Printer Working Group
Alexandria, VA
April 17/18, 1996

Attendees:

Jay Martin - Underscore
Don Wright - Lexmark
Tom Hastings - Xerox
Harry Lewis - IBM Pennant
Randy Turner - Sharp
Ron Bergman - Dataproducts
Binnur Al-Kazily - HP
Jeff Dunham - HP
Bill Wagner - DPI
Bob Setterbo, Adobe
Rick Landau, DEC
Atsushi Yuki, Kyocera
Ovidiu Rancu, Genoa

The meeting was opened at 8:30 by Binnur.

The future meeting plans were presented:

May - Chicago
June - Montreal
July - Seattle

There had been no discussion on the reflector concerning the reorganization of the group. Action was deferred to the end of the day.

A discussion of what was necessary to advance the draft MIB. A discussion on whether a specific printer MIB management application was needed. There was some belief that a simple MIB browser is all that is necessary to test the MIB implementations. Genoa discussed using their test suite as a management application that could be used to test the printers.

Randy Turner told the group the operational experience with the MIB (i.e. products that use the MIB but perhaps only with the same manufacturer's software) should be identified and that experience documented for the IETF in order to move to the proposed state.

The group began discussing the process of getting new enumerations defined. There has been a long discussion on the e-mail reflector concerning a request from DPI on a new TCP channel enumeration. One of the problems with channels is that the expectations for channels is unclear. Some want to be able to use channels to understand completely how to print to a printer while others have lower expectations for channels.

The registration requirements for an enum:
- descriptive name
- description
- published spec
- contact
- submitter

will be added to the front section of the MIB.

The group discussed the need for available documentation for an enumeration. Some thought that documentation must eventually be made available while others believe that is not a requirement. Tom Hastings proposed that an enum request should indicate whether documentation is or soon will be publicly available or if that enum represents a private value that will not be documented publicly. Another concern expressed is the lifetime of this group. How will the enum type 2 process continue after this group no longer meets? One of the values of the type 2 enum is to prevent duplication of enums.

Jay Martin moved to change the channel type enums from type 2 to type 3. The motion was seconded. The motion was tabled until after the more global channels discussion.

After a break, Randy Turner began discussing work he was doing in the area of channels. Randy discussed the "Network Services Monitoring MIB" which is currently being developed. This work was started as RFC1565. (proposal #1)

- Advantages: Open, experience in implementation, we can expand
- Disadvantages: All mandatory, URL instability

Randy then began discussing RFC1573 "Evolution of the Interfaces Group of MIB-II" and how it might be useful in our work on channels. This MIB has added a way using linked lists to deal with layers and sub-layers of protocols. (proposal #2)

- Advantages:
- Disadvantage: Difficult to implement in the management station

The group broke for lunch at noon.

The group resumed discussion on channels at 1:15PM following lunch.

Randy's third proposal was to take the best of the other MIBs and incorporate those object into the printer MIB. (proposal #3)

The group discussed the three proposals in search of a solution to the current channels group's less then perfect architecture.

Jay Martin proposed solution #4a which deprecates channel type and adds an object pointing to the entry in the Network Services Monitoring MIB table.

Proposal #4b is the same as #4a except without deprecating channel type.

Proposal #5a is pick up object called assocApplicationProtocol from the Network Services Monitoring MIB and add that object to the channel table and to deprecate the channel type.

Proposal #5b is the same as #5a except without deprecating the channel type.

Proposal #6 leaves the channel table as is except adds a port number to the channel table.

Proposal #0a - leave as is.
Proposal #0a1 - same as #0a but change channel type to type 3 enum
Proposal #0b - leave as is and clean up channel type
Proposal #0c - leave as is and add range of enums for TCP & UDP ports.

Proposal #7 - Add assocApplicationProtocol to channel table, add an index to the NSM table (if non-zero) and deprecate channel type.

By consensus, the following options were removed from consideration: #0a, #0a1, #1, #2, #4a, #4b, #5a, #5b, #6.

By a vote of 13-0, proposal #7 was accepted. Randy will be making this change in the draft.

From the last meeting the following enums for PrtChannelTypeTC were approved:
35 - IrDA
36 - Printxchange

At this meeting, the following enums were assigned for PrtChannelTypeTC:

chPortTCP = 37

chBidirPortTCP = 38
chUNPP = 39

Randy Turner reviewed the changes he has made to the RFC document:

1) Steve Waldbusser's contact information was updated.
2) Textual Conventions were created and moved to the front of the document.
3) The agreed to alert codes including generics were added.

The following review assignments were made

Alerts: Tom Hasting, Harry Lewis
Channel: J.K. Martin, Bill Wagner
Console: Rick L., Ron B
Interpreter: Ron B.
General, Covers, Localization: Bill Wagner, Bob Setterbo
Marker: Don Wright

A discussion ensued over the uses of the various MIBs related to, etc., of MIB-II. What should the printer MIB require?

Randy asked for documents from any implementors of RFC1759 describing their implementation experiences. He needs those documents by the end of May.

The group moved on to discuss the issue over the current semantics used to provide printer interpreter language and version. Ray Lutz has created a new way of describing interpreters/versions that meets the need of the MFPA specification. Should that work be incorporated into the MIB? If so how? One proposal would be to use the description screen as a place to store the MFPA format. Options:

1) prtInterpreterLangFamilyandVersion - new object and group (optional)
1a) Same as #1 only mandatory
2) new object using MFPA formatted object : eg. HP|PCL|5 or ADOBE|PS|2
3) new object except creating enums for each of the MFPI strings
4) Add a recommendation to use the MFPA format and place it in the interpreter description string.
5) To be MFPA compliant a manufacturer would consistently use the last part of the MFPA language information (e.g. 5e) as the prtInterpreterLangLevel. In addition, downlevel versions of an interpreter language would be explicitly listed in the interpreter table.

By consensus, the group agreed to #5.

Tom Hastings discussed his e-mail referencing inconsistencies in specifying ranges of certain integer objects. Some start at 0, others at 1.

If it is not accessible - doesn't matter.
If R/W the objects will be examined during the review process.

Ron Bergman presented the Virtual Printer Technology. Ron requested a new channel type for a virtual printer. Because of the potential confusion over using a channel for a virtual device and because of its similarity to chPortTCP, the group denied the request. Ron agreed to the result.

The group adjourned at 7:30.

The meeting stated on Thursday morning at 8:45 with a SENSE working session. This session lasted until 11:15. The minutes for this meeting were taken by Rick Landau and included here.

Properties in SENSE currently are strings. In fact, both the name and value are strings. As an example, "foo=bar" says that the property named foo has value bar.

(Currently these are spec'ed as display strings rather than general octet strings, though that may not be necessary. It makes the marshaling of data into CPAP-style framing easier if the character sets of the strings are slightly restricted.)

Property names are X11-style dotted hierarchical names of the form a.b.c.d, for example, Publication.Created=<timestamp>. This does allow some wildcarding of the form a.b.*, which may be useful.

There can be several Publishers for a managed entity. ISVs, for instance, may create enhanced Publishers for certain printers to provide more information about an entity. Publishers of separate Editions can be independent of each other: the vendor's default Publisher and the ISV enhanced Publisher for an entity do not have to coordinate at all.

We must be careful to define Publication as the top level entity, not the periodic issue that comes out with timely event information.

Q: Are properties named hierarchically down from the id of the object to which they belong, e.g., is the property that stores the name of Publication number 99 called "99.Name"?
A: No, the query is of the form,
"GetPublicationProperties Id=99 PropertyName=Id.Name"

Problem with the simplistic naming of properties: how do we represent multi-valued properties? Example: a printer with multiple paper trays could list the vector of possible values. (This is what CPAP currently does.) For instance,

> InputTray.Width="8.5,11,11"
> InputTray.Length="11,8.5,17"
> InputTray.Capacity="250,250,100"
> InputTray.Name="letter,letter,special"

The values represent columns of a table of properties. Use the number of commas to compute the number of rows. Yes, have to escape the commas if they are part of the value of the property; pick an escape character.

Q: Why not put the variable part into the property name, e.g., InputTray.UpperTray.Width or InputTray.Width.1?
A: Did consider this approach, but rejected for reasons we can't remember now.

If you have a property a.b.c, what is the value of a.b? A useful answer is "the length of the vectors of the items at the next level." In the above example, the value of InputTray=3.

Q: What is the relationship of the properties of separate Editions of a Publication?
A: They are independent (except for certain referential properties, such as the name of the Publication to which the Editions belong, and the inclusion of both Editions in the Publication's Editions.IdList).

Q: Are there are mandatory properties of Publications and Editions that Subscribers can always depend on to help make their choices?
A: Yes, these will be described in the specs.

The Server's Publication includes an Edition of all changes to Publication properties. A general management application can subscribe to that Edition and be notified on any property change -- and the event messages contain the properties that changed, unlike general SNMP traps.

Need shared specs for any Editions that are to be common across vendors and applications. Spec needs to include names and other identifiers, required properties, message formats, etc. (Underscore already volunteered to act as registrar for reserved names until some formal authority is found.)

Jeff Dunham started a thread in the general SNMP mail list about reliable trap delivery. The answers seemed to be the same old tune: traps should be used only to accelerate polling, can't depend on them, can't use connections for them because those will go first when network degrades. This doesn't seem to impact much on SENSE, which does not use connection-oriented protocols, anyway.

Must keep clear focus on goals of SENSE that don't relate to SNMP or DMI or network administration:

- Reliable delivery
- Dynamic registration
- Wide platform support
- Non-privileged operation (port numbers)
- Scaling to very large networks

Q: Replace clear text strings with ASN.1 BER? And use OIDs or DMI-like hierarchical identifiers within component for name space administration?
A: Well, has some advantages, but adds considerable complexity, too.
Pro: instant acceptance, base of existing code
Con: more complex, many vendors have to buy
Con: too much like SNMP
Con: supporting tools are heavy and expensive
Pro: companies have already invested in ASN.1 expertise
Pro: less tricky to internationalize

I18n a question we will really have to deal with.

Will this project have more meetings by itself, not including status reporting at these meetings?

Goal to have specs for next month's meeting. Need to have SENSE specs and the Edition specs for some of the fundamental Editions.

After the conclusion of the SENSE discussion, the group again visited the issue of the reorganization of the Printer Working Group. The group decided that unless a printer MIB working group chair was identified before Binnur's departure, the PWG chair would function as the Printer MIB working group chair. The vote was 9 to 1.

PWG
- - Don Wright, Chair

- - Jay Martin, Secretary

Printer MIB Project
- - open, chair
- - open, secretary
- - Randy Turner, MIB Editor
- - Harry Lewis , MIF Editor

SENSE Project
- - J.K. Martin, chair
- - Rick Landau , editor/secretary

JOB Monitoring Project
- - Ron Bergman, chair
- - Tom Hastings, editor
- - Harry Lewis, secretary

Bill Wagner moved to affirm the current slate by acclamation. The motion was seconded and approved unanimously. (Note: Binnur will continue as chair of the PMP through July.)

The meeting recessed for Lunch from 12:00 until 1:15PM.

Ron Bergman kicked off the afternoon session discussing the Print Job Monitoring/Resource Accounting Project. He first reviewed the list of issues for discussion:

- Charter discussed and a draft was generated:

The Printer Job Monitoring and Resource Accounting MIB Project

Statement of Charter:

This project shall develop a set of objects for monitoring the status and progress of print jobs and to obtain resource accounting data at the completion of a job. The object set will initially be developed independent from the management format. The object set development shall not preclude use of the set with either SNMP MIB, a DMI MIF, or the SENSE protocol.

Use of the object set shall not be limited to printers. To obtain complete information regarding a print job in many printing systems will require the object set to be implemented into a file server, spooler, and/or a printer management system software.

When object set is finalized, an SNMP MIB will be will developed. The object set may also be ported to other formats, such as SENSE and a DMI MIF, as an extension to this project.

Goals:

1. The job MIB is intended to provide the following information for the indicated Role Models (see Appendix D - Roles of Users in the Printer MIB draft update to RFC 1759).

User:
(U1) - A timely notification that his job has completed and where.
(U2) - The current status of the user's job (queries)
(U3) - Error and diagnostic information for jobs that did not successfully complete.
(U4) - Ability to identify the least busy printer.

Operator:
(OP1) - A presentation of the state of all the jobs in the print system.
(OP2) - Which users submitted each job.
(OP3) - What resources does each job need.
(OP4) - For which physical printers are the jobs candidates.
(OP5) - Some idea of how long each job will take.

Capacity Planner:
(C1) - How busy are printers.
(C2) - What time of day are they used.
(C3) - How long do users' jobs wait before starting to print.

Accountant:
(A1) - A record of resources used and printer usage data for charging users or groups for resources used.

2. Schedule:

Submit a request for an IETF charter for this project at the December IETF meeting.

- Terminology (discussed)
- Job Submission Model
- Platforms and Protocols
- Job Submission Protocol
- Job Ticket
- Client Software
- OS Support
- Is SNMP the right solution
- How will the Job Monitor MIB solve these issues

The group first discussed the issue of whether SNMP is an appropriate solution to this problem for end users. Tom Hastings pointed out that even if SNMP is not appropriate for end users, it is most likely appropriate for Operators.

There was a discussion on the need to develop the Job MIB outside the IETF standards track MIB process and at a future time take the MIB to the IETF and make it a standard track MIB. The group felt it should get started now and at a future time when a MIB has been created consider going to the IETF to make the MIB a standards track MIB.

In addition, Harry Lewis asked whether the group should automatically assume that a MIF will be created. The group felt that doing a MIF should be considered but not be automatic.

The group went through the information that would be provided to the various users (i.e role models) of the MIB. Ron Bergman will publish the updated items back to the mailing list.

"This is a 'satellite with camera' MIB rather than a 'satellite with missiles' MIB" per Rick Landau.

There was identified a need to better name and define the various roles that people play and the information they need. Tom Hastings will be making an attempt at defining these for the MIB.

The discussion moved to the second item: Terminology The definitions worked through will be included in the document's glossary.

Tom Hastings presented his paper on Usage Scenarios for the Job Monitoring MIB. He presented three distinct scenarios:

- Client to printer
- Client-to-server-to-printer
- Client-to-spooler-to-supervisor-to-printer

There was a concern over the more complex client server environments and the implications that OS/NOS changes would be required especially in the more complex, multi-tiered environments. The group discussed whether the MIB should be designed to support a multi-tier, multiple agenda architecture or whether the design should exclude the ability to query the intermediate nodes. There is a belief that in order for this MIB to be really useful, the OS/NOS vendors need to participate.

The rest of the identified issues were left for the next meeting. Additionally, the group felt we should set a goal as to when we should have some kind of a first pass

MIB. When should we charter with the IETF? How do we identify jobs? What information is needed to solve each of the problems identified in the charter?

The meeting adjourned at 4:05PM.