPP Event Notification Specification

Carl-Uno led a review of the IPP Notification Specification document:

ftp://ftp.pwg.org/pub/pwg/ipp/new_NOT/ipp-not-spec-000308.pdf

ISSUE 01 – Once a number of delivery solutions have been developed and evaluated, we may want to make one or several of them REQUIRED for implementation to ensure a minimum set of interoperability. Which one or ones should be REQUIRED?

The group agreed that an informative annex should be written to include the definitions of the various delivery methods. In that section, the document could also indicate that one method is REQUIRED—however, the specific method is identified in a different document—not this one.

Carl-Uno noted that items 2 and 3 in Section 7.1.2 might need some MUST statements included. This should be evaluated by the authors.

2.5 Job Progress Attributes

Ron Bergman led a review of Job Progress Attributes document:

<ftp://ftp.pwg.org/pub/pwg/ipp/meeting-pdf-000404/ipp-job-prog-attr-000202.pdf>

ISSUE 01 – Should we change the name from "sheet-collate" to "sheet-uncollate", since the absence of the attribute (and non-support of this attribute) is more likely to indicate collated sheets and so should be the 'false' value of the attribute, rather than the 'true' value?

The group agreed that no change is necessary.

ISSUE 02 – Should we change the "sheet-collate" data type from 'boolean' to 'type2 keyword' so that it could take on more values? This would also help with the name, say, "sheet-collation (type2 keyword) with values: 'uncollated' and 'collated'. The "sheet-collation-supported" (1setOf type2 keyword) would be more usual, than the unusual '1setOf boolean' also. In the future, we could define two collated values: 'multi-pass-collation' and 'output-bin-collation' to indicate which form is requested and/or supported, since some Printers MAY want to support both.

There was some confusion as to what was meant by the terms “multi-pass-collation” and “output-bin-collation.” No one at the meeting had a clear understanding of the terms. There was also some confusion as to why this term would be defined in a Job Progress document. [At one point, someone even asked if we really need this document—or whether it should be part of the IETF standards document suite.]

After much discussion, the group agreed that we should change the data type to a type2 keyword with the values “collated” and “uncollated.”

ISSUE 03 – If we change the attribute syntax to 'type2 keyword' should we have several values for the collated case now, i.e., define: 'multi-pass-collation' and 'output-bin-collation', instead of just 'collated'?

No.

ISSUE 04 – Or should the attribute syntax be 'type2 keyword' to go with "multiple-document-handling(type2 keyword)", instead of the Job MIB enum syntax?

No, keep it as an enum.

ISSUE 05 – Or should we use the IPP out-of-band 'unknown' value (see [ipp-mod] section 4.1) instead of this unknown(2) enum Job Monitoring MIB value, i.e., "job-collation-type" (type2 keyword) instead of "job-collation-type" (type2 enum)?

No – leave it consistent with the Job MIB.

2.6 IPP Event Notifications

As a review of the current status, Carl-Uno provided a short description of the four notification delivery methods currently under consideration:

- ippget: use IPP to get (poll) accumulated event notifications
- indp: use INDP over HTTP to send event notifications
- mailto: e-mail an event notification
- snmpnotify: target of URL is an SNMP notification receiver

Carl-Uno also mentioned that he was interested in the possible use of “instant messaging” as a notification delivery method. However, he was disappointed to report that the IETF is unlikely to produce a standardized method of instant messaging in any reasonable timeframe.

The Printer:

- saves each event for a fixed amount of time
- supports the IPP Get-Notifications operation
- requires clients to poll for event notifications with IPP Get-Notifications operation
- responds to the client with:
 - * one or more event notifications
 - * the recommended time to poll again
 - * the lease time of future event notifications

There were a few HTTP Issues that were raised at the IETF Plenary meeting about the ippget: delivery method:

- In order for a client to receive events as they occur, should there be an operation with a single HTTP POST where the Printer returns event notifications in multiple response-parts spread out over several minutes?
Keith Moore said that this approach might work.
- Should this operation use HTTP GET instead of HTTP POST?
No answer or recommendation was established.
- Should each response-part be a separate message body in MIME multi-part?
At the IETF Plenary meeting, it was determined that MIME multi-part should not be used for delivery notification.

- Do the lease and server recommended times make this polling mechanism a reasonable alternative to Printer-initiated event delivery methods (i.e. mailto, indp, snmpnotify)?

2.7 IPP Notify Poll

Carl-Uno led a review of the Notification Polling Method document:

ftp://ftp.pwg.org/pub/pwg/ipp/new_NOT/ipp-notify-poll-000308.pdf

ISSUE 01 – The 'ipp' is a convenient reuse of 'ipp', but does it imply the existence of a print service at each client that is not a reality?

This method will be named “ippget.”

ISSUE 02 – Now that the Get-Notification operation does not affect the Event Notifications in the Printer, it should not require write access to access them.

Correct. Issue closed.

ISSUE 03 – Is it possible for this operation to have an option that causes it to delay completing its response? It would initially returns all existing Event Notifications. Then it would return additional notifications as they occur for some period of time. The client would receive these Event Notifications as they occur. The question is whether http servers or proxies would behave in this manner so that the client would get the Event Notifications without delay after they are sent by the http server? It has been suggested that the Printer send each burst of Event Notifications be in a separate message body where each message body is part of a multipart MIME content-type.

This Issue was attributed to Harry Lewis. The group felt that if he is sufficiently interested, he should create a separate document describing an operation to achieve this additional functional capability. The Issue was removed.

ISSUE 04 – The "notification-recipient" option allows a client to group multiple Subscription-ids with a single URL. A client might decide to use the same URL for all subscriptions for a user, or it might have a separate URL for each client program. In addition a client might use an URL belonging to some other known user and let that user access Event Notifications using that URL. Is the above option useful?

During the discussion it was decided that the description should be moved earlier in the document. The group agreed that option is useful and should remain in the document.

ISSUE 05 – Does the mechanism described in the above paragraph describe a useful option that "notification-recipient" alone cannot do? Should this case be an error instead?

The group agreed to eliminate the paragraph above the Issue, and state that the client should supply at least one of the three attributes—otherwise it should be considered an error.

ISSUE 06 – Is it better if "notification-recipient" is the only way to request Event Notification? If so, this behaves more like other notification delivery methods where a recipient receives those and

only those events with its delivery URL. Furthermore, if there is a single mechanism of "notification-recipient" for a client to specify Event Notifications, a Printer can better optimize event-leases because it knows which events will be accessed together. If client can specify subscription-ids, each request can contain a different mix of subscription-ids.

Although Bob Herriot thinks it might be difficult to calculate lease times, the group agreed that this Issue should be removed (for now) without any change to the document. It was noted that more experience would be gained as implementations are developed.

2.8 The INDP Event Notification Delivery Method

Carl-Uno led a review of the INDP Event Notification Delivery Method document:

<ftp://ftp.pwg.org/pub/pwg/ipp/meeting-pdf-000404/draft-ietf-ipp-method-000229.pdf>

ISSUE 01 – Should there be an Unsupported Attributes group to return attributes that are not supported to the client?

No.

ISSUE 02 – Use the same status code space as IPP, namely:

- "successful" - 0x0000 to 0x00FF
- "informational" - 0x0100 to 0x01FF
- "redirection" - 0x0200 to 0x02FF
- "client-error" - 0x0400 to 0x04FF
- "server-error" - 0x0500 to 0x05FF

The question was raised as to whether all of these status codes are truly necessary. However, the rest of the document review was deferred due to lack of attendance by the authors.

2.9 Notification Delivery Protocol

[This document review was deferred due to lack of attendance by the authors.]

2.10 Notification over SNMP

Craig Whittle presented some slides proposing SNMP-based notification. He discussed many different Job Monitoring traps that could be mapped to IPP events. Craig noted that most printers already support SNMP (although not all of them support traps.)

Don Wright suggested that SNMP notification has limited applicability because it is usually stopped by firewalls. It was generally agreed that this could not be used as the primary (i.e. mandatory) notification scheme.

2.11 Notification via E-mail

Carl-Uno led a review of the Mailto: Notification Delivery Method document:

ftp://ftp.pwg.org/pub/pwg/ipp/new_NOT/ipp-notify-mailto-000309.pdf

ISSUE 01 – Is this SMTP terminology correct?

The paragraph should be expanded and clarified more. The group agreed that the terminology is probably not correct—but no one knew the correct terms.

ISSUE 02 – Is it a good idea to list each Subscription object attribute in this spec with the applicability to this delivery method? If yes, should all delivery method specs also do it this way?

No, it is not a good idea. No change to the document.

ISSUE 03 – What should we say about any mailto parameters, if any? For example, if you want to send over secure mail, etc.

The group agreed that the specification should state that it is up to the Printer to determine how it should deliver the e-mail. No additional parameters are required.

ISSUE 04 – Do we want to define any IPP-specific mailto parameters to this document?

No.

ISSUE 05 – Ok that we made "subscriber-user-name" be REQUIRED for the Printer to support and indicate that the client MUST supply the "requester-user-name" operation attribute when the delivery method is 'mailto:', in case the Printer does not have a more authenticated printable name?

[There was a long discussion regarding the potential problems associated with spam—and the difficulties in preventing it.] It was suggested that there should be an e-mail address included that provides a destination for “failure of delivery” messages. After more discussion, the Issue was deferred.

Carl-Uno suggested that the third paragraph in Section 10 (Security Considerations) should be reworded or eliminated.

ISSUE 06 – What if "notify-format" is 'text/plain; charset=utf-8', does that have to be sent as a mail attachment, since it isn't 'text/plain' which assumes charset=us-ascii, or can it be sent as the body of the mail message properly identified as 'text/plain; charset=us-ascii'?

This issue was deferred. Someone (perhaps Ned Freed?) needs to do more research in this area and make a recommendation. It was noted that different concerns arise for different parts of the e-mail message (e.g., subject line, from headers, message body, etc.)

ISSUE 07 – Is there any way that a Notification Recipient could reply to the message in such a way as to cancel the subscription and thereby solve the spam problem?

No.

2.12 Job and Printer Set Operations

ftp://ftp.pwg.org/pub/pwg/ipp/new_OPS/ipp-job-printer-set-ops-000323.pdf

This document will be submitted for Working Group Last Call soon. There was no review.

2.13 Set 2 and Set 3 Operations

The group questioned whether or not any vendor would actually implement these operations. There was also a general question of procedure with regard to submitting new Operations proposals. A few people are concerned that the additional operations add too much complexity to justify the expected benefit. Until someone actually (re-)confirms that the proposed operations are desirable, it was suggested that they could be deferred.

Carl-Uno will issue an e-mail soliciting comments about which operations are still of significant interest.

There is an ongoing e-mail discussion regarding a proposal for a Move-Job operation. If included, it would be added to the Set 2 Operations document.

Both the “Set 2” and “Set 3” documents will be given more descriptive titles.

3. Day 2

3.1 Interoperability Bake-Off

Carl-Uno opened the second day with a discussion about the next IPP interoperability test “Bake-off”. He said that the next bake-off is planned for a test of IPP/1.1 some time in October. This test should include the new Set operations for Print and Job and (at least) the primary Notification method.

Most of the people in the room indicated that their company is interested in participating in the Bake-off. Each company should send a message to Pete Zehler to let him know about their interest. Currently, the group is still looking for a volunteer company to host the Bake-off.

3.2 IPP “Open Source”

Carl-Uno said that he is interested in pursuing an “open source” initiative for an IPP client. He mentioned that Corel had indicated some interest in this idea several months ago, and he is still in contact with them.

Carl-Uno hopes that Xerox will be able to contribute their client for this effort. He is discussing the logistics of this activity with his company lawyers.

Representatives from Peerless, Easy Software Products (makers of “CUPS”), and a few other companies have also indicated interest in participating in this effort.

Xerox code is in Java. The CUPS product is written in C. The group will need to determine how many versions of client software would be appropriate. Carl-Uno believes that having both Linux and Windows platform versions would be desirable.

Don Wright says that the PWG does not currently have a model in place for adopting this project as an official “PWG product.” However, he feels that the organization would be interested in establishing such a precedent.

Paul Moore suggested that an “Open Source Client” BOF meeting could be arranged (possibly in the evening?) during the PWG meeting in May.

Carl-Uno will launch a discussion of this effort on the PWG IPP mail list.

3.3 IPP and UPnP Printing

Paul Moore explained that he had proposed to the UPnP Printing group that the IPP print model could be mapped to UPnP. Unfortunately, the current version of UPnP has various limitations that cannot support the full model. At the previous UPnP meeting, it was suggested that a “limited subset” of the IPP model could be applicable to the UPnP environment. Paul has recently posted an update to his original proposal on the UPnP MSN Communities message board.

Should the “Open Source” IPP client include support for UPnP discovery? This should be investigated.

It was noted that UPnP uses GENA (General Event Notification Architecture) as its event notification method. Should this be included in the set of event notification methods for IPP? Paul suggests that it is too early to make any decisions on this question. He feels that it would be better to “wait and see” until the Microsoft offering becomes more solid.

3.4 IPP Over Other Transports

Paul asked if the IPP group is interested in considering the development of IPP over other transports. He feels that infrared is an interesting consideration. He also mentioned that IEEE 1394 is another possible option.

It was generally felt that while these are interesting possibilities, the group would probably not pursue them until someone submitted a specific proposal.

3.5 Exception Attributes

ftp://ftp.pwg.org/pub/pwg/ipp/new_EXC/pwg-ipp-exceptions-model-000131.pdf

There were no Issues in this document, and no one had any comments.

3.6 Production Printing

Carl-Uno led a review of the Production Printing Attributes—Set 1 document:

ftp://ftp.pwg.org/pub/pwg/ipp/new_PPE/pwg-ipp-prod-print-set1-000207.pdf

There was a question about the appropriateness of Section 3.2.4, “insert-sheet-default attribute is not defined.” It appears to discuss an attribute that does not exist. The section should be clarified further or removed.

ISSUE 01 – Should we change the name from "collate-sheets" to "uncollated-sheets", since the absence of the attribute (and non-support of this attribute) is more likely to indicate collated sheets and so should be the 'false' value of the attribute, rather than the 'true' value?

No.

3.7 Collections

Carl-Uno led a review of the Collections Attributes Syntax document:

ftp://ftp.pwg.org/pub/pwg/ipp/new_COL/ipp-collection-attr-syntax-000331.pdf

ISSUE 01 – In attribute groups [ipp-mod] allows a Printer either (1) to reject a request with duplicate named attributes OR (2) to choose exactly one of the attributes as the one to be used. Should we REQUIRE the Printer to reject duplicate named attributes in a collection value as stated above or allow the Printer to choose one member attribute as a second alternative as we do with attribute groups?

Don suggested that the wording should be changed to: “If such a malformed request is submitted to a Printer, the Printer MAY [not MUST] reject the request with the 'client-error-bad-request' status code.” He feels that there might be the possibility for a client to get in the situation of not being able to get *anything* to print.

The group agreed that the Printer should be given the opportunity to decide whether or not it should accept or reject the request. The text should be re-written to be consistent with the Model document.

ISSUE 02 – The example contains a 1setOf collection and a nested collection, but does not contain a 1setOf member attribute. Should there be four separate examples that show a simple collection, a 1setOf member attribute, a 1setOf collection, and a nested collection?

Yes.

ISSUE 03 – Since this is intended to be a standards track document, do we also register the attribute syntax with IANA?

Yes. We should allocate a number, and after the document last call it should be registered with IANA.

Bob Herriot explained that the encoding for Collections has changed. Previously, a Collection was encoded as a sequence of groups. This approach could result in a situation where a name inside the Collection is the same as a name outside the Collection. This could cause problems for legacy (non-Collection aware) parsers.

The new encoding uses an approach that looks like a 1setOf to a legacy parser. This allows it to “skip over” the Collection as a single unknown attribute.

During the review of the document, it was suggested that the examples of this encoding should be modified to improve clarity. More examples and more comprehensive examples will be added. It was also

suggested that colored text could be used as a means of indicating levels of nesting within the Collection examples.

3.8 Mail Encodings

In response to an e-mail message about Mail encodings (ref: Yuji Sasaki message of April 4), Ned Freed has explained that RFC 1652 describes some extension that expands beyond the use of 7 bit ASCII. He suggests that IPP needs to consider the use of UTF-8 instead of UTF-7.

However, the concerns related to legacy mail systems (that do not support UTF-8) are still significant. On a practical basis, this market segment cannot be ignored. Device developers that want to ensure interoperability with many of the existing systems might still choose to support 7-bit.

IPP Meeting adjourned.