

1 | **Subj: Issues for Job Monitoring MIB, version 0.8~~51~~, dated 8/9/97/24/97**

2 | From: Tom Hastings

3 | Date: 8/8/97/5/13/97

4 | File: issues.doc

5 | This file will be our standing list of open and closed issues for the Job Monitoring MIB.
6 | The open issues are in the first section, closed issues not yet incorporated into the MIB
7 | specification are in the second section, and the closed issues that are in the current
8 | specification are in the third section. When an issue is closed, I indicate the resolution and
9 | why and move the issue to the appropriate closed section, depending on whether I've also
10 | done the editing for it or not.

11 | Revision marks show changes agreed to at the Redmond, 8/6/97 meeting~~from the 3/26/97~~
12 | ~~version of the issues list~~. When I move an issue from Open to Closed, I only show the
13 | revision marks of what changed, not the movement itself.

14 | This version of the issues list includes issues that are in the Internet Draft 0400 (same as
15 | revision 0.8~~41~~) that need to be resolved.

16 | I've updated the issues with the agreements reached at the JMP meeting, 6/26/97/4/97
17 | and the reasons behind the resolutions.

18 | If you object to any of the proposed resolutions to the issues, please send e-mail.

19 | **1. Open Issues**

20 | The following issues are open:

21 | none.

22 | **2. NoneClosed Issues - not yet reflected in the current draft**

23 | The following issues have been closed but have not been incorporated in a draft:

24 | none.

25 | **3. Closed Issues - reflected in the current draft (0.85)**

26 | The following issues have been closed and have been incorporated into the Internet Draft
27 | 0500 and version 0.8574 or earlier:

28 | Issue 1 - Should we add a standard SNMP **RowStatus** object to the **jmJobTable** and
29 | **jmAttributeTable**?

30 | Closed: No. We decided to go with simplicity. Also customers have had experience with
31 | printers that stop because some software isn't installed or running properly and so prevents
32 | the printer from working. Finally, see issue 9 resolution in which no objects are read-
33 | write, but implementers can augment this MIB with objects that allow any object in this
34 | MIB to be changed by duly authorized users. We can produce another MIB that
35 | augments this MIB in the future, if there is interest and a solution to standard means to
36 | write objects (and delete rows).

37

38 Issue 2 - If we add **RowStatus** to the **jmJobTable**, should we add a
39 **jmGeneralTableOverflowPolicy** object to the **jmGeneralGroup**?

40 Closed: No, because we did not agree to add RowStatus and we are also concerned about
41 the printer stopping for some configuration reason. Customers that care about accounting
42 should make sure that the accounting applications are running properly, perhaps even with
43 a daemon program that monitors the accounting programs.

44

45 Issue 3 - If we add **jmGeneralTableOverflowPolicy** object should it be read-write?

46 Closed: No, see issues 1, 2, and 9.

47

48 Issue 4 - Need to re-draw the job state transition diagram to add the needsAttention state

49 Tom H will make a table of all possible state transitions, so that it will mean the IETF
50 requirements for plain text in RFCs.

51 Closed: I added a job state transition table to replace the job state transition diagram. See
52 the Internet Draft 00 and version 0.71. Called the job state: 'processing-stopped'.

53

54 Issue 5 - Restore NMS having to access both the server and printer agents (Configuration
55 2b)?

56 Yes, with the following understandings: Configuration 2b will show a monitoring
57 application, a server, and a printer. The MIB will be only in the printer, but the
58 monitoring application is also monitoring the server by some other means than the Job
59 Monitoring MIB. The Job Monitoring MIB in the Printer shall have enough information
60 in it for the monitoring application to find the job in the Printer's Job Monitoring MIB that
61 it found in the server (by other means). In such cases, the server usually deletes its copy
62 of the job, but need not. This configuration covers the configuration supported by the HP
63 5si Mopier private job monitoring MIB when driven from a Novell server.

64 ACTION ITEM (Tom Hastings, Bob Pentecost): Tom draw up a new configuration 2b
65 and show to Bob before distributing it to the group.

66 Closed: I added the agreed configuration, but called it configuration 3, since it has
67 similarities to both configuration 1 and 2. See the Internet Draft 00 and version 0.71.

68

69 Issue 6 - If Configuration 2b is added to the spec, how does the monitor relate a job in the
70 server and the copy that is in the printer?

71 Closed: The new Configuration 2b does NOT have a Job Monitoring MIB in both the
72 server and printer; only in the printer.

73

74 Issue 7 - If Configuration 2b is added to the spec, add a Boolean General object that says
75 whether this Job Set requires the NMS to contact the printer's agent too?
76 Closed: No need, since the configuration won't have a Job Monitoring MIB in the server.
77

78 Issue 8 - Should we make a new **jmServerGroup** for objects needed by the
79 **serverOnly(4)** and **bothPrinterAndServer(5)** configurations?
80 Closed: No, we haven't identified any that are server only that can't be put into the
81 jmAttributeTable, so that non-server's need not implement.
82

83 Issue 9 - What objects should be read-write, so that the system administrator can set
84 policy?
85 Closed: No, for simplicity. Also implementers can augment the Job Monitoring MIB with
86 means to write any object. Add a paragraph that indicates that implementers could allow
87 monitoring applications to modify objects, by adding a private table that contains an
88 encrypted password, with date and time mixed in and set in the clear. In addition, the OID
89 of the object to be written and the new value is contained in the table.
90

91 Issue 10 - Should we add an object to specify the policy for SNMP Gets for other user's
92 jobs?
93 Closed: No, authorization and authentication are beyond SNMP and this Job Monitoring
94 MIB.
95

96 Issue 11 - Should the policy object for SNMP Gets for other user's jobs be writeable?
97 Closed: No, we didn't agree to add such an object in Issue 10 and if we had, it wouldn't
98 have been writeable (see answer to issue 9).
99

100 Issue 12 - What is the SNMPv1 and SNMPv2 error that an agent shall return if there is no
101 instrumentation for an object?
102 Closed: There is no such SNMP error. ALL uninstrumented objects in mandatory groups
103 of any MIB should always correctly return 'read-only' static values specified in 'DEFVAL'
104 clauses. 'DEFVAL' is a perfectly good SMIV2 feature intended to cover this situation.
105 Returning ANY SNMP error for ANY object in a mandatory group with a legal instance
106 qualifier (i.e., set of indices) is NOT legal in a literal reading of the SNMPv2 Protocol
107 spec (RFC 1905, page 10, in 'Get-Request PDU' handling). That's what 'shall implement
108 ALL the objects in this group' means! So add DEFVAL clauses to all objects.
109

110 Issue 13 - Why didn't the Printer MIB use this SNMP error instead of returning
111 unknown(2) enums?
112 Closed: The Printer MIB was correct to use these unknown(2) enums, instead of an
113 SNMP error.
114
115 Issue 14 - How do we add traps without adding too much network traffic?
116 Closed: For simplicity and to avoid the design problem of registering and unregistering
117 for traps, we decided not to add traps. The HP 5si Mopier private job monitoring MIB
118 has only one trap: when a job is added to the table. However, no application is using the
119 trap. Polling seems sufficient and not a problem.
120
121 Issue 15 - Should **jmGeneralQueuingAlgorithm** be writeable, so that the system
122 administrator using an NMS can change the scheduling algorithm?
123 Closed: No, we agreed to delete this object, since no implementations of job monitoring
124 with any protocol have such an object. Even Printxchange did not implement this
125 attribute, even though ISO DPA has this Printer attribute.
126
127 Issue 16 - Add **passThrough(6)** to **jmGeneralQueuingAlgorithm** for servers that just
128 pass jobs through without queuing?
129 Closed: No, none of the implementations even had the **jmGeneralQueuingAlgorithm**
130 object, so we decided to delete the entire object. So we don't need to even decide
131 whether to add a new enum value to it.
132
133 Issue 20 - OK to have added **fileName(3)** to **JmAttributeTypeTC**?
134 Closed: Yes, some implementations have both file names and document names, so we
135 need both.
136
137 Issue 21 - Change **physicalDevice(11)** to a text string, so it can be used with servers that
138 don't have the Printer MIB?
139 Closed: Have both **physicalDeviceName** and **deviceIndex** as resources in the
140 **jmAttributeTable**, so that neither, one, or the other or both can be implemented.
141
142 Closed: Issue 22 - Why not require the agent to always return FAX numbers in ASCII,
143 since it is easy to convert from Unicode to ASCII?
144 Closed: We decided to remove the fax number resource entirely, since it doesn't relate to
145 printing. When a FAX job monitoring MIB is developed to augment this Job Monitoring
146 MIB, it will need other objects, besides the FAX phone numbers. The FAX phone

147 numbers enum can be registered at that time as a type 2 enum for use with the
148 jmAttributeType object. (Had we kept this enum, and when it is registered, the data type
149 will be ASCII, rather than the coded character set of the implementation), in order to fit
150 into 63 octets. Two-octet Unicode would exceed the space since numbers with password,
151 and extensions, etc., could exceed 31 digits.

152

153 Issue 23 - Add resource item to indicate the output-bins that the job requests/uses?

154 Closed: Yes, it was in several implementations. So add outputBin as a resource enum
155 which can have multiple entries, since some jobs may actually use more than one output
156 bin.

157

158 We also agreed to add **colorantConsumed** as a resource enum and **mediumConsumed**
159 where the **jmResourceName** object would be the name of the actual colorant or medium
160 consumed, with one row per different colorant and medium.

161

162 Issue 24 - Move any resource items to the **jmJobGroup**, because monitoring applications
163 needs to access the resource frequently without having to read the entire
164 jmAttributeTable?

165 Closed: No, don't move any. In fact see issue 30, where we agreed to add a fourth index
166 to the jmAttributeTable, which makes all resources directly addressable by
167 jmAttributeTypeIndex as the third index.

168

169 Issue 26 - Which indexes shall be persistent across power off and which need not be?

170 Closed: The persistent indexes shall be: **jmJobIndex** and **jmGeneralJobSetIndex**

171

172 Issue 27 - Should **jmGeneralJobCompletedPolicy** be writeable, so that the system
173 administrator using an NMS can change the length of time that completed jobs are kept?

174 Closed: No, see answer to issue 9.

175

176 Issue 28 - Should **jmGeneralQueuingAlgorithm** be writeable, so that the system
177 administrator using an NMS can change the scheduling algorithm?

178 Closed: No, none of the implementations even had the **jmGeneralQueuingAlgorithm**
179 object, so we decided to delete the entire object. So we don't need to even decide
180 whether to make the object writeable.

181

182 Issue 29 - OK to require that **jmCompletedIndex** be monotonically increasing?

183 Closed: Yes. Also **jmQueueIndex**.

184

185 Issue 30 - Should we move any **jmJobGroup** objects to the **jmResourceGroup**?

186 Closed: We agreed to add a fourth index to **jmResourceTable**. That index is
187 **jmResourceType** enum, which will actually be the third index. The **jmResourceIndex**
188 will be moved from third to fourth position. See move8obj.doc in the contributions sub-
189 direction.

190 This makes each resource in the **jmResourceTable** directly addressable by
191 **jmResourceType**. making it as efficient to find an resource in the **jmResourceTable** as
192 to get an object in the **jmJobTable**.

193 Consequently, we agreed to move the following 9 objects to the **jmResourcesTable**:

- 194 1. **jmDeviceIndex**
- 195 2. **jmJobSourceChannel** - see Issue 37
- 196 3. **jmJobSubmissionTime**
- 197 4. **jmJobComment**
- 198 5. **jmJobTotalKOctets**
- 199 6. **jmJobKOctetsCompleted**
- 200 7. **jmJobStartedProcessingTime**
- 201 8. **jmJobCompletionTime**
- 202 9. **jmJobAccountName**

203

204 Issue 31 - Should we re-introduce **jmJobDeviceId** to handle configuration 2b?

205 Closed: We agreed to bring back configuration 2b in which the NMS has to query the
206 server by means outside the Job Monitoring MIB and the printer using the Job Monitoring
207 MIB. However, in this scenario, there is no Job Monitoring MIB in the server, so there is
208 no need to add back **jmJobDeviceId** for the job id assigned by the Printer for use in a
209 server's Job Monitoring MIB.

210

211 Issue 32 - Shouldn't we require any numeric portion of the client-side identifiers to always
212 be in the **jmJobIdNumber** object?

213 The new Job Id table that Harry proposed replaces the **jmJobIdNumber** and **jmJobIdName**
214 objects.

215

216 Issue 33 - Why have two client-side identifier objects?

217 Closed: Some job submission protocols, such as BSD LPR/LPD require two.

218

219 Issue 34 - What is the SNMPv1 and SNMPv2 error that an agent shall return if there is no
220 instrumentation for an object?

221 Closed: There is no such SNMP error. See the answer to Issue 12.

222

223 Issue 37 - Change the **jmJobSourceChannel** from an index in the Printer MIB to the
224 enum, since the server need not implement the Printer MIB?

225 Closed: No, spoolers typically know the job source channel. So keep as an index into the
226 Printer MIB for use in Printers. Since we are moving this to the Resource Table, make
227 the enum name reflect the index: `jobSourceChannelIndex`.

228

229 Issue 38 - Do we need to add the **jmJobChannelInformation** object to the new
230 `jmServerGroup` for servers that don't have a corresponding Printer MIB?

231 Closed: No, since we agreed to keep the Job Source Channel as an index into the Printer
232 MIB, the monitoring application can access the **jmJobChannelInformation** object in the
233 Printer MIB. see issue 37.

234

235 Issue 39 - ISSUE - Why not return the SNMP error ???, instead of -2, if the total K octets
236 is unknown?

237 Closed: There is no such error. Return `unknown(-2)` if not known.

238

239 Issue 41 - Is it worth rounding down **jmJobKOctetsCompleted** until the job completes
240 and then round up?

241 Closed: No, lets round up to the next higher K as with **jmJobKTotalOctets**. The only
242 time rounding down could make a difference is for a 1K job and that short a job will
243 happen so quickly that the difference between rounding up versus rounding down can not
244 be seen by people.

245

246 Issue 42 - Are `interpreters(10)`, `sheetsCompleted(14)`, `processingTime(20)` the right
247 resource items to require agents to implement?

248 Closed: We agreed only to make **sheetsCompleted(14)** mandatory, since the Printer MIB
249 requires it. All other resource items are conditionally mandatory.

250

251 Issue 43 - How can **jmResourceName** be a union of OCTET STRING, Integer32, and
252 Counter32?

253 Closed: We agreed to replace **jmResourceName** by two objects:
254 **jmResourceNameAsText** and **jmResourceNameAsType**. See `res-type.doc` in
255 `contributions` sub-directory. We also agreed that each resource fills in either
256 **jmResourceNameAsText** or **jmResourceNameAsType**, but not both. Also, except for
257 **mediaConsumed**, no resource item that fills in a **jmResourceAmount** with a count, also
258 fills in **jmResourceNameAsText** or **jmResourceNameAsType**.

259 Subsequently we agreed to combine the three objects: **jmResourceNameAsText**,
260 **jmResourceNameAsType** and **jmResourceAmount** into just two objects:
261 **jmResourceValueAsText** and **jmResourceValueAsInteger**.

262

263 Issue 44 - There was agreement that the Job Monitoring MIB needs a primary index that
264 can separate jobs into disjoint sets for purposes of scheduling. This primary index serves
265 the same purpose for the Job Monitoring MIB as the hrDeviceIndex does for the Printer
266 MIB, except that the Job Monitoring MIB did not want to require the Host Resources
267 MIB. What is the definition of Job Set? How do job sets relate to queues? Can a Printer
268 have more than one job set without having queuing? For a printer that is fed from
269 multiple external queues, are all the jobs from all those queues in a single job set in the
270 Printer?

271 Closed: See the Internet Draft 00 (and version 0.71). I added clarifications that an agent
272 in a server (configuration 2) or printer (configuration 1 or 3) will represent each queue in
273 that server or printer as a distinct job set, since the agent is representing the queue that is
274 in the server or printer that the agent is instrumenting. The agent cannot be expected to
275 represent a queue that is elsewhere (upstream or downstream). If the printer is fed from
276 multiple queues from the same or different servers, but the printer has only one queue,
277 then the agent in that printer shall represent that queue in the printer as a single job set.

278 Though a Job Set is most often representing a job queue, we're not calling the job set and
279 the job set index a queue and a queue index, since a server or printer need not implement a
280 queue and need not have any queuing. See if I succeeded in clarifying the definition of job
281 set in the draft. I added cross references back to the terminology section as well, so that
282 people who skip over the terminology or forget after reading the terminology will be
283 reminded to go back for further explanation.

284

285 Issue 45 - After fully agreeing on what a Job Set is, should the name remain Job Set, or be
286 changed to Queue, or Job Pool or Job Group to improve understandability? The new
287 **jmGeneralJobSetName** would also be changed to **jmGeneralQueueName**, or
288 **jmGeneralJobPoolName**, or **jmGeneralJobGroupName**. See the Internet Draft 00 and
289 version 0.71.

290 Closed: We agreed to leave the name as Job Set.

291

292 Issue 46 - Do we need to add a **jmGeneralJobSetName** object so that an operator can
293 determine which job set he/she is looking at. The **jmGeneralJobSetName** is
294 administratively assigned and could be (1) the queue name, (2) the server name if the
295 server has only one job set or (3) the printer name if the printer has only one job set.

296 Closed: I added the **jmGeneralJobSetName** object to the general table for the same
297 reason that we added **prtGeneralPrinterName** to the Printer MIB. I added the
298 explanation that the Job Set Name can be (1) the queue name, (2) the server name if the

299 server has only one job set or (3) the printer name if the printer has only one job set in the
300 Terminology section and under each group where the **jmJobSetIndex** is used.

301

302 Issue 47 - Should we change the name of the **jmResourceTable** to **jmAttributeTable**,
303 since some of the enums are not resources, such as **fileName** and **documentName**,
304 **jobSubmissionTime**, **jobCompletionTime**, etc.

305 Closed: I renamed the Resource Group and Table to **jmAttributeGroup** and
306 **jmAttributeTable**. See version 0.71 and the Internet Draft 00.

307

308 Issue 48 - Since **jmJobIndex** cannot be 0 according to SNMP rules for indexes, what
309 shall an agent do that is instrumenting a printer or server that uses 0 as a valid job-
310 identifier? Use largest positive integer for job 0? Or the agent map the 0 value to the
311 value that is one higher than the maximum that the server or printer uses?

312 Closed: With the Job ID table, what the values of the jmJobIndex are relative to job
313 identifiers that the server or printer generate has become less important. However, agents
314 can preserve the same values assigned by the server or printer in the jmJobIndex, by
315 mapping a zero job-identifier value to one higher than the server or printer assigns.

316

317 Issue 49 - Should we change the definition of the **jmGeneralMaxNumberOfJobs** to
318 **jmGeneralMaxJobIndex** meaning the maximum value that the **jmJobIndex** object can
319 have and the roll over to 1 happens for the next job received? Or add
320 **jmGeneralMaxJobIndex** as another object in the General table? Then the monitoring
321 application would know what the roll over limit would be. For agents that instrument
322 servers or printers that use a job identifier of 0, the actual maximum number would be one
323 more than the actual job identifier that the server or printer generates. So for LPD, the
324 value of **jmGeneralMaxJobIndex** would be 1000, not 999.

325 The current definition of **jmGeneralMaxNumberOfJobs** is:

```
326 jmGeneralMaxNumberOfJobs OBJECT-TYPE
327     SYNTAX          Integer32(0..2147483647)
328     MAX-ACCESS     read-only
329     STATUS          current
330     DESCRIPTION
331         "The maximum number of queued and completed jobs
332         that this server or print can support at the same
333         time.
334
335         The value (-1) indicating other shall indicate that
336         there is no fixed limit."
337     ::= { jmGeneralEntry 4 }
```

338 What is the purpose of **jmGeneralMaxNumberOfJobs** as currently defined?

339 Closed: No one could make a good case for this object, so we agreed to delete it. Also
340 with the new **jmGeneralOldestActiveJobIndex** and **jmGeneralNewestActiveJobIndex**,
341 an application can discover roll-over when the newest is less than the oldest.

342

343 ISSUE 50 - Should we define the **PrtInterpreterLangFamilyTC** used with the
344 **documentFormat(12)** attribute in the Job Monitoring MIB module, since the
345 **PrtInterpreterLangFamilyTC** textual convention is not yet in an RFC? Or should the
346 Job Monitoring MIB IMPORT the **PrtInterpreterLangFamilyTC** from the Printer-MIB
347 to be?

348 Closed: Since the **documentFormat(12)** is an attribute, not an object, the value is the
349 general **jmAttributeValueAsInteger** and so does not explicitly IMPORT any enum for
350 use as attributes (only as objects). So the Internet Draft 00 (and version 0.71) no longer
351 have **PrtInterpreterLangFamilyTC**, which nicely side-steps the issue that there is no
352 RFC with **PrtInterpreterLangFamilyTC** enum defined.

353

354 ISSUE 51 - Should the **jmJobCurrentState** (and **JmJobStateTC**) be changed from a
355 type 2 enum to a type 1 enum, since adding states would have serious impact on released
356 clients? Currently the IPP draft has the **job-state** and **printer-state** attributes defined as
357 type 1 enums (actually they've changed the terminology, but not the concept, to *keyword*).

358 Closed: No, we will keep as a type 2 enum, in case we need to make an addition, such as
359 follow IPP.

360

361 ISSUE 52- Should **JmJobStateReasonsTC** and **JmJobServiceTypesTC** be defined
362 using the RFC 1902 BITS built-in? **JmJobStateReasonsTC** is 540 bits, while
363 **JmJobServiceTypesTC** is only 31 bits.

364 Closed. Instead of using **BITS** which requires compilers that conform to the 19xx RFC
365 series, define the bit values as part of the DESCRIPTION in hexadecimal in each 32-bit
366 integer. Create four **jobStateReasonn** ($n=1..4$) and four **JmJobStateReasonsn** TCs.

367

368 ISSUE 53 - Should there be an object that specifies the current default coded character set
369 of the Job Monitoring MIB, so that the client can figure out how to interpret objects of
370 type OCTET STRING that are coded characters, in case the client might not be
371 configured the same as the server or printer? See Section 6 in the Internet Draft 00 and
372 revision 0.71 for the discussion of coded character sets, including the use of Unicode/ISO
373 10646.

374 Closed: No do not add such an object. Unlike the Printer MIB, the text strings in the Job
375 Monitoring MIB originate from clients, not from the server or printer. Therefore,
376 monitoring application programs are apt to be configured to support the same character
377 set and language as the job submission clients. Alternatively, a Job Monitoring agent that

378 also implements the Printer MIB could use the objects in the Printer MIB to indicate the
379 current localization of the Job Monitoring MIB.

380

381 ISSUE 54 - Should we move all of the objects in the **jmJobTable** (9 objects) into the
382 **jmAttributeTable** as enums and then specify some of them to be required for
383 implementation? What about the **jmQueueTable** (6 objects)? What about the
384 **jmCompletedTable** (3 objects)? This would reduce the number of required objects from
385 21 mandatory objects and 6 conditionally mandatory objects to just 9 mandatory objects
386 and no conditionally mandatory objects.

387 Closed: We agreed to Harry's proposal to move all of jmJobTable objects to the Attribute
388 Table, to delete the Queue and Completed Table and replace them with the jmJobIDTable
389 and the jmJobStateTable.

390

391 ISSUE 55 - Should we remove the **needsAttention** state to align with IPP (and DPA)?
392 The **printer-stopped job-state-reasons** value from IPP has already been added to the
393 **JmJobStateReasonsTC**, so that the user will be able to find out that the job that is
394 **processing** has a printer stopped.

395 Closed: No, we like the **needsAttention** state. Furthermore, it will be straightforward for
396 an agent instrumenting an IPP server or printer, to map the job and printer states,
397 including job-state-reasons and printer-state-reasons, to the Job Monitoring MIB job
398 states.

399

400 ISSUE 56 - Which of the **jmJobTable** entries that were moved to the **jmAttributeTable**
401 should be mandatory enums, if any? They were all mandatory when they were in the
402 **jmJobTable**.

- 403 1. **physicalDeviceName** (or **physicalDeviceIndex**)
- 404 2. **jobSourceChannelIndex**
- 405 3. **jobSubmissionDateAndTime** or **jobSubmissionTimeStamp**
- 406 4. **jobComment**
- 407 5. **jobKOctetsTotal**
- 408 6. **jobKOctetsCompleted**
- 409 7. **jobStartedProcessingDateAndTime** or
- 410 **jobStartedProcessingTimeStamp**
- 411 8. **jobCompletedDateAndTime** or **jobCompletedTimeStamp**
- 412 9. **jobAccountName**

413 Or should we move any mandatory attributes, such as **sheetsCompleted** back to the
414 **jmJobTable**, so that the attributes table contains no mandatory attributes, only
415 conditionally mandatory attributes and the **jmJobTable** is the place that we put the
416 mandatory information?

417 Closed: We agreed that any of the objects in the jmJobStateTable that are duplicates of
418 the jmAttributeTable, shall be mandatory. This gives the following 8 attributes as
419 mandatory:

420 **jobState**
421 **numberOfInterveningJobs**
422 **deviceAlertCode**
423 **jobKOctetsRequested**
424 **jobKOctetsCompleted**
425 **impressionsRequested**
426 **impressionsCompleted**
427 **outputBinName**
428
429

430 ISSUE 57 - OK to change **jmAttributeTypeIndex** from not-accessible to read-only, so
431 that it can be mentioned in the conformance clause where we specify that
432 **sheetsCompleted** is the only attribute that is mandatory (so far). Currently the Internet
433 Draft and version 0.71 get a compile error when attempting to mention an enum that is
434 mandatory for an object that is not listed in the conformance clause. Objects that are not-
435 accessible cannot be mentioned in the conformance clauses, but read-only can, since they
436 (**jmAttributeTypeIndex**) would also get added to the list of objects in the group.

437 Closed: No, leave **jmAttributeTypeIndex** as non-accessible. Provide comments to list
438 the attributes that are mandatory in the conformace section.

439

440 Issue 58 - OK that sides, **documentFormat**, and **physicalDeviceIndex/Name**, remain as
441 the only attributes that are both requested and used at the same time (with the same
442 instance in the **jmAttributeTable**) as suggested in Ron's EMail of 2/15, or should we
443 make four new attributes that are used/consumed and change the current ones to be
444 requested?

445 Closed: Leave as they are. The difference between requested and consumed is not very
446 great for these attributes.

447

448 Issue 59 - Add the name of the source server or client? As an object or an attribute? See
449 Bob Pentecost EMail of 2/25. Need more specific text for such proposals.

450 Closed: Agreed in principle, but need a specific specification for each.

451

452 Issue 60 - Add the "file name of the job" and a "source port object to tell which client port
453 the job came from"? As objects or attributes? See Bob Pentecost EMail of 2/25. Need
454 more specific text for such proposals.

455 ACTION ITEM (Tom, Bob): Write up a proposal and send to the DL.

456

457 Issue 61 - Need to clarify the semantics of each object and attribute with respect to
458 Configuration 1, 2, and 3. See Bob Pentecost EMail of 3/10 (HP internal review). Most
459 objects refer to the jobs as they exist in the server or printer that the agent embedded in,
460 i.e., is instrumenting. A few objects, represent information that comes from upstream

461 places in the case of configuration 1 from the client, in the case of configuration 2, the
462 client as well, and in the case of configuration 3, the server and maybe even the client as
463 well.

464 ACTION ITEM (Tom): Analyze the existing attributes to see if the semantics need
465 clarification depending on which configuration and send to the DL.

466 Closed: No more ambiguous attributes found.

467

468 Issue 62 - Harry Lewis has a proposal for a mapping table that allows a monitoring
469 application that knows a client identifier to directly address the mapping table with a single
470 get in order to find the **jmJobIndex** that the printer is using. See 3/5/97 and 3/28/97
471 EMail and <ftp.pwg.org/pub/pwg/jmp/contributions/pwgjm.pdf>. Harry will make a
472 presentation at the JMP meeting.

473 Closed: The JMP accepted Harry's proposal.

474

475 Issue 63 - Should we add attributes for inkjet plotters? See EMail from Patrick Powell, of
476 3/4/97

477 Closed: Wait until someone comes forward and wants such a value.

478

479 Issue 64 - Need to fill out Appendix A on mapping from the job submission protocols to
480 the Job Monitoring MIB for each of the three configurations.

481 Closed: Put into a separate document.

482 ACTION ITEM (all): Write up your job submission protocol mapping to the Job
483 Monitoring MIB.

484

485 Issue 65 - What Appendices should remain, which should be separate Internet Drafts
486 and/or informational RFCs and which should disappear?

487 Closed: No appendices for the Job Monitoring MIB, except for supplemental information
488 about the semantics of job states. A second appendix indicates the job submission
489 protocols that support the **jmJobSubmissionID** concept and what mechanism is used for
490 same. Put any other information into a separate informational RFC, such as mapping to
491 ISO DPA, mapping to IPP, mapping to other job submission protocols, etc.

492

493 Issue 66 - What attributes need to have Server vs. Printer values, since both may be
494 needed in a single implementation?

495 From Harry Lewis mail of 2/19. There was a need for an attribute to specify whether the
496 submission time was when the job was submitted to the server vs. the printer, since with

497 configuration 2 and 3, they could happen at quite different times. So I would think that
498 we would need another attribute.

499 Presumably a server would always have access to a date and time clock, so would *not*
500 need the TimeStamp form for the server. So should we replace:

- 501 1. **jobSubmissionDateAndTime**
- 502 2. **JobSubmissionTimeStamp**

503 with:

- 504 1. **jobSubmissionToServerDateAndTime**
- 505 2. **jobSubmissionToPrinterDateAndTime**
- 506 3. **JobSubmissionToPrinterTimeStamp**

507 If you are implementing just a printer and don't know when the job was submitted to the
508 server, you would not implement the **jobSubmissionToServerDateAndTime** attribute.

509 What other attributes do we need separate instances for the server vs.

510 the printer in order to handle configuration 2 and 3?

511 Closed: Change the two attributes to the three attributes.

512 ACTION ITEM (Tom): Send to the DL, if find other attributes that are different between
513 the server and the printer.

514

515 ISSUE 67 - Delete the three objects in the Job State table that duplicate attributes?

516 **jmJobStateKOctetsCompleted**, **jmJobStateImpressionsCompleted**, and
517 **jmJobStateAssociatedValue**?

518 An app can get all of these from the AttributeTable directly:

519 **jobStateKOctetsCompleted**, **jobStateImpressionsCompleted**, and
520 **jobStateAssociatedValue**.

521 Closed: Delete the duplicates from the Attributes Table instead.

522

523 ISSUE 68 - Delete the **Job State Group/Table** all together, since all objects are also
524 duplicated as attributes?

525 If ISSUE 67 does delete the 3 objects from the **Job State table**, then only the

526 **jmJobState** object remains. But that is also available in the Attribute Table as the
527 **jobState** attribute.

528 Closed: No, the Job State Table is useful to scan for jobs using Get Next and select the
529 desired columns.

530

531 ISSUE 69- Does order of assignment of **JmAttributeTypeTC** enums make any
532 difference?

533 Would it help if the mandatory attributes were first, so that Get Next would pick them up
534 first when getting the next conceptual row?

535 Closed: No, can't use Get Next to step through jobs. The requester can specify which
536 attributes using Get, since the agent is now required to materialize each supported
537 attribute when the job is accepted. So the application can supply a number of Gets in a
538 single PDU without fear of a error, once the application has learned which attributes the
539 agent implements.

540

541 ISSUE 70 - Add some simple general device alert TC, instead of using the Printer MIB
542 Alert Codes.

543 The **PrtAlertCodeTC** generic values are *not* much good to an end user without knowing
544 which subunit. For example, **SubUnitEmpty** isn't very informative by itself. If an
545 implementation also has the Printer MIB, then a lot more information is available, so a
546 copy of the Printer Alert isn't very useful. If the implementation doesn't have the Printer
547 MIB, then the Printer Alert codes aren't informative enough.

548 Even worse, the deviceAlertCode(10) is Mandatory, which can't be implemented, if there
549 isn't a Printer MIB also implemented.

550 Closed: No, use the alert codes.

551

552 ISSUE 71 - Are there any attributes that need to be clarified as to which apply to servers
553 and which apply to devices and which apply to either?

554 ACTION ITEM (Tom): Send to the DL, if find other attributes that are different between
555 the server and the printer.

556 Closed: no more found.

557

558 ISSUE 72 - What should happen to **jmGeneralNewestActiveJobIndex** when all the
559 active jobs complete?

560 Shall agent set it to 0 or leave it alone as a pointer to increment when the next job is
561 accepted? If it is reset to 0 (to indicate that there are no more active jobs), what
562 remembers where to put the next received job? What is remembered across power cycles,
563 so that jmJobIndex values are not immediately re-assigned upon power up? If the newest
564 active job is completed before an older one, shall the agent search back to find the newest
565 still active job and decrement **jmGeneralNewestJobIndex** to point to it? Or should this
566 object really be left alone after the newest job completes and be called
567 jmGeneralNewestJobIndex, since the newest job may no longer be active?

568 Closed: the agent shall reset it to 0 and keep an internal variable for the next row to
569 assign. That internal variable shall be persistent across power cycles. Also the agent shall
570 find the next newest active job, when the newest is canceled or completes and there are
571 still active jobs in the tables.

572

573 Issue 73 - Is there a problem with **outputBinIndex** being made mandatory?

574 If **outputBinIndex** is made mandatory, but an implementation doesn't have the Printer
575 MIB, the agent has to put 0 as the value. Should we add one more attribute:
576 **outputBinNumber**, which is just a number, not an index into the Printer MIB? If we do,
577 which should be mandatory? Just one more reason to get rid of the jmStateTable, which is
578 forcing us to pick a particular outputBin implementation and make it mandatory. If we
579 got rid of the JobState table, we could forget about making any of the 3 outputBinName,
580 outputBinNumber, or outputBinIndex attribute mandatory.

581 Closed: Don't add outputBinNumber. Don't make outputBinIndex MANDATORY.
582 Also keep **outputBinIndex** as a MULTI-ROW attribute, so don't need to add multi(-3)
583 enum value.

584

585 ISSUE 74: Collapse pairs of attributes that use Integer vs Octets value?

586 Analysis of the 78 attributes shows that there are 8 attribute pairs that could be collapsed
587 into one attribute with the implementation using either the Integer or the Octets value (or
588 both) to represent the attribute. The application would have to query both value
589 objects. But if it is using GetNext, it has to get both each time anyway. Only if it is directly
590 accessing an attribute would it have to get both values. On the other hand, it would be
591 fewer Gets than having two attributes as we have now. The 8 pairs are:

592 physicalDeviceIndex	physicalDeviceName
593 outputBinIndex	outputBinName
594 mediumRequestedType	mediumRequestedName
595 colorantRequestedIndex	colorantRequestedName
596 colorantConsumedIndex	colorantConsumedName
597 jobSubmissionToDeviceDateAndTime	jobSubmissionToDeviceTimeStamp
598 jobStartedProcessingDateAndTime	jobStartedProcessingTimeStamp
599 jobCompletedDateAndTime	jobCompletedTimeStamp

600

601 Should we collapse them into the following:

602 physicalDevice
603 outputBin
604 mediumRequested
605 colorantRequested
606 colorantConsumed
607 jobSubmissionToDeviceTime
608 jobStartedProcessingTime
609 jobCompletedTime

610

611 Should we allow an implementation to fill in both with meaningful values in a single row?

612 Closed: Yes, and allow agents to implement one, the other, or both values. Making it
613 easy for an agent to do both will make such an application that doesn't want to depend on

614 the Printer MIB being implemented, i.e., the application wants to work with all three
615 configurations.

616

617 ISSUE 75 - Should the Attribute enum values be grouped so additions could be added in
618 the appropriate section

619 When producing the first Internet-Draft, I re-arranged the Attribute enums into logical
620 groups, so that attributes would be easier to find. We now have 78 attributes, so logical
621 grouping is becoming important to make the list more understandable. Several people had
622 proposed adding attributes that were already present in the spec. Also Harry has
623 expressed the concern that any re-assignment of at least OIDs, causes problems with
624 tracking the drafts Finally, when the standard achieves proposed status, there will be
625 additional registrations. It might be helpful if the enums could be assigned to the
626 appropriate group, instead of only at the end.

627 The current logical grouping are:

628	Job State attributes	10
629	Job Identification attributes	19
630	Job Parameter attributes	7
631	Image Quality attributes (requested and used)	6
632	Job Progress attributes (requested and consumed)	7
633	Impression attributes (requested and consumed)	6
634	Page attributes (requested and consumed)	3
635	Sheet attributes (requested and consumed)	3
636	Resource attributes (requested and consumed)	7
637	Time attributes (set by server or device)	9

638 Ok to assign Job State and Job Identification in steps of 30 and the rest in steps of 20?

639 See also Issue 69. We could put the mandatory attributes first, and then group the rest as
640 above.

641 Closed: Yes, group the enum assignments.

642

643 Issue 76 - So should **jobName**, **jobOwner**, and one of **deviceNameRequested** or
644 **queueNameRequested** be made Mandatory?

645

646 When we moved attributes from the job table to the attributes table (Issue 54 and 56), we
647 didn't make any of them mandatory for an agent to implement. Should any of them be
648 made Mandatory?

649 The old job table had the following (mandatory) objects in it:

650 **jmJobName**
651 **jmJobIdName**
652 **jmJobIdNumber**
653 **jmJobServiceType**
654 **jmJobOwner**
655 **jmJobDeviceNameOrQueueRequested**

656 **jmJobCurrentState**
657 **jmJobStateReasons**
658
659 **jmJobIdName** and **jmJobIdNumber** have been replaced by **jmJobSubmissionIDIndex**
660 which is Mandatory.
661
662 **jmJobServiceType** need not be Mandatory.
663
664 Also **jmJobDeviceNameOrQueueRequested** has been made into two separate attributes:
665 **deviceNameRequested** and **queueNameRequested**, so we'd have to make either one of
666 them mandatory.
667
668 **jmJobCurrentState** is now **jobState** and is Mandatory
669
670 **jmJobStateReasons** became four attributes: **jobStateReasons1**, **jobStateReasons2**,
671 **jobStateReasons3**, and **jobStateReasons4**. None of them need to be Mandatory.
672
673 So should **jobName**, **jobOwner**, and one of **deviceNameRequested** or
674 **queueNameRequested** be made Mandatory?
675 Closed: Only **jobOwner** is made mandatory, but it will remain in the Attribute table,
676 rather than being moved to the Job table.
677
678 Issue 77 - Should **jobCompletedDateAndTime/TimeStamp** be canceled time too, or
679 add **jobCanceledDateAndTime/TimeStamp**?
680 Should we just clarify the **jobCompletedDateAndTime** and **jobCompletedTimeStamp**
681 attributes may be used for either the time that the job completes or the time that the job
682 canceled?
683 Or is it better to add two new attributes: **jobCanceledDateAndTime** and
684 **jobCanceledTimeStamp**?
685 Closed: Clarify that completed included canceled and aborted.
686
687 Issue 78 - Should the "multiplexor" **jobStateAssociatedValue(4)** attribute be removed
688 from the Job Attribute Table and the equivalent **jmJobStateAssociatedValue** object be
689 removed from the Job State table?
690 The associated values are also available as attributes in the attribute table. The application
691 has to either (1) request all 7 associated attributes or (2) first request the **jobState(3)**
692 attribute and then request the 1 pertinent attribute. Since all 7 will easily fit in a PDU
693 (minimum of 500 octets or so on all systems) and each request takes about 20 octets, so
694 you can get about 20 (5*4) attributes into a single PDU.
695 Closed: Yes. The application can request each object and/or attribute directly and it will
696 fit into a single PDU (20 objects or attributes).

697

698 Issue 79 - Should the '**printing**' state be combined into the '**processing**' state?

699 Many printers don't distinguish between '**processing**' and '**printing**', especially desktop
700 printers. For those that do, having a state change that really reflects progress, such as the
701 transition from processing to printing, is better handled as a job state reason, not as a
702 fundamental state change. Finally, since this MIB is intended for non-printing services in
703 the future, such as fax out, CD-ROM writing, fax-in, scanning, etc., it would help if one of
704 the states wasn't '**printing**'. Even IPP, only has the state of '**processing**', with a job-
705 state-reason of '**job-printing**' for those implementations that make the distinction and
706 want to go to the trouble of indicating the difference. IPP even indicates that "most
707 implementations won't bother with this nuance".

708 Closed: Yes, but the other differences between JMP and IPP need discussion with IPP.

709

710 Issue 80 - How handle IPP "**sides**" attribute?

711 The IPP "**sides**" attribute shows more than just 1 or 2. It has the following values: '**1-**
712 **sided**', '**2-sided-long-edge**' (i.e., duplex), and '**2-sided-short-edge**' (i.e., tumble).

713 Alternatives:

714 (a) Change the Job Monitoring MIB "**sides(32)**" attribute to an enum with these three
715 values.

716 (b) Add a new sides attribute, say, "**sidesType**" with these 3 enum values, and keep the
717 current "**sides(32)**" as an Integer32(1..2) (actually Integer32(-2..2), so can represent other
718 and unknown).

719 Closed: Don't include the IPP values; the agent can map them to 1 or 2.

720

721 Issue 81 - Add IPP "numberUp" attribute?

722 One of the IPP attributes that might be helpful to an administrator would be to record
723 number-up. An administrator that is bent on saving paper, might give rewards (or lower
724 charges) to users that used number-up.

725 Closed: No. Can get whether number up is being used by comparing the conditionally
726 mandatory **pagesCompleted** attribute with the **jmJobImpressionsCompleted** object.

727

728 ISSUE 82 - Change the OID assignment as Jeff Case and David Perkins suggest:

```
729 >      jobmonMIB  
730 >      jobmonMIB.1 jobmonMIBObjects  
731 >      jobmonMIB.1.1 jmGeneral  
732      jobmonMIB.1.1.1 jmGeneralTable  
733      jobmonMIB.1.1.1.1 jmGeneralEntry  
734      jobmonMIB.1.1.1.1.1 jmGeneralJobSetIndex  
735      jobmonMIB.1.1.1.1.2 jmGeneralNumberOfActiveJobs  
736      jobmonMIB.1.1.1.1.3 jmGeneralOldestActiveJobIndex  
737      jobmonMIB.1.1.1.1.4 jmGeneralNewestActiveJobIndex  
738      jobmonMIB.1.1.1.1.5 jmGeneralJobPersistence
```

```

739         jobmonMIB.1.1.1.1.6 jmGeneralAttributePersistence
740         jobmonMIB.1.1.1.1.7 jmGeneralJobSetName
741
742 >         jobmonMIB.1.2 jmJobID
743         jobmonMIB.1.1.1 jmJobIDTable
744         jobmonMIB.1.1.1.1 jmJobIDEntry
745         jobmonMIB.1.1.1.1.1 jmJobSubmissionID
746         jobmonMIB.1.1.1.1.2 jmJobIDJobSetIndex
747         jobmonMIB.1.1.1.1.3 jmJobIDJobIndex
748
749 >         jobmonMIB.1.3 jmJob
750         jobmonMIB.1.1.1 jmJobStateTable
751         jobmonMIB.1.1.1.1 jmJobStateEntry
752         jobmonMIB.1.1.1.1.1 jmJobIndex
753         jobmonMIB.1.1.1.1.2 jmJobState
754         jobmonMIB.1.1.1.1.3 jmJobStateReason1
755         jobmonMIB.1.1.1.1.4 jmNumberOfInterveningJobs
756         jobmonMIB.1.1.1.1.5 jmJobKOctetsRequested
757         jobmonMIB.1.1.1.1.6 jmJobStateKOctetsProcessed
758         jobmonMIB.1.1.1.1.7 jmJobImpressionsRequested
759         jobmonMIB.1.1.1.1.8 jmJobStateImpressionsCompleted
760         jobmonMIB.1.1.1.1.9 jmJobOwner
761
762 >         jobmonMIB.1.4 jmAttribute
763         jobmonMIB.1.1.1 jmAttributeTable
764         jobmonMIB.1.1.1.1 jmAttributeEntry
765         jobmonMIB.1.1.1.1.1 jmAttributeTypeIndex
766         jobmonMIB.1.1.1.1.2 jmAttributeInstanceIndex
767         jobmonMIB.1.1.1.1.3 jmAttributeValueAsInteger
768         jobmonMIB.1.1.1.1.4 jmAttributeValueAsOctets
769
770         jobmonMIB.2         jobmonMIBNotifications -- reserved
771
772 >         jobmonMIB.3         jobmonMIBConformance
773         jobmonMIB.3.1         jobmonMIBCompliance
774         jobmonMIB.3.2         jmMIBGroups
775         jobmonMIB.3.2.1       jmGeneralGroup
776         jobmonMIB.3.2.2       jmJobIDGroup
777         jobmonMIB.3.2.3       jmJobGroup
778         jobmonMIB.3.2.4       jmAttributeGroup
779

```

780 Correct?

781 >yes, if

- 782 > 1. jobmon is not somewhere under printmib
- 783 > 2. you don't have something like
- 784 > jobmonMIB.2 jobmonNotifications
- 785 > in which case then
- 786 > jobmonMIB.3 jobmonMIBconformance
- 787 > would move down one arc

788 Closed: Yes, including reserving an OID for traps, in case we need them in the future.

789 Closed: Add **jmGeneralJobSetIndex** to **jmGeneralTable** and **jmJobIndex** to
790 **jmJobTable** to follow SMI rules, even though previous version compiles cleanly. This
791 will avoid what David Perkins refers to as "illegal" indexing. Then have to rename the two
792 columns in the **jmJobIDTable** to not conflict.

793

794 ISSUE 83 - Can some attributes be deleted before the **jmGeneralAttributePersistence**
795 expires?

796 Harry Lewis' 5/2 e-mail suggested that some of the attributes, such as
797 "**numberOfInterveningJobs(9)**" don't even need to persist the shorter time specified by
798 **jmGeneralAttributePersistence**.

799 Closed: No. All attributes shall be instantiated at the same time and deleted at the same
800 time. Then applications can request any number of objects and attributes in a single PDU
801 and not get an error back on one that has been implemented but hasn't been put in the
802 table. The values may change at any time.

803 Revisited: Subsequently, we relaxed this requirement, so that the agent NEED NOT
804 materialize all supported attributes when the job is received.

805

806 ISSUE 84 - Change Associated Value for '**printing**' state to
807 **impressionsCompletedCurrentCopy(56)**?

808 The MIB attribute **jobStateAssociatedValue(4)** specifies that the associated value for the
809 'printing' state is the STATIC attribute: **impressionsRequested(54)**. Should we change it
810 to the DYNAMIC value: **impressionsCompletedCurrentCopy(56)**? A monitoring
811 application that wants to create a thermometer can read the STATIC
812 **impressionsRequested(54)** attribute once from the **jmAttributesTable**.

813 What about the STATIC **impressionsRequested(54)** associated value that is associated
814 with the '**processing**' state? Leave it or change it to something dynamic, like
815 **jobKOctetsCompleted(50)**?

816 Closed: Since the AssociatedValue object/attribute is being deleted, this issue is moot.
817

818 ISSUE 85 - Break the MIB into a monitoring and an accounting MIB?

819 Need to agree if the accounting MIB augments the monitoring MIB or vice versa?

820 Then need to agree on which attributes only apply to the augmenting MIB.

821 Closed: No. There are too many attributes that are used for both monitoring and
822 accounting.

823

824 ISSSUE 86 - Clarify **jobCopiesRequested(44)** vs. **documentCopiesRequested(46)**

825 In order that systems that only support one document jobs and systems that support
826 multiple documents per job, use the same attribute when the job has only one document
827 (most usual case) and multiple copies are being made, need to clarify which attribute to
828 use: **jobCopiesRequested(44)** vs. **documentCopiesRequested(46)**. Also which to use (or
829 both) when muliple copies of a job are requested when the job has multiple documents as
830 well. Need to map clearly to IPP and other job submission protocols.

831 Closed: Use **jobCopiesRequested** for single document jobs for both systems that support
832 only one documen t per job and ones that support mujltiple documents. Only use

833 **documentCopiesRequested**, when a multiple document job actually specifies that
834 individual documents are to be made copies.

835

836 ISSUE 87 - When shall the mandatory attributes appear in the **jmJobAttributesTable**?

837 Shall an agent materialize all mandatory attributes when the job is submitted, so that a
838 requester can access them all with multiple explicit Gets in a single PDU, without fear of a
839 missing object aborting the PDU? If the mandatory attributes are represented as objects in
840 the **jmJobStateTable**, then it is clear from SNMP rules that the agent shall materialize at
841 least an empty value for each mandatory object (attribute).

842 Closed: The agent can materialize all the attributes when the job is submitted, including
843 with unknown values, or the agent can materialize the attribute subsequently when the
844 attribute values become known. Management applications need to use GetNext to get
845 attributes that might not be present yet or expect SNMP error codes to be returned.

846

847 ISSUE 88 - Add **jmGeneralNumberOfJobsProcessed** object since server or printer was
848 booted.?

849 Most MIBs have some sort of utilization counter. The Job Monitoring MIB should have
850 one also. Add the object to the **jmGeneralTable**. We assume that this object SHALL
851 NOT be persistent across power cycles.

852 The DESCRIPTION is proposed to be:

```
853 jmGeneralNumberOfJobsProcessed OBJECT-TYPE
854     SYNTAX      Integer32 (0..2147483647)
855     MAX-ACCESS  read-only
856     STATUS      current
857     DESCRIPTION
858         "The number of jobs that have completed processing
859         for this job set since the server or device was
860         powered on."
861     ::= { jmGeneralEntry 7 }
```

862 Closed: Do not add **jmGeneralNumberOfJobsProcessed** object. The number of pages
863 printed is more indicative of load, than the number of jobs, since jobs can be short or long.

864

865 ISSUE 89 - Add **jmGeneralAttributesImplemented** object with bits for each attribute
866 implemented?

867 Instead of an application not knowing which attributes an implementation implements and
868 trying to discover by getting errors, or by always using Get Next, instead of Get, how
869 about adding a **jmGeneralAttributesImplemented** object to the **jmGeneralTable** that
870 has a bit for each attribute implemented.

```
871 jmGeneralAttributesImplemented OBJECT-TYPE
872     SYNTAX      OCTET STRING(SIZE(0..32))
873     MAX-ACCESS  read-only
874     STATUS      current
```

875 DESCRIPTION
876 "A bit string indicating which **JmAttributeTypeTC**
877 enum values are implemented. The value is a
878 constant independent of which bits are currently in
879 entries in the **jmAttributeTable**. The most
880 significant bit of the first octet is assigned the
881 value 0 to correspond to enum 0 (not used), the next
882 most significant bit of the first octet is assigned
883 the value 1 to correspond to enum 1 (**other**), the
884 next bit is assigned the value 2 (**unknown**), 3
885 (**jobStateReasons2**) etc. up to $32*8 - 1 = 255$ "
886 ::= { jmGeneralEntry 8 }

887 Closed: Do not add **jmGeneralAttributesImplemented** object. The representatives of
888 the applications did not think such an object would help. Their applications have to use
889 GetNext anyway, since not all supported attributes are materialized when the job is
890 accepted.

891

892 ISSUE 90 - The (MANDATORY) OCTET STRING objects should have a minimum
893 MAX size required

894 Otherwise, trivial implementations can implement too short sizes and be conforming. The
895 SNMP conformance syntax as the end of the MIB has provision for specifying the
896 minimum maximum that SHALL be implemented. The 63 in OCTET
897 STRING(SIZE(0..63)) is the maximum size and the 0 is the minimum size for an *instance*
898 returned on a Get operation. The (MANDATORY) OCTET STRING objects are with
899 suggested MAX size required:

900 **jmGeneralJobSetName** 8 octets required to be implemented

901 **jmJobSubmissionID** 32 octets required to be implemented

902 Closed: Agreed to 8 and 48 as the minimum maximum number of octets that an agent
903 MUST support. We also agreed that jmJobSubmissionID MUST be fixed length, so that
904 an application can omit trailing octets, in order to achieve a "search" capability on just the
905 more significant part of the ID.

906

907 ISSUE 91 - The MANDATORY **jobOwner** attribute should have a minimum MAX size
908 required

909 Otherwise, trivial implementations can implement too short sizes and be conforming.

910 We suggest that 16 octets of the maximum 63 be required to be implemented. Such a
911 maximum will also help with the next issue.

912 Closed: We agreed to 16 octets be the minimum maximum that an agent MUST
913 implement.

914

915 ISSUE 92 - The MANDATORY **jobOwner** attribute needs to persist as long as there is
916 job data

917 92a: The MANDATORY **jobOwner** attribute needs to persist as long as there is job
918 data. Otherwise, a client that does not have access to the **jobSubmissionID** or a system
919 that does not have a **jobSubmissionID** cannot identify the jobs in the **jmJobTable**, after
920 the agent removes the job's attributes from the **jmAttributeTable**.

921 Closed: Agreed.

922 92b: If it is agreed that the MANDATORY **jobOwner** SHALL persist for the longer
923 period of time, then it should be moved to the **jmJobTable**, where all the other
924 MANDATORY attributes have been made objects.

925 Closed: Agreed.

926 92c: Then the **jmAttributeTable** could be made OPTIONAL, so that really low end
927 printers would not need to implement the **jmAttributeTable** at all. Experience at Xerox
928 with low end non-queuing printers suggests that not requiring the **jmAttributeTable** is a
929 win. With a required maximum of only 20 octets (see previous issue), it is reasonable to
930 move the **jobOwner** to the MANDATORY **jmJobTable**.

931 Closed: No. The **jmAttributeTable** shall be MANDATORY for consistency. There
932 isn't really any difference between an implementation that doesn't implement any attributes
933 and one that doesn't implement the **jmAttributeTable**. So keep the conformance
934 requirements simpler to understand and state.

935

936 ISSUE 93 - The **jobName** and **jobSubmission[ToDevice]Time** should be
937 MANDATORY

938 The windows queue monitoring show three attributes: **jobOwner**, **jobName**, and start
939 time. In IPP, **jobName** is a MANDATORY attribute.

940 Closed: The **jobName** and **jobSubmissionTime** attributes are NOT MANDATORY.
941 They are like any other attributes that shall be implemented, if the service or device has the
942 functionality and the agent is able to access it.

943 93a: In order to work for all configurations, the **jobSubmissionToDeviceTime** should be
944 changed to **jobSubmissionTime** and be used for all three configurations: configuration 1:
945 device, configuration 2: server, and configuration 3: device. The
946 **jobSubmissionToServerTime** shall only be used in configuration 3, where
947 **jobSubmissionTime** is also used (for the device). Then **jobSubmissionTime** can be
948 made MANDATORY (since it can be used in all three configurations).

949 Closed: Agreed.

950 93b: If the **jobName** attribute is made MANDATORY then it should have a minimum
951 maximum value specified for conformance. We suggest that 12 octets is sufficient.

952 Closed: Not MANDATORY, so no minimum maximum value specified.

953 93c: If **jobName** and **jobSubmissionTime** are made MANDATORY, then they should
954 be moved to the **jmJobTable** as well, so that the **jmAttributeTable** can remain
955 OPTIONAL (see previous issue).

956 Closed: Not MANDATORY, so not moved.

957

958 ISSUE 94 - Are the 8 octet fields in the **jobSubmissionID** printable or binary?
959 The text says "8-decimal-digits". Could it be allowed to be a binary sequence number or
960 random number? Then the chances of collision are even lower.

961 Closed: printable ASCII.

962

963 ISSUE 95 - When reducing the size of the **jobSubmissionID** field from 2 to 1, the other
964 fields weren't increased by 1.

965 Closed: Fix and increase the jmJobSubmissionID index object from 32 octets to 48 octets,
966 in order to be more unique by reducing the chances of truncation.

967

968 ISSUE 96 - Add a **jobSubmissionID** format for **jobOwner**

969 The first octet would be an ASCII '4'. The next 8 would be a sequential number, and the
970 remaining 23 octets would be the low order 23 octets of the **jobOwner**.

971 Closed: Agreed.

972

973 ISSUE 97 - Add some **jobSubmissionID** formats for numeric identifiers

974 97a - Add one for POSIX user numbers

975 Closed: Agreed

976 97b - Add one for user account numbers

977 Closed: Agreed

978 97c - Add one for DTMF incoming FAX routing number

979 Closed: Agreed

980

981 ISSUE 98 - The sequence number and random number in the **jmJobSubmissionID**
982 should be the least significant field, not the most significant field

983 Then a requester can leave off the sequential number or random number in a GetNext and
984 find all of the jobs from a particular MAC address or client URL (or for a particular
985 **jobOwner**). In order to make this switch, we need to specify that when the MAC
986 address, client URL, (jobOwner, or numeric id) is shorter than 23 octets, that the field is

987 shortened, rather than being padded out to 23 octets. The least significant field is always
988 8 octets with leading zeroes, so that we don't need any delimiters between the two fields.

989 So the spec would become:

990	Format	
991	Number	Description
992	-----	-----
993	1	octets 2 to n: upto last 23 bytes of the jobName
994		attribute; n < 26
995		octets n to n+7: 8-decimal-digit random number
996		
997	2	octets 2 to n: Client MAC address; n < 26
998		octets n to n+7: 8-decimal-digit sequential
999		number
1000		
1001		
1002	3	octets 2 to n: last 23 bytes of the client URL;
1003		n < 26
1004		octets n to n+7: 8-decimal-digit sequential
1005		number
1006		
1007	..	to be registered according to procedures of a
1008		type 2 enum. See section Error! Reference source
1009		not found. on page Error! Bookmark not defined..

1010 Closed: Agreed, but extend the total length from 32 to 48 octets.

1011

1012 ISSUE 99 - The **jobSubmissionID** format '0' makes no sense

1013 99a: There are two interpretations of the text there:

1014 1. The agent only puts a single digit of '0' for the entire jobSubmissionID index - but each
1015 subsequent submission would replace the entry.

1016 2. The agent tacks on whatever it wants after the '0' to make a unique jobSubmissionID
1017 index for each job.

1018 98b: If interpretation 2 is correct, then could we require the agent to use the **jobOwner**
1019 instead, rather than leaving it unspecified how the agent fills in the entry if interpretation is
1020 correct?

1021 Closed: The agent SHALL use any of the registered formats when the submitting client
1022 does not supply a **jobSubmissionID**.

1023

1024 ISSUE 100 - The **jobSubmissionIDTable** should have **jmJobSet** and **jmJobIndex** as
1025 indexes

1026 The current formats will have collisions of **jobSubmissionIDs** occasionally. The
1027 statement on lines 1695-1697: "None-the-less, collisions could occur, but without bas
1028 consequences, since this MIB is intended to be used only for monitoring jobs, not for
1029 controlling and managing them." is incorrect, since future IETF and private MIBs are

1030 likely to have 'missles', not 'cameras'. We've had experience with a table like this
1031 (**jobClientIdTable**) for a year and a half and we added the **jmJobIndex** equivalent to the
1032 row entry that the agent adds to make sure that no entry ever gets overwritten.

1033 So we propose that the **jmJobSet** object and the **jmJobIndex** objects be added as the
1034 least significant indexes to the **jmJobSubmissionIDTable**. They are only simple integers
1035 and **jmJobSet** is likely to be 1 in most implementations.

1036 Closed: Do not add **jmJobSet** and **jmJobIndex** as indexes to the **jmJobIDTable**. If
1037 implementors want to guarrantee uniqueness, they can request to register a new format in
1038 which the agent supplies some number of the least significant octets (same as if the
1039 submitting client did not supply a **jobSubmissionID**).

1040

1041 ISSUE 101 - add **jobSubmissionID** as an attribute, so can find the ID when scanning a
1042 job or attribute table.

1043 An accounting program that wants to find the **jobSubmissionID** would have to scan the
1044 entire **jmJobSubmissionIDTable**

1045 Closed: No, if a management application really needs the **jobSubmissionID** and doesn't
1046 know what it is apriori, then that application can scan the **jmJobIDTable** looking for the
1047 **jmJobIndex** value that matches.

1048

1049 ISSUE 102 - Make a TC out of the **jobSubmissionID** formats, so can publish new ones
1050 more easily?

1051 Closed: Yes.

1052

1053 ISSUE 103 - Specify a minimum required persistence time for
1054 **jmGeneralAttributePersistence**

1055 Put the lower bound right in the ASN.1. We suggest 60 seconds as a minimum with a
1056 recommendation of 300 (5 minutes). Even add a DEFVAL of 300 as the default. The
1057 real low cost device that doesn't want to keep job information around will have a small
1058 number of jobs anyway, since how many jobs can just a device process in a minute
1059 anyway?

1060 The spec would become:

1061 SYNTAX **Integer32(60..2147483647)**

1062 MAX-ACCESS read-only

1063 STATUS current

1064 DESCRIPTION

1065 "The minimum time in seconds for this instance of
1066 the Job Set that an entry will remain in the
1067 **jmAttributeTable** after **processing** has *completed* ,
1068 i.e., the time in seconds starting when the job
1069 enters the **completed**, **canceled**, or **aborted** state.
1070 The value of this object MAY be either (1) set by

1071 the system administrator by means outside this
1072 specification or MAY be (2) fixed by the
1073 implementation, depending on implementation.
1074
1075 This value SHALL be equal to or less than the value
1076 of **jmGeneralJobPersistence**. This value SHOULD be at
1077 least 300 which gives a monitoring application five
1078 minutes in which to poll for job data."
1079 DEFVAL { 300 } -- 5 minutes
1080 ::= { jmGeneralEntry 5 }

1081 Closed: but the agreed values are 60 (one minute) as a recommendation and 15 (15
1082 seconds) as the minimum, instead of 300 and 60, respectively.

1083

1084 ISSUE 104: Add deviceAlertIndex attribute which is index into alert table?

1085 deviceAlertCode (6) needs pointer to SNMP alert table - See pg. 36. When the device is a
1086 printer, the alert code SHALL be the printer alert code. This is the current definition. But,
1087 this is not very effective when genericAlertCodes are used. An index into the alert table
1088 would provide more information (rather than just JAM, you'd know jam in Input 3, for
1089 example). Maybe this is too much info for job monitoring? But it's just as easy for the
1090 agent.

1091 Proposed new attribute:

1092 deviceAlertIndex(8) -- Integer32(0..2147483647)
1093 -- INTEGER: MULTI-ROW: The device alert table index for the device that
1094 -- the job is using. When the device is a printer, this index SHALL be the
1095 -- index into the prtAlertTable defined by the Printer MIB[1]. Whether
1096 -- this attribute is instantiated for this job when another job is
1097 -- using the device depends on implementation.

1098 Closed: No, do not add **deviceAlertIndex**. Also remove **deviceAlert(7)**, since an
1099 application can access the Printer MIB if implemented. Also it is problematic for
1100 attributes to go away during the life of the job, such as the alert code attribute. Also its
1101 not good to duplicate information between two MIBs, since the information can become
1102 stale.

1103

1104 ISSUE 105: Ok to clarify that serverAssignedJobName(22) can be all digits?

1105 What happened to **serverAssignedJobNumber** (2x) - See pg. 37

1106 We used to have **serverAssignedJobNumber**, with syntax integer. I think we combined
1107 this with **serverAssignedJobName** (22) and dropped it, but in so doing, it is not listed as
1108 Octets (only). What about the original concern that (OS/2 and perhaps other) some os's
1109 use an integer not a text string. Are we saying the integer must be converted to text?

1110 I think you are referring to the **jmJobIdNumber** attribute that Ron proposed along with
1111 jmJobIdName. We deleted both when switching over to your jmJobSubmissionID table

1112 scheme. Ron is in agreement I believe with not needing either jmJobIdName or
1113 jmJobIdNumber.

1114 So for servers that assign numbers to jobs before submitting them to devices
1115 (configuration 3), rather than names, I would suggest that having the agent converting a
1116 number that it received in the job submission protocol to a text string would mean that we
1117 could use the same attribute and that an application would not need to deal with two
1118 attributes when querying.

1119 However, we should add a sentence clarifying that the text string may be a name or a
1120 number.

1121 Closed: Agreed, that **serverAssignedJobName** can be all numbers or even a URL.
1122

1123 ISSUE 106: Should **serverAssignedJobName** persist longer too?
1124 Persistence of serverAssignedJobName - See pgs. 36, 37

1125 Two other attributes (jobOwner and jobName) mandate "long" persistence. If you read the
1126 note under serverAssignedJobName, it leads in with the same reasoning, but stops short of
1127 requiring "long" persistence. Which persistence value is serverAssignedJobName intended
1128 to follow?

1129 Closed: No.
1130

1131 ISSUE 107: Ok to remove the three IPP **timeSinceXxx** attributes?
1132 Eliminate "timeSince" Attributes - See pg. 47 This is too much work for the agent and is
1133 contrary to SNMP in that sysUpTime should do the trick. I don't mind using a
1134 JmTimeStampTC rather than sysUpTime so much, but the NMS, not the agent, should
1135 calculate the times since.

1136 Good issue. The three timeSinceJobWasSubmitted(192),
1137 timeSinceStartedProcessing(195), and timeSinceCompleted(197) were added because this
1138 is the way IPP does these time attributes.

1139 On the other hand, is easier as you point out for the agent to use the JmTimeStampTC
1140 jobSubmissionToDeviceTime(191), jobStartedProcessingTime(194), and
1141 jobCompletedTime(196) and the NMS can do the work. Also having fewer time
1142 attributes to choose from does make the NMS's job easier. I'm in favor of removing the
1143 three timeSinceXXX attributes. An agent instrumenting an IPP system will have to do a
1144 little more computation.

1145 An agent instrumenting an IPP implementation will either have access to the time that the
1146 job's state change happened and can convert to JmTimeStampTC or only has the time-
1147 since-xxx value and will have to subtract that from the time of day and subtract
1148 sysUpTime from the result to return the jmTimeStampTC value.

1149 Closed: Yes, remove the three timeSinceXxxx attributes. They can be computed from
1150 other time attributes which are also easier for the agent to deal with since the value does
1151 not change with time.

1152

1153 ISSUE 108: Add IMPLIED to **jmJobIDEntry** INDEX statement on page 61?

1154 IMPLIED/IMPLICIT - See pg. 61 The note reads "an IMPLICIT statement is NOT
1155 provided in the following INDEX clause, since it was not an SMIV2 feature. Therefore,
1156 the extra ASN.1 tag SHALL be included in the varbind in the SNMP request and the
1157 response." First, we think the terminology is IMPLIED, not IMPLICIT.

1158 Closed: No. Change jmJobSubmissionID (index) to fixed length, so that a management
1159 application can omit trailing octets and get a partial match with SNMPv1 and SNMPv2.

1160

1161 ISSUE 109: Ok for agent to supply defaults for the job attributes depending on
1162 implementation?

1163 "Requested Attribute" defaults For requested attributes like copies, toner, quality etc.
1164 what if the requested value is not passed in? Should the agent use the device default?

1165 This is a difference between ISO DPA and IPP. In DPA, the server does populate the job
1166 object with the defaults. In IPP, the server doesn't, so that the defaults are only applied if
1167 neither the requester nor the document PDL supplies the attribute.

1168 Closed: Clarify that requested attributes include values defaulted by the server or printer
1169 where they have the same semantics as if the requester had supplied them (and so are
1170 requested by the system).

1171

1172 ISSUE 110: Break **jmAttributeTable** into two tables: **jmAttributeAsIntegerTable** and
1173 **jmAttributeAsOctetsTable** as suggested by David Perkins? Then an application will
1174 need about half as many Get Next operations in order to "harvest" all of the attributes,
1175 since there won't be any zero-length string values for attributes that don't have strings and
1176 -1 or 1 values representing 'other' for attributes that don't have integer values.

1177 Closed: Do not divide the table into two tables. There is some benefit to walking the
1178 table in pairs even though most attributes are either integer or octet string, and not both.

1179

1180 ISSUE 111 (restated): How does an application determine the coded character set for the
1181 objects and attributes that the agent generates (that cannot come from the job submitting
1182 client)? Raised by David Perkins.

1183 The following 3 objects and attributes are in question:

1184 **jmGeneralJobSetName** object

1185 **processingMessage** attribute

1186 **physicalDevice** (name value) attribute

1187 Closed: Add the **JmUTF8StringTC** for these three object/attributes. See separate
1188 issue111.doc .pdf.

1189

1190 ISSUE 112 (re-stated): How does a management application determine the coded
1191 character set for the per-job objects and attributes that are returned by the agent whether
1192 submitted by the job submitting client or defaulted by the agent when the job submitting
1193 client does not supply? Raised by David Perkins.

1194 Closed: Add the JmUTF8StringTC for these objects and add a jobCodedCharSet attribute
1195 to indicate the set being used.

1196

1197 ISSUE 113: The big concern [from the Area directors] is that from the user's perspective,
1198 jobs can be submitted via serial, parallel, or network connections, and the Job MIB is only
1199 going to know about the network connections. Raised by Chris Wellens.

1200 Closed: Add additional text to clarify that the device may be directly connected over a
1201 serial or parallel port or networked to the client and/or server.

1202

1203 ISSUE 114: If nested jobs are sent to the printer and all have a JobSubmissionID
1204 attached, what does the agent do? When a spooler receives a job, it can put a banner page
1205 on the job by wrapping it inside of a job. When this occurs, there can be separate
1206 JobSubmissionID's for each job. In fact, if the printer doesn't find a JobSubmissionID on
1207 the outer job, it will assign one. When the printer gets to the inner job, it will get the true
1208 JobSubmissionID that was attached to the client's job. Raised by Bob 'Pentecost on
1209 4/17/97 and re-raised by Harry Lewis on 7/16/97.

1210 Closed: Add a reference to the Appendix B where PJJ use of the job submission ID is
1211 included and indicate in Appendix B that the later occurrence of a job submission ID is the
1212 one that the agent uses. The Appendix B text is:

1213 NOTE - Some PJJ implementations wrap a banner page as a PJJ job around a job
1214 submitted by a client. In this case, there will be two job submission ids. The outer one
1215 being the one with the banner page and the inner one being the original user's job. The
1216 agent SHALL use the last received job submission ID for the jmJobSubmissionID index,
1217 so that the original user's job submission ID will be used, not the banner page job ID.

1218

1219 ISSUE 115: If the agent changes state to CANCEL as soon as it becomes aware of the
1220 cancel command (to satisfy the end user), there may still be a page or two in the pipeline
1221 that the accounting application would miss if it noticed the state change and performed it's
1222 data collection.

1223 So, we suggest using jobStateReasons in this case.

1224 CANCEL - processingToStopPoint
1225 which progresses to

1226 CANCEL - jobCanceledByUser

1227 The question is, since these jobStateReasons are not "mandatory", how do we
1228 communicate and agree on this recommendation? In other words, how do we achieve
1229 interoperability? Raised by Harry Lewis on 7/8/97.

1230 Closed: Move the **processingToStopPoint** reason from the **jobStateReasons2** attribute
1231 to the **jmJobStateReasons1** object immediately after abortedBySystem reason and
1232 renumber the ones following. Also recommend using processingToStopPoint reason with
1233 the aborted and canceled states. The new text is:

1234 **processingToStopPoint** **0x4000**
1235 The requester has issued an operation to cancel
1236 or interrupt the job or the server/device has
1237 aborted the job but the server/device is still
1238 performing some actions on the job until a
1239 specified stop point occurs or job
1240 termination/cleanup is completed.

1241
1242 This reason is recommended to be used in
1243 conjunction with the **canceled** or **aborted** job
1244 state to indicate that the server/device is
1245 still performing some actions on the job after
1246 the job leaves the **processing** state, so that
1247 some of the jobs resources consumed counters may
1248 still be incrementing while the job is in the
1249 **canceled** or **aborted** job states.

1250
1251

1252 ISSUE 116: Delete the 'other(1)' jmJobState value? I thought we had covered all
1253 possible states with the 7 primary plus the **JmJobStateReasons**. Why do we need other?
1254 Did we not accomplish what was claimed? Raised by Ron Bergman on 7/28/97.

1255 Closed: delete the 'other(1)' jmJobState value.

1256

1257 ISSUE 117: What is the difference between Toner Economy (**tonerEconomyRequested**
1258 and **tonerEconomyUsed**) and Toner Density (**tonerDensityRequested** and
1259 **tonerDensityUsed**)? (see page #51) They appear to be identical from the description (or
1260 very very close!) Raised by Ron Bergman on 7/28/97.

1261 Closed: Cut and paste error copied the explanation from one to the other. Change the
1262 tonerEconomy to say toner economy. Economy is an enum, density is a number from 1 to
1263 100. Several printers have both, so we need both.

1264

1265 ISSUE 118: Alignment of IPP and JMP. A job monitoring MIB agent providing access
1266 to an IPP system should be able to access the same data as is created by IPP. Lengths of
1267 text characters should be the same, not different (255 vs 63, respectively). Are there any

1268 other differences that would cause problems. See ietf-ipp.doc .pdf in protomap sub-
1269 directory. Raised by Paul Moore on 8/7/97.

1270 Closed: Leave the lengths as they are. Most real IPP implementations will not exceed the
1271 63 length in practice.

1272

1273 ISSUE 119: Is the new 32-bit IPP "job-identifier" now the same as the 32-bit
1274 jmJobIndex? Is the maximum range 31-bits: 1 to 2**31-1 in both? Raised by Paul Moore
1275 on 8/7/97.

1276 Closed: Clarify that the jmJobIndex is the same as the IPP "job-identifier", if IPP keeps
1277 the job-identifier. However, push back since the IPP meeting indicates that maybe this
1278 issue isn't settled. Also there was discussion that an agent that is providing access to a
1279 device that supports multiple job submission protocols, including IPP, may have a problem
1280 using the IPP "job-identifiers", unless the device also assigns the identifiers for the other
1281 job submission protocols from the same job-identifier number space.

1282 The following text has been incorporated into section 3.6:

1283 It is recommended that agents that are providing access to servers/devices that already
1284 allocate job-identifiers for jobs as integers use the same integer value for the **jmJobIndex**.
1285 Then the jobs will have the same job identifier value as the **jmJobIndex** value, so that
1286 users viewing jobs by management applications using this MIB and applications using
1287 other protocols will see the same job identifiers for the same jobs. Agents providing
1288 access to systems that contain jobs with a job identifier of 0 SHALL map the job identifier
1289 value 0 to a **jmJobIndex** value that is one higher than the highest job identifier value that
1290 any job can have on that system. Then only job 0 will have a different job-identifier value
1291 than the job's **jmJobIndex** value.

1292 NOTE - If a server or device accepts jobs using multiple job submission protocols, it may
1293 be difficult for the agent to meet the recommendation to use the job-identifier values that
1294 the server or device assigns as the **jmJobIndex** value, unless the server/device assigns
1295 job-identifiers for each of its job submission protocols from the same job-identifier number
1296 space.

1297

1298 ISSUE 120: The management app should be able to read an object(s) to get the number
1299 of consecutive rows in a "packed" table to put into a single Get, so as to reduce network
1300 traffic and decrease the time to get the data. What about falling off the end of the job and
1301 attribute tables? Raised by Paul Moore on 8/7/97 after having this trouble with the Printer
1302 MIB Alert table.

1303 Closed: The attribute table is by definition sparse, since the fourth index is present. For
1304 the other tables, since jobs may be canceled or aborted before completing, there can be
1305 holes in the job tables, even though we require that the **jmJobIndex** be assigned
1306 incrementally on job acceptance. So we can't add any objects that can help in addition to
1307 the **jmGeneralOldestActiveJobIndex** and **jmGeneralNewestActiveJobIndex** that we

1308 already have that say where to start and end. SNMP v2 solves the problem of reducing
1309 network traffic by using bulk get.

1310

1311 ISSUE 121: **jmJobKOctetsProcessed** can be a multiple of **jmJobKOctetsRequested**,
1312 for implementations that make multiple passes over the data. However, the same
1313 difference is not specified for
1314 **jmJobImpressionsComplete/jmJobImpressionsRequested** objects and for
1315 **pagesCompleted/pagesRequested** attributes.

1316 Closed: Make **jmJobImpressionsComplete/jmJobImpressionsRequested** objects and
1317 for **pagesCompleted/pagesRequested** attributes consistent with
1318 **jmJobKOctetsProcessed/jmJobKOctetsRequested**.