

1 Subj: Pros and Cons of a separate jmJobStateTable
 2 From: Tom Hastings, Harry Lewis, and Ron Bergman
 3 Date: 5/14/97
 4 File: sepstate.doc
 5

6 The biggest issue remaining in the Job Monitoring MIB is the duplication of information in the
 7 **jmJobStateTable** and the **jmAttributeTable**. Should we get rid of the duplication? And if so, do we
 8 delete the **jmJobStateTable** or the duplicated attributes in the **jmAttributeTable**? A second issue is
 9 whether the AssociatedValue object/attribute that provides a discriminant union of values based on the
 10 job's state should be kept in either table. This paper is intended to further the discussion about this topic.

11 **1. Summary of current overlap of jmJobStateTable and jmAttributeTable**

12 The overlap between the **jmJobStateTable** and **jmAttributeTable** in the current MIB specification is
 13 summarized by the following table:

14 **Table 1**

jmJobStateTable object	corresponding jmAttributeTable attribute	Mand atory?	static/dy namic?	AssociatedVal ue state
jmJobState	jobState(3)	yes	dynamic	-
jmJobStateKOctetsCompleted	jobKOctetsCompleted(50)	yes	dynamic	-
jmJobStateImpressionsCompleted	impressionsCompleted(55)	yes	dynamic	canceled
jmJobStateAssociatedValue	jobStateAssociatedValue(4)	yes	dynamic	-
	jobStartedBeingHeldTimeSta mp(73)	no	dynamic	held
	numberOfInterveningJobs(9)	yes	dynamic	pending
	jobKOctetsRequested(48)	yes	static	processing
	impressionsRequested(54) [currentCopy(??) proposed]	yes yes	static dynamic	printing printing
	deviceAlertCode(10)	yes	dynamic	needsAttention
	outputBinIndex(34)	yes	dynamic	completed

15

16 The **jmJobStateTable** is indexed by **jmJobSetIndex** and **jmJobIndex**.

17 The **jmAttributeTable** is indexed by **jmJobSetIndex**, **jmJobIndex**, **jmAttributeTypeIndex**, and
 18 **jmAttributeInstanceIndex**.

19 **2. Summary of the Issues**

20 The issues around the above objects/attributes are (in a logical order for consideration):

21 ISSUE 68 - Delete the **Job State Group/Table** all together, since all objects are also duplicated as
 22 attributes in the **jmAttributeTable**?

23

24 Sub-issues to Issue 68 are:

25 ISSUE 68a: If we keep the **jmJobStateTable**, should we delete the attributes out of the
26 **jmAttributeTable** that already appear as objects in the **jmJobStateTable**, namely **jmJobState(3)**,
27 **jobKOctetsCompleted(50)**, and **impressionsCompleted(55)**?

28

29 ISSUE 68b: If we keep the **jmJobStateTable**, should we move the *mandatory* associated attributes (1)
30 out of the **jmAttributeTable** that the **jmJobStateAssociatedValue** object provides a convenient copy and
31 (2) into the **jmJobStateTable** as objects? Then the **jmAttributeTable** would contain only conditionally
32 mandatory attributes and the **jmAttributeTable**, itself, would change from Mandatory to Conditionally
33 Mandatory.

34 In other words, move **numberOfInterveningJobs(9)**, **jobKOctetsRequested(48)**,
35 **impressionsRequested(54)**, [or the proposed **currentCopy(??)**], **deviceAlertCode(10)**, and
36 **outputBinIndex(34)** into the **jmJobStateTable** as *mandatory* objects:
37 **jmJobStateNumberOfInterveningJobs**, **jmJobStateKOctetsRequested**,
38 **jmJobStateImpressionsRequested** [or proposed **jmJobStateCurrentCopy**],
39 **jmJobStateDeviceAlertCode**, and **jmJobStateOutputBinIndex**. (Don't move the non-mandatory
40 **jobStartedBeingHeldTimeStamp(73)**).

41

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

43 Would it help if the mandatory attributes were first, so that Get Next would pick them up first when
44 getting the next conceptual row? Does making the attribute table easier to navigate using Get Next help
45 with the decision to Issue 68 and 68b?

46

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

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

56 The current logical grouping are:

57	Job State attributes	10
58	Job Identification attributes	19
59	Job Parameter attributes	7
60	Image Quality attributes (requested and used)	6
61	Job Progress attributes (requested and consumed)	7
62	Impression attributes (requested and consumed)	6
63	Page attributes (requested and consumed)	3
64	Sheet attributes (requested and consumed)	3
65	Resource attributes (requested and consumed)	7
66	Time attributes (set by server or device)	9

67 OK to assign Job State and Job Identification in steps of 30 and the rest in steps of 20?

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

69

70 Issue 78 - Should the "multiplexor" (discriminant union?) **jobStateAssociatedValue(4)** attribute be
71 removed from the Job Attribute Table and the equivalent **jmJobStateAssociatedValue** object be removed
72 from the Job State table?

73 The associated values are also available as attributes in the attribute table. The application has to either
74 (1) request all 7 associated attributes or (2) first request the **jobState(3)** attribute and then request the 1
75 pertinent attribute. Since all 7 will easily fit in a PDU (minimum of 500 octets or so on all systems) and
76 each request takes about 20 octets, so you can get about 20 (5*4) attributes into a single PDU.

77

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

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

87

88 ISSUE 68c: If we keep the **jmJobStateAssociatedValue** object, we could just change the attributes listed
89 in ISSUE 68b from *mandatory* to *optional* and keep them only in the **jmAttributeTable**. The
90 **jmJobStateAssociatedValue** object would remain in the **jmJobStateTable** to provide access to these
91 attributes mandatorially.

92

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

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

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

98 **jmJobName**
99 **jmJobIdName**
100 **jmJobIdNumber**
101 **jmJobServiceType**
102 **jmJobOwner**
103 **jmJobDeviceNameOrQueueRequested**
104 **jmJobCurrentState**
105 **jmJobStateReasons**

106

- 107 1. **jmJobIdName** and **jmJobIdNumber** have been replaced by **jmJobSubmissionIDIndex** which is
108 Mandatory.
- 109 2. **jmJobServiceType** need not be Mandatory.
- 110 3. Also **jmJobDeviceNameOrQueueRequested** has been made into two separate attributes:
111 **deviceNameRequested** and **queueNameRequested**, so we'd have to make either one of them
112 mandatory.
- 113 4. **jmJobCurrentState** is now **jobState** and is Mandatory
- 114 5. **jmJobStateReasons** became four attributes: **jobStateReasons1**, **jobStateReasons2**,
115 **jobStateReasons3**, and **jobStateReasons4**. None of them need to be Mandatory.

116

117 So should **jobName**, **jobOwner**, and one of **deviceNameRequested** or **queueNameRequested** be made
118 Mandatory?
119

120 ISSUE 76a - If yes, then should they be put into the **jmJobStateTable**, instead of the **jmAttribute** table, if
121 Issue 68b concluded that the **jmAttributeTable** should have no mandatory attributes.
122 ISSUE 70 - Add some simple general device alert TC, instead of using the Printer MIB **Alert** Codes.
123 The **PrtAlertCodeTC** generic values are *not* much good to an end user without knowing which subunit.
124 For example, **SubUnitEmpty** isn't very informative by itself. If an implementation also has the Printer
125 MIB, then a lot more information is available, so a copy of the Printer Alert isn't very useful. If the
126 implementation doesn't have the Printer MIB, then the Printer Alert codes aren't informative enough.
127 Even worse, the deviceAlertCode(10) is Mandatory, which can't be implemented, if there isn't a Printer
128 MIB also implemented.
129

130 Issue 73 - Is there a problem with **outputBinIndex** being made mandatory?
131 If **outputBinIndex** is made mandatory, but an implementation doesn't have the Printer MIB, the agent has
132 to put 0 as the value. Should we add one more attribute: **outputBinNumber**, which is just a number, not
133 an index into the Printer MIB? If we do, which should be mandatory? Just one more reason to get rid of
134 the **jmJobStateTable**, which is forcing us to pick a particular outputBin implementation and make it
135 mandatory. If we got rid of the **jmJobStateTable**, we could forget about making any of the 3
136 outputBinName, outputBinNumber, or outputBinIndex attribute mandatory.
137 Closed: Don't add **outputBinNumber**. Just add **other**(-1), **unknown**(-2), and **multi**(-3) values and keep
138 **outputBinIndex** as mandatory. This does also means that **jmAttributeValueAsInteger** needs a lower
139 bound of -3, not -2.
140

141 ISSUE 87 - When shall an agent make the mandatory attributes appear in the **jmAttributeTable**?
142 Shall an agent materialize all mandatory attributes when the job is submitted, so that a requester can
143 access them all with multiple explicit Gets in a single PDU, without fear of a missing object aborting the
144 PDU? If the mandatory attributes are represented as objects in the **jmJobStateTable**, then it is clear from
145 SNMP rules that the agent shall materialize at least an empty value for each mandatory object (attribute).
146

147 ISSUE 83 - Can some attributes be deleted before the **jmGeneralAttributePersistence** expires?
148 Harry Lewis' 5/2 e-mail suggested that some of the attributes, such as "**numberOfInterveningJobs**(9)"
149 don't even need to persist the shorter time specified by **jmGeneralAttributePersistence**.
150 However, if we move the mandatory attributes to the **jmJobStateTable** and make them objects, then they
151 shall persist for the longer persistence specified by **jmGeneralJobPersistence**.
152

153 See the rest of the issues list for the issues that do *not* relate to the overlap objects/attributes between the
154 **jmJobStateTable** and the **jmAttributeTable**.

155 **3. Accessing the jmJobStateTable and the jmAttributeTable**

156 In order to understand the pros and cons, it seems necessary to understand how an application would use
157 Get and Get Next to get information from these two tables. We need to consider the three basic types of
158 applications: (1) a job monitoring application that is monitoring a particular job, (2) a job monitoring
159 application that is monitoring all jobs on a device or server, and (3) a job accounting or utilization

160 program. The first two kinds of applications are interested in active jobs and the third is interested in
161 inactive jobs (canceled, or completed).

162 **3.1 OID assignments to the objects**

163 In order to construct complete examples, it is helpful to use the actual OIDs that will be assigned to the
164 objects and attributes in the MIB:

```
165  
166 >      jobmonMIB  
167 >      jobmonMIB.1  jobmonMIBObjects  
168 >      jobmonMIB.1.1  jmGeneral  
169      jobmonMIB.1.1.1  jmGeneralTable  
170      jobmonMIB.1.1.1.1  jmGeneralEntry  
171      jobmonMIB.1.1.1.1.1  jmGeneralNumberOfActiveJobs  
172      jobmonMIB.1.1.1.1.2  jmGeneralOldestActiveJobIndex  
173      jobmonMIB.1.1.1.1.3  jmGeneralNewestActiveJobIndex  
174      jobmonMIB.1.1.1.1.4  jmGeneralJobPersistence  
175      jobmonMIB.1.1.1.1.5  jmGeneralAttributePersistence  
176      jobmonMIB.1.1.1.1.6  jmGeneralJobSetName  
177  
178 >      jobmonMIB.1.2  jmJobID  
179      jobmonMIB.1.1.1  jmJobIDTable  
180      jobmonMIB.1.1.1.1  jmJobIDEntry  
181      jobmonMIB.1.1.1.1.1  jmJobSubmissionIDIndex  
182      jobmonMIB.1.1.1.1.2  jmJobSetIndex  
183      jobmonMIB.1.1.1.1.3  jmJobIndex  
184  
185 >      jobmonMIB.1.3  jmJobStateG  
186      jobmonMIB.1.1.1  jmJobStateTable  
187      jobmonMIB.1.1.1.1  jmJobStateEntry  
188      jobmonMIB.1.1.1.1.1  jmJobState  
189      jobmonMIB.1.1.1.1.2  jmJobStateKOctetsCompleted  
190      jobmonMIB.1.1.1.1.3  jmJobStateImpressionsCompleted  
191      jobmonMIB.1.1.1.1.4  jmJobStateAssociatedValue  
192  
193 >      jobmonMIB.1.4  jmAttribute  
194      jobmonMIB.1.1.1  jmAttributeTable  
195      jobmonMIB.1.1.1.1  jmAttributeEntry  
196      jobmonMIB.1.1.1.1.1  jmAttributeTypeIndex  
197      jobmonMIB.1.1.1.1.2  jmAttributeInstanceIndex  
198      jobmonMIB.1.1.1.1.3  jmAttributeValueAsInteger  
199      jobmonMIB.1.1.1.1.4  jmAttributeValueAsOctets  
200  
201 >      jobmonMIB.2      jobmonMIBConformance  
202      jobmonMIB.2.1      jobmonMIBCompliance  
203      jobmonMIB.2.2      jmMIBGroups  
204      jobmonMIB.2.2.1      jmGeneralGroup  
205      jobmonMIB.2.2.2      jmJobIDGroup  
206      jobmonMIB.2.2.3      jmJobStateGroup  
207      jobmonMIB.2.2.4      jmAttributeGroup  
208
```

209 **3.2 Tables and the Get operation**

210 Recall that the OIDs for table entries consist of the OID for the entry (column) in the table, *followed by*
211 the index(es) to that entry. To get the job state object in the **jmJobStateTable** for the job with a
212 **jmJobIndex** of 1000 in job set 1, the requester must pass the following OID as a Get input parameter:

213 **jmJobState.1.1000**, i.e., jobmonMIB.1.1.1.1.1.1.1000.

214 To get the corresponding from the **jmAttributeTable**, which is the **jobState(3)** attribute, the requester
215 must pass the following OID as a Get input parameter:

216 **jmAttributeValueAsInteger.1.1000.3.1**, i.e., jobmonMIB.1.1.1.1.2.1.1000.3.1.

217 Thus an application can always get the corresponding attribute from the **jmAttributeTable** with an OID
218 that is only *two* octets longer than is required on a Get for the corresponding object **jmJobStateTable**.

219 An application can get multiple objects from the **jmJobStateTable** and can get multiple attributes from
220 the **jmAttributeTable** by supplying multiple Get operations in a single PDU.

221 If there is no such object, the Get operation returns an error (and does *not* perform any further Get
222 operations in the submitted PDU, correct?)

223 **3.3 Tables and the GetNext operation**

224 The SNMP GetNext operation returns the value of the object specified by the *next lexically higher* OID
225 from the one supplied as an input parameter. GetNext also returns that next lexically higher OID itself, so
226 that the application can pass it back as an input parameter to a subsequent GetNext in order to get the next
227 object. If there are no lexically higher objects, GetNext returns an error.

228 The OID input parameter does not need to be "fully specified". Trailing OID arcs can be omitted and they
229 shall behave as if the requester supplied 0 for those arcs.

230 For a single index table, Get Next can be used to get the "next conceptual row" in the table. GetNext must
231 be used when the agent scatters rows in a table, i.e., the table is a "sparse" table. MIB specifications can
232 specify that tables shall not be sparse. The **jmJobStateTable** is specified such that agents shall enter
233 conceptual rows such that jmJobIndex is monotonically increasing, until wrap occurs. However, because
234 jobs may be canceled, a canceled job may be removed from the middle of the table (after persisting for the
235 **jmGeneralJobPersistence** time), thereby making the **jmJobStateTable** have an empty row, i.e., be a
236 "little bit sparse". Also a system that processes jobs out of order may result in some empty rows in
237 between rows that are awaiting the **jmGeneralJobPersistence** time to expire.

238 An application can get the state of the next job after job 1000 in the **jmJobStateTable** by passing in the
239 (same) OID:

240 **jmJobState.1.1000**, i.e., jobmonMIB.1.1.1.1.1.1.1000

241 If job 1001 had been canceled, say, and the agent removed it subsequently, the agent might return the
242 state of job 1002 and the OID:

243 **jmJobState.1.1002**, i.e., jobmonMIB.1.1.1.1.1.1.1002

244 The application could copy the returned OID to the input parameter of a subsequent GetNext and the get
245 the state of the next job after 1002, and so forth.

246 If the application wanted to get more than just one object in the next conceptual row, the application could
247 supply several GetNext operations in the same PDU. So to get the **jmJobState**,
248 **jmJobStateKOctetsCompleted**, **jmJobStateImpressionsCompleted**, and **jmJobStateAssociatedValue**
249 objects from the **jmJobStateTable** for the next job after 1002, the application could pass in the following
250 four OIDs in four Get Next operations in the same PDU:

251 **jmJobState.1.1002**, i.e., jobmonMIB.1.1.1.1.1.1.1002

252 **jmJobStateKOctetsCompleted.1.1002**, i.e., jobmonMIB.1.1.1.1.2.1.1002

253 **jmJobStateImpressionsCompleted.1.1002**, i.e., jobmonMIB.1.1.1.1.3.1.1002

254 **jmJobStateAssociatedValue.1.1002**, i.e., jobmonMIB.1.1.1.1.4.1.1002

255 The agent shall *return* the next OID in each GetNextResponse for each of these inputs, which would be
256 the corresponding column in the next row, say, job 1003:

257 jmJobState.1.1003, i.e., jobmonMIB.1.1.1.1.1.1.1003
258 jmJobState**KOctetsCompleted**.1.1003, i.e., jobmonMIB.1.1.1.1.2.1.1003
259 jmJobState**ImpressionsCompleted**.1.1003, i.e., jobmonMIB.1.1.1.1.3.1.1003
260 jmJobState**AssociatedValue**.1.1003, i.e., jobmonMIB.1.1.1.1.4.1.1003

261 NOTE - An application could *not* perform the above by using a individual repeated GetNext operation
262 copying each result to the single input parameter, because GetNext increments the least significant part of
263 the OID first. Thus, each individual GetNext would get the *same* column in the next row, *not* step across
264 the columns in the same row.

265 In order to perform the equivalent of the above example in the **jmAttributeTable**, i.e., get the
266 **jobState(3)**, **jobKOctetsCompleted(50)**, **impressionsCompleted(55)**, and the
267 **jobStateAssociatedValue(4)** attributes in the **jmAttributeTable** for the next job after job with
268 **jmJobIndex** 1002, the application must first determine the next valid jmJobIndex, which cannot be done
269 by simply passing in the following OID in GetNext operation:

270 **jmAttributeValueAsInteger.1.1002.3.1**, i.e., jobmonMIB.1.1.1.1.2.1.1002.3.1

271 because the next lexically higher OID might be:

272 **jmAttributeValueAsInteger.1.1002.9.1**, i.e., jobmonMIB.1.1.1.1.2.1.1002.9.1

273 which is the **numberOfInterveningJobs(9)** attribute.

274 Instead, the application must determine what the next **jmJobIndex** value either by doing a GetNext on the
275 **jmJobStateTable** or by passing in the "incremented" partial OID that the application has incremented
276 "by hand" and shortened by removing the trailing OID arcs after the **jmJobIndex** arc:

277 **jmAttributeValueAsInteger.1.1003**, i.e., jobmonMIB.1.1.1.1.2.1.1003

278 which will return the first attribute in the next job. The **jmJobIndex** arc the comes back in that
279 GetNextResponse is the next **jmJobIndex** in the **jmJobAttributeTable**.

280 **3.4 Monitoring a single specific job**

281 When a user submits a job, the client could fire up a monitoring application that monitors the job just
282 submitted. The monitoring application needs to determine the job's **jmJobIndex** by one of several
283 methods, depending on the implementation and the configuration:

284 (1) is told the **jmJobIndex** of the job to be monitored because the server returned the job-identifier which
285 the application knows the map to **jmJobIndex** value,

286 (2) can determine the **jmJobIndex** by doing a Get supplying the OID for the **jmJobSubmissionIDIndex**
287 to the **jmJobIDTable** as follows. Suppose that the job submission id generated by the client is:
288 "12345678nnnnnnnnnn"

289 **jmJobIndex.1."12345678nnnnnnnnnn"**, i.e., jobmonMIB.1.1.1.1.3.1."12345678nnnnnnnnnn"

290 which returns the **jmJobIndex** for the job, or

291 (3) can scan the **jmAttributeTable** looking for attributes that match, such as **jobOwner(15)**,
292 **jobName(13)**, etc., though such a scan requires two probes: first to find the next **jmJobIndex** either
293 from the **jmAttributeTable** or more straightforwardly from the **jmJobStateTable**.

294 Give the jmJobIndex for the single job being monitored, the application can use direct Get operations to
295 get any objects from the **jmJobStateTable** or attributes from the **jmAttributeTable** as shown above.

296 **3.5 Monitoring all active jobs on a server or device**

297 An operator might run an application that monitors all *active* jobs on a server or device. Such an
298 application polls at some frequent enough interval to show changes, but not too frequently to bog down

299 the network or server/device. An end-user might fire up an application to monitor all jobs on a server or
300 printer, especially when searching for a "least busy printer". Here the time to find the jobs and get their
301 attributes needs to be relatively short, or the user will not want to fire up such an application.

302 With either scenario, the application has to determine the oldest active job with a Get specifying the
303 `jmJobSet=1`, and it may as well get the number of active jobs and the newest active job index in the same
304 PDU:

305 `jmGeneralNumberOfActiveJobs.1`, i.e., `jobmonMIB.1.1.1.1.1.1`

306 `jmGeneralOldestActiveJobIndex.1`, i.e., `jobmonMIB.1.1.1.1.2.1`

307 `jmGeneralNewestActiveJobIndex.1`, i.e., `jobmonMIB.1.1.1.1.3.1`

308 If the value of `jmGeneralOldestActiveJobIndex` is 0, there are no active jobs and the application updates
309 the display to show no jobs. Say the value of `jmGeneralOldestActiveJobIndex` is 2000.

310 Then the application requests, say, the four (column) objects in the `jmJobStateTable` with four Gets in a
311 single PDU as shown above for job 1000. Then the application submits four GetNext operations in the
312 same PDU for each of the four objects in the `jmJobStateTable` as described above for job 1002.

313 Finally, if there are some additional attributes that the application wishes to get, such as
314 `jobStateReasons1(5)` and `jobName(13)`, the application submits several Gets in a single PDU of the form:

315 `jmAttributeValueAsInteger.1.2000.5.1`, i.e., `jobmonMIB.1.1.1.1.2.1.2000.5.1`

316 `jmAttributeValueAsOctets.1.2000.13.1`, i.e., `jobmonMIB.1.1.1.1.2.1.2000.13.1`

317 **3.6 Accounting/Utilization application gathering data on** 318 **completed/canceled jobs**

319 The accounting or utilization application remembers the lowest `jmJobIndex` from last time. The
320 application can either get all `jmJobStateTable` objects and all `jmAttributeTable` attributes, or may get
321 only certain selected attributes.

322 To get all attributes, that application starts with the lowest `jmJobIndex` that it had on the previous poll
323 cycle and supplies a number of GetNext operations in a single PDU.

324 To get only selected attributes the application must first determine the next `jmJobIndex` by using GetNext
325 on the `jmJobStateTable`. The application may as well get the other objects from the `jmJobState` with a
326 bunch of GetNext operations in the same PDU. If the job is active, that data is probably thrown away, and
327 the application steps on to the next job. If the job is inactive (canceled or completed), then the application
328 would specify multiple Get operations in a single PDU, one for each attribute that it wished.

329 **4. Conclusions**

330 The `jmJobStateTable` is very useful because its lowest order index is `jmJobIndex`, so that any number of
331 selected objects can be obtained with multiple GetNext operations in a single PDU for the *next* job,
332 skipping over jobs that have been removed from the table. A subsequent PDU can contain multiple Get
333 operations for any attributes desired using the returned `jmJobIndex` value.

334 If the mandatory attributes are all put into the `jmJobStateTable` as objects, and not in the
335 `jmAttributeTable` as attributes, it is clear by SNMP rules that all of the mandatory objects shall be
336 instantiated at the same time when the new job row is put into the `jmJobStateTable`. Also the persistence
337 time is clearly separated by which table the information is contained.

338 The `jmAttributeTable` only contains conditionally mandatory attributes, no mandatory attributes, so that
339 the `jmAttributeTable` itself can be conditionally mandatory, thereby allowing a very small
340 implementation to only implement the `jmJobStateTable` and not the `jmAttributeTable`.