

Job-Policy Overview

Version 19.72

Job policies are used to control an end-to-end migration of an Avamar backup to a different storage medium including to a tape application, Cloud enabled storage, stand alone disk such as USB's/ JBOD or to another Avamar system when replication cannot be used. If exporting to a tape backup application, the backend target storage type used whether it be tape, disk, VTL or Cloud is transparent to ADMe any of which may be leveraged. If using cloud or object storage as a final staging/migration target, a compatible cloud gateway must be used providing the storage is presented as NFS, iSCSI, CIFS or other compatible method which an Avamar recovery session can utilize.

Job policy definitions contain a series of relevant parameters to control and influence the migration process. There are two types of job policies, **Standard** and **Batch** but the Standard is the one that performs the work of a migration while a Batch policy is used simply to automate scheduling a series of Standard policies processing them sequentially to a common migration thread. Interactive dialogues are provided for creating each type of job policy as described below.

Standard Policy Fields:

Policy Field Names	Description
Job Policy Name	User defined and limited to 10 characters or less, do not use the following special characters -\&#*(){} or space. Best practice is to name a job policy the same as that of the client group it will act on.
Environment Number	A predefined environment number used to identify the MCS thread the job will use during the migration
Client Group Name	A predefined client group name containing the client names this job will act on. A client group also defines the staging path, export or tape policy name to call and the required staging server OS type.
Backup Selection Criteria	A range of options used to limit the backup selection by controlling date search range, Plug-in type and various backup count reduction filters
Backup Migration Criteria	A range of options used to influence the migration process such as whether to stage files incrementally , perform stageonly or tapeonly phases or both and how to erase data from a Windows staging drive letter etc.

Note: The environment or thread number referenced equates to a predefined environment file where a single unique staging server name is defined. Each environment file can support a single migration at a time. To support concurrent migrations, unique environment files must be established for each and jobs assigned to them as required. The assigned client group name effectively represents the associated workload of a migration therefore, for any client backup to be migrated its client name must first be defined to a client group.

Creating a Standard Policy

Click on the **Add Job** tool icon to open the Standard job dialogue. Create a job by completing the necessary fields and appropriate options then click the **Create** button. The Selection or Migration command fields will update to contain the appropriate syntax for the options you choose but this field can also be manually edited to adjust or add as needed any additional options not available within the wizard.

Policy Field Names	Description
Job Name	<p>Enter a job name up to 10 characters. For ease of administration use a best practice convention by matching the job name to that of the assigned client group name.</p> <p>MAX job length name can be adjusted from the interactive CLI menu using option a=Admin->a=ADMe customize-></p>
Environment-#	<p>Choose the appropriate staging server thread as defined in the numerically numbered environment file.</p>
Client Group	<p>Choose the appropriate client group name to assign to this job.</p> <p>Note for simplicity the group and job names should match on another</p>
Plug-IN Type	<p>Choose the Plug-IN type the job will act on. A given job policy can act on one Plug-IN type or related plugins such as file systems for all OS's. In cases where clients have backups using different plugin types, these client names must be placed in an additional client group dedicated for the plugin type involved.</p> <p>Note: When using any Data-Base Plug-In such as Exchange, SQL, Oracle, Lotus, Sybase, SAP, SharePoint ensure your selection filter is acting on FULL backups only. Refer to the Label/Name string filter (-gname) below which can be used to further refine the backup selection based on a unique case sensitive string value as viewed from the Label/Name field within the Avamar ADMIN Recovery UI.</p> <p>Note: The use of -ndmp plugin with ADMe relies on a unique capability of Avamar allowing cross platform recoveries from an NDMP backup to a non NAS file system such as to Windows or Unix/Linux. This recovery results in the loss of ACLS and on Windows any extended file attributes.</p> <p>Note: Due to the loss of Windows extended file attributes, Windows OS based encryption or compressed files, cannot be supported due to their corresponding extended attribute being lost. This restriction does not pertain to files compressed or encrypted using applications such as WinZip or gzip etc.</p> <p>Note: If the NAS platform is Celerra/VNX and its internal parameter Backup Data Threshold is set to the default value of 90, it results in the possibility of random files automatically being compressed by VNX to help reduce the overall backup size prior to being sent to the NDMP backup solution. This scenario is incompatible with cross platform recoveries to Windows due to the loss of its extended file attribute. To be supported by ADMe, the Backup Data Threshold parameter must be set to 0 on each Celerra/VNX Data Mover and NAS file system(s) involved to ensure files remain</p>

	<p>in their uncompressed state when being backed up to Avamar. If Avamar has already performed a NDMP Lvl-0 backup prior while the threshold parameter was NOT set to 0, these historical backups are not compatible with ADMe. However, once the threshold is set to 0 and a new NDMP Lvl-0 backup performed, going forward these backups will be compatible with ADMe.</p> <p>Historical NDMP backups should not be cross platform restored or migrated using ADMe unless the above requirements were confirmed to have been met from the start.</p> <p>Refer to ETA 000495263</p>
--	--

Backup Selection Criteria

Default Date Range Search

By default, the search for eligible backups starts from the 1st of your current month up to the date the job is executed. This default suffices in many cases but there are circumstances where this behavior is not suitable and can be easily be modified as needed.

Advanced Date Range Search

Advanced date range search options are used to override the default date search range using either a **fixed** or **relative** range. A relative date range is relative to the date the policy is executed searching backwards in time for a specified number of months or days from your current date. The relative date range concept eliminates the need to use fixed date values and the need to continually modify a date within a job policy each time it is run for Tapeout scenarios. Fixed date type options are typically used for one time job executions when used for an ADS migration.

Search Date Range	Description
Days	Use to define a fixed number of days to search against rather than the default starting from the 1st of the month. Examples -nday 3 or -nday +3 where the + sign will extend the default search range by an additional 3 days into the previous month.
Fixed Date Range	Use to define a specific fixed start and/or end date range formatted as YYYY-MM-DD Example syntax, start date: -sdate yyyy-mm-dd end date: -edate yyyy-mm-dd
Relative	Use to define a date range relative to the date of a job's execution eliminating the need to modify on a monthly basis the syntax for a given date range. Examples -rday day# or day#-day# – act on the relative day or day range of the month

	<p>-rmonth # will start the search from the 1st of the relative month count # backward in time up to the 1st of your current month when the job is being executed</p> <p>-rmonthx # will start the search from the 1st of the relative month count # backward in time up to the end of that specific month only. This is useful when there's a need to select the last backup of the month -rmonthx 1</p> <p>-last or when a series of historical backups must be acted on systematically.</p>
--	---

Note: All the above date range filters can be combined with any of the backup reduction filters documented below

Backup Reduction Filters

Reduction filters are used to limit the quantity of backups selected on a per client basis, the goal being to limit selection to a single backup. These filters can be verified interactively from the CLI menu UI (Interactive Selection) to confirm in advance their suitability prior to using them within a job policy and to calculate a client group's aggregate size without inadvertently moving any data. If unfamiliar with the behavior of any reduction filter, it is advisable to use the CLI interactive capability prior to using and executing them within a job policy.

Note: When checking a group's aggregate size from the Web-UI, it defaults the reduction filter to **-last_F** looking for the last backup available for it in the current month against each client in the group. If you need to check aggregate sizes back in time you can use the interactive CLI Menu Interactive Selection option or you use the CLI directly. When using the CLI or interactive Menu always use Env-20 reserved for this purpose and will not impact any active jobs.

Backup Reduction Filters	Description
-first or -last	Select the first or last backup(s) completed on a given day within the effective date search range. If more than one backup for a client completes on the same day all will be selected. See _F and -gname options below for solutions on limiting multiple BU's per day
-first_F or -last_F	Select the very first or very last available backup within the effective date search range in cases where more than one backup completed on a given day against the same client name. Also useful when clients are located in a REPLICATE domain and BU dates do not align exactly to the original resulting in unexpected BU's on the same date.
Retention Tag -rtype	Select backups whose retention tag specifically matches a single chosen tag value -rtype yearly monthly weekly To act on more than one tag at a time, use the following approach where each desired tag can be specified and joined using a + sign -rtype yearly+monthly+weekly

Days of the Week -week	<p>Select backups which completed on a specific day of the week.</p> <p>Week Number – limits backups to a specified week of the month 0–5 and a chosen day of the week</p> <p>Examples</p> <p>0=all week instances</p> <p>-week fri_0 -last (scan all Fridays but select last Friday of the month)</p> <p>-week fri_1 (first Friday in the first week of the month)</p>
Include Backup Type -but	<p>By default, ADMe acts only on Avamar scheduled backups for selection. To override this behavior, one or more of the following options can be included to act on nonscheduled backup types. If more than one nonscheduled backup type is required, they must be included in a quoted, space separated string as shown</p> <p>Example,</p> <p>-but “cod mod”</p> <p>COD - Client on demand backups</p> <p>MOD=Made on demand backup</p> <p>NAH=Naked Add Hoc backups</p>
Label/Name String -gname	<p>Select backups where the case sensitive string value specified is found within the backup label field as seen from the Avamar Recovery UI. Normally the label field contains the Avamar group/schedule name string values but can also contain a user defined label. This option is useful to force acting on Full backups for databases or when multiple backups are performed against a client and only specific ones are required to be migrated.</p> <p>Example: -gname mystringtext</p>
Select Command	<p>This field echoes the options you’ve chosen within the wizard panel allowing you to view/edit the resulting syntax. This field can edited as required to include options not readily available from the Web-UI. Refer to ADMe User Guide or online help from the CLI menu UI for the complete suite of selection criteria. Refer also to KB 334541</p>

Backup Migration Criteria:

Migration Criteria	Description
Staging Path Suffix	<p>Insert a folder name immediately prior to where the backup data is recovered to</p> <p>Example: -path <i>your-folder-name</i></p>
Post Migration User Script	<p>Execute a user defined or canned script immediately after an export backup has completed successfully. A common use case of this call is to ensure staged data is removed reliably from a Windows staging drive letter for which ADMe provides predefined calls to quick format against the staging drive letter involved. Supported drive letters are E through Z, C & D are excluded and not supported as a staging disk with ADMe.</p> <p>Example: -upst FORMAT-E</p> <p>A user script can also be acted on which must reside on the staging server in the Avamar agent install folder under <install-folder>avs/etc/scripts. For Windows, the script name must conform to the 8.3 file name format such as myscript.bat</p>
Folder Include Exclude	<p>Include/Exclude options provide a means to limit what is staged from the backup by including or excluding top level folder names only from a backup. Use of these filters may require some trial and error therefore verification against a small backup would be a worthwhile step.</p> <p>Option to include is -data <complete-path-folder-name only> is supported and it cannot contain wild cards.</p> <p>Option to exclude is -xdata <complete or partial folder-name> is supported and wild cards * can be used within the specified path name.</p> <p>The Web-UI wizard supports specifying one of these options, however multiple calls for either option can be made by manually entering them into the Migration Command field area within the wizard. Using either option within a job policy will result in the filter being applied to all clients within the client group name being used.</p> <p>Examples:</p> <p>When the path name specified contains spaces, its entry must be enclosed in single quotes.</p> <p>-data <i>complete-path-folder-name</i> (include filter must specify the complete path) 'C:/Program Files/avs/var/clientlogs'</p> <p>-xdata <i>partial or complete-folder-name</i> (Exclude filters can specify a full or partial path and include the use of wild cards) 'Program Files/avs/*/clientlogs'</p>

	<p>If there's a need to limit the filter to specific client(s) only within a client group, refer to Client-Config tab where embedded filters can be applied at the individual client level. Highlight the client entry under Client-Config tab then right click to choose from the context menu to update either the exclude or include fields which are stored in Fld-08 and Fld-17 respectively in the client config csv file. Syntax of an embedded filter is shown here and must be strictly adhered to. Encapsulate each path name with + signs and when multiple path entries are used, each must be separated with a pipe character.</p> <p>Examples:</p> <p>+/etc/bluetooth+ (single folder entry)</p> <p>+/etc/bluetooth+ +/etc/X11+ (multiple folder entry with character)</p> <p>+C:/Program Files/avs/var/clientlogs+ (single folder entry)</p> <p>+avs/var/clientlogs+ +var/update+ (two folder entry using partial path names with a wild card supported with excludes only)</p> <p>The use of both include/exclude filters at the same time is considered undefined and will most likely not provide desired results. The path syntax varies based on whether MCCLI or avtar is being used to perform the recovery and whether include or exclude is being done but ADMe as of build 19.64i masks these differences allowing a common approach for all to be used. Refer to KB 332775 for details on specific syntax to use.</p>
Export Policy Name Override	<p>The regular month-to-month tape policy name(s) are normally specified from within each client group. This option is used in unique situations only where the migration of a given group to tape or other target requires an override when more than one export policy is needed. This is a more efficient alternative to having to create multiple (duplicate) client groups and job policies by creating only an additional job policy name and passing the custom export tape policy details from it as needed to override the one defined within the client group. The primary use of this option is when performing ADS to Avamar-DD migrations, it is not common this be used for a Tapeout use case.</p> <p>NOTE: The entire string must be enclosed in quotes as shown. Options specified within the quotes normally prefixed with either a single or double hyphen, must have the hyphens substituted with tilde characters ~.</p> <p>Example:</p> <p>-tpolicy "-p ADME -w workflow-name" (to tape)</p> <p>-tpolicy "~~expires=151 ~~vardir=C:/progra~1/avs/var3" (for ADS)</p>
Include Meta Data Listing	<p>This option is intended to be used when staging to an archive file (PAX file). It results in a separate text file containing catalog like meta data for the logical content of the archive. Given it is much smaller and human readable it can be reviewed to confirm whether an archive contains the required files</p>

	<p>prior to retrieving the archive itself from tape or cloud storage. When used for other than with an archive file, you must override the Recovery Method option to use –avtar versus mccli</p>
Custom Avtar Flags	<p>Provides a capability for passing optional avtar flags to influence a backup recovery rather than placing them into an avtar.cmd file on the staging server. For example, by default a Windows file system restore does not include files which have the system attribute set therefore, if you require these files the appropriate avatar flag can be passed here, refer to example below.</p> <p>Multiple flags may be passed within the same quoted space separated string.</p> <p>Example, -avflags “restoresystem”</p> <p>Note: Do not include the normal preceding hyphens for each avtar flag passed, ADMe will insert these automatically.</p>
Migration Command	<p>This field echoes the various options you’ve chosen within the wizard panel allowing you to view their resulting syntax. This field can be manually edited if required or to include options not readily available from the Web-UI. Refer to ADMe User Guide or online help from the CLI menu UI for the complete suite of migration criteria.</p>

Interactive Updates to Existing Standard Policies

The following field columns **Client Group Name**, **Thread Number**, **Backup Selection Criteria** and **Backup Migration Criteria** can be modified interactively after a policy has been created. To edit an existing policy field, highlight the desired row then click on one of the above-mentioned fields to be modified. Edit the field as required being sure not to remove any existing required content. When done, press enter which will display a **red triangle** in the upper left corner of the field acknowledging the change, now click on the **Save Changes** tool icon to commit the change. This can be leveraged when a permanent or temporary change is needed.

Creating a Batch Policy

Batch policies are used to automate and overcome scheduling challenges to act on a series of Standard policies sharing a common environment thread number sequentially. Rather than scheduling each Standard individually, only the Batch policy would be scheduled to process the Standard policies it contains sequentially to a common environment thread until all policies it contains complete.

To create a Batch job, use the **Add Batch Job** tool icon to open a batch job creation dialogue. Give it a name following the convention **Benv##** where **##** equals the environment number the **Standard** jobs it will contain are assigned to such as 01, 02 etc. ensuring a two-character nomenclature is used to accommodate sorting in the UI. Choose a **Fail-Action** followed by highlighting the appropriate **Standard** policy names paying attention to its prefixed square bracket value which represents its assigned environment file number. Include required jobs by holding the **Control-Key** and clicking on each job name then press **Create-Batch-Job** button.

Note: Batch policies act on series of **Standard** job policies and they **MUST** share a common environment thread number. It is not supported to include another **Batch** job name nor should it contain **Standard** jobs assigned to different environment file numbers.

Batch Policy Fields:

Policy Field Names	Description
Job Name	A batch job is to be named to reflect the environment number it is being used for. Example, Benv01 , Benv02
Fail Action	When set to STOP , it prevents remaining standard policies assigned to it from being executed when a fatal failure occurs preventing overrunning a common staging disk area in cases where the export phase fails resulting in staged data being left on the staging disk. When set to CONTINUE , allows remaining standard policies assigned to it to be initiated. This is typically only used with incremental staging or if staging to cloud enabled storage because in both cases the staging disk consumed is considered dedicated as with cloud storage the staging disk is considered infinite.

Interactive Update of an Existing Batch Policy

To modify an existing Batch policy first highlight the appropriate row then click on **Edit Batch Job** tool which will display two panels. The right panel shows existing contents and the left displays all other standard policies available. The numeric in the square brackets represents the environment or thread number assigned to a given standard policy and is the reference to ensure only policies sharing a common environment or thread number are assigned to a given Batch policy. To add/remove Standard policies, drag the highlighted policy to the opposite panel.

Cancelling an Active Batch Policy

Canceling a Batch policy is supported **ONLY** from the **CLI menu UI** and by doing so prevents any remaining Standard policies queued to it from being started. If a Standard policy belonging to it already active, it can now be cancelled without concern of the next job in the Batch being initiated.

The command sequence performed from the interactive CLI menu UI to cancel an active Batch policy:

adme -> **j**=Job Manager -> **c**=Cancel -> assigned **#** of the active Batch job to cancel

Scheduling a Job Policy

Scheduling a job policy to be run automatically is accomplished by highlighting the appropriate row then right-click to open the following context menu where you choose the desired option. Scheduling is accomplished using the Linux CRONTAB for **root** userid therefore syntax must conform to requirements and behavior of Linux CRONTAB.

A script is available under the ADMe install folder in /tools named **admesched.sh** which can be used to overcome some scheduling challenges which CRON does not support directly. For example, if there's a requirement to execute a policy the first Saturday of the month, CRON syntax does not accommodate this natively unless an explicit date is specified. The script overcomes this and provides a repetitive automated approach. A corresponding PDF is also present there and can be accessed directly from the Dashboard tab.

Scheduling Choices	Description
Schedule-Enable	Add a new scheduled entry or enable a previously disabled entry
Schedule-Disable	Disable a scheduled entry but retain its current scheduling values
Schedule-Remove	Remove an already enabled or disabled scheduled entry
Schedule-Modify	Modify schedule parameters of a currently scheduled entry

A dialog window is displayed containing the following fields:

Scheduling Parameters	Description
Time	Time of day in HH:MM based on a 24Hr clock
Date	Date of the month 1-31 (*=every or anyday)
Month	Month of the year 1-12 (*=every month)
Day	Day of the week 0-6 where 0=Sunday (*=any day of the week)

Note: An * is considered a wild card implying any value

For more frequent than once per month scheduling use the following syntax examples:

Parameter Value	Description
Date = 1+15	schedule on the 1st and 15th of the month
Date = 1-7	schedule on the 1st through the 7th each month
Month= 1+4+8+12	Schedule the months of Jan, April, Aug and December only
Month = 1-3+10-12	Schedule the months of Jan, Feb, Mar then Oct, Nov, Dec
Day = 0	Schedule only if the day is a Sunday, the above must also align with a Sunday
Day = 5+6+0	Schedule only if it's a Fri, Sat or Sun

Note: The CRON Day parameter does not behave as expected when an explicit Date is also specified versus a wild card. If a Date or Date range is specified combined with a Day of the week value, these two values behave as an OR condition therefore a match of either will result in the job being executed. Refer to the **admesched.sh** script above for an automated approach to overcome this behavior.

Deleting a Policy

To delete an existing job policy, highlighting the appropriate row and click on the **Delete Job** tool icon. Deleting a job policy removes the policy definition, its associated log, its historical logs and will remove it from the scheduler (CRON) if present. Refer to Retiring a job policy below if you prefer to maintain a historical reference as opposed to deleting it.

Retiring a Policy

To retire an existing job policy, highlight the appropriate row and click on the **Retire Job** tool icon. Retiring a policy moves both the policy and its associated logs to **/home/admin/admeadmin/retire** folder for possible future reference and removes it from the Job-Activity panel and scheduler if present.