ISSUE 1
> Figure 2 references rfc1213 for the interfaces MIB. Should this be
> updated to RFC 2863?
>
> ~~You~~Yes, I would like that.
> In fact I would prefer that in figures they do not use RFC numbers but
> MIB names or table names.
>
Any objections from the working group?
What is the impact of this change?

ISSUE 2
> 2.2.1.1. specifies I18N names versus non-I18N by the variable name.
> Shouldn't this be done by the SYNTAX clause, using T-Cs? Control codes
> are allowed IF specified in the DESCRIPTION clause. This makes it very
> hard for applications to know what to expect. Using a T-C would make
> this decision machine-parseable.
>
> The Localization table is being used to control the localization
> information used to display info. I think this is mixing the "agent" and
> application responsibilities. Since an operator console is historically
> local to the system that implements the agent, this distiction may not
> seem important, but for an application remote from the agent, it should
> be the application that determines its capabilities for display, not the
> agent. The agent should simply support UTF-8.
>
> I am also disturbed by their use of international character strings.
> In fact, in the new MIB, they now import DisplayString which they
> did NOT use in the RFC1759 version. And we've been pushing other
> MIBs to move from DisplayString to the use of UTF-8 based strings
> (Textual Conventions).
>
This issue needs to be studied to determine the impact.

ISSUE 3
> 2.2.4 recommend changing "A printer contains one or more output
> mechanisms. The Output Group in the model represents these." to "The
> Output Group in the model represents the one or more output mechanisms
> contained by a printer."
>
Agreed we should make this change. (page 11)

ISSUE 4
> 2.2.13.1 the list is inconsistent in sentence structure. clause 3 needs
> a verb.
>
Agreed. Change "(3) Alert codes reported in the Alert Table." to "(3) Alert codes are reported in the Alert Table." (page 14)

ISSUE 5
> 2.2.13.4 "must" s/b MUST if an RFC2119 must.
>
"must" appears in three places. None of these occurrences seem to be specification requirements. No change is recommended. (page 21)

ISSUE 6
> 2.4.1 the text has been changed in a way that seems to allow

> modifications in a way that is not consistent with IETF/SMI rules, which
> RFC1759 did not allow.
>
I do not understand this comment. Do we need clarification?

ISSUE 7
> 3.4.x recommends modifications to the usage of the Host resources MIB
> that may not be consistent with design of the Host resources MIB. This
> could lead to interoperability issues for applications which manage
> printers and other hosts simultaneously. This section should be reviewed
> closely by an expert in the Host resources mib design. It may be
> desirable to upgrade the Host Resources MIB printer to better reflect
> the nature of existing printers.
>
> Nor sure I read it the same way... but it is text that indeed may need
> some more review.
>
This item will require some study. If I recall correctly, the changes here were primarily clarification. (page 24–25)

ISSUE 8
> CodedCharset description removed information about SNMP encoding. Is
> this needed?
>
> Yep I agree... it contains info that seems useful
>
Here is the text that was removed. Anyone recall why it was removed? (page 30)

> The space of the coded character set enumeration has been
> divide into three regions. The first region (3-999) consists
> of coded character sets that have been standardized by some
> standard setting organization. This region is intended for
> standards that do not have subset implementations. The
> second region (1000-1999) is for the Unicode and ISO/IEC 10646
> coded character sets together with a specification of a (set
> of) sub-repetoires that may occur. The third region (>1999)
> is intended for vendor specific coded character sets.
>
> NOTE: Unicode and ISO 10646 character coded data may be
> processed and stored in either Big Endian (most significant
> octet first) or Little Endian (least significant octet
> first) order. Intel x86, VAX, and Alpha/AXP architectures are
> examples of Little Endian processor architectures.
> Furthermore, in environments where either order may occur,
> so-called Unicode BYTE ORDER MARK (BOM) character (which is
> ISO 10646 ZERO WIDTH NO BREAK SPACE), coded as FEFF in two
> octets and 0000FEFF in four octets is used at the beginning
> of the data as a signature to indicate the order of the
> following data (See ISO 10646 Annex F). Thus either
> ordering and BOM may occur in print data streams sent to the
> interpreter. However, ISO 8824/8825 (ASN.1/BER) used by
> SNMP is quite clear that Big Endian order shall be used and
> BOM shall NOT be used in transmission in the protocol.
> Transmitting Unicode in Big Endian order in SNMP should
> not prove to be a hardship for Little Endian machines,
> since SNMP ASN.1/BER requires integers to be transmitted
> in Big Endian order as well. So SNMP implementations on

Little Endian machines are already reversing the order of integers to make them Big Endian for transmission via SNMP. Also Unicode characters are usually treated as two-octet integers, not short text strings, so that it will be straightforward for Little Endian machines to reverse the order of Unicode character octets as well before transmitting them and after receiving them via the SNMP protocol.

Where a given coded character set may be known by more than one name, the most commonly known name is used as the name of the enumeration and other names are shown in the comments. The comments also indicate where to find detailed information on the coded character set and briefly characterize its relationship to other similar coded character sets.

ISSUE 9

> PrtChannelStateTC description doesn't quite match with enumerations.
> This could lead to interoperability problems if implementors interpret
> differently.
>
> I agree that it could at least use some clarification
>

I don't see what the problem is here. The text is almost identical to RFC 1759. Anyone disagree? (page 32)

ISSUE 10

> PrtChanneltypeTC seems so all-inclusive that it might not be useful.
>

The working group totally agrees with your statement. With the exception of several new enumerations and the change to a textual convention, this is identical to RFC 1759. When RFC 1759 was drafted, the group attempted to define every possible channel type. Implementation experience has shown that a number of these channel types will never be used and many will only be used only by a few vendors. To reduce some of the confusion in this set of enumerations, a number of them have been deprecated. SNMP rules do not allow a real cleanup of this group. (pages 32-40)

ISSUE 11

> The enumeration types for each TC is specified in a comment. Would it be
> better to make it part of the description to ensure the update semantics
> are considered part of the object semantics?
>

This, again, is formatted as in RFC 1759.

ISSUE 12

> Yep... but ... I also wonder why we need to define all
> these TCs. The new MIB has 31 TCs compared to 5 in RFC1759.
>
> That bothered me as well. Just because you can doesn't mean you should.
>

The embedded enumerations were carefully examined by the working group to determine which could be used in related MIBs. The additional 26 TCs were agreed to be created. Several of these TCs are NOW being used by the Finisher MIB.

ISSUE 13
> PrtOutputPageDeliveryOrientationTC expects different interpretations in
> different locales; thsi will probably lead to interoperability issues
> between applications, and could confuse users. I recommend this be
> changed to use terms that are not interpreted differently per locale.
>
The text in question was taken directly from the corresponding object in RFC 1759. (page 46)

ISSUE 14
> PrtConsoleDisableTC allows variable interpretation; this will likely
> lead to interoperability issues.
>
Working Group discussion needed. (page 50)

ISSUE 15
> PrtAlertGroupTC is declared as being both type 1 and type 2 enumeration.
> Isn't this just a type 2 enumeration? if not, I question whether this is
> consistent with SMI rules.
>
This should be just a type 1 enumeration. Values defined in other MIBs require conformance to type 1 rules in those documents. Remove "...type 2 enumerations and are..." (page 52)

ISSUE 16
> I am concerned that many objects have so many unrelated enumerations. It
> seems that it would make more sense, and be easier for applications to
> deal with, if more objects were created which represented different
> aspects of printer management. PrtAlertCodeTC is one item that
> demonstrates this use of one object where multiple objects might be a
> better choice.
>
This would be a fundamental change in the architecture of RFC 1759. In practice the alert table is conceptually very clean and practical. (pages 53-56)

ISSUE 17
> PrtGeneralEntry has had new columns added to each row; shouldn't this
> require a name change? This could be done using augments to not violate
> SMI rules, couldn't it?
>
This comment applies to the Auxiliary Sheet Group, the Administrative section, and the General Alert Table section. Should this augment the current prtGeneralEntry? (pages 62-63)

> Is it legal to change the types in an Entry to a TC?
>
> As long as the underlying data type stays the same and the semantics
> don't change, then I think it is OK.
>

ISSUE 18
> prtGeneralConfigChanges changed semantics, btu used the same name and
> OID. I am also concerned that the semantics are ambiguous - it should be
> incremented if an output tray is added, but not when an input tray is
> removed?

>
> This one is questionable? If they call out consensus, that it is a
> clarfification, then I might accept.
>
> Note that if they do not explain it that way, and if it indeed is a
> change in semantics, then ecven if they recycle at PS, it would
> still mean that they need top deprecate this object and define
> a new one with the new semantics
>
> Compare the following descriptions:
> RFC - "Counts configuration changes that change the capabilities of a
> printer, such as the addition/deletion of input/output bins, the
> addition/deletion of interpreters, or changes in media size. Such
> changes will often affect the capability of the printer to service
> certain types of print jobs."
>
> draft - "Counts configuration changes within the printer. A
> configuration change is defined to be an action that results in a change
> to any MIB object other than those that reflect status or level, or
> those that act as counters or gauges. In addition, any action that
> results in a row being added or deleted from any table in the Printer
> MIB is considered a configuration change. Such changes will often affect
> the capability of the printer to service certain types of print jobs."
>
> In the original, MIB objects and tables aren't even mentioned.
> In the draft, that is how to determine what to count.
>
> In the original, it counts "... changes that change the capabilities
> ..."
> In the draft, "Such changes will often affect the capability..."
>
> These appear to be counting different, but overlapping or related sets
> of things.
>
This object was extensively discussed by the working group and the changes were intended to clarify the use of the object. No change in semantics is intended. You have, however, pointed out a remaining ambiguity. We will change "...input tray is removed..." to "...input tray is removed to load additional paper..." This statement was not intended to imply a permanent removal of the tray. (page 58)

ISSUE 19
> prtGeneralCurrentLocalization has semantic changes.
>
This object has only a slight wording change intended as a clarification. No semantic changes were expected. (pages 58-59)

ISSUE 20
> prtGeneralReset has semantic changes.
>
Wording was added to explain the use of the enumerations. There was no intent to change the semantics. (page 59)

ISSUE 21
> prtAlertCriticalEvents and prtAlertAllEvents - the description seems to
> imply a zero-based counter, rather than an SNMP-style counter.
>

You are correct. This operation is desired to allow a management application to be able to better track printer alerts. If this value is less than the previous value observed, it would indicate that the printer has been power cycled and the alert table has been reset. If these values were not initialized, the starting value could indicate to the management app that a large number of alerts were generated on the printer since it last queried these values. This is important since printers can be frequently power cycled. (page 63)

ISSUE 22

> prtCoverIndex is ambiguous as to whether it should remain consistent
> across reboots.
>

Many modern printers can be upgraded with new features such as additional input trays and finishing units. The installation of these features almost always requires power to be off. While this does not happen often, it can happen. A management station should be aware of this possibility. (page 64)

ISSUE 23

> prtCoverStatus - check whether TC matches original semantics.
>

The only change is that "door" was changed to "cover" in enumerations 3 and 4. This was agreed to be a clarification. (page 29)

ISSUE 24

> prtLocalizationLanguage and prtLocalizationCountry - changed semantics -
> maybe, depending on ISO 639 and ISO 3166.
>

The only change here is the use of "DisplayString" in place of "OCTET STRING". Should this be reverted back to the original? (pages 65-66)

ISSUE 25

> prtStorageRefSeqNumber and prtStorageRefIndex and prtDeviceRefSeqNumber
> and prtDeviceRefIndex changed range, without changong name or OID.
>

The working group agreed that 0 would never be a valid index value and was not appropriate, so the lower value was changed to 1. If this is an illegal change we will restore the value to 0. (Pages 67-68)

ISSUE 26

> PrtInputEntry has new columns added per row.
>

Added "prtInputMediaLoadTimeout" and "prtInputNextIndex". We should add these using an augments clause. (pages 69, 76)

ISSUE 27

> prtInputMediaDimFeedDirChosen semantic differences?
>

This appears to be identical to RFC 1759. The format of the text in the Description clause has been fixed. (page 70)

ISSUE 28

> prtInputMediaName - does the REFERENCE constarin valuesa and thus change
> semantics?
>

There is no intention to change the semantics of this object. This appendix was also present in RFC 1759. This appendix lists the ISO/IEC 10175

standardized names, but there is no requirement that these must be used. (page 72)

ISSUE 29

> prtInputSerialNumber changed OCTET STRING SIZE. Applications may be
> broken by the expanded size.
>

We will change this back to the original value of 32. (page 73)

ISSUE 30

> prtInputMediaType - semantics change
>

This was believed to be a clarification. Since the object contains a string that is intended to be displayed, it was always believed that the values shown in the Description were examples. (pages 74-75)

ISSUE 31

> prtInputMediaColor - semantics change (effectively an enumeration
> extension)
>
> This one I do not see? Dave can you explain?
>
> In the RFC, there is an enumeration of fixed strings, and the ability
> for vendors to add extensions. Assuming the fixed strings enumerated by
> the other standards have some common meaning, it is important to avoid
> name space conflicts with vendor-specific names.
>
> The RFC specifies the standard list with the words "which are:", which
> implies that is is a non-changing list. Vendors could therefore feel
> free to add their own names, such as "ivory". The separation of
> namespaces is guaranteed by the fixed list of standard colors.
>
> The draft changes "which are:" to "such as:" which implies that the
> standard list may grow. If the ISO 10175 added "ivory" after a vendor
> did so, and the underlying color definition differed between the "ivory"
> of the vendor and the "ivory" of the ISO 10175 spec, then there is an
> interoperability problem due to the ambiguity caused by the name
> conflict.
>
> I suspect that for 90+% of the printing community a vendors's "ivory"
> and the standard "ivory" will be close enough that such a conflict won't
> matter. However, I assume the reason why this is even in the mib is to
> ensure the ability to control the input to a print job remotely with a
> reasonable expectation that the input media color will be a specific
> color. For a high-quality print job, the distinction between the two
> ivories may be significant.
>
> If the new text had existed at the time of PS review, it might have led
> a reviewer to request a change to the object design to better ensure
> name space diferentiation. The text as specified in the RFC has the
> necessary differentiation. The draft does not. If "such as:" was the way
> the text had read for the advancement to PS, and I had been the reviewer
> then, I think I would have questioned the issues related to name space
> collision. If the text had been "which are:", I think I would have found
> that adequate to protect against name space collision. I most certainly
> would have wanted to see "such as:" discussed in the WG.
>

> I think this is a semantic change. I am concerned about the name space
> collisions that could occur, and the potential for interoperability
> problems in the future. Whether it is enough to justify deprecation is a
> judgement call of the WG, the chairs, and the AD.
>
We will change "such as" back to "which are". (page 75)

ISSUE 32
> prtOutputMaxDimFeedDir - DimUnits versus MediaUnits ?
>
Good catch! This should be PrtMediaUnitTC. (Also applies to prtOutputMaxDimXFeedDir, prtOutputMinDimFeedDir, and prtOutputMinDimXFeedDir.) (pages 81-82)

ISSUE 33
> prtOutputMaxDimFeedDir - SYNTAX changed to TC, which includes additional
> semantics
>
Don't understand this comment. No syntax changes are in this object.

ISSUE 34
> prtOutputPageCollated and prtOutputOffsetStacking - clarificiation or
> semantics change?
>
The changes in both of these object Description clauses is intended to provide a reference to the glossary. These terms are well known in the printer industry and the change was merely to help others. (page 83)

ISSUE 35
> prtMarkerLifeCount - semantic change?
>
This change is only a clarification. CounterUnit was changed to prtMarkerCounterUnit, which is the full name of the appropriate object that defines the units. (page 85)

ISSUE 36
> prtMarkerSuppliesIndex - is "successive printer power cycles" the same
> as "successive power cycles"?
>
Yes, but "printer" should not have been removed and will be restored. (page 88)

ISSUE 37
> prtMarkerSuppliesColorantIndex - semantic change?
>
This is a clarification from the interoperability test. (page 88)

ISSUE 38
> prtMarkerColorantIndex - semantic change
>
The deleted text was not intentional and will be restored. (page 90)

ISSUE 39
> prtMarkerColorantValue - semantic change
>
This change was made to clarify that the list was not exhaustive and any additions to ISO 10175 and ISO 10180 are allowed. (page 91)

ISSUE 40
> prtMarkerColorantTonality - no range specification in syntax
>
This specification is unchanged from RFC 1759. (page 91)

ISSUE 41
> The Print Job Delivery Channel Group commented text differs from
> original - does this result in a semantic change?
>
The added text was determined by the working group to be required for clarification. No change is semantics is intended. Two examples were removed from the text since they were determined to be redundant and are now deprecated. (pages 94-96)

ISSUE 42
> prtChannelIndex - range changed; same name, OID
>
The range added was assumed to be the defined range for an index. We will remove the added range to avoid confusion. (page 97)

ISSUE 43
> prtChannelCurrentJobCntlLangIndex and prtChannelDefaultPageDescLangIndex
> - range change per decsription
>
The added text is intended to be a clarification. Implementations that participated in the interoperability test were conformant and requested the additional text to be added. (page 97)

ISSUE 44
> prtConsoleLightIndex - range change; weasel words added - stanadrd or
> not?
>
Again this is a clarification. See issue 22. (page 106)

ISSUE 45
> prtConsoleOnTime and prtConsoleOffTime - semantic change
>
This again is a clarification. The original description had no way of indicating the light was on and not blinking. (pages 106-107)

ISSUE 46
> prtAlertIndex - semantic change to range and access-type
>
The range added was assumed to be the defined range for an index. We will remove the added range to avoid confusion. The syntax was changed to read-only since the traps return this oid and it was our understanding that the oid could not be not-accessible. Should we revert? (page 108)

ISSUE 47
> prtAlertLocation - semantic change
>
This change is a clarification from the interoperability test. (page 110)

ISSUE 48
> printerV1Alert - should this be object-type or pbect-identity?
>

This is unchanged from RFC 1759. (page 110)

ISSUE 49
> prtGeneralGroup, prtChannelGroup, prtAlertTableGroup - changed objects
> required for conformance
>
??? (page 111?)

ISSUE 50
> prtAlertTimeGroup - needs official status of deprecated, not just
> comment.
>
Will change as suggested. (page 120)

ISSUE 51
> References - refers to internet drafts and out-of-print books.
>
This is a consequence of this document being in the queue for almost 5 years.
This section will be updated as appropriate. (page 122)

> I'd hope the authors can find decent responses to the above.
>
> In addition I found:
>

ISSUE 52
- You must mention in abstract that this doc obsoletes RFC1759
- These days we want the MIB boilerplate as found on the
  www.ops.ietf.org web page. You can see many other MIB
  documents that have it.
- I am missing an appendix with list of changes (I know some other
  AD will worry about it).

The above will be corrected.

> - This MIB cannot advance to DS. They for example add a set of
>   objects, and so they need to recycle for that.
> - If the semantics of some of these objects and TCs have changed,
>   then the names (and OIDs for objects) need to change also. The old
>   ones then need to be deprecated or obsoleted.
> - There seem so many changes that I wonder if we can do this without
>   resurrecting the PRINTMIB WG. Or is the external PWG group
>   trying to do this work. In that case, we need like a 4-week IETF Last
>   Call (even to recycle at PS).... but we need to see answers first.
>

ISSUE 53
> - I see SMICng compile error/warning:
>   W: f(printmib.mi2), (4494,1) Item "printerV2Alert" is not contained in
>   any group defined in the current module
>   This needs to be fixed.
>
> - Need the LAST-UPDATED and REVISION clauses to use 4-digit
>   years as specified in RFC2578.
>
> - They need to add REVISION clauses. And the changes also need to
>   be documented in such a clause. See RFC2863 as an example.

The above will be corrected.

ISSUE 54

```
> - quite a set of requirements and or compliance related statements
>   are made throught the document. I believe that a lot of those can
>   actually be captured in the MODULE-COMPLIANCE statement
>   (maybe we need multiple such statements).
>   For example, if you have a pre-req that the system group of MIB-II
>   is implememented, then you can add to the MODULE-COMPLIANCE
>   this clause:
>
>         MODULE SNMPv2-MIB
>                   MANDATORY-GROUPS { systemGroup }
>
```

The above will be corrected.

ISSUE 55

```
> - You cannot ADD objects to a MODULE-COMPLIANCE statement
>   after the MIB is published as PS. So if you want to add new
>   objects (even if you recycle at PS) then you must add another
>   new MODULE-COMPLIANCE statement. You can deprecate
>   or obsolete the old one.
>
```

This will be corrected.

```
> > Some of your comments could be responded to with "we claim they are
> > clarifications". I guess I'd be willing to accept that, but in order to
> > do so I would want to see an explicit call for consensus that such is
> > indeed agreed upon.
>
> Yes, some may simply be clarifications.
>
> >
```

ISSUE 56

Who's going to do all this work ;-)

Comments compiled by:
Ron Bergman
Hitachi Koki Imaging Solutions