

Module 9: Performing Cloning and Staging

Upon completion of this module, you should be able to:

- Describe the cloning process
- Configure automatic save set cloning
- Perform manual save set and volume cloning
- Perform scheduled clone operations
- Describe save set staging
- Perform manual staging of a save set
- Configure a NetWorker stage resource



This module focuses on how to configure NetWorker resources to perform automatic cloning and staging. You also learn how to perform manual cloning and staging.

Module 9: Performing Cloning and Staging

Lesson 1 : Performing Cloning

During this lesson the following topics are covered:

- Difference between cloning and staging
- Configuring automatic cloning
- Performing manual clone operations
- Perform scheduled clone operations
- Clone-controlled replication

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 2

This lesson covers the procedures for performing cloning in the NetWorker environment including configuring automatic, manual, and scheduled clone operations. Additionally, the process for configuring clone-controlled replication is discussed.

Cloning and Staging Save Sets

Task	Description	Method
Cloning	Copies save sets or a volume of save sets to another volume belonging to a clone pool	<code>nsrclone</code> NetWorker Administration
Staging	Moves save sets to another volume	<code>nsrstage</code> NetWorker Administration

Detailed information about NetWorker cloning can be found in the *NetWorker Cloning Integration Guide*.

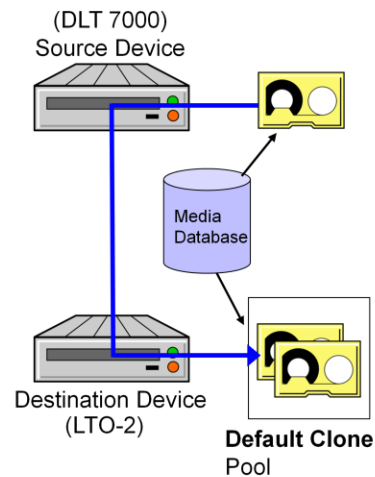


NetWorker provides the ability to further manage and protect save sets and volumes through the use of cloning and staging.

The basic description of these features and methods of performing the tasks are shown in the slide.

Cloning Overview

- Cloning creates copies of save sets, to enhance data protection.
- Cloning can be done either manually or automatically.
- Two devices are required to clone a save set.
- A save set can be cloned to a different type of media.
- Cloning of a save set may require multiple source or destination volumes.
- No volume may contain more than one instance of a save set.
- The destination volume(s) must belong to a clone pool.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 4

Cloning allows you to create identical copies of save sets to be used in case of damage to the original media or for offsite storage.

You can clone save sets either manually or automatically. To clone save sets manually, use the `nsrclone` CLI or NetWorker Administration. To clone save sets automatically, set the **Clones** attribute in a group resource to **Yes**. If cloning is enabled, all save sets in the group are cloned. Additionally, a clone job can be created to automatically clone savesets after they have been backed up using specified criteria.

Cloning involves simultaneously reading from a source volume and writing to a destination volume. Therefore two devices are required. These devices need not be of the same type. For example, you can clone from a DLT 7000 to an LTO-2 or from a file type device to an AIT-3.

Save sets are always completely cloned. Thus, if a save set begins on one volume and continues (spans) onto one or more additional volumes, each of the source volumes will be mounted and read during the clone operation. Conversely, if the destination volume becomes full during a clone operation, another volume from the same pool must be made available for the cloning to continue.

Concurrent clone, backup, and recovery operations can be performed on the same device at the same time when using advanced file type or Data Domain devices.

No volume may contain more than one instance (copy) of a save set. This eliminates the possibility of losing multiple instances of a save set if a single volume becomes damaged.

Since backup data cannot be mixed with clone data on a volume, it is required that the destination volume belong to a clone pool.

<https://t.me/learningnets>

Enabling Automatic Cloning Group Resource

The screenshot shows the 'Group Properties - Default' dialog box with the 'Setup' tab selected. The 'Clones' section is highlighted with a red box. A callout box with a red border and a blue arrow points to the 'Clones' checkbox and the 'Clone pool' dropdown menu. The callout text reads: 'Check the box and select the clone pool'. The 'Clones' section includes a checked 'Clones' checkbox, a 'Clone mode' dropdown menu set to 'Start on group completion', and a 'Clone pool' dropdown menu set to 'Default Clone'. Other sections visible include 'Identity', 'Setup', 'Status', 'Output', and 'Snapshots'.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 5

Automatic cloning is performed during backup operations by enabling the **Clones** attribute of a NetWorker group resource. The **Clone pool** attribute is used to select a destination pool other than **Default Clone**.

The **Clone Mode** option is discussed on the next slide.

When `savegrp` is executed and cloning is enabled, save sets are backed up and cloned in the following order:

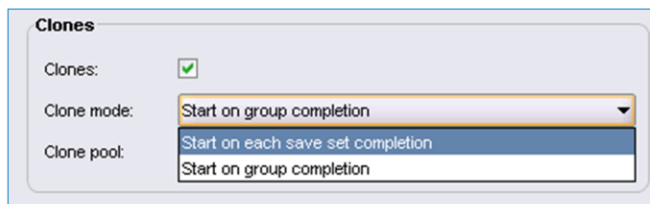
1. All save sets, with the exception of the bootstrap, are backed up.
2. All save sets that were backed up are cloned to a volume from the specified clone pool.
3. The bootstrap save set is backed up.
4. The bootstrap save set is cloned.

Note: Although client data may be sent to several pools during the backup, all cloned save sets are directed to a single clone pool.

Automatic Group Cloning Mode

Cloning Mode in NetWorker 8.1 and higher – 2 Options:

- ▶ Start on Save Set completion
 - ▶ Allows for faster time to DR readiness by beginning cloning as soon as save set is complete
- ▶ Start on Save group completion
 - ▶ Traditional automatic cloning mode
 - ▶ Starts cloning when group is completed



Clones

Clones:

Clone mode: Start on group completion

Clone pool: Start on each save set completion
Start on group completion

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 6

Immediate cloning was introduced in NetWorker 8.1. It enables cloning to begin as soon as a save set as part of a group has finished being backed up. This means that the cloning process can begin almost concurrently with the backup process, speeding up the time to DR readiness by up to 90%.

Previous versions of NetWorker required that the entire group finish its backup before a clone operation could be initiated. These groups can often be very large, therefore, immediate cloning helps support Recovery Point Objectives. This enables a faster time to Disaster Recovery readiness.

The default option with NetWorker 8.1 and higher is **Start on each save set completion**.

Cloning a Save Set: nsrclone

`nsrclone [-options] -S ssid ...`

```

C:\>nsrclone -v -S 2496202421
Obtaining media database information on server leg1-win2
80470:nsrclone: Following volumes are needed for cloning
80471:nsrclone: 11PM.001.RO <Regular>
5874:nsrclone: Automatically copying save sets(s) to other volume(s)
79634:nsrclone:
Starting Regular cloning operation...
6217:nsrclone: ...from storage node: leg1-win2
81542:nsrclone: Successfully cloned all requested Regular save sets (with new cloneid)
79629:nsrclone: Clones were written to the following volume(s) for Regular save sets:
DFE03M03

C:\>nsrclone -v -S 2496202421 -r "volume,ssid,name<20>,cloneid,client"
volume      ssid      name      clone id  client
11PM.001    2496202421 C:\WUTemp 1305020085 leg1-win2
11PM.001    1976109659 C:\WUTemp 1305021019 leg1-win2
11PM.001    1909089493 C:\WUTemp 1305109717 leg1-win2
11PM.001.RO 2496202421 C:\WUTemp 1305020084 leg1-win2
11PM.001.RO 1976109659 C:\WUTemp 1305021018 leg1-win2
11PM.001.RO 1909089493 C:\WUTemp 1305109716 leg1-win2
DFE02400    1305194401 C:\WUTemp 1305194401 leg1-win2
DFE03M01    751621286 C:\WUTemp 1305269414 leg1-win2
DFE03M03    2496202421 C:\WUTemp 1305444610 leg1-win2
    
```



The `nsrclone` command is the CLI used to clone save sets and volumes.

When the `-S` option is used, a list of save set IDs must be specified. If the `-S` option is not used, arguments following any options must be NetWorker volume names.

`nsrclone(1m)` syntax:

`nsrclone [options] -S ssid ... | volume ...`

where `ssid` is a save set to clone; `volume` is a volume containing save sets to clone. Note that `ssid/cloneid` may also be used to specify which save set with multiple copies to use as a source. Additional information including a full list of the command options can be found in the *NetWorker Command Reference Guide*, or the *NetWorker Cloning Integration Guide*.

Option	Description
<code>-b pool</code>	Write clone data to <code>pool</code> , which must be a backup clone pool.
<code>-f file</code>	Read <code>file</code> , which contains the list of SSIDs (<code>-S</code>) or volumes to clone. If <code>file</code> is '-', the list is read from standard input (STDIN).
<code>-v</code>	Generate verbose output. By default, only the completion of the operation is reported.

Table 9-1: nsrclone Options

Cloning a Volume : nsrclone

`nsrclone [-options] volume ...`

```

C:\>mminfo -av DFE03M03
volume      type      client      date      time      size      ssid      fl  lvl name
DFE03M03    LTO      Ultrium    leg1-sun2 5/10/2011 5:30:51 AM 84 MB 2512979365 cb full /var
DFE03M03    LTO      Ultrium    leg1-win2 4/26/2011 9:40:25 AM 5185 MB 4206284334 ca manual C:\

C:\>nsrclone -v DFE03M03
Obtaining media database information on server leg1-win2
80470:nsrclone: Following volumes are needed for cloning
80471:nsrclone:      DFE03M03 (Regular)
5874:nsrclone: Automatically copying save sets(s) to other volume(s)
79634:nsrclone:
Starting Regular cloning operation...
6217:nsrclone:   ...from storage node: Leg1-sun2
81542:nsrclone: Successfully cloned all requested Regular save sets (with new cloneid)
                2512979365/1305441372
79629:nsrclone: Clones were written to the following volume(s) for Regular save sets:
                DFE03M04

C:\>mminfo -av DFE03M04
volume      type      client      date      time      size      ssid      fl  lvl name
DFE03M04    LTO      Ultrium    leg1-sun2 5/10/2011 5:30:51 AM 84 MB 2512979365 cb full /var
    
```

List save sets on the volume

Clone the volume

Verify all save sets were cloned to a clone volume

When cloning a volume, it is not a byte-by-byte copy. Only save sets that begin on the volume are cloned. If a save set begins on the volume and spans to one or more additional volumes, each of those volumes will be mounted and read. Thus, to clone a volume really means to clone, in their entirety, all save sets beginning on the volume.

Multiple volumes can be specified on the command-line. The `-f` option of the `nsrclone` command can be used to specify a file (or standard input) containing a list of volumes to clone. When using an input file, each volume must be on a line by itself.

Note: The first flag associated with a save set indicates which part of the save set is stored on a volume. This flag can be displayed with the `mminfo -v` command and is also displayed when viewing the save sets for a volume in the **Volume Save Sets** window in NetWorker Administration **Media**. Values for the first flag are:

- **c:** save set is completely contained on this volume.
- **h:** save set spans volumes and the head is contained on this volume.
- **m:** save set spans volumes and a middle section is contained on this volume.
- **t:** the tail section of a spanning save set is contained on this volume.

Volume Save Sets			
Volume: nwwindows.emc.edu.002			
Client	Save Set ▲	SSID	Flags
nwwindows....	index:winclient.emc.edu	3815460117	cr
nwwindows....	index:nwwindows.emc.edu	3798682985	cr
winclient.em...	C:\WUTemp\Java	3865791730	cb
winclient.em...	C:\WUTemp\inf	3882568943	cb

<https://t.me/learningnets>

Manually Cloning a Save Set : NetWorker Admin

1. Specify selection criteria.

2. Perform query by changing tabs.

3. Select save sets to clone.

4. Right-click

5. Select pool if desired.

Client	Save Set	SSID	Clone ID	Level	Status	Type	Media Type	Clone Retention Time	Checkpoint ID
leg1-win2	C:\MUTemp	1305194401	1305108717	incr	browsable	dir		5/1/12 11:58:59 PM	
leg1-win2	C:\MUTemp	1305194401	1305108716	incr	browsable	dir		5/1/12 11:58:59 PM	
leg1-win2	C:\MUTemp	1305194401				LTO Ultrium DFEB03A01 (S.D., 11PM	718	4	5/1/12 5:59:59 AM 5/1/12 11:58:59 PM
leg1-win2	C:\MUTemp	751621				LTO Ultrium DFEB03A01 (S.D., Def.	418	0	5/1/12 2:50:12 AM 5/1/12 11:58:59 PM
leg1-win2	C:\MUTemp	751621				LTO Ultrium DFEB03A01 (S.D., 11PM	418	0	5/1/12 2:50:12 AM 5/1/12 11:58:59 PM

When using NetWorker Administration to perform save set cloning, click **Save Sets** in the left pane of the *Media* window. In the right pane, specify the save set characteristics of those save sets you want to clone. Change to the *Save Set List* tab to perform the query and report matching save sets.

After specifying save set characteristics and performing the query, select any or all of the matching save sets, right-click any selected save set to display a drop-down menu, and select **Clone**. A pop-up window is displayed from which you can specify the destination pool. Upon clicking **OK**, a final verification window is displayed before cloning begins.

Manually Cloning a Volume : NetWorker Admin

The screenshot shows the NetWorker Admin interface. On the left, the 'Tape Volumes' folder is selected under 'Leg1-win2'. The main window displays a table of tape volumes:

Volume Name	Barcode	Used	% Used	Mode	Expiration	Pool	Location
DFE00L00	DFE00L00	8113 MB	8%	appendable	manual	Default	Moved to 2nd Shelf of
DFE00L01	DFE00L01	7808 KB	0.1%	appendable	5/13/12	11PM	STK@2.0.0
DFE02400	DFE02400	4930 MB	5%	appendable	5/13/12	11PM	rd-Leg1-sun2:STK@2.5.3
DFE02401	DFE02401			appendable	5/11/12	Default Clone	rd-Leg1-sun2:STK@2.5.3
DFE0241				appendable	5/12/12	Default	rd-Leg1-sun2:STK@2.5.3
DFE024				appendable	5/11/12	Default	rd-Leg1-sun2:STK@2.5.3
DFE03M				appendable		Cloud Clone	
FE03M				appendable	5/13/12	11PM	
FE03M				appendable	5/12/12	Default Clone	
FE03M				appendable	5/12/12	Default Clone	

Three callouts indicate the steps:

1. Select volume to clone. (Points to the 'Tape Volumes' folder in the sidebar)
2. Right-click. (Points to the context menu for a selected volume)
3. Select pool if desired. (Points to the 'Target Clone Media Pool' dropdown in the 'Clone' dialog box)

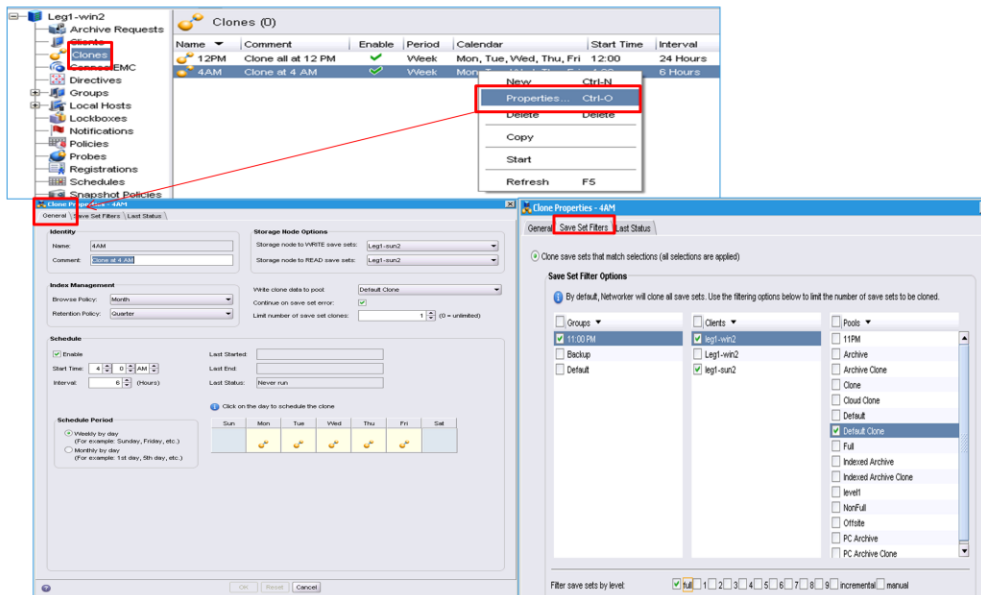
The 'Clone' dialog box shows 'Source volumes to clone:' and 'Target Clone Media Pool: Automatic Selection'. The 'Question' dialog box asks: 'Would you like to clone selected volume(s) to the media pool: Automatic Selection?' with 'Yes' and 'No' buttons.

EMC²
Copyright © 2013 EMC Corporation. All Rights Reserved. Module 9: Performing Cloning and Staging 10

Cloning of a volume is done in the **Media > Volumes** window. Selecting one or more source volumes and right-clicking any of them produces a drop-down menu from which you select **Clone**. A pop-up window is displayed that allows you to choose the backup clone pool that the destination volume(s) must belong to.

As previously mentioned, cloning a volume clones all save sets beginning on the volume and may require mounting additional source volumes to read save sets which continue on other volumes.

Scheduling Clone Operations



Copyright © 2013 EMC Corporation. All Rights Reserved.

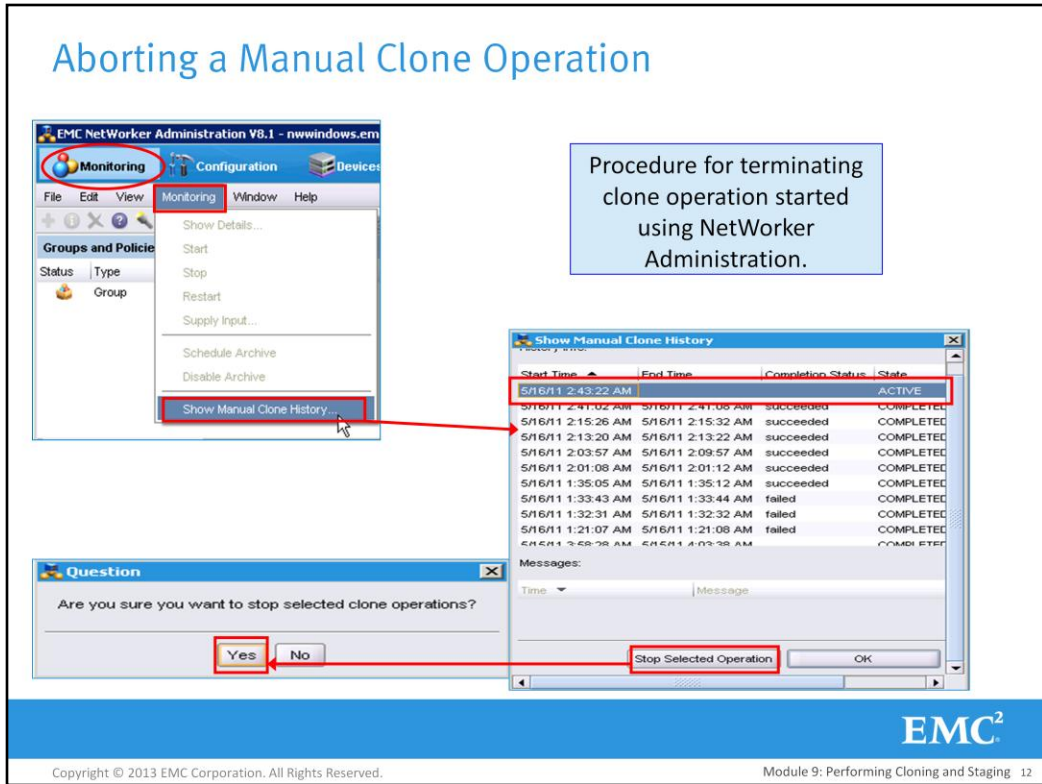
Module 9: Performing Cloning and Staging 11

Cloning operations can be scheduled to run automatically, independent of backup operations. Setting up a scheduled clone operation is done from the **Configuration** window by creating a **clone** resource. The clone operation can be scheduled to occur at a specified time either on one or more days of the week or on one or more days of the month. Select the storage node and pool to be used to write the clone save sets.

The data to be cloned is specified using the options on the **Save Set Filters** tab including filtering by group, client, pool, level and/or save set name. Clicking **Preview Save Set Selection** brings up a box listing the save sets to be cloned. You can view the status of scheduled clone jobs from the **Monitoring** window. A scheduled clone job can also be manually started at any time without affecting the regularly scheduled start time.

<https://t.me/learningnets>

Aborting a Manual Clone Operation



The slide shows the procedure to follow to abort a manual cloning of either save sets or volumes using NetWorker Administration.

In the slide, the **Clone Sessions** screen in the **Monitoring** window shows a clone operation is in progress. If you want to stop the clone operation, you could abort the clone operation by using the **Monitoring** menu, selecting **Show Manual Clone History**, selecting the active clone operation in the pop-up window, and clicking the **Stop Selected Operation** button.

You can also right-click a clone operation from the *Clone Sessions* screen in the *Monitoring* window and select **Stop Manual Clone Operation** to abort the operation.

Cloning Save Sets Across Storage Nodes

The screenshot shows the NetWorker Client Properties window for 'leg1-win2'. The 'Configuration' tab is active, showing the 'Clone storage nodes' attribute set to 'Leg1-Sun2'. A red box highlights this attribute. A red arrow points from a text box to the 'Clone storage nodes' attribute.

In this example, when Leg1-win2 supports the read side of a clone operation, the clone copy is written to a volume on storage node Leg1-sun2.

The diagram below illustrates the setup: a Server (Leg1-win2) is connected to a Storage Node (Leg1-sun2). A blue arrow indicates data flow from the server to the storage node.

EMC²
 Copyright © 2013 EMC Corporation. All Rights Reserved. Module 9: Performing Cloning and Staging 13

When doing cloning, the storage nodes to be used for the clone read operation (of the source volume) and the write operation (to the destination volume) are based on the following prioritized criteria. For reading the clone data:

1. If the volume to be read is mounted, the storage node of that device is used.
2. Storage nodes in the **Recover storage nodes** attribute of the NetWorker server's client resource having access to the volume.
3. Storage nodes in the **Storage nodes** attribute of the NetWorker server having access to the volume.
4. Storage node in the **Read hostname** attribute of the jukebox resource, or if this is empty, storage nodes on which a device in the library is configured.

For writing the clone data, the **Clone storage nodes** attribute of the client resource of the storage node performing the read operation is used. Next in priority is the **Clone storage nodes** attribute of the client resource of the NetWorker server. Otherwise, the **Storage nodes** attribute of the NetWorker server is used.

Each of the above attributes may contain multiple storage node host names, which are also listed by priority. If a storage node is not available, the next storage node on the list is contacted. For more information on cloning and storage node selection, see the *EMC NetWorker Administration Guide*.

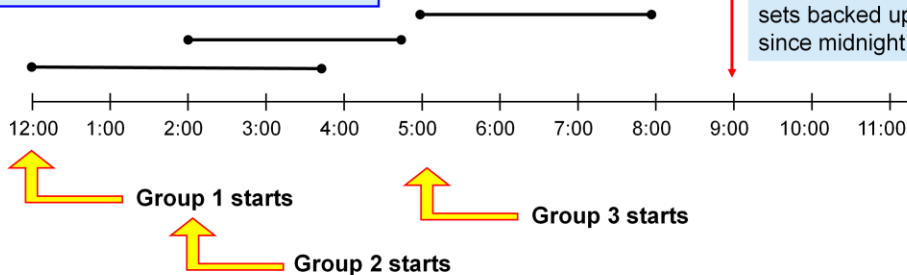
nsrclone : Using the -t and -e options

```
nsrclone -S -t start_time -e end_time
```

- Save sets backed up within specified time range are cloned.
- Must be used with -S option
- Either or both of -t and -e options may be used
- -c, -C, -g, -l options extend save set selection capabilities

```
nsrclone -S -t today
```

Clone all save sets backed up since midnight



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 14

The `-t start_time` option causes `nsrclone` to automatically determine which save sets have been backed up since `start_time` (based upon the *savetime* value) and clone them. `start_time` can be specified using any `nsr_getdate(3)` format. By default, all save sets backed up since `start_time` are cloned. To specify a time range, the `-e end_time` option can be used to specify the end time of the range. If `-e end_time` is used, the default value of `start_time` is "`end_time - 24 hours`".

Options `-c client_name`, `-C less_than_copies_in_pool`, `-g group_name`, `-l level` can be used with the `-t` or `-e` option to extend save set selection capabilities. Also, `-N saveset_name` allows for selection on save set name.

Examples

- Clone all save sets backed up since 1:00 a.m. this morning:

```
nsrclone -S -t "01:00"
```
- Clone all save sets backed up in the last 24 hours with backup level full and group Default:

```
nsrclone -S -e now -l full -g Default ("now" is a valid nsr_getdate format)
```
- Clone all save sets backed up between 9:00 p.m. yesterday and 8:00 a.m. this morning:

```
nsrclone -S -t "yesterday 21:00" -e "08:00"
```

Note: See the *NetWorker Command Reference Guide* for more information on command options as well as additional valid time formats.

<https://t.me/learningnets>

Clone Browse and Retention

- Browse and retention policy for clone data can be specified independently:
 - ▶ Clone data can be browsed and recovered after the browse and retention policies of the original save set have expired.
 - ▶ Backup (non-clone) volumes can be recycled more frequently than the clone volumes.
- To specify the clone retention time:
 - ▶ Run `nsrclone` with the `-y` option
 - ▶ Run `nsrmm` with the `-e` option
 - ▶ Configure a retention policy for a clone pool
 - ▶ Set in a clone resource
- To specify the clone browse time:
 - ▶ Run `nsrclone` with the `-w` option
 - ▶ Set in a clone resource



Each instance of a save set has its own clone browse and retention time which is tracked in the save set record of the media database. Browse and retention times for clone data can be extended beyond that of the original save set, enabling browsing and recovery of clone data after the original save sets have expired.

If the clone instance is written to a pool having a retention policy, the retention time of that save set instance is determined by the pool's retention policy instead of the client's retention policy. A different clone retention time can also be set using the `-y retent_time` option with `nsrclone` and with the `nsrmm -e` command. Setting the clone's retention to a longer period than the client's retention allows the clone to remain recoverable even after the original backup is no longer retained. Note that retention specified from the command line overrides the retention policy for the clone pool.

The browse period for a clone can be extended with the `-w` option of `nsrclone` when creating a clone save set. Note that the browse period is left unchanged if the save set's browse date is later or if the new time has already passed. This option requires the `-y` retention option and must not be greater than the retention time.

Important: The date on which a volume becomes recyclable is determined by the clone retention times of save set instances on the volume, not by the save set retention times. For example, if 10/17/13 was the longest save set retention time on a volume and the longest clone retention time on the volume was 1/1/014, the volume would not become recyclable until 1/1/14.

Clone Retention Policy Attribute – Pool Resource

The screenshot displays the 'Create Media Pool' wizard in the Veeam Backup & Replication console. The wizard is in the 'Configuration' tab. Red boxes and arrows highlight key settings:

- Media Pools List:** A 'New' button is highlighted with a red box and arrow.
- Basic Tab:** The 'Basic' tab is selected. The 'Name' field contains '11 PM'.
- Configuration Tab:** The 'Pool type' dropdown is set to 'Backup Clone'. The 'Store index entries' checkbox is checked. The 'Retention policy' dropdown is set to 'Decade'.

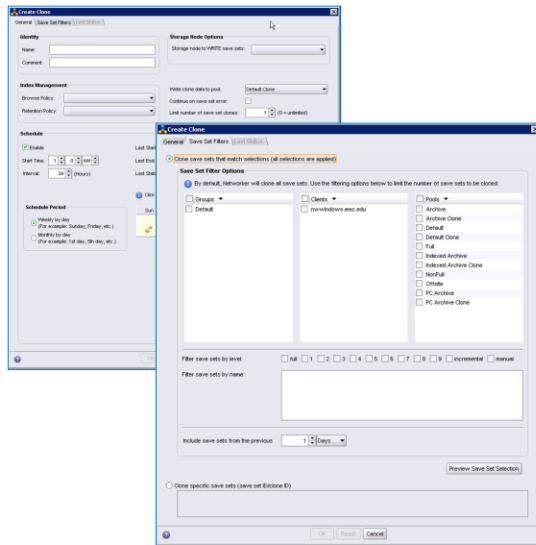
Callout boxes provide instructions:

- 'Select pool type' points to the 'Pool type' dropdown.
- 'Deselect **Store index entries** if configuring a Backup Clone pool' points to the 'Store index entries' checkbox.
- 'Select pool retention policy' points to the 'Retention policy' dropdown.

EMC²
Copyright © 2013 EMC Corporation. All Rights Reserved. Module 9: Performing Cloning and Staging 16

The slide shows how to set the **Retention policy** attribute in the pool resource. When creating a backup clone pool, it is necessary to deselect the **Store index entries** attribute. This is because duplicate CFI entries cannot be created during a clone operation.

Clone Controlled Replication



- NetWorker managed cloning leveraging Data Domain replication
- Cloning between Data Domain devices
- Managed and tracked in NetWorker indexes
- Transfers only deduplicated and compressed data
- Transparent during NetWorker cloning
- NetWorker storage node is not in the replication path

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 17


When cloning is taking place from one Data Domain system to another Data Domain system with NetWorker 7.6 SP1 and higher, NetWorker is able to detect that both the clone source and destination are Data Domain systems and automatically initiate Data Domain file copy replication using the DD Boost interface. Data Domain managed file replication is performed transparently during network cloning, this is known as clone controlled replication.

A benefit of Data Domain managed file replication is that only unique, compressed data is replicated between the two Data Domain systems, which can reduce the replication workload by 90% – 99%.

Even though data transmitted via clone controlled replication uses Data Domain replication to transmit the data, it is still tracked within the NetWorker databases as if it were a traditional clone job.

<https://t.me/learningnets>

Cloning to a Cloud



- Placing backup data on a cloud provides offsite backup storage without the complexity of managing tapes.
- Use NetWorker cloning to clone backup data to a cloud device.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.


Module 9: Performing Cloning and Staging 18

With the NetWorker Cloud Backup Option, copies of backup data can be stored on internet-based storage as an alternative to sending tapes offsite. This provides a tape-less offsite storage solution, eliminating the complex requirements of managing tapes.

Cloning backup data to a cloud complements backing up to disk. In the example shown on the slide, backups are first written to disk. Then, the backup data is cloned to a volume on an Atmos cloud device. The original backup data is retained on disk only as long as required for short term recovery operations; data on cloud storage is retained for a longer period of time according to business requirements for long term/offsite storage.

Clone Reporting within the NMC

- Reporting on clone processes can be performed within the NetWorker Management Console **Reports** task pane.
- Four different types of clone reports can be generated:
 - ▶ Server Summary
 - ▶ Clone Details
 - ▶ Save Set Details
 - ▶ Clone Summary over Time
- Clone records are gathered by the Console server every 12 hours.
- Reports include basic and drill down reports.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 19

Reporting on clone jobs can be achieved through the clone reports available in the NetWorker Management Console **Reports** task pane. These reports allow you to view the history of automatic and scheduled clone operations performed by the NetWorker servers for any server version 7.6 SP2 and later.

Four different types of clone reports can be generated: server summary, clone details, save set details, and clone summary over time.

Be aware that clone reports may not be up-to-date because clone records are gathered by the Console server every 12 hours.

The NetWorker server report category includes basic and drill down reports for each selected NetWorker server within the enterprise. The configuration tab allows you to limit the scope of the report that was selected.

For further information on the clone reports available within NetWorker Management Console, refer to the *EMC NetWorker Administration Guide*.

Lab Exercise 9-1: Configure a Scheduled Clone Operation



This lab covers procedures for creating scheduled clone jobs within EMC NetWorker.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 20

In this lab, you will:

- Configure a clone resource
- Run the scheduled clone

<https://t.me/learningnets>

Module 9: Performing Cloning and Staging

Lesson 1 : Summary

During this lesson the following topics were covered:

- Difference between cloning and staging
- Configuring automatic cloning
- Performing manual clone operations
- Perform scheduled clone operations
- Clone-controlled replication

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 21

This lesson covered the procedures for performing cloning in the NetWorker environment including configuring automatic, manual, and scheduled clone operations. Additionally, the process for configuring clone-controlled replication is discussed.

Module 9: Performing Cloning and Staging

Lesson 2 : Performing NetWorker Staging

During this lesson the following topics are covered:

- Configuring automatic staging
- Configuring manual staging

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

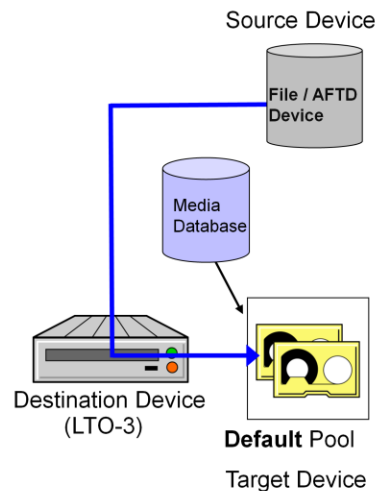
Module 9: Performing Cloning and Staging

22

This lesson covers the procedures for configuring automatic and manual staging of data in NetWorker.

Staging Save Sets – Overview

- Staging moves a save set from one volume to another.
- Two devices are required to perform staging.
- A save set can be staged to a different type of media.
- Multiple source or destination volumes may be required to stage a save set.
- After a save set is moved, its original media database record is updated.
- The destination volume can belong to any pool.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 23

Staging a save set moves it from one storage volume to another.

Like cloning, staging requires two devices, one or more source volumes, and one or more destination volumes.

When a save set is staged, it is actually cloned, resulting in an additional instance (copy) of the save set being tracked in the media database save set record. Upon successful completion of the clone operation, the information pertaining to the original instance (copy) of the save set is removed from the save set record.

If the save set being staged is on tape, it remains on the tape until the tape is relabeled. If the save set being staged is on a file or adv_file type device, it is immediately deleted from the device/volume (directory).

Unlike cloning, destination volumes do not have to belong to a clone pool.

Staging is often used to move save sets from file and adv_file devices to long term media such as tape. This allows the most recent backups to be written to and recovered from disk, then moved to tape to free space for subsequent backups. Staging is also used to remove non-recyclable save sets from an otherwise recyclable volume.

Staging Save Sets : nsrstage

```
nsrstage -m [ -options ] -S ssid[/cloneid ] ...
```

```
C:\>nsrinfo -q "ssid=1909089493" -r "copies.name(10),ssid,cloneid,volume"
copies name      ssid      clone id volume
1 C:\NUTemp 1909089493 1305526690 DFE82402

C:\>nsrstage -m -h Default -S 1909089493
80470:nsrstage: Following volumes are needed for cloning
80471:nsrstage: DFE82402 (Regular)
Successfully deleted original clone 1305526690 of save set 1909089493 from media database.
Recovering space from volume 600802858 failed with the error 'Space can only be recovered from
Refer to the NetWorker log for details.
80855:nsrstage: Space can only be recovered from Data Domain, adv_file and file type devices.

C:\>nsrinfo -q "ssid=1976109659" -r "copies.name(10),ssid,cloneid,volume"
copies name      ssid      clone id volume
2 C:\NUTemp 1976109659 1305021019 11PM.001
2 C:\NUTemp 1976109659 1305021018 11PM.001.R0

C:\>nsrstage -m -h Default -S "1976109659/1305021019"
80470:nsrstage: Following volumes are needed for cloning
80471:nsrstage: 11PM.001 (Regular)
Successfully deleted original clone 1305021019 of save set 1976109659 from media database.
Successfully deleted AFD's companion clone 1305021018 of save set 1976109659 from media database with
Completed recover space operation for volume 3066625111
Refer to the NetWorker log for any failures.

C:\>nsrinfo -q "ssid=1976109659" -r "copies.name(10),ssid,cloneid,volume"
copies name      ssid      clone id volume
1 C:\NUTemp 1976109659 1305526690 DFE82402

C:\>nsrstage -m -h Default -S 1976109659
80470:nsrstage: Following volumes are needed for cloning
80471:nsrstage: DFE82402 (Regular)
Successfully deleted original clone 1305526690 of save set 1976109659 from media database.
Recovering space from volume 600802858 failed with the error 'Space can only be recovered from Data Domain, adv_file and file type devices.
Refer to the NetWorker log for details.
80855:nsrstage: Space can only be recovered from Data Domain, adv_file and file type devices.

C:\>nsrinfo -q "ssid=1976109659" -r "copies.name(10),ssid,cloneid,volume"
copies name      ssid      clone id volume
1 C:\NUTemp 1976109659 1305526692 DFE82403

C:\>_
```

Use *ssid/cloneid* to stage only copy from volume 11PM.001

Specifying only *ssid* causes all copies (clones) to be deleted except the staged copy.

nsrstage is the command line utility used to stage save sets.

nsrstage syntax:

```
nsrstage [ -options ] -m -S ssid[/cloneid ] ...
```

-m is a required option to stage (move) save sets and -S *ssid* specifies which save set(s) to stage. The optional */cloneid* is for save sets with more than one instance (copy), to identify the instance of the save set to stage. If an instance is not specified, all instances except for the staged copy are deleted from the media database. Table 9-2 lists several other nsrstage options.

Note: See the *NetWorker Command Reference Guide* for more information and specific command options..


Option	Description
-b <i>pool</i>	Stage save set to a volume in <i>pool</i> , which can be any pool.
-f <i>file</i>	Read <i>file</i> , which contains the list of SSIDs to stage. If <i>file</i> is '-', the list is read from standard input (STDIN).
-C -V <i>volume</i>	Clean. Check the file type device supporting the specified volume and delete any aborted save sets or save sets not tracked in the media database.

Table 9-2: nsrstage Options

<https://t.me/learningnets>

Staging Save Sets Automatically

- A NetWorker stage resource monitors selected file and adv_file type devices to automatically move (stage) save sets based on configured criteria.
- Allows you to specify the maximum number of days or hours a save set may reside on disk before it must be staged.
- Uses a high water mark to specify at what percentage of file system fullness staging should start, regardless of save set age.
- Continues staging save sets until a low water mark is reached.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 25

A NetWorker stage resource is used to monitor selected file and adv_file type devices and automatically stage save sets from the device's volume to other media when the volume becomes too full.

Automatic save set staging is designed to move data from file/adv_file type devices to tape. Staging allows you to perform backups to disk, potentially maximizing backup performance, and later move the save sets to tape.

Staging prevents the file/adv_file type device from becoming full by periodically checking the following:

- How long each save set has been on the file type device. Save sets are staged after a specified number of days or hours, regardless of how full the volume (file system) is.
- The percentage fullness of the file system on which the file/adv_file type device directory resides. Save sets are staged when the file system reaches a certain percentage of utilization (the *high water mark*), regardless of a save set's age. Once staging begins, it continues until the file system utilization has decreased to the specified *low water mark*.

NetWorker Stage Resource

Configuration -> Staging -> New

Water Marks:
High=% full when staging starts
Low=% full when staging stops

Save set movement order once the high water mark is reached.

Maximum amount of time save sets reside on disk

How often to check devices for save sets not in media database

How often to check for Max storage period and High Water Mark

Resource must be enabled to use

Disk type devices to monitor

Select destination pool

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved. Module 9: Performing Cloning and Staging 26

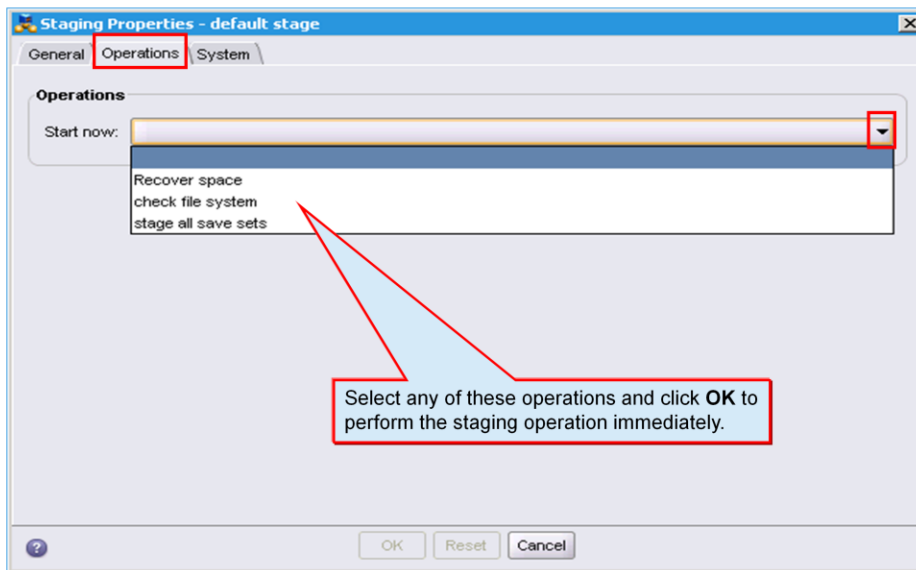
A NetWorker stage resource is used to monitor and manage selected disk type devices. There is one preconfigured stage resource, **default stage**, having the default attribute values shown in the slide.

Attribute	Description
Devices	File type device(s) to monitor.
Destination pool	Pool to which destination volumes must belong.
High Water Mark(%) Low Water Mark (%)	Percentage fullness of the file system containing the file device at which staging will begin and end.
Save set selection	Type of save set (oldest, newest, largest, smallest) that will be staged first.
Max storage period	Maximum number of hours or days that a save set can reside on a file device before it is staged.
Recover space interval	Length of time (hours or minutes) between checks to determine if save sets exist that are not tracked in the media database. These save sets are removed. Aborted save sets are also removed.
File system check interval	Length of time (hours or minutes) between checks to determine if the high water mark or max storage period has been reached.

Table 9-3: Stage Resource Attributes

<https://t.me/learningnets>

Manual Staging Using the Administration Window



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 27

The *Operations* tab of the stage resource allows you to perform manual staging. After selecting and performing any of the operations, the **Start now** attribute is returned to a null value.

Choose **Recover space** to immediately perform a recover space operation.

Select **check file system** to perform an immediate check of the fullness of the file system(s) to determine whether the high-water mark has been reached, thereby requiring automatic staging.

After selecting **stage all save sets** and clicking **OK**, all save sets residing on all devices managed by the stage resource will be staged.

Lab Exercise 9-2: Configure Automatic Staging



In this lab, you will configure automatic staging.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 28

In this lab, you will configure a stage resource to stage the data backed up to the advanced file type devices.

<https://t.me/learningnets>

Module 9: Performing Cloning and Staging

Lesson 2 : Summary

During this lesson the following topics were covered:

- Configuring automatic staging
- Configuring manual staging

The EMC logo is located in the bottom right corner of the slide. It consists of the letters 'EMC' in a bold, white, sans-serif font, with a small superscript '2' to the right of the 'C'. The logo is set against a solid blue background that spans the width of the slide.

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 29

This lesson covered the procedures for configuring automatic and manual staging of data in NetWorker.

Module 9: Summary

Key points covered in this module include:

- The description of the cloning process
- Automatic save set cloning
- Manual save set and volume cloning
- Scheduled cloning
- Save set staging
- Manual staging of a save set
- Configuring a NetWorker stage resource

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 9: Performing Cloning and Staging 30

This module covered how to configure NetWorker resources to perform automatic cloning and staging. We also looked at how to perform manual cloning and staging.

<https://t.me/learningnets>