



IBM Domino Plug-in Operations

Snap Creator Framework

NetApp
February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snap-creator-framework/domino-ops/concept_backup_operations_using_domino_plug_in.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- IBM Domino Plug-in Operations 1
 - IBM Domino plug-in overview 1
 - IBM Domino backup and recovery workflow 3
 - Preparing for IBM Domino backup and restore 4
 - Creating a backup and restore configuration 8
 - Backing up databases 19
 - Restoring databases 21
 - Viewing job status and logs 29
 - Troubleshooting Domino plug-in-specific errors 31
 - Where to go next 34

IBM Domino Plug-in Operations

You can configure and use the IBM Domino plug-in for Snap Creator 4.3.3 to back up and restore IBM Domino databases.

IBM Domino plug-in overview

The IBM Domino plug-in for the Snap Creator Framework offers a complete backup and recovery solution for Domino databases on NetApp storage. With the IBM Domino plug-in, you can back up databases efficiently and restore them as needed without taking database servers offline. The plug-in uses IBM-provided APIs to ensure application consistency.

With key NetApp data protection technologies tightly integrated in the Snap Creator Framework, you can use the IBM Domino plug-in to:

- Create application-consistent Snapshot copies on primary storage
- Replicate Snapshot copies to secondary storage for disaster recovery and archiving

Integrated technologies include Snapshot, SnapMirror, and SnapVault.

Understanding IBM Domino plug-in backup operations

Unlike other backup technologies, the IBM Domino plug-in allows data to be written to the database during backup operations. It ensures that the backup is consistent with application data by recording any changes to the database during backup operations in a temporary directory called changeinfo. At restore time, it applies these changes to the database after restoring data from the Snapshot copy.

During backup operations, database files are said to be in backup mode. Although the Snap Creator Framework refers to these files as “quiesced,” that is not the case. Data continues to be written to the files while the Snapshot copy is being made.

High-level steps are as follows:

1. List databases, templates, and mailboxes in the volumes to be backed up.
2. Put files in backup mode, one database at a time.
3. Check for inconsistent or corrupted databases.



You can force the plug-in to proceed with backup operations despite errors, as described in step [task_using_the_gui_to_create_a_configuration_file.md#STEP_AA41331683A24598B7845367CB967F99](#).

4. Make a Snapshot copy of each volume.
5. If Domino transaction logging is in use, archive copies of Domino transaction logs for use in up-to-the-minute restore operations.
6. Take files out of backup mode, one database at a time.

7. Record changes to the database since files were put in backup mode in the changeinfo directory.

Understanding IBM Domino plug-in restore operations

The IBM Domino plug-in restores database files from the Snapshot copy of the volume on which the database is stored. When the Snapshot restore operation is complete, the plug-in applies any changes recorded in the changeinfo directory.

High-level steps are as follows:

1. Restore database files from Snapshot copies.
2. Apply changes recorded in changeinfo directory.
3. Apply transaction log information, if available.

Restore options

The plug-in offers a wide variety of restore options:

Restore type	Description
Volume restore (point-in-time)	Restores the entire volume to its state at the time of the Snapshot copy.
Volume restore (up-to-the-minute)	Restores the entire volume to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to the most recent copy.
Volume restore (selectable up-to-the-minute)	Restores the entire volume to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to a specified time.
Single-file restore (point-in-time)	Restores a single file to its state at the time of the Snapshot copy (NFS only).
Single-file restore (up-to-the-minute)	Restores a single file to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to the most recent copy (NFS only).
Single-file restore (selectable up-to-the-minute)	Restores a single file to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to a specified time (NFS only).

Note: For up-to-the-minute and selectable up-to-the-minute restores, Domino transaction logging must be enabled.


Transaction logs can only be played forward. Selecting a time before the Snapshot copy was created causes a restore error.

IBM Domino backup and recovery workflow

Before you can create backups with the IBM Domino plug-in, you need to install the Snap Creator Server and Agent software and provision NetApp storage. If you plan to replicate Snapshot copies to secondary storage for disaster recovery and archiving, you need to set up SnapMirror and SnapVault relationships.



Especially in a SAN environment, you might want to use SnapDrive on the Snap Creator Agent host to back up databases or to mount Snapshot copies for single-file restores. For more information, see [Adding commands to the backup and restore configuration](#) and [Performing a single-file restore in a SAN environment](#).



Preparing for IBM Domino backup and restore

Before you deploy the IBM Domino plug-in, make sure that your storage system and hosts meet minimum resource requirements. You also need to configure storage system layouts for databases, and optionally set up SnapMirror and SnapVault relationships.

For Snap Creator Server and Agent installation requirements, see the [Snap Creator Framework 4.1.2 Installation Guide](#). Pay particular attention to the IBM Domino preinstallation requirements for the Agent host:

- On UNIX hosts, you must create symbolic links to IBM Domino shared object files.
- On Windows hosts, you must add the IBM Domino installation path to the PATH environment variable.

Storage layout requirements

A typical IBM Domino environment has at least three Domino volumes, one each for Domino data, Domino transaction logs, and the plug-in changeinfo directory. Many sites also have volumes for Domino DAOS and for view rebuilds.

The IBM Domino plug-in uses the changeinfo directory for changes recorded during backup operations and for copies of transaction logs used in up-to-the-minute restore operations. It is a best practice to store the changeinfo directory on a separate volume, to avoid inadvertently overwriting the information and to make it easier to back up.

You may also find it useful to have separate volumes for Domino DAOS (if it is enabled) and for view rebuilds. When Domino rebuilds a view (for example, when a user opens a view whose index has been deleted or when updll --R is run), it may generate temporary files to sort the data for rapid view rebuilding.

By default, these temporary files are located in the system's temporary folder or in the Domino data folder. IBM recommends changing the location of the temporary files to a different drive to distribute disk I/O and to ensure adequate space to rebuild views. To change the temporary folder used for view rebuilds, add the View_Rebuild_Dir setting to the notes.ini file.

The following table shows the preferred volume layout:

Volume	Contents	Notes
Volume 1	Domino data	FC, SAS, or SSD drives preferred.
Volume 2	Domino transaction logs	FC, SAS, or SSD drives preferred.
Volume 3	changeinfo	Stores changes recorded during backup operations and copies of transaction logs for use in up-to-the-minute restore operations.
Volume 4	View rebuild	Optional. Stores temp files created during index updates. Can use RAM disk. Add View_Rebuild_Dir setting to notes.ini file.

Volume	Contents	Notes
Volume 5	DAOS repository	Optional. Contains .dlo files from DAOS. Low I/O requirements make this a good candidate for SATA drives.



In virtual environments, guest-mounted disks are preferred.

SnapMirror and SnapVault setup

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. SnapVault is archiving technology, designed for disk-to-disk Snapshot copy replication for standards compliance and other governance-related purposes.

Before you can use Snap Creator with these products, you need to configure a data-protection relationship between the source and destination volumes, then initialize the relationship.



The procedures in this section describe how to set up replication relationships in clustered Data ONTAP. You can find information about setting up these relationships in Data ONTAP operating in 7-Mode in the .

Preparing storage systems for SnapMirror replication

Before you can use to mirror Snapshot copies, you need to configure a data-protection relationship between the source and destination volumes, then initialize the relationship. Upon initialization, SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks that it references to the destination volume. It also transfers any other, less recent Snapshot copies on the source volume to the destination volume.

- You must be a cluster administrator.
- For Snapshot copy verification on the destination volume, the source and destination Storage Virtual Machines (SVMs) must have a management LIF as well as a data LIF.

The management LIF must have the same DNS name as the SVM. Set the management LIF role to data, the protocol to none, and the firewall policy to mgmt.

You can use the Data ONTAP command-line interface (CLI) or OnCommand System Manager to create a SnapMirror relationship. The following procedure documents CLI usage.



If you are storing database files and transaction logs on different volumes, you must create relationships between the source and destination volumes for the database files and between the source and destination volumes for the transaction logs.

The following illustration shows the procedure for initializing a SnapMirror relationship:

1. Identify the destination cluster.
2. On the destination cluster, use the volume create command with the -typeDP option to create a SnapMirror destination volume that is either the same or greater in size than the source volume.



The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2 GB destination volume named dstvolB in SVM2 on the aggregate node01_aggr:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate  
node01_aggr -type DP  
-size 2GB
```

3. On the destination SVM, use the snapmirror create command with the -type DP parameter to create a SnapMirror relationship.

The DP type defines the relationship as a SnapMirror relationship.

The following command creates a SnapMirror relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2, and assigns the default SnapMirror policy DPDefault:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path  
SVM2:dstvolB  
-type DP
```



Do not define a mirror schedule for the SnapMirror relationship. does that for you when you create a backup schedule.

If you do not want to use the default SnapMirror policy, you can invoke the snapmirror policy create command to define a SnapMirror policy.

4. Use the snapmirror initialize command to initialize the relationship.

The initialization process performs a baseline transfer to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume. It also transfers any other Snapshot copies on the source volume to the destination volume.

The following command initializes the relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```


Preparing storage systems for SnapVault replication

Before you can use to perform disk-to-disk backup replication, you need to configure a data-protection relationship between the source and destination volumes, then initialize the relationship. On initialization, SnapVault makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

- You must be a cluster administrator.

You can use the Data ONTAP command-line interface (CLI) or OnCommand System Manager to create SnapVault relationships. The following procedure documents CLI usage.



If you are storing database files and transaction logs on different volumes, you must create relationships between the source and destination volumes for the database files and between the source and destination volumes for the transaction logs.

The following illustration shows the procedure for initializing a SnapVault relationship:

1. Identify the destination cluster.
2. On the destination cluster, use the volume create command with the -typeDP option to create a SnapVault destination volume that is the same size as or larger than the source volume.



The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2 GB destination volume named dstvolB in SVM2 on the aggregate node01_aggr:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate  
node01_aggr -type DP  
-size 2GB
```

3. On the destination SVM, use the snapmirror policy create command to create a SnapVault policy.

The following command creates the SVM-wide policy SVM1-vault:

```
SVM2::> snapmirror policy create -vserver SVM2 -policy SVM1-vault
```



Do not define a cron schedule or Snapshot copy policy for the SnapVault relationship. does that for you when you create a backup schedule.

4. Use the snapmirror create command with the -type XDP parameter and the -policy parameter to create a SnapVault relationship and assign a vault policy.

The XDP type defines the relationship as a SnapVault relationship.

The following command creates a SnapVault relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2, and assigns the policy SVM1-vault:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path  
SVM2:dstvolB  
-type XDP -policy SVM1-vault
```

5. Use the snapmirror initialize command to initialize the relationship.

The initialization process performs a baseline transfer to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

The following command initializes the relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Creating a backup and restore configuration

You specify most of the information you need for backup and restore operations in a configuration file. The configuration file is the input for Snap Creator backup and restore operations.

You can create as many configuration files as you need, but you can specify only a single configuration file as input for an operation. You can use either the Snap Creator GUI or the Snap Creator CLI to create configuration files.



This guide shows how to use the GUI to create configuration files. For information on how to use the CLI, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

You organize configuration files in Snap Creator profiles. Profiles and configuration files reside in the \engine\configs subdirectory in the Snap Creator installation directory:

- On UNIX hosts, the directory to which the installation file was extracted
- On Windows hosts, C:\Program Files\NetApp\Snap_Creator_Framework\scServer4.1.x, by default

Related information

[Logging in to the Snap Creator GUI](#) [Creating a configuration file](#) [Adding commands to the backup and restore configuration](#) [Cleaning up archived copies of Domino transaction logs](#) [Backing up the changeinfo directory](#)

Logging into the Snap Creator GUI

You use a web browser to log in to the Snap Creator GUI. Check with your administrator for the Snap Creator Server IP address, port, and login credentials.

1. Enter the following URL in your web browser: `https://server_name:port`, where:
 - `server_name` is the name or IP address of the Snap Creator Server
 - `port` is the port for the Snap Creator Server (8443, by default)
2. Enter the login credentials for the Snap Creator Server.

The Snap Creator GUI is displayed in the browser. If you are launching the GUI for the first time, a welcome message appears on screen. Click **OK** to dismiss the message.

Creating a configuration file

Before you can back up IBM Domino databases, you need to create a configuration file. The configuration file is the input for Snap Creator backup and restore operations.

Only alphanumeric and underscore characters are allowed in configuration and profile names. Names cannot start with a number.



Many sites use a profile-wide or site-wide global configuration file to apply storage system or VMware credential information to backup and restore configurations. For details about how to create a global configuration file, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

1. In the Profiles and Configurations pane, click **Add Profile**.
2. In the New Profile dialog box, enter the name of the profile, and then click **OK**.

Creating a new profile automatically launches the Snap Creator Framework Configuration wizard. To add a new configuration to an existing profile, right-click the profile, and then select **New Configuration**.

Throughout the Configuration wizard, click **Next** to advance to the next page and click **Back** to return to the previous page.

3. On the Configuration page, enter the name of the configuration, and then specify whether you want to enable password encryption.

Password encryption is enabled by default to prevent passwords from being displayed in clear text in the configuration file.



Hover over a field in the wizard to display a tool tip.

4. On the Plug-in Type page, select **Application plug-in**.
5. On the Application Plug-ins page, select **IBM Domino**.
6. On the Plug-in parameters page, specify IBM Domino backup configuration details:



The examples in the following table are for UNIX environments. You can view examples of the settings for Windows in the screenshot that follows the table.

GUI parameter	CLI parameter	Description
Domino Data Directory	DOMINO_DATA_PATH	The path to the Domino data directory. For example, /Domino/data.

GUI parameter	CLI parameter	Description
notes.ini Path	DOMINO_INI_PATH	The path to the notes.ini file. For example, /Domino/data/notes.ini.
Change info Directory	DOMINO_CHANGE_INFO_PATH	The path to changeinfo files. For example, /changeinfo. Important: Make sure that you store changeinfo files on a volume other than the Domino data volume.
Backup Type	DOMINO_DATABASE_TYPE	<p>The database files to be backed up. Select a value from the GUI drop-down, or specify the corresponding number in the CLI command:</p> <ul style="list-style-type: none"> • Everything (recommended) (CLI = 0) • *.box files (CLI= 1) • *.nsf, *.nsg, and *.nsh files (CLI = 2) • *.ntf files (CLI = 3)
IBM Domino installation directory	LOTUS	The path to the directory where Domino binary files are installed. In Linux, for example, /opt/ibm/domino (assuming the default installation path). In Windows, the parent of the Domino Shared Objects directory. Note: This is not the path to IBM Notes, formerly Lotus Notes.
Path to Domino shared objects	Notes_ExecDirectory	The path that contains Domino shared object files (.so or .dll). For example, /opt/ibm/domino/notes/latest/linux.
Data Restore Directory	DOMINO_RESTORE_DATA_PATH	For volume restores (assuming you are restoring to the same location from which you took the backup), the Domino data directory path. For single-file restores (NFS only), a location on the same volume as the Domino Data Directory. For example, /Domino/data/restore.

GUI parameter	CLI parameter	Description
Validate Domino Data Directory	VALIDATE_ DOMINO_ DATA_ PATH	<p>If the Domino environment has multiple mount points, some mount points might be on NetApp storage, while others might not be. Select a value from the GUI drop-down, or specify the corresponding value in the CLI command:</p> <ul style="list-style-type: none"> • Yes backs up files on the Domino data directory path only. • No backs up all files in the Domino environment regardless of their location on disk.

The following example shows how you might complete the fields in a Windows environment:

Domino Data Directory:	<input type="text" value="F:\Domino\data"/>
notes.ini Path:	<input type="text" value="C:\Lotus\Domino\notes.ini"/>
Change info Directory:	<input type="text" value="I:\changeinfo"/>
Backup Type:	<input type="text" value="Everything"/> ▼
IBM Domino installation directory:	<input type="text" value="C:\Lotus"/>
Path to Domino shared objects:	<input type="text" value="C:\Lotus\Domino"/>
Data Restore Directory:	<input type="text" value="F:\Domino\data"/>
Validate Domino Data Directory:	<input type="text" value="Yes"/> ▼

7. On the Agent Configuration page, specify Snap Creator Agent connection information:

For this field...	Do this...
IP/DNS	Enter the IP address or DNS host name of the Snap Creator Agent host.
Port	If you are not using the default port for the Snap Creator Agent (9090), enter the port number.
Timeout (secs)	Leave the default.

8. When you are satisfied with your entries, click **Test agent connection** to verify the connection to the Agent.



If the Agent is not responding, verify the Agent details and confirm that host name resolution is working correctly.

9. On the Storage Connection settings page, specify connection information for the Storage Virtual Machine (SVM, formerly known as Vserver) on the primary storage system:

For this field...	Do this...
Transport	Select the transport protocol for communications with the SVM: HTTP or HTTPS.
Controller/Vserver Port	If you are not using the default port for the SVM (80 for HTTP, 443 for HTTPS), enter the port number.

Note: For information about how to use an OnCommand proxy, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

10. On the Controller/Vserver Credentials page, specify the credentials for the SVM on the primary storage system:

For this field...	Do this...
Controller/Vserver IP or Name	Enter the IP address or DNS host name of the SVM host.
Controller/Vserver User	Enter the user name for the SVM host.
Controller/Vserver Password	Enter the password for the SVM host.

Important: If you are planning to replicate Snapshot copies to a SnapMirror or SnapVault destination, the name of the SVM you enter in this step must match exactly the name of the SVM you used when you created the SnapMirror or SnapVault relationship. If you specified a fully qualified domain name when you created the relationship, you must specify a fully qualified domain name in this step, regardless of whether SnapCreator can find the SVM with the information you provide. Case is significant.

You can use the `snapmirror show` command to check the name of the SVM on the primary storage system: `snapmirror show -destination-path destination_SVM:destination_volume` where `destination_SVM_name` is the name of the SVM on the destination system and `destination_volume` is the volume. For more information about creating SnapMirror and SnapVault relationships, see [SnapMirror and SnapVault setup](#).

When you click **Next**, the Controller/Vserver Volumes window is displayed.

11. In the Controller/Vserver Volumes window, specify the volumes to be backed up by dragging and dropping from the list of available volumes in the left pane to the list of volumes to be backed up in the right pane, and then click **Save**.

The specified volumes are displayed on the Controller/Vserver Credentials page.



If you plan to back up the changeinfo directory, you must configure the volume containing the directory as a metadata volume, as described in [Backing up the changeinfo directory](#). This option tells the IBM Domino plug-in to create a Snapshot copy of the changeinfo volume *after* creating the Snapshot copy for database files.

12. On the Controller/Vserver Credentials page, click **Add** if you want to specify SVM details and volumes to be backed up for another primary storage system.
13. On the Snapshot details page, specify Snapshot configuration information:

For this field...	Do this...
Snapshot Copy Name	Enter the name of the Snapshot copy. Tip: Click Allow Duplicate Snapshot Copy Name if you want to reuse Snapshot copy names across configuration files.
Snapshot Copy Label	Enter descriptive text for the Snapshot copy.
Policy Type	Click Use Policy , and then select the built-in backup policies you want to make available for this configuration. After you select a policy, click in the Retention cell to specify how many Snapshot copies with that policy type you want to retain. Note: For information about how to use policy objects, see the Snap Creator Framework 4.1.2 Administration Guide .
Prevent Snapshot Copy Deletion	Specify Yes only if you do not want Snap Creator to automatically delete Snapshot copies that exceed the number of copies to be retained. Note: Specifying Yes might cause you to exceed the number of supported Snapshot copies per volume.
Policy Retention Age	Specify the number of days you want to retain Snapshot copies that exceed the number of copies to be retained. You can specify a retention age per policy type by entering policy type:age, for example, daily:15.
Naming Convention	Leave the default.

The configuration specified in the following example performs a daily backup and retains four Snapshot copies:

Snapshot copy Name:

Snapshot copy Label:

Policy Type: ☒ Use Policy ☐ Use Policy Object

Snapshot copy Policies		
Enable Policy	Policy Name	Retention
<input type="checkbox"/>	hourly	0
<input checked="" type="checkbox"/>	daily	4
<input type="checkbox"/>	weekly	0
<input type="checkbox"/>	monthly	0

Prevent Snapshot copy Deletion:

Policy Retention Age:

Naming Convention: ☐ Recent ☒ Timestamp

14. On the Snapshot details continued page, set **Ignore Application Errors** to Yes if you want to force the backup operation to proceed even if one or more databases are in an inconsistent or corrupted state.

You should ignore the remaining fields.



A Domino environment might consist of hundreds or thousands of databases. If even a single database is in an inconsistent or corrupted state, the backup will fail. Enabling **Ignore Application Errors** allows the backup to continue.

15. On the Data Protection page, specify whether you want to perform optional Snapshot copy replication to secondary storage:
 - a. Click **SnapMirror** to mirror Snapshot copies.

The policy for mirrored Snapshot copies is the same as the policy for primary Snapshot copies.
 - b. Click **SnapVault** to archive Snapshot copies.
 - c. Specify the policy for archived Snapshot copies.
 - d. The instructions are in the following step [13](#).
 - e. In **SnapVault wait time**, enter the number of minutes you want Snap Creator to wait for the SnapVault operation to complete.
 - f. You need to have set up SnapMirror and SnapVault relationships before performing replication to secondary storage. For more information, see [SnapMirror and SnapVault setup](#).
16. On the Data Protection Volumes page, click **Add**, and then select the SVM for the primary storage system.

When you click **Next**, the Data Protection Volume Selection window is displayed.

17. In the Data Protection Volume Selection window, specify the source volumes to be replicated by dragging and dropping from the list of available volumes in the left pane to the list of volumes in the SnapMirror and/or SnapVault areas in the right pane, and then click **Save**.

The specified volumes are displayed on the Data Protection Volumes page.

18. On the Data Protection Volumes page, click **Add** if you want to specify SVM details and volumes to be replicated for another primary storage system.
19. On the Data protection relationships page, specify the credentials for the SVM on the SnapMirror and/or SnapVault destination systems.
20. If you prefer to use NetAppOnCommand Unified Manager APIs instead of Data ONTAP APIs for Snapshot copies and SnapMirror/SnapVault updates, complete the fields on the DFM/OnCommand Settings page:
 - a. Click **Operations Manager Console Alert** if you want to receive Unified Manager alerts, and then enter the required connection information for the Unified Manager virtual machine.
 - b. Click **NetApp Management Console data protection capability** if you are using the NetApp Management Console data protection feature for 7-Mode SnapVault replication, and then enter the required connection information for the Unified Manager virtual machine.
21. Review the summary, and then click **Finish**.

Snap Creator lists the configuration file below the specified profile in the Profiles and Configurations pane. You can edit the configuration by selecting the configuration file and clicking the appropriate tab in the Configuration Content pane. You can rename the configuration by clicking **Rename** in the right-click menu. You can delete the configuration by clicking **Delete** in the right-click menu.

Adding commands to the backup and restore configuration

Occasionally, a backup or restore operation might need to run a command that is not available in Snap Creator. On Windows hosts using FC or iSCSI, for example, you might prefer to use SnapDrive commands for backups instead of native Snap Creator commands, to ensure file system consistency at the OS level.

Where to run supplemental commands

By default, supplemental commands run on the Snap Creator Agent. You can run a supplemental command on the Server by prepending the command with the text **SERVER**, followed by a colon (:). For example, **SERVER:C:\myscript.bat**.

If you want to run a supplemental command on the Agent, you must explicitly permit the command to be executed by including it in the `allowed_commands.config` file on the Agent. There is no requirement to allow supplemental commands on the Server.

Where to add supplemental commands

You can use either the Snap Creator GUI or the Snap Creator CLI to add commands to your backup and restore configuration.



This guide shows how to use the GUI to add commands. For information on how to use the CLI, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

Related information

Allowing supplemental commands on the Snap Creator Agent

If you want to run a supplemental command on the Snap Creator Agent, you must explicitly permit the command to be executed by including it in the `allowed_commands.config` file on the Agent.

`allowed_commands.config` is located in the `etc` subdirectory of the Snap Creator Agent installation directory. For example:

- On UNIX hosts, `/install/path/scAgent4.1.x/etc/allowed_commands.config`
- On Windows hosts, `C:\Program Files\NetApp\Snap_Creator_Framework\scAgent4.1.x\etc\allowed_commands.config`

To allow supplemental commands on the Agent, open `allowed_commands.config` in an editor. Enter each command on its own line, exactly as you would enter the command at a command prompt. Case is significant. Make sure to specify the fully qualified pathname. Enclose the pathname in quotation marks if it contains spaces. For example:

```
"C:\Program Files\NetApp\SnapDrive\sdcli.exe"  
myscript.bat
```

Restart the Agent for the changes to take effect.



For security reasons, you should not use a wildcard entry (*) to allow all commands.

Adding supplemental commands to the configuration

You can add supplemental commands to your configuration after you create the configuration file. You can run the commands on the Snap Creator Server or Agent.

To run a supplemental command on the Agent, you should already have added the command to the `allowed_commands.config` file on the Agent.

By default, supplemental commands run on the Agent. You can run a supplemental command on the Server by prepending the command with the text `SERVER`, followed by a colon (:). For example, `SERVER:C:\myscript.bat`. There is no requirement to allow the command on the Server.

You can use either the Snap Creator GUI or the Snap Creator CLI to add commands to your backup and restore configuration.



This guide shows how to use the GUI to add commands. For information on how to use the CLI, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

1. In the Profiles and Configurations pane, select the configuration file and click the **Commands** tab.

The following command types are available:

Command type	Description
Application Quiesce	For non-plug-in use only. Ignore this command type.
Application Un-Quiesce	For non-plug-in use only. Ignore this command type.
Pre Exit	Use this type to run commands that tell Snap Creator what to do after a fatal error.
Snapshot Copy Create	Use this type to run your own Snapshot copy create commands.



Only a subset of the supported commands are available in the GUI. For information on the APP, ARCHIVE, MOUNT/UNMOUNT, and POST commands, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

- On the **Commands** tab, scroll to the section for the command type you want to use and click **Add**.
- Type over the text **Add Command Here** to enter the command.

To use SnapDrive for Windows to create a backup, you might enter the following in the section for the Snapshot Copy Create command type:

```
"C:\Program Files\NetApp\SnapDrive\sdcli.exe" snap create -s %SNAME-
%SNAP_TYPE_%SNAP_TIME -D F G H
```

where:

- "C:\Program Files\NetApp\SnapDrive\sdcli.exe" is the default installation path for the SnapDrive CLI.
 - %SNAME-%SNAP_TYPE_%SNAP_TIME specifies that the Snapshot copy name be formed from the configuration file name, the Snapshot policy (hourly, daily, or monthly), and the time and date the Snapshot copy was created.
 - D is the switch and F G H are the mount points for the data to be backed up. TIP: To run the command on the Server, prepend the command with the text SERVER, followed by a colon (:).
- When you are satisfied with your entry, press **Enter**.
 - Repeat steps [#STEP_3F9C83DD05D84075AD0277213FD67C56](#) through [#STEP_FBEDD7AE105F42138EB35315EA9FA7CC](#) to add another command.
 - Click **Save** in the Configuration Content pane to save your changes.

Cleaning up archived copies of Domino transaction logs

If you are using Domino transaction logging, the IBM Domino plug-in archives copies of Domino transaction logs to the changeinfo directory for use in up-to-the-minute restore operations. You need to enable transaction log management before Snap Creator will clean up the copies of the transaction logs you no longer need.

The Domino plug-in stores changeinfo files and archived copies of Domino transaction logs in the changeinfo

directory. The plug-in automatically cleans up .info files based on the retention policy for the configuration. It does not clean up unneeded copies of transaction logs unless you explicitly enable transaction log management in Snap Creator.

In setting up transaction log management, make sure to strike an appropriate balance between the number of archived copies of transaction logs you want to retain for up-to-the-minute restore operations and the size of the volume containing the changeinfo directory.



Do not set the Archive Log directory parameter to the path of Domino transaction logs. Doing so can damage your Domino environment.

1. In the Profiles and Configurations pane, select the configuration file and click the **Archive Log Management** tab.
2. In the **Archive Log Management** tab, specify the following values for the transaction log management settings:

Parameter	Setting	Description
Archive Log Enable	Y	Enables transaction log cleanup.
Archive Log Retention	Enter a value greater than or equal to the number of days Snapshot copies are retained.	The number of days to retain transaction logs.
Archive Log Directory	changeinfo/logs	The directory to clean up. Important: Do not specify the path of Domino transaction logs. Doing so can damage your Domino environment.
Archive Log Extension		The extension for copies of Domino transaction log files.
Archive Log Recursive Search	N	The Snap Creator plug-in puts .txn files in the logs subdirectory of the changeinfo directory, so a recursive search is not required.

3. Click **Save** in the Configuration Content pane to save your changes.

Backing up the changeinfo directory

If you plan to back up the changeinfo directory, you must configure the volume containing the directory as a metadata volume. Doing so tells the IBM Domino plug-in to create a Snapshot copy of the changeinfo volume *after* creating the Snapshot copy for database files.

You should already have identified the changeinfo volume as a volume to be backed up when you created the configuration file. For more information, see step [task using the gui to create a configuration](#).

You use the `META_DATA_VOLUME` parameter in the configuration file to identify the changeinfo volume as a metadata volume. The parameter is not available in the GUI.

1. Open the configuration file in a text editor.

Configuration files reside in the Snap Creator installation directory, in a subdirectory named for the profile. On a Windows host, for example, `C:\Program Files\NetApp\Snap_Creator_Framework\scServer4.1.x\engine\configs\Domino\Fuji15.conf`, where Domino is the name of the profile and Fuji15.conf is the name of the configuration file.

2. Enter the name of the changeinfo volume in the `META_DATA_VOLUME` parameter.

The changeinfo volume should already be listed in the `VOLUMES` parameter.

The following example shows how to configure Fuji15_DomChangeInfo as a metadata volume:

```
VOLUMES=domino_vserver:Fuji15_DomDAOS,Fuji15_DomChangeInfo,  
Fuji15_DomTxn,Fuji15_DomData  
META_DATA_VOLUME=domino_vserver:Fuji15_DomChangeInfo
```

Backing up databases

You can back up databases on demand or on a schedule. You can use either the Snap Creator GUI or the Snap Creator CLI to back up databases.



You can use the GUI to back up databases. For information on how to use the CLI, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

Related information

- [Backing up databases on demand](#)
- [Scheduling backups](#)

Backing up databases on demand

You should back up your databases as soon as they are available in NetApp storage. You can schedule recurring backups after the initial backup.

You should have created the configuration file for the backup, as described in [Creating a configuration file](#).

1. In the Profiles and Configurations pane, select the configuration file and click **Actions > Backup**.
2. In the Additional Parameters dialog box, select the policy for the backup job in the **Policy** drop-down and click **OK**.



You must have configured the policy in the configuration file.

Snap Creator starts the backup job. Job information is displayed in the Console pane.

Scheduling backups

You can schedule recurring backup jobs after performing the initial backup.

You should have created the configuration file for the backup, as described in [Creating a configuration file](#).

1. In the Profiles and Configurations pane, select the configuration file and click **Management > Schedules** at the top of the Snap Creator Framework window.
2. In the Jobs pane, click **Create**
3. In the New Job dialog, specify the job information:

For this field...	Do this...
Job Name	Enter the name of the job.
Start Date	Use the calendar control to select the date on which the schedule should start.
Active	Select this box to activate the schedule. Deselect the box to deactivate the schedule.
Profile	Select the profile for the configuration file.
Configuration	Select the configuration file.
Action	Select backup.
Policy	Select the policy for the backup job. Note: You must have configured the policy in the configuration file.

For this field...	Do this...
Frequency	Select the frequency with which the job should run. Depending on your choice, new fields are displayed where you can specify the minute, hour, and day to run the job. If you want to use a cron job, select cron and enter the cron command.

The following example shows how to schedule a backup job to run every day at midnight:

4. Click **Save**.

The scheduled job is listed on the Schedules tab in the Jobs pane. Select the job to run, edit, or delete it.

Restoring databases

The IBM Domino plug-in restores database files from the Snapshot copy of the volume on which the database is stored. When the Snapshot restore is complete, the plug-in applies any changes recorded in the changeinfo directory.

Understanding SnapMirror and SnapVault restore operations

You cannot use Snap Creator to restore a SnapMirror copy of the data. If you want to restore a mirror copy, you need to break the mirror first, then manually restore the data. For more information, see the [ONTAP 9 Volume Disaster Recovery Express Guide](#).

Only point-in-time volume restores are available when you restore a SnapVault copy from secondary storage. You cannot perform an up-to-the-minute volume restore or a single-file restore. You can work around this

limitation by first restoring from secondary storage to primary storage. From primary storage, you can then perform an up-to-the-minute volume restore or a single-file restore as needed.

Understanding destination directories


Volume restores and single-file restores typically have different destination directories. You specify the destination directory in the Data Restore Directory parameter in the backup and restore configuration file (see [step task_using_the_gui_to_create_a_configuration_file.md#STEP_E6C507729B3647FD8B2E8D0818F22D3B](#)):

- For a volume restore, you need to set Data Restore Directory to the Domino data directory.
- For a single-file restore, you need to set Data Restore Directory to a subdirectory on the Domino data volume.

This means that if you alternate restore types, you need to modify Data Restore Directory to point to the correct setting for each restore type. Otherwise, change information will not be applied correctly, and the data will be inconsistent with the database state you want to restore to. Of course, you can create different configurations for volume and single-file restores if you prefer.

Where to run a restore

With one exception, you can use either the Snap Creator GUI or the Snap Creator CLI to restore databases. For single-file restores in a SAN environment, you must use the CLI, as described in [Performing a single-file restore in a SAN environment](#).



This guide shows how to use the CLI only for single-file restores in a SAN environment. Otherwise, it shows how to use the GUI to restore databases. For information on how to use the CLI, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

Related information

[Performing a volume restore](#) [Performing a single-file restore](#) [Performing a single-file restore in a SAN environment](#)

Performing a volume restore

You can use the IBM Domino plug-in to perform point-in-time, up-to-the-minute, or selectable up-to-the-minute volume restores for all major storage types.

Make sure to:

- Set the Data Restore Directory in the configuration file to the Domino data directory.
- Stop the Domino server.

Volume restore options are as follows:

Restore type	Description
Volume restore (point-in-time)	Restores the entire volume to its state at the time of the Snapshot copy.

Restore type	Description
Volume restore (up-to-the-minute)	Restores the entire volume to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to the most recent copy.
Volume restore (selectable up-to-the-minute)	Restores the entire volume to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to a specified time.



For up-to-the-minute and selectable point-in-time restores, Domino transaction logging must be enabled.

Transaction logs can only be played forward. Selecting a time before the Snapshot copy was created causes a restore error.

1. In the Profiles and Configurations pane, select the configuration file and click **Actions > Restore**.

Selecting **Restore** automatically launches the Snap Creator Framework Restore wizard. Throughout the Restore wizard, click **Next** to advance to the next page. Click **Back** to return to the previous page.



If you are restoring from a SnapVault copy, you are prompted to restore from primary or secondary storage. Your selection depends on how you are handling the limitation described in [Restoring databases](#).

2. On the Restore details page, specify the restore details:

For this field...	Do this...
Controller/Vserver name	Select the SVM on the storage system you want to restore from.
Restore volume name	Select the volume you want to restore from.
Policy	Select the policy for the backup you want to restore from.
Snap Creator Snapshot Copies/All Snapshot Copies	Leave the default.
Restore Snapshot copy name	Select the Snapshot copy you want to restore from.
Restore type	Select Volume Restore .

3. On the following page, specify the volume restore type:

For this field...	Do this...
Restore Type	Click the volume restore type. If you choose Selectable up the minute , new fields are displayed where you can specify the date and time to which you want to play forward Domino transaction logs. Click Ignore Validation if the time on the Snap Creator Server does not match the time on the storage controller.
Disable Replication	Select this box to set a new Domino replica ID for restored database files.

- Review the summary and click **Finish**.
- In the Volume Restore dialog, click **OK**.

Snap Creator starts the restore job. Job information is displayed in the Console pane.

Restart the Domino server when the restore operation is complete.

Performing a single-file restore

You can use the IBM Domino plug-in to perform point-in-time, up-to-the-minute, or selectable up-to-the-minute single-file restores for databases in NFS environments.

You must set the Data Restore Directory parameter in the configuration file to the correct subdirectory on the Domino data volume.



You do not have to stop the Domino server for a single-file restore.

Single-file restore options are as follows:

Restore type	Description
Single-file restore (point-in-time)	Restores a single file to its state at the time of the Snapshot copy (NFS only).
Single-file restore (up-to-the-minute)	Restores a single file to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to the most recent copy (NFS only).
Single-file restore (selectable up-to-the-minute)	Restores a single file to its state at the time of the Snapshot copy, then plays forward Domino transaction logs to a specified time (NFS only).



For up-to-the-minute and selectable point-in-time restores, Domino transaction logging must be enabled.

Transaction logs can only be played forward. Selecting a time before the Snapshot copy was created causes a

restore error.

1. In the Profiles and Configurations pane, select the configuration file, and then click **Actions > Restore**.

Selecting **Restore** automatically launches the Snap Creator Framework Restore wizard. Throughout the Restore wizard, click **Next** to advance to the next page and click **Back** to return to the previous page.



If you are restoring from a SnapVault copy, you are prompted to restore from primary or secondary storage. Your selection depends on how you are handling the limitation described in [concept_domino_database_restore_overview.md#GUID-4D864E27-DE54-43BF-9B9F-EF2F240F65D9](#).

2. On the Restore details page, specify the restore details:

For this field...	Do this...
Controller/Vserver name	Select the SVM on the storage system you want to restore from.
Restore volume name	Select the volume you want to restore from.
Policy	Select the policy for the backup you want to restore from.
Snap Creator Snapshot Copies/All Snapshot Copies	Leave the default.
Restore Snapshot copy name	Select the Snapshot copy you want to restore from.
Restore type	Select Single File Restore .

3. On the following page, select the file you want to restore.

The file is displayed with its full path name in the **Source path** field.



Copy the path name to your clipboard for later use.

4. On the following page, select the directory to which you want to restore the file.

The name of the parent subdirectory for the file must be identical to the name of the parent subdirectory in the source path. If, for example, the source path is /domino/data/mail/user1.nsf, the restore path is /domino/data/restore/mail/user1.nsf.

The directory is displayed with its full path name in the **Destination path** field.

5. On the following page, specify the single-file restore type:

For this field...	Do this...
Restore Type	Click the single-file restore type. If you choose Selectable up the minute , new fields are displayed where you can specify the date and time to which you want to play forward Domino transaction logs. Click Ignore Validation if the time on the Snap Creator Server does not match the time on the storage controller.
Disable Replication	Select this box to set a new Domino replica ID for the restored database file.

- Review the summary, and then click **Finish**.
- In the Single File Restore dialog, click **OK**.

Snap Creator starts the restore job. Job information is displayed in the Console pane.

Performing a single-file restore in a SAN environment

For single-file restores in a SAN environment, you can use SnapDrive to mount the Snapshot copy on the Snap Creator Agent host. You can then copy the file you want to restore to the Data Restore Directory and use a Snap Creator custom action to complete restore processing.

Related information

[Using SnapDrive to mount a Snapshot copy](#)

[Copying the file to the Data Restore Directory](#)

[Using a Snap Creator custom action to complete restore processing](#)

[Disconnecting the Snapshot Copy](#)

Using SnapDrive to mount a Snapshot copy

You can use SnapDrive to mount the Snapshot copy for the file you want to restore. You can then manually copy the file to the Data Restore Directory.

The following procedure shows how to use SnapDrive for Windows to mount a Snapshot copy on a Windows Server 2008 host. SnapDrive for Windows and the Snap Creator Agent must be running on the host.

- In Server Manager, click **Storage > SnapDrive > physical_server_name > Disks**.
- In the Disk Identification pane, navigate to the Snapshot copy that contains the copy of the file you want to restore and choose **Connect Disk** in the right-click menu.

The Connect Disk wizard opens.

Throughout the Connect Disk wizard, click **Next** to advance to the next page. Click **Back** to return to the previous page.



Make a note of the name of the Snapshot copy. You will need the name when you copy the file to the Data Restore Directory.


3. On the Provide a Storage System Name, Lun Path and Name page, optionally enter descriptive text for the LUN.
4. On the Select a LUN Type page, leave the default.
5. On the Select LUN Properties page, select the drive letter or mount point for the Snapshot copy.
6. On the iSCSI initiator/FC host bus adapter (HBA) page, select the iSCSI initiator or FC host bus adapter (HBA).
7. On the Management Type page, select the management type you want to use to map to the LUN: automatic or manual.
8. Click **Finish** to mount the Snapshot copy to the specified mount point.

Copying the file to the Data Restore Directory

After you have mounted the Snapshot copy for the file you want to restore, you can copy the file to the Data Restore Directory specified in the backup and restore configuration file.

1. Copy the file to be restored from the mounted Snapshot copy.
2. Paste the file in the Data Restore Directory specified in the backup and restore configuration file.

The following example shows how you might copy the database file t04.nsf from the Snapshot copy mounted on the R: drive to the Data Restore Directory specified in the configuration file.



Using a Snap Creator custom action to complete restore processing

You can use a Snap Creator custom action to complete processing of restores originated

outside of Snap Creator.

You can perform this action in the Snap Creator CLI only.

1. In the Snap Creator CLI, enter the following command: `snapcreator.exe ---action custom --server server --port port --user user --passwd password --profile profile --config config --params snapname=snapshot_name datapath=datapath restoretype=u2m | su2m restoretime="MM/DD/YYYY HH:MM:SS" disablereplication=Y|N`

where

- server is the host name or IP address of the Snap Creator Server.
- port is the port number where the Snap Creator Server is running.
- user is the name of the Snap Creator user.
- password is the Snap Creator user's password.



You can omit the `--passwd` option if you do not want to type the password in clear text. Snap Creator will prompt for the password when you execute the command

- profile is the name of the Snap Creator profile.
- config is the name of the Snap Creator configuration file.
- snapshot_name is the name of the Snapshot copy from which you want to restore the file.
- datapath is the pathname of the directory to which you want to restore the file.
- restoretype is the type of restore to be performed:
 - u2m performs an up-to-the-minute restore.
 - su2m performs a selectable up-to-the-minute restore. Specify the date and time to which you want to play forward Domino transaction logs in MM/DD/YYYY HH:MM:SS format in the `--restoretime` option. For example, 01/23/2014 12:45:00.
- disablereplication disables Domino replication:
 - Enter N (default) to allow Domino replication. The plug-in uses an existing Domino replica ID for the restored database file.



An alternative method to allow Domino replication is to leave this parameter out of the command.

- Enter Y to disable Domino replication. The plug-in sets a new Domino replica ID for the restored database file. NOTE: Appending `--verbose` to the command will show additional information about the restore operation. Otherwise, you are immediately returned to the command prompt when the command finishes executing.

```
snapcreator.exe --action custom --server Tampico --port 8443
--user admin --profile Domino --config Nantes
--params snapname=Nantes-Test_Policy_20140123121459
datapath=F:\Domino\data\restore\mail\t04.nsf
restoretype=su2m restoretime="01/23/2014 12:45:00"
disablereplication=y
```

Disconnecting the Snapshot Copy

You need to disconnect the mounted Snapshot copy when the restore operation is complete.

The following procedure shows how to disconnect a mounted Snapshot copy on a Windows Server 2008 host. SnapDrive for Windows and the Snap Creator Agent must be running on the host.

1. In Server Manager, click **Storage > SnapDrive > physical_server_name > Disks**.
2. In the Disk Identification pane, select the mounted disk and click **Disconnect Disk** in the right-click menu.
3. In the Disconnect Disk dialog, click **Yes**.

Viewing job status and logs

You can view the status of a job in the Snap Creator Job Monitor. You can view a log of job actions in the Console pane while the job is running, or in the Out Log for the job. A number of specialized logs also are available.

You can view the following logs in the GUI:

Log	Description
Out	Provide verbose-level information for a job.
Debug	Provides verbose-level information plus information useful in debugging.
GUI	Provides information about Snap Creator GUI actions.
Stderr	Provides information about errors written to standard error (Reports pane only).
Error	Provides all error information (Reports pane only).

You can view all the logs in a text editor if you prefer not to use the Job Monitor or Reports pane. Logs reside in the installation directory for the Snap Creator Server.

Related information

[Viewing job status and logs in the Job Monitor](#)





[Viewing logs in the Reports pane](#)

[Using scdump to gather logs in a compressed format](#)

Viewing job status and logs in the Job Monitor

The Snap Creator Job Monitor is a dashboard-like interface for Snap Creator jobs. You can use the Job Monitor to view job status and download logs.

The Job Monitor indicates job status as follows:

-  indicates a job in progress.
 -  indicates a job that has completed successfully.
 -  indicates a job that completed with errors.
 -  indicates a job that failed.
1. In the Profiles and Configurations pane, select the configuration file and click **Management > Job Monitor** at the top of the Snap Creator Framework window.

The Job Monitor pane is displayed, where you can view job status.

2. Select a job, then click **Download Logs > Log_Type**, where Log_Type is one of the following:
 - **Out Logs** provide verbose-level information for a job.
 - **Debug Logs** provides verbose-level information plus information useful in debugging.
 - **GUI Logs** provides information about Snap Creator GUI actions. The GUI log is not specific to the job that is selected.
3. In the download dialog, click **Open** or **Save** as required.

Viewing logs in the Reports pane

You can view logs in the Snap Creator Reports pane, in the same format used to display them in the Console pane. You can view error logs as well as standard logs in the Reports pane.

1. At the top of the Snap Creator Framework window, click **Reports > Logs**.
2. In the Reports pane, select the profile, configuration file, log type, and log file, where log types are as follows:
 - **Out Logs** provide verbose-level information for a job.
 - **Debug Logs** provides verbose-level information plus information useful in debugging.
 - **Stderr Logs** provides information about errors written to standard error.
 - **Error Logs** provides all error information. **Note:** Depending on the job size, choosing **Debug Log** can cause the browser to become slow or unresponsive.
3. Click:
 - **Download Selected Log** to download the selected log.
 - **Download GUI Logs** to download the GUI logs. In the download dialog, click **Open** or **Save** as required.

Using scdump to gather logs in a compressed format

You can use the Snap Creatorscdump action to gather logs into a compressed file archive. You can then send the archive to technical support or a Snap Creator expert for review.

You should have created a backup and restore configuration file.

An scdump includes the following files:

File	Description
scdump.txt	Contains the Snap Creator Agent version and host operating system, the Data ONTAP release, and the Domino version.
engine.log	Contains information on Snap Creator workflow engine operations.
gui.log	Contains information on Snap Creator GUI operations and messages. Multiple logs may be included.
sc_server.log	Contains detailed information on Snap Creator Server operations. Multiple logs may be included.
out logs	Contain verbose-level information for jobs. Output logs for all configurations in the selected profile are included.
debug logs	Contain verbose-level information plus information useful in debugging. Debug logs for all configurations in the selected profile are included.
error logs	Contain error information for logs. Error logs for all configurations in the selected profile are included.
configuration files	All configuration files in the selected profile are included, including global configuration files.



Although a configuration file is the input for scdump, the utility gathers information for all configurations in the parent profile.

1. In the Profiles and Configurations pane, select a configuration file and click **Actions** > **scdump**.

scdump gathers logs and other files into a compressed file archive.

2. In the download dialog, click **Open** or **Save** as required.

Troubleshooting Domino plug-in-specific errors

Common Snap Creator errors fall into two categories: generic Snap Creator errors and Domino plug-in-specific errors. Domino plug-in-specific errors typically involve preinstallation issues, bad databases, or problems with a custom restore action.



For information about generic Snap Creator errors, see the [Snap Creator Framework 4.1.2 Administration Guide](#).

Domino plug-in is not supported on this platform

You typically receive this error when the prerequisites for the Domino plug-in have not been met. Generally, the error message is similar to the following:

```
Failed to load plug-in domino. Reason: ERROR: [ltd-00009] Domino plug-in
couldn't be loaded. Reasons could be 1) Domino plug-in is not supported
on this platform.
```

The following error message has also been known to appear:

```
ERROR: [tampico:9090(4.1.x)] SCF-00038: Application quiesce for
plug-in [domino] failed with exit code [99], continuing with backup.
```

As well as this one:

```
ERROR: [scf-00094] Application environment set for plug-in domino failed
[ERROR: [scf-00060] unknown application object [app] / application domino
at snapcreator.pl line 6410.
```

These error messages appear if one of the following conditions is not met:

- Preinstallation steps were not performed or were performed incorrectly:
 - On UNIX hosts, you must create symbolic links to IBM Domino shared object files.
 - On Windows hosts, you must add the IBM Domino installation path to the PATH environment variable. For more information, see the [Snap Creator Framework 4.1.2 Installation Guide](#).
- The bit level of Domino is different from the bit level of Snap Creator. For example, if you are using 32-bit Domino, you should use 32-bit Snap Creator.

Snap Creator backup fails because of a bad database

By default, a backup operation will fail if one or more databases are in an inconsistent or corrupted state. You can force the plug-in to proceed with backup operations despite errors, as described in step [task using the gui to create a configuration file](#).

Domino restore using custom action fails

Restore errors typically occur when restore paths are not set correctly. Check the paths for extra characters and backslashes. Also, verify that the Data Restore Directory is set correctly.

Errors are similar to the following:

```
##### Custom plug-in action #####
[Thu Apr 19 16:05:55 2012] DEBUG: Restoring to current time.
[Thu Apr 19 16:05:55 2012] DEBUG: Calling dominocore::postRestoreToTime
with
arguments(C:\Lotus\Domino\notes.ini,F:\Domino\data\,H:\changeinfo\fujil5-
daily_20120419130836,-1,F:\Domino\data\restore\mail\dadams.nsf,UP-TO-THE-
MINUTE,H:\changeinfo\logs\
[Thu Apr 19 16:05:55 2012] TRACE: Notes API initialized
Applying Change info for:F:\Domino\data\restore\mail\dadams.nsf
Error: unable to open file 'H:\changeinfo\fujil5-
daily_20120419130836/restore_mail_dadams.nsf.info', for path
'F:\Domino\data\restore\mail\dadams.nsf'.
Error:5114
NSFRecoverDatabases failed for
database:F:\Domino\data\restore\mail\dadams.nsf
Recovery Manager: Recovery only supported for Backup Files.ERROR_CODE:5114
[Thu Apr 19 16:05:55 2012] ERROR: [ltd-00008] Restoring databases finished
with errors
[Thu Apr 19 16:05:55 2012] ERROR: [scf-00154] Action custom for plugin
domino failed with exit code 1, Exiting!
[Thu Apr 19 16:05:55 2012] DEBUG: Exiting with error code - 2
```

All Domino plug-in-specific errors

The following table lists the IBM Domino plug-in-specific error messages in reference format:

Error code	Error message	Description/resolution
ltd-00001	Domino plug-in cannot work with SNAP_TIMESTAMP_ONLY = N. SNAP_TIMESTAMP_ONLY should be set to Y.	For more information, see step task_using_the_gui_to_create_a_configuration_file.md#STEP_81795CF9D6294AC891BC3D0CE4827CA3 .
ltd-00002	Quiescing databases finished with errors	Domino plug-in was unable to put all databases in backup mode. Check the logs to find the exact reason for the error or run Snap Creator in debug mode to find the error.
ltd-00003	Unquiescing databases finished with errors	Domino plug-in was unable to take all databases out of backup mode. Check the logs to find the exact reason for the error or run Snap Creator in debug mode to find the error.

Error code	Error message	Description/resolution
ltd-00004	Discovering databases failed	An application error caused application discovery to fail. Check the configuration and application settings. You can disable automatic discovery by setting APP_AUTO_DISCOVERY=N and commenting out VALIDATE_VOLUMES.
ltd-00005	Collection of operating system information failed - \$@	An error collecting operating system information caused the scdump action to fail. Check the logs and try running the command manually.
ltd-00006	Collection of SnapDrive information failed - \$@	An error collecting SnapDrive information caused the scdump action to fail. Check the logs and try running the command manually.
ltd-00008	Restoring databases finished with errors	Domino plug-in was unable to restore all the databases successfully. Check the logs to find the exact reason for the error or run Snap Creator in debug mode to find the error.
ltd-00009	Domino plug-in cannot be run as root	Reasons could be: <ul style="list-style-type: none"> • Domino plug-in is not supported on this platform. • Prerequisites for the Domino plug-in to run are not satisfied. For more information, see Domino plug-in is not supported on this platform .
ltd-00010	Errors encountered while opening databases	Check the logs to find the exact reason for the error or run Snap Creator in debug mode to find the error.

Where to go next

You can find more information about Snap Creator, including release-specific information, on the NetApp Support Site.

- [Snap Creator Framework 4.3.3 Installation Guide](#)

Describes how to install the Snap Creator Agent and Snap Creator Server. The Snap Creator Agent installation includes the IBM Domino plug-in.

- [Snap Creator Framework 4.3.3 Administration Guide](#)

Describes how to administer the Snap Creator Framework after installation is complete.

- [Snap Creator Framework 4.3 Release Notes](#)

Describes new features of, important cautions for, known problems with, and limitations of Snap Creator Framework 4.1.1.

- [SnapDrive 7.1 for Windows Installation Guide](#)

Describes how to install SnapDrive for Windows.

- [ONTAP 9 Cluster Peering Express Guide](#)

Describes how to quickly configure peer relationships between clusters and Storage Virtual Machines (SVMs).

- [ONTAP 9 Volume Disaster Recovery Preparation Express Guide](#)

Describes how to prepare a storage system for SnapMirror replication.

- [ONTAP 9 Volume Disaster Recovery Express Guide](#)

Describes how to prepare a storage system for SnapMirror recovery.

- [ONTAP 9 Volume Backup Using SnapVault Express Guide](#)

Describes how to prepare a storage system for SnapVault replication.

- [NetApp Technical Report 3917: Best Practices for Domino on NetApp](#)

Describes best practices for the IBM Domino plug-in.

- [Technical Report 3588: Integrating Domino 8.5 for Windows 2008 with a NetApp Storage System](#)

Describes how to integrate IBM Domino 8.5 for Windows 2008 with a NetApp storage system.

- [Technical Report 3691: Integrating IBM Domino 8.5 for Linux with a NetApp Storage System](#)

Describes how to integrate IBM Domino 8.0.2 for Linux with a NetApp storage system.

- [Technical Report 3723: Deploying Domino 8.0.2 in an AIX and NetApp Environment Using NFS](#)

Describes how to integrate IBM Domino 8.0.2 for AIX with a NetApp storage system.

- [Snap Creator Framework Discussions](#)

Enables you to connect with peers, ask questions, exchange ideas, find resources, and share Snap Creator best practices.

- [NetApp Video: SnapCreatorTV](#)

Enables you to view videos that demonstrate key Snap Creator technologies.

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.