PWG Job MIB Explanations

07 April 1997
Revision 1.00

Harry Lewis

Phone: 303.924.5337

harryl@us.ibm.com

File: jmdoc.doc
Printed: 04/07/97 4:24 PM

Table Of Contents

REVISION HISTORY

revision 1.00 - 04 April 1997

- Derived from IBM prototyping of PWG MIB and PWG Austin April meeting discussion.

Job MIB Specification

1. Overview

The IETF standard Printer MIB is now in use by several printer vendors as a means of facilitating printer management via SNMP. Many printers contain SNMP agents which respond to software ranging from hi-end NMS applications like Netview and OpenView to printer specific management applications and self-configuring print drivers such as those provided by individual printer companies.

The next step in SNMP management of print is to address the print job itself. This explanation represents the culmination of over 1 year of work by the PWG on Job MIB development and is intended to aid in the design and implementation of solutions based on the PWG Job Monitoring MIB.

1.1 Job MIB Roles

The Job MIB is intended to serve the needs of various applications that provide information based on a specific Role.

1.1.1 User

♦ Timely notification that a job has completed
♦ The current state of a user's job
♦ Job progress (pages complete)
♦ What output bin the completed job is in
♦ Error information
♦ Help in selecting the least busy printer.

1.1.2 Administrator

♦♦ State of all the jobs in the printer.
♦♦ Which users submitted each job.
♦♦ Resources used by each job.
♦♦ Some idea of how long each job will take.

1.1.3 Accounting

– A record of resources used on a particular print job

1.1.4 Capacity and Supplies planning

♦♦ How busy are printers.
♦♦ What time of day are they used.
♦♦ How long are users waiting for their jobs to print.
♦♦ Supplies usage rates

1.2 Job MIB Model

The Job MIB will provide for printers that can contain more than one job at a time, but still be usable for low end printers that can only process a single job at a time.

As with the Printer MIB (RFC1759), a subset of the ISO Abstract Syntax Notation One (ISO ASN.1) will be used to define the Job MIB syntax which will consist of a tree architecture defined to organize all available information. Each piece of information in the tree is a labeled node. Each labeled node contains the following:

1. An object identifier
2. A short text description

The object identifier (OID) is a series of integers separated by periods that identifies the node and denotes the exact traversal of the ASN.1 tree. The text description describes the labeled node. A labeled node can have subtrees containing other labeled nodes. If the labeled node has no subtrees (a leaf node), it contains a value and is known as an object.

The Job MIB will initially reside in the Experimental tree, starting at 54.105 and will contain information in 4 related groups:

General - Contains policy related to persistence of entries and queuing

Job Information - Contains information used to uniquely identify each job yet rapidly index it's information.

State - Tracks job state, progress and order of execution

Attribute - Indicates resources required and used for each job

An Agent module within the printer will manage the Job MIB objects. The Agent will respond to SNMP requests from network print job management applications that want to read data from Job MIB objects. Data in the MIB can be set in the following ways:

- Objects that are defined at build time and will never change. For example, jmGeneralJobPersistence - the time in seconds that the Job State and Job ID table entries will remain.

- •• Objects provided by dynamic modules within the Printer OS. For example, when a print job starts, a sheet is pushed to the exit tray or a print job ends, the Printer OS will notify the Agent of these job related events.

2. Reliability

Reliability of the Job monitoring solution will be obtained as the result of overall system behavior, as opposed to being guaranteed by one particular component. The Job MIB agent, in the printer, will track every job from the time it is received by the printer until it has completed printing or has been canceled. Agent assigned job ID's (referred to as jmJobIndex) will be increasing and sequential in nature, although sparse entries may occur as a result of jobs being canceled, power interruptions or other management anomalies. Job Complete transitions shall not be missed by the agent, except for possible power interruptions. If any jobs are lost due to a power cycle, the job monitoring application will deduce this when it detects a power on event with print jobs outstanding but finds no record of them in the job ID or job State tables.

Since job monitoring via SNMP relies on polling, overall data integrity will be enhanced by table entry latency. Life of a table entry is a function of the rate of printing and table size. Based on a hypothetical scenario of a 18ppm printer printing 1 page jobs, each job containing 10 managed resources, the following *minimum* table sizes are recommended:

Table	Size in Rows	Appx. Storage
Job ID	60	2.5K Bytes
Job State	60	1K Bytes
Job Attributes	300	14K Bytes

These recommended minimums will result in appx 3 minutes for the job monitor to obtain the job ID, allows 3 minutes of non-polling (after the job ID has been obtained) to determine job completion and provides 100 seconds for the accounting application to collect data before a job is removed from the attribute table. Note that persistence of job ID and State information should always exceed that of accounting entries for a particular job. This allows the accounting application to assure the attribute information pertains to a completed job.

These table storage calculations are approximate and do not address overall storage required for the Job MIB agent.

3. Use of NVRAM

Total Job MIB NVRAM storage stands at 4 bytes in the Controller. The 4 bytes of NVRAM are used to insure monotonically increasing values for jmJobIndex.

4. Print Submission Languages

4.1 PJL and Postscript

Many printer controllers, today, have interpreter architectures rooted in some dialect of HP's PJL or Adobe's Postscript language. Whichever print submission datastream is used, for robust print job monitoring, objects described in the Job MIB should be supported.

It would be ideal if these two predominant print management datastreams were extended to accommodate any new Job monitoring objects that need to be associated with the job at submission time. Proprietary extensions to these datastreams may be required as a lesser alternative.

4.2 Internet Printing Protocol

Several forms of internet print submission are under development. One, called IPP (Internet Printing Protocol) by the IETF Printer Working Group, may require future protocol updates to the printer. IPP could may provide it's own job monitoring features or could serve as a standard, platform independent, print job submission protocol for jobs that will be monitored using SNMP and the Job MIB. IPP and the Job MIB may coexist and should strive for compatibility.

4.3 "Raw" Print jobs

Some jobs will reach the printer via submission protocols which are not capable of assigning a Job Submission ID. In this case, the agent will assign a Job Submission ID with format 0 but management will be limited to accounting applications and queue views, not specific job state determination.

5. Print Job MIB Objects

The Print Job monitoring MIB contains 4 groups of objects. These are the General group, Job ID Table, Job State Table and Job Attribute Group.

The following sections define each of the groups in the Print Job monitoring MIB. The general format of the definition is to list the object ID and description, the storage requirements for the object, and whether or not the object may be directly modified by the SNMP management application. All objects in the PWG Job monitoring MIB are wither read-only ("r/o"), or inaccessible ("na"), i.e. used internally by the NIC or Controller operating system.

Objects that are defined as Counter32, Integer32 and Integer are all 32 bit integers. The objects defined as Integer are Enumerations, defined by the relevant standards, or by the Job monitoring MIB itself.

5.1 Job General Group

The Job **General Printer** Group reflects local policy pertaining to how long jobs are kept in the Job ID Table, the Job State Table and how frequently job events will be reported There is only one Job General Printer Group per printer.

1 jmGeneral 16 bytes per Job Set

1 jmJobSetIndex Integer32 r/o
A running index of Job Set instances supported - hardcoded to (1) for printer implementations. This object is included for compatibility with potential server implementation managing jobs for more than 1 printer.

2 jmGeneralJobPersistence Integer32 r/o
The minimum time, in seconds, that an entry will remain in the Job ID and Job State tables. This value should be equal to or greater than the value for jmGeneralAttributePersistence.

3 jmGeneralAttributePersistence Integer32 r/o
The minimum time, in seconds, that an entry will remain in the Attribute table.

4 jmGeneralLowestActiveIndex Integer32 r/o
This value indicates the dividing line between jobs that have been completed and/or canceled vs. jobs that are considered in the active path for printing..

5.2 Job ID Group

The Job ID group serves as a means to get the value of jmJobIndex if, for some reason, or on some platforms, this cannot be obtained by any other means. This could be necessary on platforms such as Netware which does not currently facilitate an integrated client server implementation. Since other identifying parameters (Sever name, Queue name etc.), may appear in the attribute table as resource types, the Job ID table is not intended for use in correlating these characteristics for accounting purposes - only for client job tracking from start to finish. The unique jmJobSubmissionID will not appear in the attribute table.

Every job will have an entry in the JobID table.

2 jmJobIDGroup 40 bytes per Job

1 jmJobIDTable na

A table of print jobs that the printer is engaged in or aware of.

1 jmJobIDEntry **na**

1 ***jmJobSubmissionID*** ***String31*** ***r/o***

An index to this table. An ASCII (human readable) text string, the contents of which are specified by the SubmissionID format byte. Typical contents are 8 digit host generated random number followed by MAC address or last 22 bytes of the job Name. For any job which does not contain a Job Submission ID the printer will assign one using Submission ID Format 0.

2 JobSetIndex **Integer32** **r/o**

Hardcoded to (1) for printers.

3 jmJobIndex **Integer32** **r/o**

A sequential index assigned to each entry in the Job ID Table. This value is used to index the Job State and Attribute tables.

5.2.1 Job Submission ID formats

Format (Byte 0)	Part-1	Part-2
0	Bytes 1-30 = Printer specific ID (intended to be ASCII representation of jmJobIndex)	n/a
1	Bytes 1-8 = 8 digit number (intended to be random)	Bytes 9-30 = Last 21 bytes of Job Name
2	Bytes 1-8 = 8 digit number (intended to be sequential)	Bytes 9-20 = MAC address
3	Bytes 1-8 - 8 digit number (intended to be sequential)	Bytes 9-30 = Last 21 bytes of Printer URL

The notion behind the Job Submission ID is that of uniqueness among a set of print job submission clients with integrated print job monitoring. The Job Submission ID is made up of a pair of identifiers, typically, a number (represented as ASCII) and a string. Since the Job Submission ID only needs to be unique among the limited set of clients for the limited duration while the print jobs from these clients are resident in the printer, an 8 digit randomly generated number is probably sufficient in itself. We have further guaranteed uniqueness, however, by combining the random number with most, if not all, of the job name. Thus, if Client A submits the same job twice, it is very unlikely to assign the same number, if Client B happens to assign the same number, it is very unlikely that it will be associated with the same job name.

When part-2 of the Job Submission ID is, itself, guaranteed to be unique (as in the MAC address or URL of the client), then a sequential number should suffice (and may be easier for the client to manage).

To give the client options, we have registered the 3 methods above. Other methods may be registered as needed. For example, one method could contain Server and Queue name, together, combined with a job number.

5.3 Job State Table

The Job **State** table shows the current state and relevant metrics of each job, from start to finish in the printer, and also, for some duration post completion. The size of this table is adjusted for each printer implementation according to processing speed, print speed and ability to queue jobs in order to provide "buffer" for the job monitoring applications to assure they can obtain knowledge of job completion.

3 jmStateGroup 20 bytes per Job

1 jmStateTable na
A table of print jobs and their current state.

1 jmStateEntry na

1 jmJobSetIndex Integer32 r/o
An index to this table. Hardcoded to (1) for printers.

2 jmJobIndex Integer32 r/o
An index to this table. A unique value used by the printer to identify each job. The index can be obtained from the jmJobID Table.

3 jmJobState Integer r/o
Current state of the job (printing, complete etc.). See "Job States and Associated Values" for a table of enumerations.

4 jmJobKOctetsCompleted Integer r/o
Amount of job currently processed - rounded up to the next higher K.
This value would be compared to the baseline to gauge job progress.
If total Octets is known by the printer, it will appear as a resource in the attribute table.

5 jmJobImpressionsCompleted Integer r/o

Number of impressions currently “stacked” for this job. Value for this object may increment in multiples of 2, 3, 4, 6, 8 etc. if the job is printing in duplex and/or n-up.

6 **jmStateValue** **Integer** **r/o**
An auxiliary value of type specific to each job state.

5.3.1 Job States Transition Diagram

The following diagram shows all possible job state transitions.

Job State Diagram

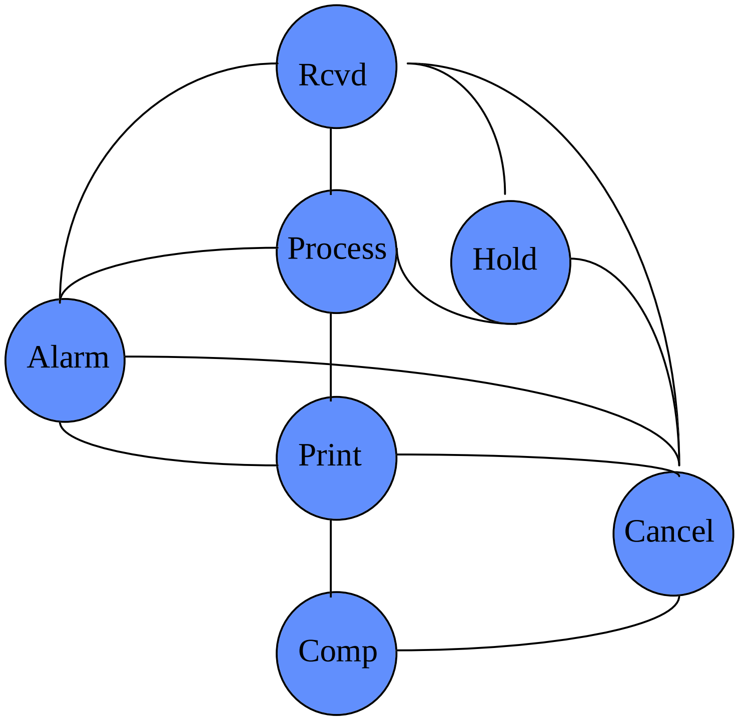

5.3.2 Job States Transition Table

The following shows job state transitions in table form.

	Processing	Printing	Completed	Alarm	Canceled	Held
Received	Yes			Yes	Yes	Yes
Processing		Yes		Yes	Yes	
Printing			Yes	Yes	Yes	

Completed						
Alarm	Yes	Yes			Yes	
Canceled			Yes			
Held	Yes				Yes	

5.3.3 Job States and Associated Values

The meaning of jmStateValue will vary depending on the Job State.

Job State Enumeration

Job State	Symbolic Name	Description	Associated Value
1	other		
2	unknown		
3	pending	Job is buffered or queued in the printer	Queue Position
4	processing	Job is being “ripped” and sent to the imaging mechanism	Total Octets
5	printing	Job is printing	TotalImpressions
6	held		TimeTicks
7	canceled		ImpressionsCompleted
8	alarm	An intervention condition that prevents printing	Alarm Type •• RFC1759 alerts
9	completed	Last page of job has been “stacked”	Bin number •• -1 Unknown •• -2 Multi •• prtOutput Index
10	retained		TimeTicks

5.4 Job Attribute Table

The resource attribute table is used by accounting and supplies tracking applications to gather information about total resource usage once the jobs have been completed. The accounting application must monitor the Job State table to determine when a job has completed. The accounting application can determine the jmJobIndex of completed jobs from the Job State table before collecting information from the Job Attribute table.

4 jmAttributeGroup 48 bytes per Resource

1 jmAttributeTable na

A table of accounting information indexed by job and resource.

1 jmAttributeEntry na

1 jmJobSetIndex Integer32 r/o

An index into this table. Hardcoded to (1) for printers.

2 jmJobIndex Integer32 r/o

An index into this table. Identifies the print job this attribute pertains to.

3 jmAttrubuteType Integer r/o

An index into this table. A category of attribute being either requested or consumed. There can be more than one attribute of the same general type or category used by each job. There is a separate row in the Attribute table for each new attribute type.

4 jmAttributeIndex Integer32 r/o

An index into this table. A job can contain more than one instance of the same attribute type. An example would be a job that uses 2 PDL's or several types of media. This index distinguishes different attributes of the same type.

5 jmAttributeValueAsInteger Integer32 r/o

The quantity of this attribute, either requested or consumed, in Attribute Units. If not used, the agent shall return a (-1) signifying other as in the Printer MIB.

6 jmAttributeValueAsOctet String31 r/o

The quantity of this attribute, either requested or consumed, as a text string.

When jm**AttributeValueAsOctet** isn't used, the agent shall return a zero length string.

5.4.1 Attribute Type Enumeration

Value	Symbolic Name	Value as Integer	Value as Text	PJL ?
1	other			
2	unknown			
3	fileName		File Name of Job	
4	Reserved			
5	jobAccountName		May be encrypted	
6	Reserved			
7	Reserved			
8	jobSourceChannelIndex	Index of row in Printer MIB		
9	outputBin	Index of row in Printer MIB		

10	sides	1 or 2		
11	documentFormat *	prtLangTC		
12	physicalDeviceIndex	hrDeviceIndex		
13	physicalDeviceName		Printer Name	
14	jobCopiesRequested	-2 Unknown		
15	jobCopiesCompleted	-2 Unknown		
16	Reserved			
17	Reserved			
18	jobTotalKOctets	Rounded Up		
19	jobKOctetsCompleted	Rounded Up		
20	Reserved			
21	Reserved			
22	Reserved			
23	Reserved			
24	Reserved			
25	Reserved			
26	pagesRequested	Logical Pages		
27	pagesCompleted	Logical Pages		
28	Reserved			
29	sheetsRequested	x		
30	sheetsCompleted	x		
31	Reserved			
32	mediumRequested *		Name of Medium	
33	mediaConsumed *	Sheets	Name of Medium	
34	colorantRequested *		Name of Colorant	
35	colorantConsumed *	Amount	Name of Colorant	
36	jobSubmissionDateTime		DateTime	
37	jobSubmissionTick	Ticks		
38	Reserved			
39	jobStartTick			
40	Reserved			
41	jobCompletedTick	Ticks		
42	Reserved			
43	jobSource	other - 1 unknown - 2 NetWare Qserver - 3 AppleTalk - 4 OS/2 - 5 NT - 6 AIX - 7 Other Unix - 8 AS/400 - 9		
44	jobServerName			
45	jobQueueName			
46	jobServerJobName			
47	jobServerJobNumber			
48	finishingUsed *			
49	jobQualityRequested			
50	jobQualityUsed			
51	tonerEconomyRequested			

52	tonerEconomyUsed			
53	tonerDensityRequested			
54	tonerDensityUsed			
55	finalState			
56	jmJobType *	Other - 1 Unknown - 2 Print - 4 Scan - 8 faxIn - 16 faxOut - 32 getFile - 64 putFile - 128 mailList - 256		
57	jobCopiesCollated	other - 1 unknown - 2 yes - 3 no - 4		
58	jmJobOwner		Owner’s name	

* Means this resource may require an index.