



job commands

ONTAP 9.3 commands

NetApp
February 11, 2024

Table of Contents

- job commands 1
 - job delete 1
 - job pause 1
 - job resume 2
 - job show-bynode 3
 - job show-cluster 4
 - job show-completed 7
 - job show 9
 - job stop 12
 - job unclaim 13
 - job watch-progress 13
- job history commands 14
- job initstate commands 17
- job private commands 19
- job schedule commands 28

job commands

job delete

Delete a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job delete` command deletes a job. Use the [job show](#) command to view a list of running jobs that can be deleted.

Parameters

-id <integer> - Job ID

The numeric ID of the job you want to delete. A job ID is a positive integer.

[-vserver <vserver name>] - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example deletes the job that has ID 99:

```
cluster1::> job delete -id 99
```

Related Links

- [job show](#)

job pause

Pause a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job pause` command pauses a job. Use the [job resume](#) command to resume a paused job. Use the [job show](#) command to view a list of running jobs that can be paused.

Parameters

-id <integer> - Job ID

The numeric ID of the job you want to pause. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example pauses the job that has ID 183:

```
cluster1::> job pause -id 183
```

Related Links

- [job resume](#)
- [job show](#)

job resume

Resume a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job resume` command resumes a job that was previously paused by using the [job pause](#) command. Use the [job show](#) command to view a list of paused jobs that can be resumed.

Parameters

`-id <integer>` - Job ID

The numeric ID of the paused job to be resumed. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example resumes the paused job that has ID 183:

```
cluster1::> job resume -id 183
```

Related Links

- [job pause](#)
- [job show](#)

job show-bynode

Display a list of jobs by node

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job show-bynode` command displays information about jobs on a per-node basis. The command output depends on the parameters specified with the command. If no parameters are specified, the command displays information about all jobs in the cluster that are currently owned by a node.

To display detailed information about a specific job, run the command with the `-id` parameter. The detailed view includes all of the default information plus additional items.

You can specify additional parameters to display only information that matches the values you specify for those parameters. For example, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

Use this parameter to display information only about the jobs that are associated with the node you specify.

[-id <integer>] - Job ID

Use this parameter to display information only about the jobs that match the ID or range of IDs you specify.

[-vserver <vserver name>] - Owning Vserver

Use this parameter with the name of a Vserver to display only jobs that are owned by that Vserver.

[-name <text>] - Name

Use this parameter to display information only about the jobs that match the job name you specify.

[-description <text>] - Description

Use this parameter to display information only about the jobs that match the description you specify.

[-affinity {Cluster|Node}] - Affinity

Use this parameter with an affinity value to display only jobs that match the affinity you specify.

[-username <text>] - User Name

Use this parameter with a username to display only jobs that are associated with that user.

Examples

The following example displays information about all jobs on a per-node basis:

```
node::> job show-bynode
```

Node	Job ID	Name	Owning Vserver	Affinity
node0	1501	log-rotation	node-vserver	Cluster
		Descr:logrotation job		
node1	1498	log-rotation	node-vserver	Cluster
		Descr:logrotation job		
node2	1499	log-rotation	node-vserver	Cluster
		Descr:logrotation job		
node3	1500	log-rotation	node-vserver	Cluster
		Descr:logrotation job		

job show-cluster

Display a list of cluster jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job show-cluster` command displays information about cluster-affiliated jobs. The command output depends on the parameters specified with the command. If no parameters are specified, the command displays information about all cluster-affiliated jobs.

To display detailed information about a specific job, run the command with the `-id` parameter. The detailed view includes all of the default information plus additional items.

You can specify additional parameters to display only information that matches the values you specify for those parameters. For example, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

[`-instance`] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-id <integer>`] - Job ID

Use this parameter to display information only about the jobs that match the ID or range of IDs you specify.

[`-vserver <vserver name>`] - Owning Vserver

Use this parameter with the name of a Vserver to display only jobs that are owned by that Vserver.

[`-name <text>`] - Name

Use this parameter to display information only about the jobs that match the job name you specify.

[`-description <text>`] - Description

Use this parameter to display information only about the jobs that match the description you specify.

[`-priority {Low|Medium|High|Exclusive}`] - Priority

Use this parameter to display information only about the jobs that match the priority you specify.

[`-node <nodename>`] - Node

Use this parameter to display information only about the jobs that are associated with the node you specify.

[`-affinity {Cluster|Node}`] - Affinity

Use this parameter with an affinity value to display only jobs that match the affinity you specify.

[`-schedule <job_schedule>`] - Schedule

Use this parameter to display information only about the jobs that run on the schedule you specify.

[`-queuetime <MM/DD HH:MM:SS>`] - Queue Time

Use this parameter to display information only about the jobs that match the queue time you specify.

[`-starttime <MM/DD HH:MM:SS>`] - Start Time

Use this parameter to display information only about the jobs that match the start time you specify.

[`-endtime <MM/DD HH:MM:SS>`] - End Time

Use this parameter to display information only about the jobs that match the end time you specify.

[`-dropdeadtime <MM/DD HH:MM:SS>`] - Drop-dead Time

Use this parameter to display information only about the jobs that match the final timeout time you specify.

[`-restarted {true|false}`] - Restarted?

Use this parameter to display information only about the jobs that match the restart value you specify.

[`-state`

{`Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant`}] - State

Use this parameter to display information only about the jobs that match the job state you specify.

[-code <integer>] - Status Code

Use this parameter to display information only about the jobs that match the status code you specify.

[-completion <text>] - Completion String

Use this parameter to display information only about the jobs that match the completion text you specify.

[-jobtype <text>] - Job Type

Use this parameter to display information only about the jobs that match the job type you specify.

[-category <text>] - Job Category

Use this parameter to display information only about the jobs that match the job category you specify.

[-uuid <UUID>] - UUID

Use this parameter to display information only about the jobs that match the UUID you specify.

[-username <text>] - User Name

Use this parameter with a username to display only jobs that are associated with the user you specify.

Examples

The following example displays information about all cluster-affiliated jobs:

```
cluster1::> job show-cluster
```

Job ID	Name	Owning Vserver	Node	State
305	Auto_Mirror	node-vserver	-	Running
6202	mirror-03_10	node-vserver	-	Queued
	Descr:Auto mirror			
6203	mirror-04_10	node-vserver	-	Queued
	Descr:Auto mirror			
6204	mirror-01_10	node-vserver	-	Queued
	Descr:Auto mirror			
6205	mirror-02_10	node-vserver	-	Queued
	Descr:Auto mirror			
6206	mirror-05_10	node-vserver	-	Queued
	Descr:Auto mirror			

job show-completed

Display a list of completed jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job show-completed` command displays information about completed jobs. The command output depends on the parameters you specify with the command. If you do not use any parameters, the command displays information about all completed jobs.

To display detailed information about a specific job, run the command with the `-id` parameter. The detailed view includes all of the default information plus additional items.

You can specify additional parameters to display only information that matches those parameters. For instance, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-id <integer>] - Job ID

Use this parameter to display information only about the jobs that match the ID or range of IDs you specify.

[-vserver <vserver name>] - Owing Vserver

Use this parameter with the name of a Vserver to display only jobs that are owned by that Vserver.

[-name <text>] - Name

Use this parameter to display information only about the jobs that match the name you specify.

[-description <text>] - Description

Use this parameter to display information only about the jobs that match the description you specify.

[-priority {Low|Medium|High|Exclusive}] - Priority

Use this parameter to display information only about the jobs that match the priority you specify.

[-node <nodename>] - Node

Use this parameter to display information only about the jobs that are associated with the node you specify.

[-affinity {Cluster|Node}] - Affinity

Use this parameter with an affinity value to display only jobs that match the affinity you specify.

[-schedule <job_schedule>] - Schedule

If you use this parameter, the command displays information only about the jobs that have the schedule you specify.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time

If you use this parameter, the command displays information only about the jobs that have the queue time you specify.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Use this parameter to display information only about the jobs that have the start time you specify.

[-endtime <MM/DD HH:MM:SS>] - End Time

Use this parameter to display information only about the jobs that have the end time you specify.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time

Use this parameter to display information only about the jobs that time out at the time you specify.

[-restarted {true|false}] - Restarted?

Use this parameter to display information only about the jobs that match the restart value you specify.

[-state

{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State

Use this parameter to display information only about the jobs that match the job state you specify.

[-code <integer>] - Status Code

Use this parameter to display information only about the jobs that match the status code you specify.

[-completion <text>] - Completion String

Use this parameter to display information only about the jobs that match the completion text you specify.

[-jobtype <text>] - Job Type

Use this parameter to display information only about the jobs that match the job type you specify.

[-category <text>] - Job Category

Use this parameter to display information only about the jobs that match the job category you specify.

[-uuid <UUID>] - UUID

Use this parameter to display information only about the jobs that match the UUID you specify.

[-username <text>] - User Name

Use this parameter with a username to display only jobs that are associated with that user.

Examples

The following example displays information about all completed jobs:

```
node::> job show-completed
```

Job ID	Name	Owning Vserver	End Time	Code	Completion
305	Auto_Mirror	node-vserver	10/10 08:07:05	0	Succeeded
6202	mirror-03_10	node-vserver	10/10 11:10:07	0	
6203	mirror-04_10	node-vserver	10/10 12:10:09	0	
6204	mirror-01_10	node-vserver	10/10 09:10:03	0	
6205	mirror-02_10	node-vserver	10/10 10:10:08	0	
6206	mirror-05_10	node-vserver	10/10 05:10:04	0	

job show

Display a list of jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job show` command displays information about jobs. By default, the command displays information about all current jobs.

To display detailed information about a specific job, run the command with the `-id` parameter.

You can specify additional parameters to select information that matches the values you specify for those parameters. For example, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-inprogress]

Displays the job ID, the job name, the owning Vserver, and the progress of the job.

| **[-jobstate]**

Displays information about each job's state, including the queue state, whether the job was restarted, and when the job has completely timed out.

| **[-sched]**

Displays the job ID, the job name, the owning Vserver, and the schedule on which the job runs.

| **[-times]**

Displays the job ID, the job name, the owning Vserver, the time when the job was last queued, the time when the job was last started, and the time when the job most recently ended.

| **[-type]**

Displays the job ID, the job name, the job type, and the job category.

| **[-jobuuid] (privilege: advanced)**

Displays the job ID, the job name, the owning Vserver, and the job UUID.

| **[-instance] }**

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-id <integer>] - Job ID

Selects the jobs that match the ID or range of IDs that you specify.

[-vserver <vserver name>] - Owning Vserver

Selects jobs that are owned by the specified Vserver.

[-name <text>] - Name

Selects the jobs that match this parameter value.

[-description <text>] - Description

Selects the jobs that match this parameter value.

[-priority {Low|Medium|High|Exclusive}] - Priority

Selects the jobs that match this parameter value.

[-node <nodename>] - Node

Selects the jobs that match this parameter value.

[-affinity {Cluster|Node}] - Affinity

Selects the jobs that match this parameter value.

[-schedule <job_schedule>] - Schedule

Selects the jobs that match this parameter value.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time

Selects the jobs that match this parameter value.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Selects the jobs that match this parameter value.

[-endtime <MM/DD HH:MM:SS>] - End Time

Selects the jobs that match this parameter value.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time

Selects the jobs that match this parameter value.

[-restarted {true|false}] - Restarted?

Selects the jobs that match this parameter value.

[-state

{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State

Selects the jobs that match this parameter value.

[-code <integer>] - Status Code

Selects the jobs that match this parameter value.

[-completion <text>] - Completion String

Selects the jobs that match this parameter value.

[-jobtype <text>] - Job Type

Selects the jobs that match this parameter value.

[-category <text>] - Job Category

Selects the jobs that match this parameter value.

[-uuid <UUID>] - UUID (privilege: advanced)

Selects the jobs that match this parameter value.

[-progress <text>] - Execution Progress

Selects the jobs that match this parameter value.

[-username <text>] - User Name

Selects the jobs that match this parameter value.

[-restart-is-delayed-by-module <text>] - Restart Is Delayed by Module

Selects jobs which are or were delayed by the specified module during the restart. For example:
MCC_SWITCHBACK

Examples

The following example displays information about all jobs on the node named node1:

```
cluster1::> job show -node node1
```

Job ID	Name	Owning Vserver	Node	State
308114	mirror-daily-3587206	node-vserver	node1	Running
	Descr:Auto-replicate to 1 mirror(s)			
308115	mirror-daily-3618985	node-vserver	node1	Running
	Descr:Auto-replicate to 1 mirror(s)			
308116	mirror-daily-3619010	node-vserver	node1	Queued
	Descr:Auto-replicate to 1 mirror(s)			
308117	mirror-daily-3749547	node-vserver	node1	Queued
	Descr:Auto-replicate to 1 mirror(s)			

4 entries were displayed.

job stop

Stop a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job stop` command stops a running job. A stopped job cannot be resumed. Use the [job pause](#) command to pause a job so that you can later resume it. Use the [job show](#) command to view a list of running jobs.

Parameters

-id <integer> - Job ID

The numeric ID of the job to stop. A job ID is a positive integer.

[-vserver <vserver name>] - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example stops the job that has ID 101:

```
cluster1::> job stop -id 101
```

Related Links

- [job pause](#)
- [job show](#)

job unclaim

Unclaim a cluster job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job unclaim` command causes a cluster-affiliated job that is owned by an unavailable node to be unclaimed by that node. Another node in the cluster can then take ownership of the job. Use the [job show-cluster](#) command to obtain a list of cluster-affiliated jobs.

Parameters

-id <integer> - Job ID (privilege: advanced)

Use this parameter to specify the ID number of the job to unclaim.

[-vserver <vserver name>] - Owning Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example shows how to unclaim the cluster-affiliated job with the ID 27 that is owned by the Vserver `vs1`:

```
cluster1::*> job unclaim -vserver vs1 -id 27
```

Related Links

- [job show-cluster](#)

job watch-progress

Watch the progress of a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job watch-progress` command displays the progress of a job, and periodically updates that display. You can specify the frequency of the updates.

Parameters

-id <integer> - Job ID

Use this parameter to specify the numeric ID of the job to monitor.

[-vserver <vserver name>] - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

[-interval <integer>] - Refresh Interval (seconds)

Use this parameter to specify the number of seconds between updates.

Examples

The following example show how to monitor the progress of the job that has ID 222 on Vserver `vs0` . The progress display updates every 3 seconds.

```
cluster1::> job watch-progress -vserver vs0 -id 222 -interval 3
```

job history commands

job history show

Display a history of jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job history show` command displays a history of completed jobs with newer entries displayed first. You can specify optional parameters to select information about job history items that match only those parameters. For example, to display information about jobs that were completed on February 27 at noon, run the command with `-endtime "02/27 12:00:00"` .

Parameters

{ [-fields <fieldname>, ...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

Selects the completed jobs that match this parameter value.

[-record <Sequence Number>] - Record ID

Selects the completed jobs that match the record ID or range of record IDs you specify. Note that record IDs are unique for each node, not for the cluster as a whole. As a result, there can be two records with the same record ID within the cluster.

[-vserver <vserver name>] - Owning Vserver

Selects the completed jobs that are owned by the Vserver you specify.

[-id <integer>] - Job ID

Selects the completed jobs that match this parameter value.

[-endtime <MM/DD HH:MM:SS>] - End Time

Selects jobs that completed at the time you specify. This parameter is most useful when used with a range of times.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Selects completed jobs that were started at the time you specify. This parameter is most useful when used with a range of times.

[-name <text>] - Name

Selects the completed jobs that match this parameter value.

[-description <text>] - Description

Selects the completed jobs that match this parameter value.

[-code <integer>] - Status Code

Selects the completed jobs that match this parameter value. Each job defines its own status codes. The completion text is more informative, but support technicians may request this numeric code.

[-progress <text>] - Progress String

Selects the completed jobs that match this parameter value.

[-completion <text>] - Completion String

Selects the completed jobs that match this parameter value.

[-jobuuid <UUID>] - Job UUID (privilege: advanced)

Selects the completed jobs that match this parameter value.

[-event-type {Idle|Running|Succeeded|Failed|Paused|Stopped|Deleted|Error}] - Event Type

Selects the completed jobs that match this parameter value.

[-event-time <MM/DD HH:MM:SS>] - Event Time

Selects the completed jobs that match this parameter value. This parameter is most useful when used with a range of times.

[-error-code <integer>] - Job Manager Error Code

Selects the completed jobs that match this parameter value.

[-error-text <text>] - Job Manager Error Text

Selects the completed jobs that match this parameter value.

[-username <text>] - User Name

Selects the completed jobs that match this parameter value.

Examples

The following example displays information about all completed jobs:

```
cluster1::> job history show
```

Time	Node	Owning Vserver	Name	Event
08/23 08:58:24	node1	node1-vs	Vol Create	Succeeded
Description: Create testvol Completion: Successful				
08/23 08:58:22	node1	node1-vs	Vol Create	Running
Description: Create testvol				
08/22 08:16:36	node1	node1-vs	CLUSTER BACKUP AUTO weekly	Succeeded
Description: Cluster Backup Job				
08/22 08:15:49	node1	node1-vs	CLUSTER BACKUP AUTO weekly	Running
Description: Cluster Backup Job				
08/22 08:15:08	node1	node1-vs	CLUSTER BACKUP AUTO weekly	Idle
Description: Cluster Backup Job				
08/22 08:15:03	node1	node1-vs	CLUSTER BACKUP AUTO weekly	Running
Description: Cluster Backup Job				

6 entries were displayed.

The following example shows how to use a range with the "endtime" parameter to select only the events that ended between 8:15 and 8:16 on August 22nd.

```
cluster1::> job history show -endtime "08/22 08:15:00".. "08/22 08:16:00"

```

Time	Node	Owning Vserver	Name	Event
08/22 08:15:49	node1	node1-vs	CLUSTER BACKUP AUTO	weekly Running
4 Description: Cluster Backup Job				
08/22 08:15:08	node1	node1-vs	CLUSTER BACKUP AUTO	weekly Idle
4 Description: Cluster Backup Job				
08/22 08:15:03	node1	node1-vs	CLUSTER BACKUP AUTO	weekly Running
4 Description: Cluster Backup Job				

3 entries were displayed.

job initstate commands

job initstate show

Display init state for job managers

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `job initstate show` command displays information about the initialization states of job-manager processes.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the nodes that match this parameter value.

`[-process <process_name>]` - Process Name (privilege: advanced)

Selects the nodes that match this parameter value.

`[-initialized {true|false}]` - Initialized? (privilege: advanced)

Selects the nodes that match this parameter value (`true` means initialized; `false` means not initialized).

`[-cache-root <text>]` - Cache Root (privilege: advanced)

Selects the nodes that match this parameter value.

`[-siteid <UUID>]` - Site ID (privilege: advanced)

Selects the nodes that match this parameter value.

`[-hp-threads <integer>]` - High Priority Threads (privilege: advanced)

Selects the nodes that have the number of high-priority threads you specify.

`[-mp-threads <integer>]` - Medium Priority Threads (privilege: advanced)

Selects the nodes that have the number of medium-priority threads you specify.

`[-lp-threads <integer>]` - Low Priority Threads (privilege: advanced)

Selects the nodes that have the number of low-priority threads you specify.

`[-tx-interval <integer>]` - Transaction Interval (privilege: advanced)

Selects the nodes that have the number of seconds you specify as their transaction interval.

`[-initmsg <text>]` - Initialization Message (privilege: advanced)

Selects the nodes that match this parameter value.

`[-thread-initmsg <text>]` - Thread Initialization Message (privilege: advanced)

Selects the nodes that match this parameter value. The thread initialization message contains information about thread status. If there is no information to communicate, this message is empty.

`[-recovery-enabled {true|false}]` - Job Failover Enabled? (privilege: advanced)

Selects the nodes that match this parameter value (`true` means enabled, `false` means not enabled).

`[-ex-threads <integer>]` - Exclusive Priority Threads (privilege: advanced)

Selects the nodes that match this parameter value.

Examples

The following example shows how to display general job-manager initialization-state information for a cluster.

```
cluster1::*> job initstate show
```

Node	Process	Init?	HP Thr	MP Thr	LP Thr	EX Thr	TX Int	Failover?
node1	mgwd	true	2	3	5	8	300	true
node2	mgwd	true	2	3	5	8	300	true

2 entries were displayed.

The following example shows how to display detailed job-manager initialization-state information for a node named node0 .

```
cluster1::*> job initstate show -instance -node node0
```

```
Node: node0
Process Name: mgwd
Initialized?: true
Cache Root: /mroot/jm_cache
Site ID: 824e8f7d-f49-1d9-84af-00423b7352
High Priority Threads: 2
Medium Priority Threads: 3
Low Priority Threads: 5
Transaction Interval: 300
Initialization Message: Initialized
Are Threads Running?: -
Job Failover Enabled?: true
Exclusive Priority Threads: 8
```

job private commands

job private delete

Delete a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private delete` command deletes a private job. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

If you use this command on a job that does not support the delete operation, the command returns an error message.

Use the [job private show](#) command to view a list of private jobs that can be deleted.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node with which the private job is associated.

-id <integer> - Job ID (privilege: advanced)

Use this parameter to specify the numeric ID of the private job to be deleted. A job ID is a positive integer.

[-vserver <vserver name>] - Owing Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example shows how to delete the job that has ID 273 from the node named `node2` :

```
cluster1::*> job private delete -node node2 -id 273
```

Related Links

- [job private show](#)

job private pause

Pause a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private pause` command pauses a private job. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

If you use this command to pause a job that does not support it, the command returns an error message.

Use the [job private resume](#) command to resume a paused private job.

Use the [job private show](#) command to view a list of private jobs.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node with which the private job is associated.

-id <integer> - Job ID (privilege: advanced)

Use this parameter to specify the numeric ID of the paused private job to be paused. A job ID is a positive integer.

[-vserver <vserver name>] - Owing Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example pauses the private job that has ID 99 on the node `node1` :

```
cluster1::*> jobs private pause -node node1 -id 99
```

Related Links

- [job private resume](#)
- [job private show](#)

job private resume

Resume a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private resume` command resumes a private job that was paused by using the [job private pause](#) command. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

Use the [job private show](#) command to view a list of paused private jobs that can be resumed.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node with which the paused private job is associated.

-id <integer> - Job ID (privilege: advanced)

Use this parameter to specify the numeric ID of the paused private job to be resumed. A job ID is a positive integer.

[-vserver <vserver name>] - Owing Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example resumes the paused private job that has ID 99 on a node named `node2` :

```
cluster1::*> job private resume -node node2 -id 99
```

Related Links

- [job private pause](#)
- [job private show](#)

job private show-completed

Display a list of completed jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private show-completed` command displays information about completed private jobs. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node (privilege: advanced)

Use this parameter to display information only about completed jobs that are associated with the node you specify.

[-id <integer>] - Job ID (privilege: advanced)

Use this parameter to display information only about completed jobs that have the ID you specify.

[-vserver <vserver name>] - Owning Vserver (privilege: advanced)

Use this parameter to display only completed jobs that are owned by the Vserver you specify.

[-name <text>] - Name (privilege: advanced)

Use this parameter to display information only about completed jobs that have the name you specify.

[-description <text>] - Description (privilege: advanced)

Use this parameter to display information only about completed jobs that have the description you specify.

[-priority {Low|Medium|High|Exclusive}] - Priority (privilege: advanced)

Use this parameter to display information only about completed jobs that have the priority you specify.

[-schedule <job_schedule>] - Schedule (privilege: advanced)

Use this parameter to display information only about completed jobs that have the schedule you specify.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the queue time you specify.

[-starttime <MM/DD HH:MM:SS>] - Start Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the start time you specify.

[-endtime <MM/DD HH:MM:SS>] - End Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the end time you specify.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the final timeout time you specify.

[-restarted {true|false}] - Restarted? (privilege: advanced)

Use this parameter to display information only about completed jobs that have the restart value you specify.

[-state

{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State (privilege: advanced)

Use this parameter to display information only about completed jobs that have the job state you specify.

[-code <integer>] - Status Code (privilege: advanced)

Use this parameter to display information only about completed jobs that have the status code you specify.

[-completion <text>] - Completion String (privilege: advanced)

Use this parameter to display information only about completed jobs that have the completion text you specify.

[-jobtype <text>] - Job Type (privilege: advanced)

Use this parameter to display information only about completed jobs that have the job type you specify.

[-category <text>] - Job Category (privilege: advanced)

Use this parameter to display information only about completed jobs that have the job category you specify.

[-uuid <UUID>] - UUID (privilege: advanced)

Use this parameter to display information only about completed jobs that have the UUID you specify.

[-username <text>] - User Name (privilege: advanced)

Use this parameter to display information only about completed jobs that are associated with the user you specify.

Examples

The following example shows how to display information about all completed private jobs on the node named node1 :

```
cluster1::*> job private show-completed -node node1
Node: node1
```

Job ID	Name	Owning Vserver	End Time	Code	Completion String
1	sync task	node1	02/17 15:03:23	0	
2	load_balancing	node1	02/17 16:29:28	0	DONE_VIF_STATS
3	snap-hourly	node1	02/17 16:05:00	0	
4	snap-daily	node1	02/17 00:10:00	0	
5	snap-weekly	node1	02/13 00:15:00	0	
8	Cross-Cluster Manager	node1	02/17 16:27:27	0	complete
9	reconcile service policy	node1	02/17 15:03:12	0	

7 entries were displayed.

job private show

Display a list of jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private show` command displays information about private jobs. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-inprogress] (privilege: advanced)

Displays the job ID, name, owning Vserver, and progress of each private job.

| [-jobstate] (privilege: advanced)

Displays information about each private job's state, including the queue state, whether the job was restarted, and when the job has timed out.

| [-jobuuid] (privilege: advanced)

Displays the ID, name, owning Vserver, and UUID of each private job.

| [-sched] (privilege: advanced)

Displays the job ID, name, owning Vserver, and run schedule of each private job.

| [-times] (privilege: advanced)

Displays the queue time, start time, and end time of each private job.

| [-type] (privilege: advanced)

Displays the type and category of each private job.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the private jobs that match this parameter value. .

[-id <integer>] - Job ID (privilege: advanced)

Selects the private jobs that match the ID or range of IDs that you specify.

[-vserver <vserver name>] - Owning Vserver (privilege: advanced)

Selects the private jobs that match this parameter value.

[-name <text>] - Name (privilege: advanced)

Selects the private jobs that match this parameter value.

[-description <text>] - Description (privilege: advanced)

Selects the private jobs that match this parameter value.

[-priority {Low|Medium|High|Exclusive}] - Priority (privilege: advanced)

Selects the private jobs that match this parameter value.

[-schedule <job_schedule>] - Schedule (privilege: advanced)

Selects the private jobs that match this parameter value.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-starttime <MM/DD HH:MM:SS>] - Start Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-endtime <MM/DD HH:MM:SS>] - End Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-restarted {true|false}] - Restarted? (privilege: advanced)

Selects the private jobs that match this parameter value.

[-state

{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State (privilege: advanced)

Selects the private jobs that match this parameter value.

[-code <integer>] - Status Code (privilege: advanced)

Selects the private jobs that match this parameter value.

[-completion <text>] - Completion String (privilege: advanced)

Selects the private jobs that match this parameter value.

[-jobtype <text>] - Job Type (privilege: advanced)

Selects the private jobs that match this parameter value.

[-category <text>] - Job Category (privilege: advanced)

Selects the private jobs that match this parameter value.

[-uuid <UUID>] - UUID (privilege: advanced)

Selects the private jobs that match this parameter value.

[-progress <text>] - Execution Progress (privilege: advanced)

Selects the private jobs that match this parameter value.

[-username <text>] - User Name (privilege: advanced)

Selects the private jobs that match this parameter value.

Examples

The following example displays information about all private jobs on the local node:

```
cluster1::*> job private show -node local
Node: node1

Job ID Name                                Owing
      Vserver                               State
-----
3      snap-hourly                         cluster1  Queued
      Description: Auto-Snapshot
4      snap-daily                         cluster1  Queued
      Description: Auto-Snapshot
5      snap-weekly                        cluster1  Queued
      Description: Auto-Snapshot
6      sync task                          cluster1  Queued
      Description: sync task
7      ldap-certs                         cluster1  Queued
      Description: ldap resync
5 entries were displayed.
```

job private stop

Stop a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private stop` command stops a running private job. A private job is a job that is associated with a specific node and does not use cluster facilities. A stopped job cannot be restarted.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

This specifies the node on which the job is running.

-id <integer> - Job ID (privilege: advanced)

This specifies the numeric ID of the job that is to be stopped.

[-vserver <vserver name>] - Owing Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example stops a private job with the ID 416 on a node named node0:

```
cluster1::*> job private stop -node node0 -id 416
```

job private watch-progress

Watch the progress of a job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `job private watch-progress` command displays and periodically updates the progress of a private job. A private job is a job that is associated with a specific node and does not use cluster facilities. You can specify the frequency of the progress updates.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

This specifies the node on which the job is running.

-id <integer> - Job ID (privilege: advanced)

This specifies the numeric ID of the job whose progress is to be monitored.

[-vserver <vserver name>] - Owing Vserver (privilege: advanced)

Use this parameter to specify the Vserver with which the paused private job is associated.

[-interval <integer>] - Refresh Interval (seconds) (privilege: advanced)

This optionally specifies, in seconds, the frequency of the updates.

Examples

The following example monitors the progress of the private job that has ID 127 on a node named node1. The progress is updated every 2 seconds.

```
cluster1::*> job private watch-progress -node node1 -id 127 -interval 2
Queued
```

job schedule commands

job schedule delete

Delete a schedule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule delete` command deletes a schedule. Use the [job schedule show](#) command to display all current schedules.

You cannot delete any schedules that are in use by jobs. Use the [job schedule show-jobs](#) command to display jobs by schedule.

You cannot delete any schedules that are referenced by:

- Volume Snapshot copy policy entries
- SnapMirror entries
- SIS policy entries
- configuration backup settings

You must remove all references to a schedule before you can delete it. If you attempt to delete a schedule that is referenced, an error message will list which entries reference the schedule you want to delete. Use the `show` command for each of the items listed by the error message to display which entries reference the schedule. You may need to use the `-instance` parameter to display more detail.

Parameters

`[-cluster <Cluster name>]` - Cluster

This parameter specifies the name of the cluster on which you want to delete a schedule. By default, the schedule is deleted from the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

`-name <text>` - Schedule Name

Use this parameter with the name of an existing schedule to specify the schedule you want to delete.

Examples

The following example deletes a schedule named overnightbackup:

```
cluster1::> job schedule delete -name overnightbackup
```

Related Links

- [job schedule show](#)
- [job schedule show-jobs](#)

job schedule show-jobs

Display the list of jobs by schedule

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `job schedule show-jobs` command displays information about jobs that are associated with schedules.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-name <text>] - Schedule Name

Use this parameter to display information only about the jobs that are associated with the schedule you specify.

[-affinity {Cluster|Node}] - Cluster / Node

Use this parameter to display information only about the jobs that match the affinity value you specify.

[-owner <text>] - Owner

Use this parameter to display information only about the jobs that are owned by the nodes you specify.

[-jobid <integer>] - ID

Use this parameter to display information only about the jobs that match the ID or range of IDs that you specify.

[-jobname <text>] - Job Name

Use this parameter to display information only about the jobs that match the name you specify.

Examples

The following example shows information about schedules that are associated with jobs:

```
cluster1::> job schedule show-jobs
```

Name	Type	Owner	Job ID	Job Name
hourly	Cluster	-	98644	mirror-hourly
weeklylog	Node	node0	1501	log-rotation
weeklylog	Node	node1	1498	log-rotation
weeklylog	Node	node2	1499	log-rotation
weeklylog	Node	node3	1500	log-rotation

5 entries were displayed.

job schedule show

Display a list of available schedules

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule show` command displays information about schedules.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-cluster <Cluster name>] - Cluster

Selects the schedules that match this parameter value.

[-name <text>] - Schedule Name

Selects the schedules that match this parameter value.

[-type {cron|interval|builtin}] - Schedule Type

Selects the schedules that match this parameter value.

[-description <text>] - Description

Selects the schedules that match this parameter value.

Examples

The following example displays information about all schedules:


```
cluster1::> job schedule show
Cluster      Name      Type      Description
-----
cluster1
      5min      cron
@:00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
      daily      cron      @0:10
      hourly      cron      @:05
      monthly      cron      1@0:20
      weekly      cron      Sun@0:15
```

job schedule cron create

Create a cron schedule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule cron create` command creates a cron schedule. A cron schedule, like a UNIX cron job, runs at a specified time. You can also specify months, days of the month, or days of the week on which the schedule will run.

If you specify values for both days of the month and days of the week, they are considered independently. For example, a cron schedule with the day specification Friday, 13 runs every Friday and on the 13th day of each month, not just on every Friday the 13th.

Parameters

[-cluster <Cluster name>] - Cluster

This parameter specifies the name of the cluster on which you want to create a cron schedule. By default, the schedule is created on the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

-name <text> - Name

Use this parameter to specify the name of the cron schedule that you want to create.

[-month <cron_month>,...] - Month

Use this parameter to specify months in which the schedule runs. Valid values are January, February, March, April, May, June, July, August, September, October, November, December, and all. Specify "all" to run the schedule every month.

[-dayofweek <cron_dayofweek>,...] - Day of Week

Use this parameter to specify days of the week on which the schedule runs. Valid values are Sunday, Monday, Tuesday, Thursday, Friday, and Saturday, and all. Specify "all" to run the schedule every day.

[`-day <cron_dayofmonth>,...`] - Day

Use this parameter to specify days of the month on which the schedule runs. Valid values range from 1 to 31.

[`-hour <cron_hour>,...`] - Hour

Use this parameter to specify the hours value of the time of day at which the schedule runs. Valid values range from 0 (midnight) to 23 (11:00 p.m.). Specify "all" to run the schedule every hour.

`-minute <cron_minute>,...` - Minute

Use this parameter to specify the minutes portion of the time of day at which the schedule runs. Valid values range from 0 to 59.

Examples

The following example creates a cron schedule named `weekendcron` that runs on weekend days (Saturday and Sunday) at 3:00 a.m.

```
cluster1::> job schedule cron create -name weekendcron -dayofweek
"Saturday, Sunday" -hour 3 -minute 0
```

job schedule cron delete**Delete a cron schedule**

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule cron delete` command deletes a cron schedule. Use the [job schedule cron show](#) command to display all current cron schedules.

You cannot delete any cron schedules that are associated with jobs. Use the [job schedule show-jobs](#) command to display jobs by schedule.

Parameters**[`-cluster <Cluster name>`] - Cluster**

This parameter specifies the name of the cluster on which you want to delete a cron schedule. By default, the schedule is deleted from the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

`-name <text>` - Name

Use this parameter with the name of an existing cron schedule to specify the cron schedule that you want to delete.

Examples

The following example deletes a cron schedule named `midnightcron`:

```
cluster1::> job schedule cron delete -name midnightcron
```

Related Links

- [job schedule cron show](#)
- [job schedule show-jobs](#)

job schedule cron modify

Modify a cron schedule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule cron modify` command modifies a cron schedule. A cron schedule, like a UNIX cron job, runs at a specified time. You can also specify months, days of the month, or days of the week on which the schedule runs. Use the [job schedule cron show](#) command to display all current cron schedules. See the documentation for [job schedule cron show](#) for more information about how cron schedules work.

Modifying one parameter of a cron schedule does not affect the other parameters. For example, if cron schedule is set to run at 3:15 AM, and you modify the "hour" parameter to 4, the schedule's new time will be 4:15am. To clear a parameter of the schedule's interval, you must explicitly set that portion to "0" or "-". Some parameters can also be set to "all".

Parameters

[-cluster <Cluster name>] - Cluster

Use this parameter to specify the cluster of an existing cron schedule you want to modify. The local cluster is provided as the default value. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

-name <text> - Name

Use this parameter with the name of an existing cron schedule to specify the cron schedule you want to modify.

[-month <cron_month>,...] - Month

Use this parameter to specify a new "month" value for the cron schedule. Valid values are January, February, March, April, May, June, July, August, September, October, November, December, or all. Specify "all" to run the schedule every month.

[-dayofweek <cron_dayofweek>,...] - Day of Week

Use this parameter to specify a new "day of week" value for the cron schedule. Valid values include Sunday, Monday, Tuesday, Thursday, Friday, Saturday, or all. Specify "all" to run the schedule every day.

[-day <cron_dayofmonth>,...] - Day

Use this parameter to specify a new "day of month" value for the cron schedule. Valid values range from 1 to 31.

[-hour <cron_hour>,...] - Hour

Use this parameter to specify a new "hour of the day" value for the cron schedule. Valid values range from 0 (midnight) to 23 (11:00 p.m.), Specify "all" to run the schedule every hour.

[-minute <cron_minute>,...] - Minute

Use this parameter to specify a new "minute of the hour" value for the cron schedule. Valid values range from 0 to 59.

Examples

The following example modifies a cron schedule named weekendcron so that it runs at 3:15 a.m.:

```
cluster1::> job schedule cron modify -name weekendcron -hour 3 -minute 15
```

Related Links

- [job schedule cron show](#)

job schedule cron show

Show cron schedules

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule cron show` command displays information about cron schedules. A cron schedule runs a job at a specified time on specified days.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-cluster <Cluster name>] - Cluster

Selects the cron schedules that match this parameter value.

[-name <text>] - Name

Selects the cron schedules that match this parameter value.

[-month <cron_month>,...] - Month

Selects the cron schedules that match this parameter value. Valid values are January , February , March , April , May , June , July , August , September , October , November , December , or all .

[-dayofweek <cron_dayofweek>,...] - Day of Week

Selects the cron schedules that match this parameter value. Valid values include Sunday , Monday , Tuesday , Wednesday , Thursday , Friday , Saturday , or all .

[-day <cron_dayofmonth>,...] - Day

Selects the cron schedules that match this parameter value. Valid values range from 1 to 31 .

[-hour <cron_hour>,...] - Hour

Selects the cron schedules that match this parameter value.

[-minute <cron_minute>,...] - Minute

Selects the cron schedules that match the minute or range of minutes that you specify.

[-description <text>] - Description

Selects the cron schedules that match this parameter value.

Examples

The following example displays information about all current cron schedules:

```
cluster1:> job schedule cron show
Cluster      Name      Description
-----
cluster1
            5min      @:00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
            8hour     @2:15,10:15,18:15
            weekly    Sun@0:15
```

The following example displays information about the cron schedule named weekly:

```
cluster1:> job schedule cron show -name weekly -instance
Cluster: cluster1
      Name: weekly
      Month: -
Day of Week: Sunday
      Day: -
      Hour: 0
      Minute: 15
Description: Sun@0:15
```

job schedule interval create

Create a schedule that runs on an interval

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule interval create` creates an interval schedule. An interval schedule runs jobs at specified intervals after the previous job finishes. For instance, if a job uses an interval schedule of 12 hours and takes 30 minutes to complete, the job runs at the following times:

- Day one at 8:00 a.m. (the job's initial run)
- Day one at 8:30 p.m.
- Day two at 9:00 a.m.
- Day two at 9:30 p.m.

Each of the numerical parameters of the interval must be a whole number. These parameters can be used individually, or combined to define complex time values. For example, use a value of 1 day, 12 hours to create an interval of 1.5 days.

Large parameter values are converted into larger units. For example, if you create a schedule with an interval of 36 hours, the `job schedule interval show` command will display it with an interval of 1 day 12 hours.

Parameters

`[-cluster <Cluster name>] - Cluster`

This parameter specifies the name of the cluster on which you want to create an interval schedule. By default, the schedule is created on the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

`-name <text> - Name`

Use this parameter to specify the name of the interval schedule you want to create.

`[-days <integer>] - Days`

Use this parameter to specify the "days" portion of the schedule's interval. A day is one calendar day.

`[-hours <integer>] - Hours`

Use this parameter to specify the "hours" portion of the schedule's interval.

`[-minutes <integer>] - Minutes`

Use this parameter to specify the "minutes" portion of the schedule's interval.

`[-seconds <integer>] - Seconds`

Use this parameter to specify the "seconds" portion of the schedule's interval.

Examples

The following example creates an interval schedule named `rollingdaily` that runs six hours after the completion of the previous occurrence of the job:

```
cluster1::> job schedule interval create -name rollingdaily -hours 6
```

Related Links

- [job schedule interval show](#)

job schedule interval delete

Delete an interval schedule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule interval delete` command deletes an interval schedule. Use the [job schedule interval show](#) command to display all current interval schedules.

You cannot delete interval schedules that are currently being run. Use the [job schedule show-jobs](#) command to display jobs by schedule.

Parameters

[`-cluster <Cluster name>`] - Cluster

This parameter specifies the name of the cluster on which you want to delete an interval schedule. By default, the schedule is deleted from the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

`-name <text>` - Name

Use this parameter with the name of an existing interval schedule to specify the interval schedule you want to delete.

Examples

The following example deletes an interval schedule named `rollingdaily`:

```
cluster1::> job schedule interval delete -name rollingdaily
```

Related Links

- [job schedule interval show](#)
- [job schedule show-jobs](#)

job schedule interval modify

Modify an interval schedule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule interval modify` command modifies an interval schedule. An interval schedule runs jobs at a specified interval after the previous job finishes. Use the [job schedule interval show](#) command to display all current interval schedules. See the documentation of [job schedule interval show](#) for more

information on how interval schedules work.

Modifying one parameter of a schedule's interval does not affect the other parameters. For example, if a schedule's interval is 1 day 12 hours, and you modify the "hours" parameter to 16, the schedule's new interval is 1 day 16 hours. To clear a parameter of the schedule's interval, you must explicitly set that parameter to "0" or "-".

Parameters

[-cluster <Cluster name>] - Cluster

Use this parameter to specify the cluster of an existing interval schedule you want to modify. The local cluster is provided as the default value. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

-name <text> - Name

Use this parameter with the name of an existing interval schedule to specify the interval schedule you want to modify.

[-days <integer>] - Days

Use this parameter to specify a different "days" value for the schedule's interval.

[-hours <integer>] - Hours

Use this parameter to specify a different "hours" value for the schedule's interval.

[-minutes <integer>] - Minutes

Use this parameter to specify a different "minutes" value for the schedule's interval.

[-seconds <integer>] - Seconds

Use this parameter to specify a different "seconds" value for the schedule's interval.

Examples

The following example sets the schedule named rollingdaily to run every eight hours:

```
cluster1::> job schedule interval modify -name rollingdaily -hours 8
```

Related Links

- [job schedule interval show](#)

job schedule interval show

Show interval schedules

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `job schedule interval show` command displays information about interval schedules.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-cluster <Cluster name>] - Cluster

Selects the interval schedules that match this parameter value.

[-name <text>] - Name

Selects the interval schedules that match this parameter value.

[-days <integer>] - Days

Selects the interval schedules that match the day value or range of values you specify.

[-hours <integer>] - Hours

Selects the interval schedules that match the hour value or range of values you specify.

[-minutes <integer>] - Minutes

Selects the interval schedules that match the minute value or range of values you specify.

[-seconds <integer>] - Seconds

Selects the interval schedules that match the second value or range of values you specify.

[-description <text>] - Description

Selects the interval schedules that match the description you specify.

Examples

The following example displays information about all interval schedules:

```
cluster1::> job schedule interval show
Cluster      Name      Description
-----
cluster1
              rollingdaily
              Every 8h
```

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.