



Sphere 3D

Snap Enterprise Data Replicator

Administrator Guide

For Appliances Running SnapServer®



©2008-15 Overland Storage, Inc. All rights reserved.

Overland®, Overland Storage®, ARCvault®, DynamicRAID®, GuardianOS®, NEO®, NEO Series®, PowerLoader®, Protection OS®, RAINcloud®, REO®, REO 4000®, REO Series®, Snap Appliance®, Snap Care® (EU only), SnapSAN®, SnapScale®, SnapScale X2®, SnapServer®, StorAssure®, Ultamus®, VR2®, and XchangeNOW® are registered trademarks of Overland Storage, Inc.

Tandberg Data®, AccuGuard®, AccuVault®, DPS1000 Series®, DPS1100®, DPS1200®, DPS2000®, Magnum®, QuikStation®, QuikStor®, RDX®, RDXPRO®, StorageLibrary®, StorageLoader®, Tandberg SecureService®, Tandberg StorageLibrary®, and VXA® are registered trademarks of Tandberg Data, Inc.

Desktop Cloud Orchestrator® and V3® are registered trademarks of Sphere 3D Corp.

RapidRebuild™, SnapExpansion XSR™, SnapScale X4™, SnapServer DX Series™, SnapServer XSD Series™, SnapServer XSD 40™, SnapServer XSR Series™, SnapServer XSR 40™, SnapServer XSR 120™, and SnapServer Manager™ are trademarks of Overland Storage, Inc.

BizNAS™, QuadPak™, and RDX+™ are trademarks of Tandberg Data, Inc.

Exosphere™, G-Series™, Glassware 2.0™, and SnapCLOUD™ are trademarks of Sphere 3D Corp.

All other brand names or trademarks are the property of their respective owners.

The names of companies and individuals used in examples are fictitious and intended to illustrate the use of the software. Any resemblance to actual companies or individuals, whether past or present, is coincidental.

PROPRIETARY NOTICE

All information contained in or disclosed by this document is considered proprietary by Sphere 3D Corp. By accepting this material the recipient agrees that this material and the information contained therein are held in confidence and in trust and will not be used, reproduced in whole or in part, nor its contents revealed to others, except to meet the purpose for which it was delivered. It is understood that no right is conveyed to reproduce or have reproduced any item herein disclosed without express permission from Sphere 3D Corp.

Sphere 3D Corp. provides this manual as is, without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Sphere 3D Corp. may make improvements or changes in the product(s) or programs described in this manual at any time. These changes will be incorporated in new editions of this publication.

Sphere 3D Corp. assumes no responsibility for the accuracy, completeness, sufficiency, or usefulness of this manual, nor for any problem that might arise from the use of the information in this manual.

FW 7.6

Sphere 3D Corp.
9112 Spectrum Center Boulevard
San Diego, CA 92123 USA

TEL 1.800.729.8725

1.858.571.5555

FAX 1.858.571.3664

www.sphere3d.com

Overview

The Snap EDR Management Console and Agents automate the secure flow of data between systems. This technology provides ready-to-use advanced file transfer and security features, and a scalable and flexible framework for remote data movement and management.

There are two products in the Snap EDR family: Snap EDR Express and Snap EDR. Each agent requires a separate license, and the licenses are sold by product, available from the SnapServer Web site at

<http://www.snapservers.com/Products/EDR.shtml>

For more information on licensing, see [Licensing Information](#) and [Manage License Keys](#).

Snap EDR consists of a Management Console and a collection of Agents.

The Management Console is installed on a central SnapServer and coordinates and logs the data transfer activities carried out by the distributed Agents. An Agent is also installed on the Management Console host. The Agents are responsible for the actual transfer of data.

Snap EDR Express supports replication between only two GuardianOS SnapServers, one of which is the Agent installed by default on the Management Console.

Snap EDR

Snap EDR enables you to do the following:

- Use the Replicate solution to schedule a job to replicate files between any two systems including Windows, Linux, and Mac Agents.
- Use the Distribute solution to schedule a job where files are transferred from one source host to one or more target hosts.
- Use the Aggregate solution to schedule a job to transfer files from multiple hosts to a single target host.
- Use the Remote Backup solution to schedule a job to back up data from remote hosts to a central host.
- Use the Remote Restore solution to schedule a job to restore backup data from a central storage location, on a per host basis, to the remote hosts from which the data was originally retrieved.

Snap EDR allows these transfers between Windows, Linux, Mac, and/or GuardianOS 5.0 and higher systems.

Snap EDR Express

Snap EDR Express supports replication between only two GuardianOS 5.0 and higher SnapServers, one of which is the Agent installed by default on the Management Console, and includes the Replicate tool. The Replicate tool enables users to schedule a job to replicate files between two systems. Snap EDR Express does not include the Distribute, Aggregate, Remote Backup, or Remote Restore tools, nor is it available for Windows, Linux, or Mac Agents.

Audience and Purpose

This document is intended for experienced information technology (IT) personnel. It provides information and procedures describing the use of the Snap Enterprise Data Replicator (EDR) Management Console to perform the following tasks:

- Configure, schedule, and manage the Distribute, Aggregate, Replicate, Remote Backup, and Remote Restore data management tools
- Generate reports on data transfer activities
- Perform Management Console maintenance tasks, including viewing and trimming the database logs and upgrading the different Snap Solutions
- Configure Agent administration

The procedures in this document also apply to Snap EDR Express, except for the procedures that describe how to use the Distribute, Aggregate, Remote Backup, and Remote Restore data management tools. These tools are not available with Snap EDR Express.

Licensing Information

Snap EDR installs with a 45-day trial period, after which it is necessary to obtain a license key from <http://www.snapserver.com/Products/EDR.shtml>

(or through your SnapServer distributor) to continue use of Snap EDR's functionality.

There are two licenses you can choose to purchase:

- Snap EDR Express allows only replication between two GuardianOS 5.0 or higher SnapServers, one of which is the Agent installed by default on the Management Console. This product does not include the Aggregate, Distribute, Remote Backup, or Remote Restore tools.
- Snap EDR includes the Aggregate, Distribute, Replicate, Remote Backup, and Remote Restore tools, licensed for a specific number of agents on Windows, Linux, Mac, or GuardianOS 5.0 and higher.

Technical Support

You can get additional technical support information on the Contact Support web page at:

<http://docs.overlandstorage.com/support>




Japanese Voluntary Control Council for Interference (VCCI)

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 **VCCI— A**

(Translation: This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case, the user may be required to take corrective actions.)

Conventions

This document exercises several alerts and typographical conventions.

Convention	Description & Usage
 WARNING WARNUNG	<p>A <i>Warning</i> contains information concerning personal safety. Failure to follow directions in the Warning could result in bodily harm or death.</p> <p>Eine <i>Warnung</i> enthält Informationen zur persönlichen Sicherheit. Das Nichtbeachten der Anweisungen in der Warnung kann zu Verletzungen oder zum Tod führen.</p> <p>Un <i>avertissement</i> contient des informations relatives à la sécurité personnelle. Ignorer les instructions dans l'avertissement peut entraîner des lésions corporelles ou la mort.</p>
ADVERTISSEMENT	
 CAUTION	<p>A <i>Caution</i> contains information that the user needs to know to avoid damaging or permanently deleting data or causing physical damage to the hardware or system.</p>
 IMPORTANT	<p>An <i>Important</i> note is a type of note that provides information essential to the completion of a task or that can impact the product and its function.</p>
Item_name	Words in this special boldface font indicate the names of buttons or page names found in the Web Management Interface.
Ctrl-Alt-r	This type of format details the keys you press simultaneously. In this example, hold down the Ctrl and Alt keys and press the r key.
NOTE	A Note indicates neutral or positive information that emphasizes or supplements important points of the main text. A note supplies information that may apply only in special cases, for example, memory limitations or details that apply to specific program versions.
Menu Flow Indicator (>)	<p>Words with a greater than sign between them indicate the flow of actions to accomplish a task.</p> <p>For example, Setup > User > Password indicates that you should click the Setup tab, then the User secondary tab, and finally the Password button to accomplish a task.</p>
<i>Courier Italic</i>	A variable for which you must substitute a value.
Courier Bold	Commands you enter in a command-line interface (CLI).

Information contained in this guide has been reviewed for accuracy, but not for product warranty because of the various environments, operating systems, or settings involved. Information and specifications may change without notice.

Contents

Preface

Snap EDR	3
Snap EDR Express	4
Audience and Purpose	4
Licensing Information	4
Conventions	5

Chapter 1: Management Console Install

Install and Configure the Snap EDR Management Console	4
Upgrade or Reinstall the Management Console	5
Upgrade the Management Console	5
Reinstall the Management Console	5
Management Console Overview	6
Configure the Dashboard	7
Change the Bandwidth of a Running Job	7
Add a Widget to the Dashboard	8
Change the Name of a Dashboard Widget	8
Configure a Widget	8
Change the Layout of a Widget	9
Remove a Widget	9

Chapter 2: Agent Administration

System Requirements for Snap EDR Agents	10
Install and Configure a Snap EDR Agent on a SnapServer	11
Upgrade or Reinstall a Snap EDR Agent on a SnapServer	11
Upgrade a Snap EDR Agent	11
Reinstall a Snap EDR Agent	12
Download the Snap EDR Agent Information File	12
Install the Agent Software on Windows-Based Systems	13
Prerequisites	13
Installation Procedure	13
Install the Agent Software on Linux Systems	14
Prerequisites	14
Installation Procedure	15
Install the Agent Software on Mac Systems	15
Prerequisites	15
Installation Procedure	16
Verify that the Agent is Properly Configured	16
Unsuccessful Agent Installation	16
Agent Groups	17
Create an Agent Group	17

Remove Agents from an Agent Group	17
Edit an Agent Group	17
Delete an Agent Group	17
Manage License Keys	18
Add a License Key	18
Delete a License Key	18
Revoke an Agent's Certificate	19
Uninstall Agents	20
Uninstall an Agent from a SnapServer	20
Uninstall an Agent from Windows	20
Uninstall an Agent from Linux	20
Uninstall an Agent from Mac	21
Configuration Notes	21
DNS and Name Resolution	21
Port Requirements	21
Re-Sync with Management Console Button	21
Certificate Validity Inconsistent Between Management Console and Agents	22
Warning Messages	22

Chapter 3: Data Management Tools

Aggregate Data Management Tool	23
Default Functionality	23
Create an Aggregate Job	24
Distribute Data Management Tool	30
Default Functionality	30
Create a Distribute Job	31
Replicate Data Management Tool	38
Default Functionality	38
Create a Replicate Job	38
Remote Backup Data Management Tool	48
Create a Remote Backup Job	48
File Backup Details	55
File Transfer Options	55
Application Backup Details	58
User Permissions Required	58
SystemState Backup	58
Remote Restore Data Management Tool	58
Create a Remote Restore Job	59
How Data to Restore is Resolved	65
Source Host Selection	66
Target Host Selection	66
Restore Versions	66
Specified vs. Full Restore	66
Restore Data Location on Target Host	67
Application Data Restore	67
File Transfer Options	67

Chapter 4: Manage Jobs

Manage Jobs Using the Job List Summary Screen	69
Manage Jobs Using the Job List Detail Screen	70
Check a Job's Running Status	70

View All Runs for a Job	70
View Job Statistics	71
Statistics	73
Interval Statistics	75
Perform a Task on More than One Job	76
Update Bandwidth Across Multiple Jobs	76
Updating Variables Across Multiple Jobs	77

Chapter 5: Reports

Report Types	78
Stats Summary Report	79
Report Fields	80
Detail Summary Report	80
Custom Query Report	82
Generated Report Categories	83
Generate a Report from a Template	84
Schedule a Report from a Template	84

Chapter 6: Maintenance

Trim Manager Database Logs	85
View Transfer Logs	87
Upgrade Applications	88
Upgrade an Application	88
Uninstall Applications	88
View Application Summary Information	88

Appendix A: Best Practices

Remote Backup Best Practices	89
Size the Solution	89
Regular Backup Volume and Frequency	90
Synchronize Backup Volume and Frequency	90
Required Target Agents	90
Simultaneous Running Jobs	91
Transfer Read/Write Issues	91
Changing the SnapServer Name or IP	91
Encrypted Files	91
File Ownership Transfer	91
Target Directory for Backups	92
Source Directories Specified	92
Target Host High Speed Link to Attached Storage	92
Synchronized Vs. Regular Backup	92
Mixed Environment and Preserve File Properties	93
Troubleshooting	93
Remote Restore Best Practices	93
Job Scheduling	93
File Ownership Transfer	94
Use Case Data Set	94

Appendix B: Exit Codes

Management Console Install

Installing and configuring the Snap EDR Management Console involves the procedures outlined in this section.


NOTE: Snap EDR 7.2 supports the following browsers: Internet Explorer 6.0 or higher, Firefox 2.0 or higher (Windows, Linux, Mac), and Safari 2.0 or higher.

Install and Configure the Snap EDR Management Console

Snap EDR is preinstalled on your SnapServer. You should configure only one Management Console to handle all of the Agents in your Snap EDR system.

1. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
2. Click the **Snap EDR link** in the Site Map (under Extras).
3. Click the **Configure as Management Console** option.

Once the configuration is complete, a screen appears with the following options.

Button	Description
Stop Service	Stops all Snap EDR services.
Restart Service	Restarts all Snap EDR services.
	 CAUTION: Use only if you have encountered a problem, and customer support advises you to restart the service. Any jobs currently running will stop and will not resume when you restart the service.
Disable Service on System Boot	By default, when a user reboots the SnapServer, services automatically restart. Select Disable Service on System Boot if you do not want the Snap EDR service to start up upon reboot. Note that the Enable Service on System Boot option appears when the disable service option is selected. This allows users to turn on the enable option.
Uninstall Service	Removes all of the components of the Snap EDR service.
Click Here to configure jobs	Clicking this link opens the Management Console UI where users can schedule jobs.

Upgrade or Reinstall the Management Console

NOTE: GuardianOS 5.0 or higher is required to install Snap EDR 7.2.

Upgrade the Management Console

1. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
2. Navigate to the **Maintenance > OS Update** screen.
3. Click the **Check for Updates** button and follow the instructions.
4. Once you have downloaded the latest update, use the **Browse** button to select the installation file you just downloaded.
5. Click the **Update** button.
6. Once the installation is applied, you will be prompted to reboot the server. Click **Reboot** to complete the installation process.
7. Once the system has rebooted, click the Snap EDR link in the Site Map (under Extras).
8. Click the **Complete Update** button to configure Snap EDR.

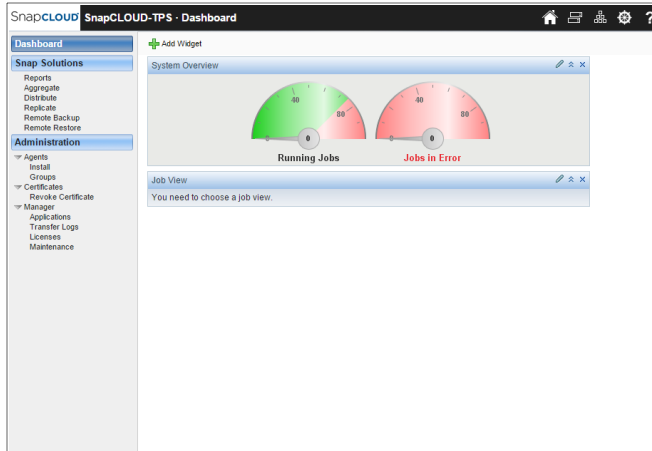
Reinstall the Management Console

If you have uninstalled Snap EDR and need to reinstall the application:

1. Download the latest installation file (**SnapEDR.gsu**) from:
<http://www.snapsrver.com/support/>
2. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
3. Navigate to the **Maintenance > OS Update** screen.
4. Use the **Browse** button to select the installation file you just downloaded.
5. Click the **Update** button.
6. Once the installation is applied, you will be prompted to reboot the server. Click **Reboot** to complete the installation process.
7. Once the system has rebooted, click the Snap EDR link in the Site Map (under Extras).
8. Click the **Configure as Management Console** option to configure Snap EDR.

Management Console Overview

The web-based Management Console enables users to manage the rules that govern data transfer activity, and deploy the Agent software on Mac, Linux, and Windows platforms. (Note that with Snap EDR Express, you can deploy the Agent software only on two GuardianOS Agents, one of which is the Agent installed by default on the Management Console. Also, only the Replicate tool appears in the Management Console with Snap EDR Express.)



By navigating menus and sub-menus, the Console displays the features and functions administrators utilize to create an automated data transfer system. The following table lists the menus that appear in the Console.

Menu	Description
Dashboard	<p>Click Dashboard to display your selected widgets in the main panel. Click Add Widget to add the following widgets to your Dashboard:</p> <ul style="list-style-type: none"> • System Overview: The Running Jobs gauge indicates the number of jobs that are currently running (the needle points to the number, and displays the number in its center). The green area of the Running Jobs gauge shows the upper range of job load for the CPU. The closer the needle is to the red area of the gauge, the more likely it is that the CPU will be straining. The Jobs in Error gauge displays the number of jobs in error. • Job View: Displays information about jobs (either running or in error), including job name, group, percent complete, job state (e.g., running, stopped, etc.), and provides action icons that allow the user to modify the job. • Click the Refresh tab to specify how often the information displayed in the widget is updated (every 15 seconds, 30 seconds, 45 seconds, 1 minute, or 2 minutes), or to turn off auto-refresh. <p>For information about configuring the Dashboard, see Configure the Dashboard.</p>

Menu	Description
Snap Solutions	<p>Click Snap Solutions to display the job data management tool options.</p> <p>NOTE: Aggregate, Distribute, Remote Backup, and Remote Restore do not appear with Snap EDR Express.</p> <ul style="list-style-type: none"> • Reports: Create customized reports on data transfer activity. • Aggregate: Schedule a job to transfer files from multiple hosts to a single target host. • Distribute: Schedule a job where files are transferred from one source host to one or more target hosts. • Replicate: Schedule a job to replicate files between any two systems including Windows, Linux, and Mac Agents. • Remote Backup: Schedule a job to back up data from remote hosts to a central host. • Remote Restore: Schedule a job to restore backup data from a central storage location, on a per host basis, to the remote hosts from which the data was originally retrieved. <p>For more information about the different job management data tools, see Data Management Tools. For information about creating Reports, see Reports.</p>
Administration	<p>Click Administration to perform the following administrative activities:</p> <ul style="list-style-type: none"> • Agent Install: Install an Agent on Windows, Linux, and Mac platforms. • Groups: Create an Agent group (collection of individual agents). • Certificates • Revoke Certificate: Revoke an agent's certificate (Does not apply to Snap EDR Express.) • Manager Applications: Install patches and upgrades to your individual Snap. <p>For more information, see Agent Administration.</p>
Solutions.	<ul style="list-style-type: none"> • Transfer Logs: Display the transfer logs generated by the Management Console during each transfer. • Licenses: Add license keys for additional agents or features. • Log Maintenance: Configure custom parameters for the scheduled log maintenance job, including amount of time to keep logs, threshold size, and scheduling parameters. <p>For more information, see Maintenance and Manage License Keys.</p>

Configure the Dashboard

The dashboard is customizable to better meet your needs. For example, you can adjust the transfer rates to optimize network loads or add multiple widgets in order to monitor multiple categories, such as the jobs running or any job errors.

Change the Bandwidth of a Running Job

The Management Console allows you to change the bandwidth of a running job. You may want to increase or decrease the bandwidth to speed up or slow down the transfer, due to network load issues during the job run.

When a job is running, the icon for changing the bandwidth appears in both the job summary screen and the Job View dashboard widget.

To specify the bandwidth throttle for a running job, follow these steps:

1. Log in to the Management Console.

2. Navigate to a screen that shows the running job whose bandwidth you want to modify. Choose the scheduled job list screen or the **Dashboard > Job View > Running Jobs** screen.
3. Click the **Set Bandwidth Limit** icon.
4. Specify the bandwidth limit in bytes, kilobytes or megabytes per second to a maximum of whatever the CPU can handle, or a percentage of a selected network connection (for example, 75% of 128 Kbps).

NOTE: Bandwidth throttles may also be employed by other network devices and policies (such as, QoS); therefore, a Snap EDR bandwidth throttle (or target maximum) may not be achievable. If you are having difficulty achieving a particular bandwidth target, ensure that other policies are not impacting your ability to reach the desired throughput.

5. Click **OK**.

Add a Widget to the Dashboard

There is no limit to the number of widgets you can add to the dashboard (except for the screen space available to display them). For example, you may want to have multiple Job View widgets in order to monitor both the Jobs in Error and the Running Jobs.

To add a widget to the dashboard, follow these steps:

1. Log in to the Management Console.
2. Click **Dashboard** (if the Dashboard area is not displayed).
3. Click **Add Widget**. Icons for the widgets you can add appear.
4. Click the widget you want to add to the dashboard. The widget appears in the dashboard.

Change the Name of a Dashboard Widget

You can change the name of any of the dashboard widgets. This is useful if you have more than one instance of the same type of widget. To change the name of a dashboard widget, double-click its name. A field appears that allows you to type in a new name. After typing the name, press **Enter** or click elsewhere on the screen.

Configure a Widget

The dashboard widgets are configurable. Users can change the type of information that the widget displays to better suit their needs by clicking the **Edit** icon and specifying the changes. Both widgets allow you to change the following:

- Refresh rate (how often the information displayed in the widget is updated).
- Widget name (double-click its name. A field appears that allows you to type in a new name. After typing the name, press **Enter** or click elsewhere on the screen.)
- Links to the underlying job view from which the information is displayed (click the **Edit** icon to select the **Running Jobs** or **Jobs in Error** job view).

To configure a dashboard widget, follow these steps:

1. Log in to the Management Console.
2. Click **Dashboard** (if the Dashboard area is not displayed).
3. Click the **Edit** icon on the widget you want to configure.
4. Change the values (most appear in a drop-down list).

5. To change the refresh rate for a job completion and job view widget, click the **Refresh** tab and select a value from the drop-down list.
6. Click **Save**.

Change the Layout of a Widget

You can change the location of the widgets simply by clicking the widget and dragging it to a different part of the screen. The dashboard page displays widgets in two columns. Clicking the mouse between the columns and dragging allows you to change the width of the columns. Note that the smallest that one of the columns can be is 250 pixels.

Remove a Widget

To remove a widget, click in the 'x' icon in the upper right corner of the widget. Note that you are not prompted when you delete a widget. If this widget was configured, you will lose this information.

Agent administration involves the following tasks:

- [Install and Configure a Snap EDR Agent on a SnapServer](#)
- [Upgrade or Reinstall a Snap EDR Agent on a SnapServer](#)
- [Install the Agent Software on Windows-Based Systems](#)
- [Install the Agent Software on Linux Systems](#)
- [Install the Agent Software on Mac Systems](#)
- [Agent Groups](#)
- [Manage License Keys](#)
- [Revoke an Agent's Certificate](#)
- [Uninstall Agents](#)

NOTE: With Snap EDR Express, you can install only two Agents, one of which is the Agent installed by default on the Management Console, and the other must be installed on another GuardianOS machine (see [System Requirements for Snap EDR Agents](#)).

System Requirements for Snap EDR Agents

Before installing an Agent, check your system to ensure that it meets the system requirements for Windows, Linux, or Mac Agents.

Item	Description
Operating System*	Solaris 8, 9, 10 Sparc; RedHat Enterprise Linux 4.x, 5.x SuSE Linux Enterprise Server 10.x MacOS-X 10.4, 10.5 (PowerPC and Intel); Windows 2000 Server (with SP4 or higher), Windows XP (with SP2 or higher), Windows 2003 Server, Windows 2003 R2 Server, Windows 7†, Windows Vista†, Windows 2008 Server†, Windows 2012†, and Windows 8†.
Web Browser	Internet Explorer 6.0 or higher, Firefox 2.0 or higher (Windows, Linux, Mac), Safari 2.0 or higher
System Memory	512 MB or more
Disk Space	Windows systems: 5 MB or more Linux systems: 10 MB or more Mac systems: 10 MB or more Installation directory: 50 MB (Windows); 100 MB, 30MB free in /tmp on Linux and Mac platforms
Network Connection	100 Mbit/sec. Ethernet or faster

* Operating systems that were supported in Snap EDR 5.2.2 but are not listed in this table are supported through backwards compatibility in Snap EDR 7.2.


† For Windows Vista, Windows 7, Windows 2008, Windows 2012, and Windows 8, User Account Control (UAC) and Data Execution Prevention (DEP) must be disabled. In addition, SystemState and Virtual Shadow Copy Service are not supported.

Install and Configure a Snap EDR Agent on a SnapServer

Snap EDR is preinstalled on your SnapServer. To configure your SnapServer as a Snap EDR Agent:

1. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
2. Click the Snap EDR link in the Site Map (under Extras).
3. Type the Name or IP address of the Management Console in the Name/IP of Management Console field. Note that the name of the Management Console must be resolvable from the Agent.
4. Click Configure as an Agent.

Once the configuration is complete, a screen appears with a number of options, as described in the following table.

Button	Description
Stop Service	Stops all replicator services
Restart Service	Restarts all Snap EDR services.
	 CAUTION: Use only if you have encountered a problem, and customer support advises you to restart the service. Any jobs currently running will stop and will not resume when you restart the service.
Disable Service on System Boot	By default, when a user reboots the SnapServer, services automatically restart. Select Disable Service on System Boot if you do not want Snap EDR service to start up upon reboot. Note that when the disable service option is selected, the Enable Service on System Boot option appears, allowing users to turn on the disabled option.
Uninstall Service	Removes all the components of the Snap EDR Agent.

Upgrade or Reinstall a Snap EDR Agent on a SnapServer

Topics include:

- [Upgrade a Snap EDR Agent](#)
- [Reinstall a Snap EDR Agent](#)
- [Download the Snap EDR Agent Information File](#)

Upgrade a Snap EDR Agent

1. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
2. Navigate to the **Maintenance > OS Update** screen.
3. Click the **Check for Updates** button and follow the instructions.

4. Once you have downloaded the latest update, use the **Browse** button to select the installation file you just downloaded.
5. Click the **Update** button.
6. Once the installation is applied, you will be prompted to reboot the server. Click **Reboot** to complete the installation process.

Reinstall a Snap EDR Agent

If you have uninstalled Snap EDR and need to reinstall the application:

1. Download the latest installation file (**SnapEDR.gsu**) from <http://www.snapserver.com/support/>
2. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
3. Navigate to the **Maintenance > OS Update** screen.
4. Use the **Browse** button to select the installation file you just downloaded.
5. Click the **Update** button.
6. Once the installation is applied, you will be prompted to reboot the server. Click **Reboot** to complete the installation process.
7. Once the system has rebooted, click the **Snap EDR** link in the Site Map (under Extras).
8. Type the name or IP address of the Management Console in the **Name/IP of Management Console** field.

NOTE: The name of the Management Console must be resolvable from the Agent.

9. Click **Configure as an Agent**.

Download the Snap EDR Agent Information File

Before installing an Agent on Windows, Linux, or Mac, you must first download the Snap EDR Agent Information file (called **sigsetup.inf**).

NOTE: This procedure does not apply to Snap EDR Express, or installing an Agent on a SnapServer.

NOTE: Before installing an Agent, check the table in [System Requirements for Snap EDR Agents](#) for the system requirements for Windows, Linux, and Mac Agents.

To download the Snap EDR Agent installation information file to a Windows, Linux, or Mac environment:

1. From the computer on which you want to install the agent, access the Snap EDR Management Console:
 - a. Log into the SnapServer as Administrator.
 - b. In the Site Map under Extras, click **Snap EDR**.
2. Click **Administration > Agents > Install**.
3. Click the **Click here** link.
4. Choose **Save Target As**.

5. Follow the instructions in the dialog to download the information file to your Agent software directory (for example, `/tmp/snapagent` for Linux-based hosts, `C:\temp\snapagent` for Windows-based hosts, or `Desktop/Snapagent` for Mac-based hosts).
6. Download the latest installation file from <http://www.snapserver.com/support/>.
7. Navigate to the installation file for your operating system using the `ReadMeFirst.html` file.
8. Follow the instructions in the dialog to download the setup program to the same Agent software directory you used earlier for the installation information file (for example, `/tmp/snapagent` for Linux-based hosts, `C:\temp\snapagent` for Windows-based hosts, or `Desktop/Snapagent` for Mac-based hosts). See the Windows, Linux, and Mac installation procedures in the following sections: [Install the Agent Software on Windows-Based Systems](#), [Install the Agent Software on Linux Systems](#), or [Install the Agent Software on Mac Systems](#).

You must place the installation information file and the setup software in the same directory.

NOTE: If your browser opens the information file instead of displaying the dialog, navigate back to the Installation Downloads page, right-click [Click here](#) and then save the file in your Agent software directory.

Install the Agent Software on Windows-Based Systems

NOTE: This installation does not apply to Snap EDR Express.

Prerequisites

- Make sure that the host system has a valid fully qualified domain name, system time and date, and time zone. The date, time, and time zone **MUST** be accurate before installing the Agent. The digital certificates for the Agents use the fully qualified host name and have their certificate validity period determined by time and date.
- Hosts using the Agent software should have IP addresses that do not change (for example, statically assigned IP addresses, or DHCP address reservations). In general, changes to IP addresses require reconfiguration of the Agents.
- You must have already installed and configured the Snap EDR Management Console on a SnapServer.
- Make sure you have downloaded the Snap EDR Agent Installation Information file, as described in [Upgrade or Reinstall a Snap EDR Agent on a SnapServer](#).
- If you are installing the Agent on Windows Vista, Windows 7, Windows 2008 Server, Windows 2012 Server, or Windows 8, make sure User Access Control (UAC) is disabled and that the Data Execution Prevention (DEP) security settings are only enabled for essential Windows programs and services. See [Disable User Access Control](#) and [Disable Data Execution Prevention](#).

Installation Procedure

To install the Agent software on Windows:

1. Log in to your system as Administrator or as a user with equivalent administration privileges.

2. In the folder where you downloaded the Agent software and installation configuration file (for example, `C:\temp\snapagent`), double-click the downloaded executable file (for example, `sig_client_x86-wnt.exe`).
The Welcome to Snap Enterprise Data Replicator screen appears.
3. Fill in the information on the screens, using **Back** and **Next** to move through the setup program screens.
Select the **Accept** radio button to accept the license agreement. If the setup program does not find the `sigsetup.inf` file in the installation directory, you are prompted to specify its location.
The Agent connects to the Management Console to request a digitally-signed certificate.
4. To exit the setup program, click **Finish**.

Disable User Access Control

This applies to Windows Vista, Windows 7, Windows 2008 Server and Windows 8 installations.

1. Click **Start > Settings > Control Panel**.
2. Double-click **User Accounts**.
3. Select **User Accounts**.
4. Make sure that the **User Access** option is turned off.
5. If necessary, reboot the computer.

Disable Data Execution Prevention

This applies to Windows Vista, Windows 7, Windows 2008 Server, Windows 2012, and Windows 8 installations.

1. Click **Start > Settings > Control Panel**.
2. Double-click **System** and choose the **Advanced** tab.
3. In the Performance area, click **Settings**.
4. Click the **Data Execution Prevention** tab.
5. Make sure DEP is enabled for essential Windows programs and services only, and click **OK**.
6. Reboot the computer.

Install the Agent Software on Linux Systems

NOTE: This installation does not apply to Snap EDR Express.

Prerequisites

- Before proceeding with a Linux installation, make sure that your system has a valid fully qualified host name, system time and date, and time zone. The date, time, and time zone **MUST** be accurate before installing the Agent. The digital certificates for the Agents use the fully qualified host name and have their certificate validity period determined by time and date.

- You must have already installed and configured the Snap EDR Management Console on a SnapServer.
- Make sure you have downloaded the Snap EDR Agent Installation Information file as described in [Upgrade a Snap EDR Agent](#) or [Reinstall a Snap EDR Agent](#).

Installation Procedure

The setup program installs the Agent software using native packaging systems such as pkgadd on Solaris systems and RPM (RedHat Package Manager) on Linux systems.

To install the Agent software on Linux:

1. Log in to your host system as a user with root privileges (for example, root).
2. Navigate to the location where the Agent software was downloaded.

```
cd /tmp/snapagent
```
3. Uncompress the downloaded file.

```
gunzip sig_client_i686-linux.tar.Z
```
4. Untar the downloaded file.

```
tar -xvf sig_client_i686-linux.tar
```
5. Run the program using the following command syntax:

```
./sigsetup
```

The Welcome to Snap Enterprise Data Replicator Setup screen appears.
6. Follow the instructions in the setup program screens to select installation options, and navigate through the program screens by typing N (next).

You need to type **A (Accept)** to accept the license agreement. If the setup program does not find the `sigsetup.inf` file in the installation directory, you are prompted to specify its location.

The setup program completes the certificate generation process.
7. You can view any installation errors in the Agent installer log file:

```
/tmp/sigsetup.log
```
8. Type **Y** to see the Readme file, or type **N** to exit the setup program.

Install the Agent Software on Mac Systems

NOTE: This installation does not apply to Snap EDR Express.

Prerequisites

- Make sure that the host system has a valid fully qualified domain name, system time and date, and time zone. The date, time, and time zone **MUST** be accurate before installing the Agent. The digital certificates for the Agents use the fully qualified host name and have their certificate validity period determined by time and date.
- Hosts using the Agent software should have IP addresses that do not change (for example, statically assigned IP addresses, or DHCP address reservations). In general, changes to IP addresses require reconfiguration of the Agents.
- You must have already installed and configured the Snap EDR Management Console on a SnapServer.

Installation Procedure

To install the Agent software on a Mac:

1. Log in to your system as a user with administrator privileges.
2. In the folder where you downloaded the Agent software and installation configuration file, double-click the downloaded executable file.
`sig_client_macintosh-OS_X-10.dmg`
A new volume (**SigAgentVol**) appears on the desktop.
3. Double-click the volume and double-click setup.
A password prompt appears.
4. Enter your password and click **OK**.
The Welcome to Snap Enterprise Data Replicator screen appears.
5. Fill in the information on the screens, using **Back** and **Next** to move through the setup program screens.
Select the **Accept** radio button to accept the license agreement. If the setup program does not find the sigsetup.inf file in the installation directory, you are prompted to specify its location.
The Agent connects to the Management Console to request a digitally-signed certificate.
6. To exit the setup program, click **Finish**.

Verify that the Agent is Properly Configured

To verify that the Agent is properly configured, run a simple replication job of a small set of data. For information on how to run a replication job, see [Create a Replicate Job](#).

NOTE: You need to have at least two Agents installed to run a replication job.

Unsuccessful Agent Installation

In some instances, the Agent installation might fail. The following table describes why an Agent installation might have failed, and possible resolutions to the problem.

Problem	Resolution
An Agent may have already been installed and then uninstalled without having the certificate revoked before re-installation.	Revoke the certificate for the host and rerun the installation. Proceed to Revoke an Agent's Certificate for instructions on revoking the certificate.
The host may not have connectivity to the network.	Correct connectivity problem and rerun the installation.
The Management Console may not be reachable from the host (due to routing problems, a physical problem, the Management Console being down, etc.).	Since the Management Console coordinates and logs the transfer activities carried out by the distributed Agents, a Snap EDR Agent must be able to connect to the Management Console. Correct the reachability problem and rerun the installation.
There may be insufficient drive space on the host for the installation files.	Correct the drive space problem and rerun the installation.

NOTE: If an Agent installation fails, the installer will generate a <hostname>.req.pem file, and save it to the <installation_directory>dds/security directory. This file can be deleted as it is not used by Snap EDR after an installation failure.

Agent Groups

Creating an Agent Group is useful when a large number of Agents are needed for data transfers.

An Agent Group is a logical collection of Agents that jobs can use in place of individual Agents. When a job uses an Agent Group, the controlling Agent sends data to or receives data from each Agent in the group. Agent Groups appear in the **Groups** submenu.

Create an Agent Group

To create an Agent Group:

1. From the Management Console, click **Administration > Agents > Groups**.
2. Enter the name of the Agent Group and provide a description of the group, if desired. Click **Add**.
3. Associate existing Agents and Agent Groups with the group you've just created by moving Agents from the **Available Agents** column to the **Selected Agents and Agent Groups** column.

NOTE: You can press <Shift>+click to choose multiple consecutive Agents/Agent Groups or <Ctrl>+click (Apple key on Mac) to choose multiple non-consecutive Agents/Agent Groups from the list.

Remove Agents from an Agent Group

To remove Agents/Agent Groups from the list of those associated with the Agent Group, click the item in the Selected Agents and Agent Groups column, and click the arrow to move the selection to the Available Agents and Agent Groups column.

Edit an Agent Group

You can change the name of the Agent Group or the description of the group using the Edit feature.

Delete an Agent Group

To delete an Agent Group, click the red 'X' next to the group.



CAUTION: If you delete an Agent Group that is part of a scheduled job, the job will fail when it runs (since the Agent Group that was defined as a source or target no longer exists).

Manage License Keys

Snap EDR installs with a 45-day trial period, after which it is necessary to obtain a license key from <http://www.snapserver.com/Products/EDR.shtml> (or through your SnapServer distributor) to continue use of Snap EDR's functionality.

There are two licenses you can choose to purchase:

- Snap EDR Express allows only replication between two GuardianOS 5.0 or higher SnapServers, one of which is the Agent installed by default on the Management Console. It includes only the Replicate solution and is not available for Windows, Linux, or Mac agents.
- Snap EDR includes the Aggregate, Distribute, Replicate, Remote Backup, and Remote Restore solutions, licensed for a specific number of agents on Windows, Linux, Mac, or GuardianOS 5.0 and higher.

One of the first things users need to do after logging in is to license the software with the number of agents and applications they have purchased.

The license screen displays a list of the components for which the user has purchased a license, as well as the number of agents associated with the component, the expiration date, license key, date the license was added, status, and a delete icon.

The Feature License Usage table displays the number of licenses associated with a component, how many are in use and their status.

Add a License Key

To add a license key:

1. Log in to the Management Console.
2. Click Administration > Manager > Licenses.
3. Click Add License Keys.
4. Type the license key(s) into the field. Separate multiple keys with a space or place each key on a separate line.
5. Click Add Keys.

If you type an invalid key an error message appears.

Delete a License Key

NOTE: Deleting a license key means that you will no longer be able to use the component associated with the license.

To delete a license key:

1. Log in to the Management Console.
2. Click **Administration > Manager > Licenses**.
3. Click the delete icon ("x") beside the license key you want to remove.
You are prompted to confirm deletion of the key.
4. Click **OK**.

Revoke an Agent's Certificate

The Certificate Authority (CA) running on the Snap EDR Management Console issues a digital certificate when you install the Agent software. Snap EDR uses the digital certificate for the identification and creation of secure communication channels among the data transfer Agents and the Management Console.

You need to revoke an Agent's certificate when you delete an Agent, otherwise it will not be possible to reinstall the same Agent should you wish to in the future. If you delete an Agent without revoking its certificate and then attempt to reinstall the Agent, an error message appears that indicates the Agent already has a valid certificate. Receipt of the certificate is part of the configuration of the Agent during the installation; when an unrevoked certificate exists, the Management Console generates the error message and does not send a certificate to the installer. With no certificate received during the installation, the Agent cannot be properly configured and is not recognized by the Management Console as part of the Snap EDR system.

Revoking an Agent's certificate does not remove the Agent software from the Agent. For information on uninstalling the Agent software, see [Uninstall an Agent from a SnapServer](#), [Uninstall an Agent from Windows](#), [Uninstall an Agent from Linux](#), or [Uninstall an Agent from Mac](#).

Once you revoke an Agent's certificate, the only way the Agent can take part in transfers again is to re-install the Agent.

To revoke an Agent's certificate:

1. Access the Snap EDR Management Console, then click **Administration > Certificates > Revoke Certificate**.

A list of Agents to which the user has access appears.

2. In the display area, click the Agent whose certificate you want to revoke.
3. Make sure a check appears in the **Also remove Agent from database** box if you want to permanently delete the Agent from the system.

Leaving the box unchecked means that the Agent is still identified in the database (along with any jobs with which it is associated). Make sure you uncheck the box if you are planning to reinstall the same Agent (for example, your Agent machine had a technical problem and needed to have Snap EDR reinstalled). Keeping the Agent in the database while you reinstall the Snap EDR Agent software on it means that the Agent will still appear in any jobs in which it is involved.



CAUTION: If you delete the Agent from the database, it will also be removed from any jobs with which it is associated. You would then have to manually add the re-installed Agent to all the jobs with which you want it associated.

However, if you are not planning to reinstall the Agent you are deleting, you may want to remove it from the database. If you do not delete it from the database, it will be included in the Agent counts when evaluating license keys. License keys specify how many Snap EDR Agents you can have in your Snap EDR installation. If the deleted Agent still appears in the database, then you will be unable to install a new agent to take its place (since the Management Console believes the agent still exists and that you have installed your Agent quota as specified in the license key).

4. Click **Revoke Certificate**.

Revoking an Agent's certificate adds the Agent to a list of revoked certificates that Agents check periodically.

Uninstall Agents

Either before or after uninstalling an Agent, you must revoke the Agent's certificate. If you do not revoke the Agent's certificate, you will be unable to reinstall the agent at a future date, since the Management Console believes the agent still has a valid certificate. For more information on revoking an Agent certificate, see [Revoke an Agent's Certificate](#).

Uninstall an Agent from a SnapServer

To uninstall a Snap EDR Agent:

1. Connect to the browser-based Administration Tool for the SnapServer, logging in as the administrator.
2. Click the **Snap EDR** link in the Site Map (under Extras).
3. Click **Uninstall Service**.
A prompt appears to confirm the uninstall.
4. Click **OK**.

NOTE: The uninstall removes all of the components associated with the software, including the configuration screen. Therefore you will not see an OK or successful uninstallation message screen once the uninstall completes.

5. Once the uninstall is complete, close the window.

Uninstall an Agent from Windows

To uninstall the Agent from Windows:

1. Choose **Start > Settings > Control Panel**.
2. Choose **Add/Remove Programs**.
3. Locate the **Snap Enterprise Data Replicator** entry in the list, and click on it.
4. Click **Change/Remove**.
5. In the Welcome screen, click **Next**.
6. In the Software Detected screen, click the **Uninstall** button and click **Next**.
7. In the Uninstallation Complete screen, click **Finish**.
8. Close the Add/Remove Programs screen.

Uninstall an Agent from Linux

To uninstall the Agent from a Linux system:

1. Log in as root.
2. Change directories to where the Agent software is installed.
`/usr/snap/dds/bin`
3. Type `./siguninstall`.
The uninstall opening screen appears.
4. Type **N** (for next).
The Uninstall program searches for installed packages.
5. Select the packages you want to uninstall (in this case, the Agent) and type **N**.

6. Verify the directories associated with the Agent software and type **N**.
The summary screen appears.
7. Confirm the information and type **N**.
A prompt appears to confirm that you want to uninstall.
8. Type **Y** (for yes).
Eventually a message appears that the uninstall is complete.
9. Press **Enter**.
A message appears that the uninstallation is running a cleanup.
10. Press any key to finish the uninstall and return to the shell.

Uninstall an Agent from Mac

To uninstall the Agent from a Mac, follow these steps:

1. Log in to your system as a user with administrator privileges.
2. In the Applications area, open the Signiant folder.
3. Double-click the uninstall icon.
A password prompt appears.
4. Enter the password and click **OK**.
5. In the Welcome Screen, click **Next**, and click **Next** again.
You are prompted if you are ready to begin uninstallation.
6. Click **Yes**.
7. Click **Finish** when the uninstall is complete.

Configuration Notes

IT personnel charged with administering data replication and migration tasks with Snap EDR must consider the following information.

DNS and Name Resolution

If there is no DNS in the customer's network, host files can be added on the SnapServers.

Port Requirements

The ports Snap EDR uses to route data traffic are 49221 and 443. To accomplish data replication or migration this port must be enabled along the network path required for the machines to properly transmit the data. If you have installed a firewall, installing a Snap EDR Agent does not damage the firewall, but you will need to manually add port 49221 through the firewall to communicate to the Snap EDR Agent.

Re-Sync with Management Console Button

In rare circumstances, a Management Console name or address change fails to propagate to the Agents due to network or server outages. The Re-Sync with Management Console option resynchronizes the Agent with the Management Console and all other Agents associated with

the Management Console's IP address. The Re-Sync option is needed when the Management Console can no longer communicate with an Agent after a name change. This option can be found on the Snap EDR Agent post-configuration screen.

Certificate Validity Inconsistent Between Management Console and Agents

By default, the Snap EDR installation provides a certificate that is valid for one year for the Management Console and two years for the Agents.

Warning Messages

Some warning messages that appear in the Maintenance Logs may also contain the word "error". These messages are simply warnings and do not indicate that an error has occurred. Due to formatting issues, some errors do not appear in the Event Log.

Snap EDR includes five Data Management Tools to automate the secure flow of data between systems:

- The [Aggregate Data Management Tool](#) allows you to schedule a job to transfer files from multiple hosts to a single target host.
- The [Distribute Data Management Tool](#) allows you to schedule a job where files are transferred from one source machine to one or more target machines.
- The [Replicate Data Management Tool](#) allows you to schedule a job to transfer files between two systems.
- The [Remote Backup Data Management Tool](#) allows you to backup data from remote hosts to a central host.
- The [Remote Restore Data Management Tool](#) allows you to transfer backup data from a central storage location to the remote hosts from which the data was originally retrieved.

This chapter provides detailed descriptions of each of the data management tools and procedures for creating jobs using them.

NOTE: Snap EDR Express includes only the Replicate data management tool. For information on this tool, see [Replicate Data Management Tool](#).

Aggregate Data Management Tool

NOTE: This tool is not available for Snap EDR Express.

The Aggregate data management tool retrieves files from multiple source Agents to a single target Agent. The files transferred from the source Agents may be stored on a per Agent basis, or in a single flat namespace on the target Agent's storage device (either local disk or NAS). This is referred to as a "Pull" file transfer.

Default Functionality

The Aggregate data management tool has the following default functionality:

- File transfer from one or many Agents to single source Agent
- Cross platform support (transfers can go from one or more GuardianOS, Windows, Linux, or Mac Agents to a single GuardianOS, Windows, Linux, or Mac Agent)
- Store files from source Agents in a single flat namespace, or in a separate namespace on a per source Agent basis
- Option to exclude files and sub-directories from the transfer
- Option to transfer changed bytes only, or entire file contents

Create an Aggregate Job

To schedule an Aggregate job:

1. From the Management Console, click **Snap Solutions > Aggregate**.
2. Click **Add a Job**.

If this is the first time you are creating an aggregate job, when you click **Aggregate**, the scheduling screen appears. You do not have to click **Add a Job**.

3. Enter information into the various fields. The following table describes the fields.

NOTE: Note You cannot use the '|' (pipe) character in a Windows directory path or errors will result.

Field	Description
Job Name	A unique name for this job run.
Source Agent(s)	Selection
Source Agent(s)	The Agent(s) from which you are transferring the data. Press <Shift>+click to select multiple consecutive agents or <Ctrl> +click to select multiple non-consecutive Agents.
Directory and File Options	
Source Directory	<p>The directory from which you want to transfer the files.</p> <p>NOTE: This repository folder must have 'sharing' and correct permissions applied. To do so, right-click on the directory to be backed up, select the Properties tab and click in the Share this folder radio button. Click the Permissions button and set the read/write privileges.</p> <p>The source directory specified may be in the following formats:</p> <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • Linux root drives; e.g., /home • Mac root drives; e.g., /Users <p>Example entry:e:\databasefiles</p> <p>Use the Browse button to select directories. Click through to the desired directory and click choose this directory. Leave the field empty to use the default Agent directory.</p>
Include File Names/Types	<p>Allows users to specify which files are transferred by filtering on the names to include. Use a comma to separate multiple filters.</p> <p>For example, if you type *.doc, *.ppt, in the field, the transfer will include files with the doc and power point extensions. However, *.txt files, for example, would be excluded.</p>
Exclude FileNames/Types	<p>Allows users to specify which files are not transferred by filtering on the names to exclude. Use a comma to separate multiple filters.</p> <p>For example, if you type *.doc, *.ppt, in the field, the transfer will not include files with the doc and power point extensions. However, *.txt files, for example, would be included.</p>
Exclude All Sub-Directories	<p>Choose Yes to specify that all sub-directories be excluded. Values in the Exclude Sub-Directories field are ignored. Choose No to define specific sub-directories for exclusion.</p>

Field	Description
Exclude Sub-Directories	<p>Subdirectories may be excluded from the data being backed up by specifying them in this field. Multiple entries must be separated with a comma. When the job runs, all directories that match those specified in the exclude subdirectories field will be excluded. Note that normal behavior is to exclude subdirectories that match regardless of where they appear in the directory path. Using the anchoring expression ("@") changes this behavior to anchor the exclude directory path.</p> <p>For example, if a user specifies a source directory of C:\data\docs, and an exclude directory of temp, any subdirectories called temp will be excluded, even those nested within another subdirectory. (For example the subdirectory C:\data\docs\publish\release\temp would be excluded as well.)</p> <p>Users who want to exclude a directory only at a certain level can use the @ symbol to anchor the exclude directory path at the starting source directory level. For example, specifying @temp in the above example would mean that the C:\data\docs\temp directory would be excluded, but the C:\data\docs\publish\release\temp directory would be included.</p> <p>If the user instead wanted to make sure just the C:\data\docs\publish\release\temp directory was excluded, but wanted to have a source directory path of c:\data\docs, the user would need to type C:\data\docs in the source directory field and @publish\release\temp in the exclude field.</p> <p>Special characters allow users to make use of pattern matching on the directory path. You must escape special characters with a backslash to be matched literally. Characters include the following:</p> <ul style="list-style-type: none"> • * (matches zero or more characters) • ? (matches any single character) • [...] (matches any one of the enclosed characters - for example, [ch] would match the characters "c" or "h") <p>A pair of digit, lowercase or uppercase characters separated by a hyphen '-' denotes a range set and matches any character sorted in the range. If the first character following the '[' is '^' or '!', then any character not enclosed is matched.</p> <p>Use commas to specify multiple distinct patterns as an alternative to using multiple option specifications. Note that these options do not enable/disable directory traversal.</p>
Target Agent Selection	
Target Agent	The Agent to which you are transferring the data.
Target Directory	<p>Specifies the directory to store the data from the source agents. If the directory does not exist it will be created. Note that this repository folder must have 'sharing' and correct permissions applied. To do so, right-click the directory to be backed up, select the Properties tab and click in the Share this folder radio button. Click the Permissions button and set the read/write privileges.. If specified, data for each Agent specified is stored on a per Agent basis under this directory. Otherwise data from all Agents are stored in this directory. The directory specified should be an absolute path name. Directories specified may be in the following format:</p> <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • Linux root drives; e.g., /home • Mac root drives; e.g., /Users <p>Example entry:e:\remotebackups</p> <p>Use the Browse button to select directories. Click through to the desired directory and click choose this directory.</p>

Field	Description
Folder Per Source	When selected, a separate folder for each source Agent is created in the specified target directory. All data transferred from that Agent appears in that folder.
Windows Volume Shadow Copy Service Options	
Use Volume Shadow Copy Service	Specifies that a Windows Volume Shadow Copy of the source data is automatically created before the transfer. The transfer is performed from the shadow copy and then the shadow copy is released. This avoids transfer issues with open files. This option does not apply to Windows 2000 agents. NOTE: If you select VSS and you are using a 5.2 source Agent, the job will still run, but no volume shadow copy will be used (a warning appears in the log).
Shadow Copy Creation Timeout	Specifies the maximum number of seconds to wait for shadow copy creation before proceeding in non-shadow mode. The recommended minimum value is 60 seconds. The recommended maximum is 600 seconds.

Field	Description
File Transfer Options	
File Ownership Transfer	<p>Specifies whether files maintain the source user ID and group ownership after they are transferred, and what method is used. Choose from the following options to preserve ownership:</p> <p>GuardianOS - When transferring between two GuardianOS machines, Permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Mac, Linux, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for Permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).</p> <p>Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its Permissions from a folder on the source is transferred to a folder on the target that has different Permissions. However, the source folder is not transferred. In this case, the file's inherited Permissions on the target will come from the target folder into which it is transferred and not from the source folder.</p> <p>NOTE: Windows attributes (such as Read Only) are not preserved in transfers between Windows and GuardianOS machines. Access Control Lists (ACLs) can be preserved, but attributes cannot.</p> <p>GuardianOS mode is the default mode of transfer.</p> <p>Windows - Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups—these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.</p> <p>Unix - Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX/Linux/Mac hosts.</p> <p>The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.</p> <p>Off - Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.</p>
Transfer File Differences Only	<p>If set to Yes, only changed bytes of files will be transferred, not the entire file. Typically used in low bandwidth situations. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps).</p>
Use Compression	<p>If set to Yes, the source Agent compresses each file before sending it. The files will be uncompressed automatically on the target Agent(s). The degree of compression depends on the type of data you are transferring. The following are typical rates of compression for different types of data:</p> <p>Plain text: 70-95% TIFF images: 20-40% Binary files: 0-5%</p> <p>The default value is No.</p> <p>Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps). This option is NOT recommended for LAN or high speed links.</p>

Field	Description
Delivery Mode	<p>Specifies the mode to transfer files:</p> <ul style="list-style-type: none"> • Fast - Do not create temporary work files. • Normal - Use temporary work files during transfer. • Log File Name - Log file names transferred. Logs are stored on the Manager. • Certify File Content - Create agent certified delivery log of files transferred.
Action If "File in Use"	<p>Allows users to specify whether a file that is in use during a transfer is skipped or generates an error message. If Error is selected, and a file is in use during the transfer, the transfer fails and generates an error. If Skip is selected, and a file is in use during the transfer, the transfer continues, skipping the file, and issuing a warning. Files are considered "in use" if, for example, it is not possible to copy them using Windows Explorer (on Windows), or the file is actively executing (on Linux).</p> <p>Using Volume Shadow Copy results in fewer cases of open files being encountered.</p>
Bandwidth Throttle	<p>Click Add a new throttle to limit this job to the amount of bandwidth specified. In addition, click the plus sign ("+") to add additional throttles. (Click the "x" to delete a throttle.)</p> <p>Specifies the rate at which the source will make the data available to the network, regardless of how fast the machine is sending. Bandwidth Throttle is the rate at which the source reads the data from disk and puts it in the send queue.</p> <p>NOTE: Bandwidth limiting is done on each stream connection, so a value specified here is passed to the controlling Agent for each template executed, and divided among the streams as follows:</p> <ol style="list-style-type: none"> 1 It is divided by the number of streams the Agent will use for each remote Agent (typically four). 2 It is further divided by the potential number of concurrent remote Agents. This will be the lesser of the maximum allowed number of concurrent Agents, and the number of remote Agents specified in the template. 3 Note that you must use 24-hour time format to specify start and end times (i.e., enter 2:00 pm as 14). Click the clock icon to display 12-hour clock options from which you can choose and click OK. The value will be converted to 24-hour time. You can also specify days of the week to which this bandwidth limit applies. <p>Once a job has started, all bandwidth throttles are applied at the times based on the Daylight Savings Time (DST) in effect when the job started. If DST changes while the job is running, bandwidth time of day changes may be off by the time change value (plus or minus an hour) after the time change.</p> <p>Note that bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.</p>

Field	Description
Encryption Level	<p>Allows users to specify the encryption level from the following values: High, Medium, Low, or No Encryption - signed. The default value is High. Note that mutual authentication is always used regardless of the encryption level specified.</p> <p>The No encryption - signed option transfers unencrypted (plain text) data, but includes the SSL protocol's message digest calculation and signing to ensure data stream integrity.</p> <p>Note Encryption incompatibilities exist with transfers that involve version 5.2.2 build 1194 and higher (including version 7.2) and version 5.2.2 build 1177 and earlier agents.</p> <p>With Snap EDR version 5.2.2 build 1194 and higher, Snap EDR was enhanced to support AES-256 and AES-128 bit encryption. We recommend upgrading your Management Console and all of your agents. If all of the Snap EDR nodes are upgraded, the system will automatically use the "high" (AES-256) encryption when the replication jobs have been set to "high" (the default).</p> <p>If you choose not to upgrade all of the Snap EDR nodes, you must change all of the Snap EDR jobs to reflect "no" encryption; otherwise, the jobs will fail since the earlier versions of Snap EDR agents do not support the improved encryption.</p> <p>If running Snap EDR version 5.2.2 build 1194 or later, these encryption settings map to these encryption levels:</p> <ul style="list-style-type: none"> • High - AES-256 • Medium - AES-128 • Low - RC4-40
Job Log Detail Level	The type of logging information for this job. Choose from Error , Warn (warning), Info (Information) or Debug . Debug provides the greatest level of detail while Error provides the least.
Scheduling Parameters	
Job Start Date/Time	The date and time at which you want the job to run.
Frequency	
Once	Run the job only once at start time.
Hourly, Daily, Weekly, Monthly, Yearly	Run the job once every selected interval.
Monthend	Run the job once every last day of the month.
None	Use none when you have a job that needs to be scheduled, but is run only at irregular, user-defined times.
Every x interval	Run the job every x interval, the first of which is the "start job at" time. Intervals may be minutes, hours, days, weeks, months or years.
X<day> of the month	Run a job on a certain day every month.

Field	Description
Time Zone	<p>Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time.</p> <p>Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere.</p> <p>Note that the time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.</p>
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	<p>Text that appears in the subject field of the report e-mail message.</p> <p>NOTE: To receive email notifications in Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification.</p>
SNMP Trap Notification	
Send Traps On	<p>Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither.</p> <p>NOTE: To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.</p>

4. Click **Add Job**.

The new job appears in the list.

Distribute Data Management Tool

NOTE: This tool is not available for Snap EDR Express.

The Distribute data management tool allows users to schedule a job where files are transferred from one source machine to one or more target machines.

Default Functionality

The Distribute data management tool has the following default functionality:

- File transfer from one source to one or many target Agents

- Cross platform support
- Option to exclude files from transfer
- Option to exclude subdirectories from transfer
- Option to synchronize target directory structure with source directory structure
- Option to transfer changed bytes only, or entire file contents

By default, any files which are in use during the transfer will be skipped without error. This allows users to transfer common files from a central repository. Users can set information in the following categories:

- Source Agent selection
- Target Agent selection
- Directory and file selection
- Windows Volume Shadow Copy Service options
- File transfer options (such as file compression, certified delivery, source and target synchronization, etc.)
- Job options (such as bandwidth throttle, encryption, and job log detail levels)
- Scheduling parameters
- E-mail notification
- SNMP trap notification

Create a Distribute Job

To schedule a Distribute job:

1. From the Management Console, click **Snap Solutions > Distribute**.
2. Click **Add a Job**.

If this is the first time you are creating a distribute job, when you click **Distribute**, the scheduling screen appears. You do not have to click **Add a Job**.

3. Enter information into the various fields. The following table describes the fields.

NOTE: Note You cannot use the “|” (pipe) character in a Windows directory path or errors will result.

Field	Description
Job Name	A unique name for this job run.
Source Agent Selection	
Source Agent	The Agent from which you want to transfer the files.
Source File/ Directory	The directory from which you want to transfer files. This repository folder must have 'sharing' and correct permissions applied. To do so, right-click the directory to be backed up, select the Properties tab and click the Share this folder radio button. Click the Permissions button and set the read/write privileges. The source directory specified may be in the following format: <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • Linux root drives; e.g., /home • Mac root drives; e.g., /UsersExample entry: e:\databasefiles Use the Browse button to select directories. Navigate to the desired directory and click choose this directory.
Directory and File Options	

Field	Description
Include File Names/Types	Allows users to specify which files are transferred by filtering on the names to include. Use a comma to separate multiple filters. For example, if you type *.doc, *.ppt, in the field, the transfer will include files with the doc and power point extensions. However, *.txt files, for example, would be excluded.
Exclude File Names/Types	Allows users to specify which files are transferred by filtering on the names or types to exclude. Separate multiple filters by commas. For example, if you type *.doc, *.ppt, in the field, the transfer will not include files with the doc and power point extensions. However, *.txt files, for example, would be included.
Exclude all Sub Directories	Choose Yes to specify that all subdirectories be excluded. Values in the Exclude Subdirectories field are ignored. Choose No to define specific subdirectories for exclusion.
Exclude Sub Directories	Subdirectories may be excluded from the data being backed up by specifying them in this field. Multiple entries must be separated with a comma. When the job runs, all directories that match those specified in the exclude sub-directories field will be excluded. Note that normal behavior is to exclude sub-directories that match regardless of where they appear in the directory path. Using the anchoring expression ("@") changes this behavior to anchor the exclude directory path. For example, if a user specifies a source directory of C:\data\docs, and an exclude directory of temp, any subdirectories called temp will be excluded, even those nested within another subdirectory. (For example the subdirectory C:\data\docs\publish\release\temp would be excluded as well.) Users who want to exclude a directory only at a certain level can use the @ symbol to anchor the exclude directory path at the starting source directory level. For example, specifying @temp in the above example would mean that the C:\data\docs\temp directory would be excluded, but the C:\data\docs\publish\release\temp directory would be included. If the user instead wanted to make sure just the C:\data\docs\publish\release\temp directory was excluded, but wanted to have a source directory path of c:\data\docs, the user would need to type C:\data\docs in the source directory field and @publish\release\temp in the exclude field. Special characters allow users to make use of pattern matching on the directory path. You must escape special characters with a backslash to be matched literally. Characters include the following: * (matches zero or more characters) ? (matches any single character) [...] (matches any one of the enclosed characters - for example, [ch] would match the characters "c" or "h") A pair of digit, lowercase or uppercase characters separated by a hyphen '-' denotes a range set and matches any character sorted in the range. If the first character following the '[' is '^' or '!', then any character not enclosed is matched. Use commas to specify multiple distinct patterns as an alternative to using multiple option specifications. Note that these options do not enable/ disable directory traversal.
Target Agent(s) Selection	
Target Agent(s)	The Agent(s) to which you want to transfer the files. Press shift +click to select multiple consecutive agents or ctrl +click to select multiple non- consecutive Agents.

Field	Description
Target Directory	<p>Specifies the directory in which to store the data from the source Agent. If the directory does not exist it will be created. This repository folder must have 'sharing' and correct permissions applied. To do so, right-click the directory to be backed up, select the Properties tab and click the Share this folder radio button. Click the Permissions button and set the read/ write privileges. The directory specified should be an absolute path name. The directory specified may be in the following formats:</p> <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • UNIX root drives; e.g., /home • Mac root drives; e.g., /Users Example entry:e:\restorefiles <p>Use the Browse button to select directories. Click through to the desired directory and click choose this directory. Leave the field empty to use the default Agent directory.</p>
Windows Volume Shadow Copy Service Options	
Use Volume Shadow Copy Service	<p>Specifies that a Windows Volume Shadow Copy of the source data is automatically created before the transfer. The transfer is performed from the shadow copy and then the shadow copy is released. This avoids transfer issues with open files. This option does not apply to Windows 2000 agents.</p> <p>NOTE: If you select VSS and you are using a 5.2 source Agent, the job will still run, but no volume shadow copy will be used (a warning appears in the log).</p>
Shadow Copy Creation Timeout	<p>Specifies the maximum number of seconds to wait for shadow copy creation before proceeding in non-shadow mode. The recommended minimum value is 60 seconds. The recommended maximum is 600 seconds.</p>

Field	Description
File Transfer Options	
File Ownership Transfer	<p>Specifies whether files maintain the source user ID and group ownership after they are transferred, and what method is used. Choose from the following options to preserve ownership:</p> <p>GuardianOS - When transferring between two GuardianOS machines, Permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Mac, Linux, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for Permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).</p> <p>Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its Permissions from a folder on the source is transferred to a folder on the target that has different Permissions. However, the source folder is not transferred. In this case, the file's inherited Permissions on the target will come from the target folder into which it is transferred and not from the source folder.</p> <p>NOTE: Windows attributes (such as Read Only) are not preserved in transfers between Windows and GuardianOS machines. Access Control Lists (ACLs) can be preserved, but attributes cannot. GuardianOS mode is the default mode of transfer.</p> <p>Windows - Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.</p> <p>Unix - Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX/Linux/ Mac hosts. The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.</p> <p>Off - Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by <i>root</i> on UNIX, <i>NT Authority/System</i> on Windows and <i>admin</i> on GuardianOS.</p>
Synchronize Source and Targets	<p>Specifies whether the Agent creates identical directory structures on the source and target nodes. If set to Yes, the Agent deletes any files in the target directory structure that do not have a corresponding file in the source directory structure, and transfers any files from the source that do not have a corresponding file in the target directory structure. The default value is No.</p>

Field	Description
Transfer File Differences Only	If set to Yes , only changed bytes of files will be transferred, not the entire file. Typically used in low bandwidth situations. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps).
Use Compression	<p>If set to Yes, the source Agent compresses each file before sending it. The files will be uncompressed automatically on the target Agent(s). The degree of compression depends on the type of data you are transferring. The following are typical rates of compression for different types of data:</p> <ul style="list-style-type: none"> • Plain text: 70-95% • TIFF images: 20-40% • Binary files: 0-5% <p>The default value is No. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps). This option is NOT recommended for LAN or high speed links.</p>
Delivery Mode	Specifies the mode to transfer files: Fast - Do not create temporary work files. Normal - Use temporary work files during transfer. Log File Name - Log file names transferred. Logs are stored on the Manager. Certify File Content - Create agent certified delivery log of files transferred.
Action If "File in Use"	Allows users to specify whether a file that is in use during a transfer is skipped or generates an error message. If Error is selected, and a file is in use during the transfer, the transfer fails and generates an error. If Skip is selected, and a file is in use during the transfer, the transfer continues, skipping the file, and issuing a warning. Using Volume Shadow Copy results in fewer cases of open files being encountered.
Ensure Source Directory is Mounted	Allows users to specify whether the job checks if the source directory is mounted before proceeding. If a user chooses Yes , the directory is checked against the current list of mounted directories. If the directory is not found in the list, the job aborts with an error. If you select No , the job runs without checking whether or not the directory is mounted. This option applies to only UNIX-based agents. If this option is enabled for heterogeneous transfers (Windows and UNIX), Windows agents are ignored.
Ensure Target Directory is Mounted	Allows users to specify whether the job checks if the target directory is mounted before proceeding. If a user chooses Yes , the directory is checked against the current list of mounted directories. If the directory is not found in the list, the job aborts with an error. If you select No , the job runs without checking whether or not the directory is mounted. This option applies to only UNIX-based agents. If this option is enabled for heterogeneous transfers (Windows and UNIX), Windows agents are ignored.

Field	Description
Job Options	
Bandwidth Throttle	<p>Click Add a new throttle to limit this job to the amount of bandwidth specified. In addition, click the plus sign ("+") to add additional throttles. (Click the "x" to delete a throttle.) Specifies the rate at which the source will make the data available to the network, regardless of how fast the machine is sending. Bandwidth Throttle is the rate at which the source reads the data from disk and puts it in the send queue. Note that bandwidth limiting is done on each stream connection, so a value specified here is passed to the controlling agent for each template executed, and divided among the streams as follows:</p> <ol style="list-style-type: none"> 1 It is divided by the number of streams the Agent will use for each remote Agent (typically four). 2 It is further divided by the potential number of concurrent remote Agents. This will be the lesser of the maximum allowed number of concurrent Agents, and the number of remote Agents specified in the template. 3 Note that you must use 24-hour time format to specify start and end times (i.e., enter 2:00 pm as 14). Click the clock icon to display 12-hour clock options from which you can choose and click OK. The value will be converted to 24-hour time. You can also specify days of the week to which this bandwidth limit applies. Once a job has started, all bandwidth throttles are applied at the times based on the Daylight Savings Time (DST) in effect when the job started. If DST changes while the job is running, bandwidth time of day changes may be off by the time change value (plus or minus an hour) after the time change. Note that bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.
Encryption Level	<p>Allows users to specify the encryption level from the following values: High, Medium, Low, or No Encryption - signed. The default value is High. Note that mutual authentication is always used regardless of the encryption level specified. The No encryption - signed option transfers unencrypted (plain text) data, but includes the SSL protocol's message digest calculation and signing to ensure data stream integrity. Note Encryption incompatibilities exist with transfers that involve version 5.2.2 build 1194 and higher (including version 7.2) and version 5.2.2 build 1177 and earlier agents. With Snap EDR version 5.2.2 build 1194 and higher, Snap EDR was enhanced to support AES-256 and AES-128 bit encryption. We recommend upgrading your Management Console and all of your agents. If all of the Snap EDR nodes are upgraded, the system will automatically use the "high" (AES-256) encryption when the replication jobs have been set to "high" (the default). If you choose not to upgrade all of the Snap EDR nodes, you must change all of the Snap EDR jobs to reflect "no" encryption; otherwise, the jobs will fail since the earlier versions of Snap EDR agents do not support the improved encryption. If running Snap EDR version 5.2.2 build 1194 or later, these encryption settings map to these encryption levels:</p> <ul style="list-style-type: none"> • High - AES-256 • Medium - AES-128 • Low - RC4-40
Job Log Detail Level	The type of logging information for this job. Choose from Error, Warn (warning), Info (Information) or Debug. Debug provides the greatest level of detail while Error provides the least.
Scheduling Parameters	
Job Start Date/ Time	The date and time at which you want the job to run.

Field	Description
Frequency	
None	Use none when you have a job that needs to be scheduled, but is run only at irregular, user-defined times.
Hourly, Daily, Weekly, Monthly, Yearly	Run the job once every selected interval.
Monthend	Run the job once every last day of the month
Every x interval	Run the job every x interval, the first of which is the "start job at" time. Intervals may be minutes, hours, days, weeks, months or years.
X<day> of the month	Run a job on a certain day every month.
Time Zone	Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time. Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere. Note that the time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	Text that appears in the subject field of the report e-mail message. NOTE: To receive email notifications for Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification .
SNMP Trap Notification	
Send Traps On	Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither. Note To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.

4. Click **Add Job**.

The new job appears in the list.

Replicate Data Management Tool

The Replicate data management tool allows users to schedule a simple one-to-one transfer where files are transferred from one source machine to one target machine.

Default Functionality

The Replicate data management tool has the following default functionality:

- File transfer from one source to one target host
- Cross platform support
- Option to exclude files from transfer
- Option to exclude directories from transfer
- Option to synchronize target host directory structure with source host directory structure
- Option to transfer changed bytes only, or entire file contents

By default, any files which are in use during the transfer will be skipped without error. This allows users to transfer common files from a central repository.

During a data transfer, the target Agent indicates that it is receiving data by the presence of `#work#{filename}` and `#check#{filename}` files. These represent a file that is partially received, and a file and the checkpoint data for an incoming file.

In the case of an incremental data transfer, the target Agent will generate checksum information for the file. (A checksum is a numeric value used to verify the integrity of a block of data.) The target Agent will send this checksum information back to the source. When byte level differences are encountered, the source will begin sending the differences to the target and the target will start creating a new file by making a copy of the "older" file and applying the byte level differences.

Users can set information in the following categories:

- Source host selection
- Directory and file options
- Target Agent selection
- Windows Volume Shadow Copy Service options
- File transfer options
- Job options
- Scheduling parameters
- E-mail notification
- SNMP trap notification

Create a Replicate Job

To schedule a Replicate job:

1. From the Management Console, click **Snap Solutions > Replicate**.
2. Click **Add a Job**.

If this is the first time you are creating a replicate job, when you click Replicate, the scheduling screen appears. You do not have to click Add a Job.

3. Enter information into the various fields. The following table describes the fields:

NOTE: You cannot use the '|' (pipe) character in a Windows directory path or errors will result.

Field	Description
Job Name	A unique name for this job run.
Source Agent Selection	
Source Agent	The Agent from which you want to transfer the files.
Source Directory	Directory from which you want to transfer files. The source directory specified may be in the following format: <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • UNIX/Linux root drives; e.g., /home • Mac root drives; e.g., /UsersExample entry:e:\databasefiles
Directory and File Options	
Include File Names/Types	Allows users to specify which files are transferred by filtering on the names to include. Use a comma to separate multiple filters. For example, if you type *.doc, *.ppt, in the field, the transfer will include files with the doc and power point extensions. However, *.txt files, for example, would be excluded.
Exclude File Names/Types	Allows users to specify which files are transferred by filtering on the names to exclude. Use a comma to separate multiple filters. For example, if you type *.doc, *.ppt, in the field, the transfer will include files with the doc and power point extensions. However, *.txt files, for example, would be excluded.
Exclude all Sub-Directories	Choose Yes to specify that all subdirectories be excluded. Values in the Exclude Subdirectories field are ignored. Choose No to define specific subdirectories for exclusion.

Field	Description
Exclude Sub Directories	<p>Subdirectories may be excluded from the data being backed up by specifying them in this field. Multiple entries must be separated with a comma. When the job runs, all directories that match those specified in the exclude subdirectories field will be excluded. Note that normal behavior is to exclude subdirectories that match regardless of where they appear in the directory path. Using the anchoring expression ("@") changes this behavior to anchor the exclude directory path.</p> <p>For example, if a user specifies a source directory of <code>C:\data\docs</code>, and an exclude directory of <code>temp</code>, any subdirectories called <code>temp</code> will be excluded, even those nested within another subdirectory. (For example the subdirectory <code>C:\data\docs\publish\release\temp</code> would be excluded as well.)</p> <p>Users who want to exclude a directory only at a certain level can use the @ symbol to anchor the exclude directory path at the starting source directory level. For example, specifying <code>@temp</code> in the above example would mean that the <code>C:\data\docs\temp</code> directory would be excluded, but the <code>C:\data\docs\publish\release\temp</code> directory would be included.</p> <p>If the user instead wanted to make sure just the <code>C:\data\docs\publish\release\temp</code> directory was excluded, but wanted to have a source directory path of <code>c:\data\docs</code>, the user would need to type <code>C:\data\docs</code> in the source directory field and <code>@publish\release\temp</code> in the exclude field.</p> <p>Special characters allow users to make use of pattern matching on the directory path. You must escape special characters with a backslash to be matched literally. Characters include the following:</p> <ul style="list-style-type: none"> * (matches zero or more characters) ? (matches any single character) [...] (matches any one of the enclosed characters - for example, [ch] would match the characters "c" or "h") <p>A pair of digit, lowercase or uppercase characters separated by a hyphen '-' denotes a range set and matches any character sorted in the range. If the first character following the '[' is '^' or ' ', then any character not enclosed is matched.</p> <p>Use commas to specify multiple distinct patterns as an alternative to using multiple option specifications. Note that these options do not enable/ disable directory traversal.</p>
Target Agent Selection	
Target Agent	The Agent to which you want to transfer the files.
Target Directory	<p>Specifies the directory to store the data from the source Agent. If the directory does not exist it will be created. This repository folder must have 'sharing' and correct permissions applied. To do so, right-click the directory to be backed up, select the Properties tab and click the Share this folder radio button. Click the Permissions button and set the read/write privileges. The directory specified should be an absolute path name. The directory specified may be in the following formats:</p> <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • UNIX/Linux root drives; e.g., /home • Mac root drives; e.g., /Users <p>Example entry: <code>e:\restorefiles</code> Use the Browse button to select directories. Leave the field empty to use the default Agent directory.</p>
Windows Volume Shadow Copy Service Options	
Use Volume Shadow Copy Service	<p>Specifies that a Windows Volume Shadow Copy of the source data is automatically created before the transfer. The transfer is performed from the shadow copy and then the shadow copy is released. This avoids transfer issues with open files. This option does not apply to Windows 2000 Agents. Using Volume Shadow Copy results in fewer cases of open files being encountered.</p> <p>NOTE: If you select VSS and you are using a 5.2 source Agent, the job will still run, but no volume shadow copy will be used (a warning appears in the log).</p>

Field	Description
Shadow Copy Creation Timeout	Specifies the maximum number of seconds to wait for shadow copy creation before proceeding in non-shadow mode. The recommended minimum value is 60 seconds. The recommended maximum is 600 seconds.
File Transfer Options	
File Ownership Transfer	<p>Specifies whether files maintain the source user ID and group ownership after they are transferred, and what method is used. Choose from the following options to preserve ownership:</p> <p>GuardianOS - When transferring between two GuardianOS machines, permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Linux, Mac, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).</p> <p>Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its permissions from a folder on the source is transferred to a folder on the target that has different permissions. However, the source folder is not transferred. In this case, the file's inherited permissions on the target will come from the target folder into which it is transferred and not from the source folder.</p> <p>NOTE: Windows attributes (such as Read Only) are not preserved in transfers between Windows and GuardianOS machines. Access Control Lists (ACLs) can be preserved, but attributes cannot. GuardianOS mode is the default mode of transfer.</p> <p>Windows - Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.</p> <p>Unix - Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX/Linux/ Mac hosts.</p> <p>The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.</p> <p>Off - Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.</p>
Synchronize Source and Targets	Specifies whether the Agent creates identical directory structures on the source and target nodes. If set to Yes, the Agent deletes any files in the target directory structure that do not have a corresponding file in the source directory structure, and transfers any files from the source that do not have a corresponding file in the target directory structure. The default value is No.
Transfer File Differences Only	If set to Yes, only changed bytes of files will be transferred, not the entire file. Typically used in low bandwidth situations. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps).
Use Compression	<p>If set to Yes, the source Agent compresses each file before sending it. The files will be uncompressed automatically on the target Agent(s). The degree of compression depends on the type of data you are transferring. The following are typical rates of compression for different types of data:</p> <ul style="list-style-type: none"> • Plain text: 70-95% • TIFF images: 20-40% • Binary files: 0-5% <p>The default value is No. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps). This option is NOT recommended for LAN or high speed links.</p>

Field	Description
Delivery Mode	<p>Specifies the mode to transfer files:</p> <ul style="list-style-type: none"> • Fast - Do not create temporary work files. • Normal - Use temporary work files during transfer. • Log File Name - Log file names transferred. • Logs are stored on the Manager. • Certify File Content - Create Agent certified delivery log of files transferred.
Action If "File in Use"	<p>Allows users to specify whether a file that is in use during a transfer is skipped or generates an error message. If Error is selected, and a file is in use during the transfer, the transfer fails and generates an error. If Skip is selected, and a file is in use during the transfer, the transfer continues, skipping the file, and issuing a warning.</p> <p>Using Volume Shadow Copy results in fewer cases of open files being encountered.</p>
Ensure source directory is mounted	<p>Allows users to specify whether the job checks if the source directory is mounted before proceeding. If a user chooses Yes, the directory is checked against the current list of mounted directories. If the directory is not found in the list, the job aborts with an error. If you select No, the job runs without checking whether or not the directory is mounted.</p> <p>NOTE: This option applies to only UNIX-based agents. If this option is enabled for heterogeneous transfers (Windows and UNIX), Windows agents are ignored.</p>
Ensure target directory is mounted	<p>Allows users to specify whether the job checks if the target directory is mounted before proceeding. If a user chooses Yes, the directory is checked against the current list of mounted directories. If the directory is not found in the list, the job aborts with an error. If you select No, the job runs without checking whether or not the directory is mounted. Note This option applies to only UNIX-based agents. If this option is enabled for heterogeneous transfers (Windows and UNIX), Windows agents are ignored.</p>
Job Options	
Bandwidth Throttle	<p>Click Add a new throttle to limit this job to the amount of bandwidth specified. In addition, click the plus sign ("+") to add additional throttles. (Click the "x" to delete a throttle.)</p> <p>Specifies the rate at which the source will make the data available to the network, regardless of how fast the machine is sending. Bandwidth Throttle is the rate at which the source reads the data from disk and puts it in the send queue. Note that bandwidth limiting is done on each stream connection, so a value specified here is passed to the controlling agent for each template executed, and divided among the streams as follows:</p> <ol style="list-style-type: none"> 1 It is divided by the number of streams the agent will use for each remote agent (typically four). 2 It is further divided by the potential number of concurrent remote agents. This will be the lesser of the maximum allowed number of concurrent agents, and the number of remote agents specified in the template. 3 Note that you must use 24-hour time format to specify start and end times (i.e., enter 2:00 pm as 14). Click the clock icon to display 12-hour clock options from which you can choose and click OK. The value will be converted to 24-hour time. You can also specify days of the week to which this bandwidth limit applies. <p>Once a job has started, all bandwidth throttles are applied at the times based on the Daylight Savings Time (DST) in effect when the job started. If DST changes while the job is running, bandwidth time of day changes may be off by the time change value (plus or minus an hour) after the time change.</p> <p>NOTE: Bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.</p>

Field	Description
Encryption Level	<p>Allows users to specify the encryption level from the following values: High, Medium, Low, or No Encryption - signed. The default value is High.</p> <p>NOTE: Mutual authentication is always used regardless of the encryption level specified.</p> <p>The No encryption - signed option transfers unencrypted (plain text) data, but includes the SSL protocol's message digest calculation and signing to ensure data stream integrity.</p> <p>NOTE: Encryption incompatibilities exist with transfers that involve version 5.2.2 build 1194 and higher (including version 7.2) and version 5.2.2 build 1177 and earlier agents.</p> <p>With Snap EDR version 5.2.2 build 1194 and higher, Snap EDR was enhanced to support AES-256 and AES-128 bit encryption. We recommend upgrading your Management Console and all of your agents. If all of the Snap EDR nodes are upgraded, the system will automatically use the "high" (AES-256) encryption when the replication jobs have been set to "high" (the default).</p> <p>If you choose not to upgrade all of the Snap EDR nodes, you must change all of the Snap EDR jobs to reflect "no" encryption; otherwise, the jobs will fail since the earlier versions of Snap EDR agents do not support the improved encryption.</p> <p>If running Snap EDR version 5.2.2 build 1194 or later, these encryption settings map to these encryption levels:</p> <ul style="list-style-type: none"> • High - AES-256 • Medium - AES-128 • Low - RC4-40
Job Log Detail Level	The type of logging information for this job. Choose from Error, Warn (warning), Info (Information) or Debug. Debug provides the greatest level of detail while Error provides the least.
Scheduling Parameters	
Job Start Date/Time	The date and time at which you want the job to run.
Frequency	
Once	Run the job only once at start time.
Hourly, Daily, Weekly, Monthly, Yearly	Run the job once every selected interval.
Monthend	Run the job once every last day of the month.
None	Use none when you have a job that needs to be scheduled, but is run only at irregular, user-defined times.
Every x interval	Run the job every x interval, the first of which is the "start job at" time. Intervals may be minutes, hours, days, weeks, months or years.
X<day> of the month	Run a job on a certain day every month.
Time Zone	<p>Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time.</p> <p>Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere.</p> <p>NOTE: The time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.</p>

Field	Description
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	<p>Text that appears in the subject field of the report e-mail message.</p> <p>NOTE: To receive email notifications for Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification.</p>
SNMP Trap Notification	<p>Send Traps On Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither.</p> <p>NOTE: To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.</p>
Target Agent Selection	
Target Agent	The Agent to which you want to transfer the files.
Target Directory	<p>Specifies the directory to store the data from the source Agent. If the directory does not exist it will be created. This repository folder must have 'sharing' and correct permissions applied. To do so, right-click the directory to be backed up, select the Properties tab and click the Share this folder radio button. Click the Permissions button and set the read/write privileges. The directory specified should be an absolute path name. The directory specified may be in the following formats:</p> <ul style="list-style-type: none"> • SnapServer shares; e.g., /shares/SHARE1 • Windows root drives; e.g., C:\ApplicationData • UNIX/Linux root drives; e.g., /home • Mac root drives; e.g., /Users <p>Example entry:e:\restorefiles</p> <p>Use the Browse button to select directories. Leave the field empty to use the default Agent directory.</p>
Windows Volume Shadow Copy Service Options	
Use Volume Shadow Copy Service	<p>Specifies that a Windows Volume Shadow Copy of the source data is automatically created before the transfer. The transfer is performed from the shadow copy and then the shadow copy is released. This avoids transfer issues with open files. This option does not apply to Windows 2000 Agents. Using Volume Shadow Copy results in fewer cases of open files being encountered.</p> <p>NOTE: If you select VSS and you are using a 5.2 source Agent, the job will still run, but no volume shadow copy will be used (a warning appears in the log).</p>
Shadow Copy Creation Timeout	Specifies the maximum number of seconds to wait for shadow copy creation before proceeding in non-shadow mode. The recommended minimum value is 60 seconds. The recommended maximum is 600 seconds.

Field	Description
File Transfer Options	
File Ownership Transfer	<p>Specifies whether files maintain the source user ID and group ownership after they are transferred, and what method is used. Choose from the following options to preserve ownership:</p> <p>GuardianOS - When transferring between two GuardianOS machines, permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Linux, Mac, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1). Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its permissions from a folder on the source is transferred to a folder on the target that has different permissions. However, the source folder is not transferred. In this case, the file's inherited permissions on the target will come from the target folder into which it is transferred and not from the source folder. Note Windows attributes (such as Read Only) are not preserved in transfers between Windows and GuardianOS machines. Access Control Lists (ACLs) can be preserved, but attributes cannot. GuardianOS mode is the default mode of transfer.</p> <p>Windows - Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.</p> <p>Unix - Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX/Linux/ Mac hosts. The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards. Off - Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.</p>
Synchronize Source and Targets	Specifies whether the Agent creates identical directory structures on the source and target nodes. If set to Yes, the Agent deletes any files in the target directory structure that do not have a corresponding file in the source directory structure, and transfers any files from the source that do not have a corresponding file in the target directory structure. The default value is No.
Transfer File Differences Only	If set to Yes, only changed bytes of files will be transferred, not the entire file. Typically used in low bandwidth situations. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps).
Use Compression	If set to Yes, the source Agent compresses each file before sending it. The files will be uncompressed automatically on the target Agent(s). The degree of compression depends on the type of data you are transferring. The following are typical rates of compression for different types of data: Plain text: 70-95% TIFF images: 20-40% Binary files: 0-5% The default value is No. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps). This option is NOT recommended for LAN or high speed links.
Delivery Mode	Specifies the mode to transfer files: Fast - Do not create temporary work files. Normal - Use temporary work files during transfer. Log File Name - Log file names transferred. Logs are stored on the Manager. Certify File Content - Create Agent certified delivery log of files transferred.
Action If "File in Use"	Allows users to specify whether a file that is in use during a transfer is skipped or generates an error message. If Error is selected, and a file is in use during the transfer, the transfer fails and generates an error. If Skip is selected, and a file is in use during the transfer, the transfer continues, skipping the file, and issuing a warning. Using Volume Shadow Copy results in fewer cases of open files being encountered.

Field	Description
Ensure source directory is mounted	Allows users to specify whether the job checks if the source directory is mounted before proceeding. If a user chooses Yes, the directory is checked against the current list of mounted directories. If the directory is not found in the list, the job aborts with an error. If you select No, the job runs without checking whether or not the directory is mounted. Note This option applies to only UNIX-based agents. If this option is enabled for heterogeneous transfers (Windows and UNIX), Windows agents are ignored.
Ensure target directory is mounted	Allows users to specify whether the job checks if the target directory is mounted before proceeding. If a user chooses Yes, the directory is checked against the current list of mounted directories. If the directory is not found in the list, the job aborts with an error. If you select No, the job runs without checking whether or not the directory is mounted. Note This option applies to only UNIX-based agents. If this option is enabled for heterogeneous transfers (Windows and UNIX), Windows agents are ignored.
Job Options	
Bandwidth Throttle	<p>Click Add a new throttle to limit this job to the amount of bandwidth specified. In addition, click the plus sign "+" to add additional throttles. (Click the "x" to delete a throttle.) Specifies the rate at which the source will make the data available to the network, regardless of how fast the machine is sending. Bandwidth Throttle is the rate at which the source reads the data from disk and puts it in the send queue. Note that bandwidth limiting is done on each stream connection, so a value specified here is passed to the controlling agent for each template executed, and divided among the streams as follows:</p> <ol style="list-style-type: none"> 1 It is divided by the number of streams the agent will use for each remote agent (typically four). 2 It is further divided by the potential number of concurrent remote agents. This will be the lesser of the maximum allowed number of concurrent agents, and the number of remote agents specified in the template. 3 Note that you must use 24-hour time format to specify start and end times (i.e., enter 2:00 pm as 14). Click the clock icon to display 12-hour clock options from which you can choose and click OK. The value will be converted to 24-hour time. You can also specify days of the week to which this bandwidth limit applies. <p>Once a job has started, all bandwidth throttles are applied at the times based on the Daylight Savings Time (DST) in effect when the job started. If DST changes while the job is running, bandwidth time of day changes may be off by the time change value (plus or minus an hour) after the time change. Note that bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.</p>
Encryption Level	<p>Allows users to specify the encryption level from the following values: High, Medium, Low, or No Encryption - signed. The default value is High. Note that mutual authentication is always used regardless of the encryption level specified. The No encryption - signed option transfers unencrypted (plain text) data, but includes the SSL protocol's message digest calculation and signing to ensure data stream integrity. Note Encryption incompatibilities exist with transfers that involve version 5.2.2 build 1194 and higher (including version 7.2) and version 5.2.2 build 1177 and earlier agents. With Snap EDR version 5.2.2 build 1194 and higher, Snap EDR was enhanced to support AES-256 and AES-128 bit encryption. We recommend upgrading your Management Console and all of your agents.</p> <p>If all of the Snap EDR nodes are upgraded, the system will automatically use the High (AES-256) encryption when the replication jobs have been set to High (the default). If you choose not to upgrade all of the Snap EDR nodes, you must change all of the Snap EDR jobs to reflect no encryption; otherwise, the jobs will fail since the earlier versions of Snap EDR agents do not support the improved encryption. If running Snap EDR version 5.2.2 build 1194 or later, these encryption settings map to these encryption levels:</p> <ul style="list-style-type: none"> • High - AES-256 • Medium - AES-128 • Low - RC4-40

Field	Description
Job Log Detail Level	The type of logging information for this job. Choose from Error, Warn (warning), Info (Information) or Debug. Debug provides the greatest level of detail while Error provides the least.
Scheduling Parameters	
Job Start Date/Time	The date and time at which you want the job to run.
Frequency	
Once	Run the job only once at start time.
Hourly, Daily, Weekly, Monthly, Yearly	Run the job once every selected interval.
Monthend	Run the job once every last day of the month.
None	Use none when you have a job that needs to be scheduled, but is run only at irregular, user-defined times.
Every x interval	Run the job every x interval, the first of which is the "start job at" time. Intervals may be minutes, hours, days, weeks, months or years.
X<day> of the month	Run a job on a certain day every month.
Time Zone	<p>Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time.</p> <p>Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere.</p> <p>Note that the time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.</p>
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	Text that appears in the subject field of the report e-mail message. Note To receive email notifications for Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification.
SNMP Trap Notification	Send Traps On Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither. Note To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.

4. Click Add Job.

The new job appears in the list.

Remote Backup Data Management Tool

NOTE: This tool is not available for Snap EDR Express.

The Remote Backup solution provides a method for backing up data from remote hosts to a central host. Data may also be backed up on a specific host to a different storage device. An application interface with Microsoft System State backup allows the backup data set to include data from this application.

The following are the features of the Remote Backup solution:

- Allow data backup from multiple hosts to a central location
- Backup data from heterogeneous host platforms to a single central location
- Specify multiple directories and mount points
- Configure number of backups to retain
- Backup and store only the files that have changed to reduce disk usage on the target host attached storage
- File backups take advantage of incremental changes in files, reducing the amount of bandwidth utilized during backup
- Preservation of file ownership permissions
- Apply file ownership changes to files at the central location, even if the file contents have not changed
- Full scheduling functionality for automating frequency of backups
- Centralized control for viewing job status, controlling job execution, and managing hosts involved in backup
- Manage the amount of bandwidth the backup job will utilize
- Ability to synchronize the directory structure and files between the remote hosts and the target host's locally-attached storage
- Network fault tolerance during backups
- Microsoft System State backup integration
- Full flexibility to allow for customization of solution to meet administrator needs

For information about Remote Backup recommended configurations and best practices, see [Best Practices](#).

The Restore solution is a complementary solution to the Remote Backup solution. For more information on the Restore solution, see [Remote Restore Data Management Tool](#).

Create a Remote Backup Job

To schedule a Remote Backup job:

1. From the Management Console, click **Snap Solutions > Remote Backup**.
2. Click **Add a Job**.

If this is the first time you are creating a remote backup job, when you click **Remote Backup**, the scheduling screen appears. You do not have to click **Add a Job**.

3. Enter information into the various fields. The following table describes the fields.

NOTE: You cannot use the '|' (pipe) character in a Windows directory path or errors will result.

Field	Description
Job Name	A unique name for this job run.
Backup Repository	
Repository Host	Select a target host to store the backup data from the source hosts.
Repository Base Directory	<p>The target directory to store the remote backup data specifies the base directory to store the backup data. If the directory does not exist it will be created. Data for each host specified is stored on a per host basis under this directory.</p> <p>A minimum of one directory should be specified for this field. The directory specified should be a fully qualified path name. Directories specified may be in the following format:</p> <p>Guardian OS root drives, e.g., /home</p> <p>Example Entry: /remotebackups</p>
Version Directories to Maintain	<p>Specifies the number of backup copies of data to maintain per host, to a maximum of 30. When a change in a file is detected between the central repository's latest version and the source host, before the file is transferred from the source host to replace the latest version, a copy is stored in the backup directories for that host. Every backup run in which changes are detected on the source host counts as a backup copy.</p> <p>Once the specified number of backup copies (directories) has been reached, the oldest backup copy (directory) will be removed.</p> <p>Specifying a value of 0 indicates only the latest version will be maintained in the target host storage and no backup copies will be maintained.</p>
Log Versions to Maintain	Specifies the number of copies of the job log to maintain, to a maximum of 100. (Each time the job runs, it generates a log file.) Once the maximum value is reached, the oldest version of the log is deleted.
Hosts to Back Up	
Source Hosts	Select one or more source hosts from which to backup data. Press shift+click to choose multiple consecutive hosts or ctrl+click to choose multiple non-consecutive hosts.
General Directory and File Options	
Source Directories	<p>The Source Directories to backup specifies the directories to backup on the source hosts. The directories specified must be fully qualified path names. A comma must separate multiple directories.</p> <p>This field must be specified, unless only application data is being backed up as part of the scheduled job. These directories should exist on all of the hosts specified as source hosts for the job. Directories specified may be in the following format:</p> <p>Windows root drives, e.g., C:\ApplicationData</p> <p>UNIX/Linux root drives, e.g., /home</p> <p>Mac root drives, e.g., /Users</p> <p>Example entry: e:\databasefiles,c:\userdata</p>
Include File Names/Types	Files may be included in the data being backed up by specifying one or more file masks in the files to include field. A comma must separate multiple data masks. Example Entry *.jpg,userdata1.doc

Field	Description
Exclude File Names/Types	<p>Files may be excluded in the data being backed up by specifying one or more file masks in the files to exclude field. A comma must separate multiple data masks.</p> <p>Example Entry: *.jpg,userdata1.doc</p>
Exclude Sub Directories	<p>Subdirectories may be excluded from the data being backed up by specifying them in this field. Multiple entries must be separated with a comma. When the job runs, all directories that match those specified in the Exclude Subdirectories field will be excluded. Note that normal behavior is to exclude subdirectories that match regardless of where they appear in the directory path. Using the anchoring expression ("@")changes this behavior to anchor the exclude directory path.</p> <p>For example, if a user specifies a source directory of C:\data\docs, and an exclude directory of temp, any subdirectories called temp will be excluded, even those nested within another sub-directory. (For example the subdirectory C:\data\docs\publish\release\temp would be excluded as well.)</p> <p>Users who want to exclude a directory only at a certain level can use the @ symbol to anchor the exclude directory path at the starting source directory level. For example, specifying @temp in the above example would mean that the C:\data\docs\temp directory would be excluded, but the C:\data\docs\publish\release\temp directory would be included.</p> <p>If the user instead wanted to make sure just the C:\data\docs\publish\release\temp directory was excluded, but wanted to have a source directory path of c:\data\docs, the user would need to type C:\data\docs in the source directory field and @publish\release\temp in the exclude field.</p> <p>Special characters allow users to make use of pattern matching on the directory path. You must escape special characters with a backslash to be matched literally. Characters include the following:</p> <ul style="list-style-type: none"> * (matches zero or more characters) ? (matches any single character) [...] (matches any one of the enclosed characters - for example, [ch] would match the characters "c" or "h") <p>A pair of digit, lowercase or uppercase characters separated by a hyphen '-' denotes a range set and matches any character sorted in the range. If the first character following the '[' is '^' or '!', then any character not enclosed is matched.</p> <p>Use commas to specify multiple distinct patterns as an alternative to using multiple option specifications. Note that these options do not enable/disable directory traversal.</p>
Backup Empty Directories	Specifies whether or not directories that do not include any files are backed up.
Windows Directory and File Options	
Exclude Program Files Folder	When set to Yes, the Program Files folder is not included in the backup.
Exclude Windows Folder	When set to Yes, the Windows system folder is not included in the backup.
Exclude System Volume Information Folder	When set to Yes, the System Volume Information folder is not included in the backup.
Exclude Recycle Bin	When set to Yes, the Recycle bin is not included in the backup.

Field	Description
Windows Volume Shadow Copy Service Options	
Use Volume Shadow Copy Service	<p>Specifies that a Windows Volume Shadow Copy of the source data is automatically created before the transfer. The transfer is performed from the shadow copy and then the shadow copy is released. This avoids transfer issues with open files. This option does not apply to Windows 2000 Agents.</p> <p>NOTE: If you select VSS and you are using a 5.2 source Agent, the job will still run, but no volume shadow copy will be used (a warning appears in the log).</p>
Shadow Copy Creation Timeout	<p>Specifies the maximum number of seconds to wait for shadow copy creation before proceeding in non-shadow mode. The recommended minimum value is 60 seconds. The recommended maximum is 600 seconds.</p>
Windows SystemState Backup	
Back up SystemState	<p>Specifies whether to run Microsoft SystemState backup or not as part of the backup job. This option is ignored on non-Windows hosts. Application Backup Details describes the behaviors of the Microsoft SystemState backup.</p>
Directory to Export Data	<p>If SystemState backup has been specified, a directory to export the SystemState data must be specified. The SystemState data is exported to disk on the source, and then transferred to the target host storage as part of the backup procedure. The directory specified must have enough space to hold the exported SystemState data. Once the transfer is complete, the data file is deleted from the source host.</p>
File Transfer Options	
Transfer File Differences Only	<p>If set to Yes, only changed bytes of files will be transferred, not the entire file. Typically used in low bandwidth situations.</p> <p>Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps).</p>
Remove Deleted Files from Version 'Latest'	<p>When set to Yes, any files deleted on the source host will either be moved to a date/time stamped version directory in the backup repository (if 'Version Directories to Maintain' is greater than zero) or simply deleted (if 'Version Directories to Maintain' is zero). When set to No, any files deleted on the source host will remain in the "latest" folder in the backup repository until such time as another backup is run with "Remove Deleted Files from Version 'Latest'" set to "Yes". This option is provided as a performance feature, since the synchronization overhead and time involved in determining files to delete can be significant in very large backups. You can specify the number of file versions to keep (up to 100) in the Version Directories to Maintain field.</p> <p>NOTE: If the path to back up is greater than 256 characters, and versioning is specified, the old file is not placed in the versions directory. The changed file is still transferred, but the older version of the file is not saved.</p>
Log Skipped Files	<p>Specifies that any files that are skipped be noted in the log file.</p>
Use Compression	<p>If set to Yes, the source agent compresses each file before sending it. The files will be uncompressed automatically on the target agent(s). The degree of compression depends on the type of data you are transferring. The following are typical rates of compression for different types of data:</p> <ul style="list-style-type: none"> • Plain text: 70-95% • TIFF images: 20-40% • Binary files: 0-5% <p>The default value is No. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps). This option is NOT recommended for LAN or high speed links.</p>

Field	Description
File Ownership Transfer	<p>Specifies whether files maintain the source user ID and group ownership after they are transferred, and what method is used. Choose from the following options to preserve ownership:</p> <p>GuardianOS - When transferring between two GuardianOS machines, permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Linux, Mac, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).</p> <p>Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its permissions from a folder on the source is transferred to a folder on the target that has different permissions. However, the source folder is not transferred. In this case, the file's inherited permissions on the target will come from the target folder into which it is transferred and not from the source folder.</p> <p>NOTE: Windows attributes (such as Read Only) are not preserved in transfers between Windows and GuardianOS machines. Access Control Lists (ACLs) can be preserved, but attributes cannot.</p> <p>GuardianOS mode is the default mode of transfer.</p> <p>Windows - Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.</p> <p>Unix - Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX/Linux/Mac hosts.</p> <p>The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.</p> <p>Off - Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.</p>

Field	Description
Job Options	
Bandwidth Throttle	<p>Click Add a new throttle to limit this job to the amount of bandwidth specified. In addition, click the plus sign ("+") to add additional throttles. (Click the "x" to delete a throttle.)</p> <p>Specifies the rate at which the source will make the data available to the network, regardless of how fast the machine is sending. Bandwidth Throttle is the rate at which the source reads the data from disk and puts it in the send queue.</p> <p>Note that bandwidth limiting is done on each stream connection, so a value specified here is passed to the controlling agent for each template executed, and divided among the streams as follows:</p> <ol style="list-style-type: none"> 1 It is divided by the number of streams the agent will use for each remote agent (typically four). 2 It is further divided by the potential number of concurrent remote agents. This will be the lesser of the maximum allowed number of concurrent agents, and the number of remote agents specified in the template. 3 Note that you must use 24-hour time format to specify start and end times (i.e., enter 2:00 pm as 14). Click the clock icon to display 12-hour clock options from which you can choose and click OK. The value will be converted to 24-hour time. You can also specify days of the week to which this bandwidth limit applies. <p>Once a job has started, all bandwidth throttles are applied at the times based on the Daylight Savings Time (DST) in effect when the job started. If DST changes while the job is running, bandwidth time of day changes may be off by the time change value (plus or minus an hour) after the time change.</p> <p>Note that bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.</p>

Field	Description
Encryption Level	<p>Allows users to specify the encryption level from the following values: High, Medium, Low, or No Encryption - signed. The default value is High. Note that mutual authentication is always used regardless of the encryption level specified.</p> <p>The No encryption - signed option transfers unencrypted (plain text) data, but includes the SSL protocol's message digest calculation and signing to ensure data stream integrity.</p> <p>NOTE: Encryption incompatibilities exist with transfers that involve version 5.2.2 build 1194 and higher (including version 7.2) and version 5.2.2 build 1177 and earlier agents.</p> <p>With Snap EDR version 5.2.2 build 1194 and higher, Snap EDR was enhanced to support AES-256 and AES-128 bit encryption. We recommend upgrading your Management Console and all of your agents. If all of the Snap EDR nodes are upgraded, the system will automatically use the "high" (AES-256) encryption when the replication jobs have been set to "high" (the default).</p> <p>If you choose not to upgrade all of the Snap EDR nodes, you must change all of the Snap EDR jobs to reflect "no" encryption; otherwise, the jobs will fail since the earlier versions of Snap EDR agents do not support the improved encryption.</p> <p>If running Snap EDR version 5.2.2 build 1194 or later, these encryption settings map to these encryption levels:</p> <ul style="list-style-type: none"> • High - AES-256 • Medium - AES-128 • Low - RC4-40
Job Log Detail Level	The type of logging information for this job. Choose from Error , Warn (warning), Info (Information) or Debug . Debug provides the greatest level of detail while Error provides the least.
Scheduling Parameters	
Job Start Date/ Time	The date and time at which you want the job to run.
Frequency	
None	Use none when you have a job that needs to be scheduled, but is run only at irregular, user-defined times.
Hourly, Daily, Weekly, Monthly, Yearly	Run the job once every selected interval.
Monthend	Run the job once every last day of the month.
Once	Run the job only once at start time.
Days of Week to Run	By default, the Days of Week to Run area specifies that backup jobs run every day. Unchecking the box beside a specific day or days allows users to specify days of the week on which the job should not run. For example, you may schedule a job to run on April 15, 2008 at 10:00 AM, with a frequency of daily, but you do not want the job to run on Saturday and Sunday. Removing the check from the Saturday and Sunday days of week boxes means the job will run at the specified time, Monday to Friday inclusive. The day specified is relative to the time zone where the Management Console is located.

Field	Description
Time Zone	<p>Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time.</p> <p>Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere.</p> <p>Note that the time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.</p>
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	<p>Text that appears in the subject field of the report e-mail message.</p> <p>NOTE: To receive email notifications for Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification.</p>
SNMP Trap Notification	
Send Traps On	<p>Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither.</p> <p>NOTE: To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.</p>

4. Click **Add Job**.

The new job appears in the list.

File Backup Details

This section describes how backup data is stored on the target host. It also describes how backup versions are maintained.

File Transfer Options

The files are transferred between the source hosts and target host using the following file transfer options:

1. Files are considered for transfer if the file size and/or date and time are not equal between the source and target file system.

2. By default, the files are transferred using incremental transfer. This reduces the amount of data that is sent because only the changed portions of the files are transferred. Users can turn this feature off when scheduling the job.
3. File ownership preservation is set to "Windows" by default. Selecting one of the file ownership options when scheduling is recommended:
 - **GuardianOS**—When transferring between two GuardianOS machines, Permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Mac, Linux, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for Permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).

Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its Permissions from a folder on the source is transferred to a folder on the target that has different Permissions. However, the source folder is not transferred. In this case, the file's inherited Permissions on the target will come from the target folder into which it is transferred and not from the source folder.

Note that Windows attributes (such as read only, and so on) are not preserved in transfers between Windows and GoS machines. Access Control Lists can be preserved, but attributes cannot.

GuardianOS mode is the default mode of transfer.

- **Windows**—Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.
 - **UNIX**—Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX, Windows and Mac, and Unix and Mac hosts.
- The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.
- **Off**—Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.
4. Synchronize source host and target host. When this option is enabled, the source directories and files will be mirrored with the corresponding directories on the target for the specified source host.

Data Layout in Target Host Storage Location

The backup data is stored on the target host on a per-host basis. A separate folder for each source host specified in the Remote Backup job will be created under the specified target directory. The directory per source host will be named with the hostname of the source host.

An example of the directory layout on the target host would be as follows: User specified target directory as /BackupData

User specified source hosts in company.com, host1, host2, and host3.

```
/BackupData/  
  host1.company.com  
  host2.company.com  
  host3.company.com
```

The latest versions of the files backed up from the source hosts are always maintained in a directory called latest under the source host directory. The directory structure under the latest directory depends on the original location of the data on the source host. Files from Windows drives are stored under the latest\drives directory. Files from GuardianOS/UNIX/Linux systems are stored under the latest\root directory.

The following is an example of the directory layout:

```
/BackupData/host1.company.com/latest/  
  drives  
  root
```

The drive letters under latest drives directory are converted to directories. For example C:\userdata is converted to ..\latest\c\userdata.

To permit easy identification, application data backed up from the source hosts is stored in a different location. The data is stored under the latest directory in a subdirectory called applicationdata. A separate subdirectory under applicationdata is created for each application backed up. Microsoft System State data is stored in a subdirectory called systemstatedata. An example of the directory layout is shown below:

```
/BackupData/host1.company.com/latest/applicationdata/systemstatedata
```

Backup Directories in the Target Host Storage Location

During a backup, if files have been modified on the source host, they will be transferred to the target host and stored on a per-host basis in the latest directory as specified above. Prior to the version of the file being replaced in the latest directory from the source host, the file is copied to a backup directory. The backup directory is named with the date/time stamp when the backup started. If no files have been modified, or if the archive version number has been set to zero, a backup copy of the file will not be created.

The files in the backup directory have the same directory structure as the files in the latest directory.

The following is an example of the backup directory structure:

```
/BackupData/host1.company.com/040105171500/applicationdata/  
  systemstatedata  
  \drives\e\dir1\file1.txt
```

In the above example, during the backup from host 1, the data had been modified and the file e:\dir1\file1.txt was also modified. Both files were copied from the host 1 latest directory and placed in the host 1 backup directory 040105171500. The host 1 latest directory now contains the latest file version from host 1.

The data in the backup directories contains older copies of the files from the latest directory PRIOR to the time backup was run. This is important to consider when restoring files from a particular backup date.

When restoring the files, the first general rule is “get the file latest files,” which would be restored from the latest directory. This would restore the most current file version from the target host storage location for the specified host.

The second general rule when restoring files is “get the version prior to N backup,” where N is when a backup occurred. In this case, the user wants to restore the version of the file from the target storage for a specified host prior to a certain backup date.

Once the number of backup copies to maintain has been reached, the next backup performed will cause the oldest backup directory to be deleted.

Application Backup Details

This section describes the specifics of the Microsoft SystemState application backup that can be configured as part of the Remote Backup job.

User Permissions Required

The source user specified for the backup job must be added to the “Backup Operators Group” on the LOCAL machine, per source host running the SystemState backup.

SystemState Backup

The SystemState backup exports a file to disk on the source host specified with the relevant system information, such as Windows registry, boot files, SYSVOL directory files, etc. The changes are transferred to the target host storage, and then the file on the source host is deleted.

Remote Restore Data Management Tool

NOTE: This tool is not available for Snap EDR Express.

The Remote Restore tool transfers backup data from a central storage location, on a per host basis, to the remote hosts from which the data was originally retrieved. The backup data in the central storage location must have been previously backed up with the Snap EDR Remote Backup solution. The following are the features of the Restore solution:

- Restore backup data from the central location to one or many hosts
- Restore backup data from the central location to heterogeneous host platforms
- Restore from any one of the backup copies stored in the central location
- Restore entire backup or specified data sets
- Restore backup data from a particular host to one or many other hosts
- Preserve file ownership permissions
- Centralized control for viewing job status, controlling job execution, managing hosts involved in backup
- Manage the amount of bandwidth the backup job will utilize
- Network fault tolerance during backups
- Restore Microsoft System State backup data
- Full flexibility to allow for customization of solution to meet customers needs

For information about Remote Restore recommended configurations and best practices, see [Best Practices](#).

Create a Remote Restore Job

To schedule a Remote Restore job:

1. From the Management Console, click **Snap Solutions > Remote Restore**.
2. Click **Add a Job**.

If this is the first time you are creating a Remote Restore job, when you click **Remote Restore**, the scheduling screen appears. You do not have to click **Add a Job**.

3. Enter information into the various fields. The following table describes the fields:

NOTE: You cannot use the '|' (pipe) character in a Windows directory path or errors will result.

Field	Description
Job Name	A unique name for this job run.
Backup Repository	
Repository Host	The backup repository host for a Restore job will be the same Agent specified as the Target Agent during the Remote Backup Job. This is the Agent where the backup data resides. The hosts that appear are only ones to which the user has access.
Repository Base Directory	The root directory where the backup data is stored is the location of the base directory in the central storage location where the backup data resides for all hosts. This is the same directory as the target directory specified for a Remote Backup job. This field must be specified. The directory specified may be in the following format: Guardian OS, Unix/Linux root drives, e.g., /remotebackups Example entry: /remotebackups The Remote Backup job would have specified "/remotebackups" as the target directory.
Host to Restore	
Restore Hosts	The host whose data you want to restore. This host must have been involved in a Remote Backup job as a source host.
Restore From Version	
Restore from Backup Version	Specifies where to look for the data to restore. For each host specified in a Backup job, there is "the latest" directory with the most recent backup data. There may also be one or many backup copies with older delta versions of the backup data stored as well. Choose from Latest , Specified Date , or Last Backup Before Specified Date .
Specified Date	The date of the backup version from which you want to restore. Use this in conjunction with the "Specified Date" or "Last Backup Before Specified Date" options selected in the Restore From Backup Version field.
General Restore Options	

Field	Description
Restore Type	The restore type to perform is described as follows: <ul style="list-style-type: none"> • Full - Restore all the data from the specified version of backup. • Specified Directories and File Data Only - Restore only the data in the "Directories to Restore" field. • Specified Application Data Only - Do not restore any user files, just the backup of the exported application data.
Restore To Host	The host to which you want to restore the backup. If you choose "Alternate Restore Host" you must also specify the host in the "Alternate Restore Host" field (see below).
Alternate Restore Host	You can redirect a restore to an existing, configured Agent. Choose a host from the drop down list.
Directory and File Data Restore Options	
Directories to Restore	By default all directories are restored. This field allows you to specify which directories you want to include in the restore. Separate multiple entries with a comma. Note that this field does not support pattern matching.
Include File Names/Types	By default, all file names and types backed up are restored. This field allows you to specify file names or types to include in the data being restored. A comma must separate multiple data masks.
Exclude File Names/Types	By default all file names and types backed up are restored. This field allows you to specify file names or types to exclude from the data being restored. A comma must separate multiple data masks.

Field	Description
Exclude Sub Directories	<p>Subdirectories may be excluded from the data being backed up by specifying them in this field. Multiple entries must be separated with a comma. When the job runs, all directories that match those specified in the exclude subdirectories field will be excluded. Note that normal behavior is to exclude subdirectories that match regardless of where they appear in the directory path. Using the anchoring expression ("@") changes this behavior to anchor the exclude directory path.</p> <p>For example, if a user specifies a source directory of C:\data\docs, and an exclude temp directory of temp, any sub-directories called temp will be excluded, even those nested within another subdirectory. (For example the subdirectory C:\data\docs\publish\release\temp would be excluded as well.)Users who want to exclude a directory only at a certain level can use the @ symbol to anchor the exclude directory path at the starting source directory level. For example, specifying @temp in the above example would mean that the C:\data\docs\temp directory would be excluded, but the C:\data\docs\publish\release\temp directory would be included.</p> <p>If the user instead wanted to make sure just the C:\data\docs\publish\release\temp directory was excluded, but wanted to have a source directory path of c:\data\docs, the user would need to type C:\data\docs in the source directory field and @publish\release\temp in the exclude field.</p> <p>Special characters allow users to make use of pattern matching on the directory path. You must escape special characters with a backslash to be matched literally. Characters include the following:</p> <ul style="list-style-type: none"> * (matches zero or more characters) ? (matches any single character) [...] (matches any one of the enclosed characters - for example, [ch] would match the characters "c" or "h" <p>A pair of digit, lowercase or uppercase characters separated by a hyphen '-' denotes a range set and matches any character sorted in the range. If the first character following the '[' is '^' or '!', then any character not enclosed is matched.</p> <p>Use commas to specify multiple distinct patterns as an alternative to using multiple option specifications. Note that these options do not enable/disable directory traversal.</p>
Restore to Directory	<p>The directory to which to restore the data. If you choose Alternate Base Directory you must specify the directory name in the Alternate Base Directory field. If the specified directory does not exist, it will be created.</p>
Alternate Base Directory	<p>The alternate restore directory field specifies the base directory to which to restore the data. If the directory does not exist it will be created. If not specified, the data will be restored to its original location. For example, the data to be restored is as follows:</p> <p>c:\userdata\file1.doc</p> <p>If the root directory on the target is specified as c:\restoredata, the data will be restored back to the host in the directory c:\restoredata\c\userdata\file1.doc. If the root directory on the target host is not specified, the file would be restored to its original location, c:\userdata\file1.doc.</p> <p>The directory specified should be a fully qualified path name. The directory specified may be in the following format:</p> <p>Windows root drives, e.g., C:\ApplicationData UNIX/Linux root drives, e.g., /home Mac root drives, e.g., /Users</p>
Application Data Restore Options	

Field	Description
Application Data To Restore	By default, the Restore job restores specified user files. Backup copies may be stored, on a per host basis, for exported application data during a Backup job. To restore a version of the backed up application data, select which application data to restore. For the Restore job to complete successfully, the specified application data must exist in the backup version specified.
Application Data Restore Directory	If you choose application data to restore, you must specify a directory in which to restore the application data on the target host. If the directory does not exist, it will be created. The directory specified must have enough space to hold the exported application data.
File Transfer Options	
Use Compression	<p>If set to Yes, the source agent compresses each file before sending it. The files will be uncompressed automatically on the target agent(s). The degree of compression depends on the type of data you are transferring. The following are typical rates of compression for different types of data:</p> <ul style="list-style-type: none"> • Plain text: 70-95% • TIFF images: 20-40% • Binary files: 0-5% <p>The default value is No. Choose this if you are running over a low speed WAN link (e.g., less than 3 Mbps). This option is NOT recommended for LAN or high speed links.</p>

Field	Description
File Ownership Transfer	<p>Specifies whether files maintain the source user ID and group ownership after they are transferred, and what method is used. Choose from the following options to preserve ownership:</p> <p>GuardianOS - When transferring between two GuardianOS machines, permissions are extracted and set using the GuardianOS routines. When transferring between Windows, Linux, Mac, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).</p> <p>Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its permissions from a folder on the source is transferred to a folder on the target that has different permissions. However, the source folder is not transferred. In this case, the file's inherited permissions on the target will come from the target folder into which it is transferred and not from the source folder.</p> <p>NOTE: Windows attributes (such as Read Only) are not preserved in transfers between Windows and GuardianOS machines. Access Control Lists (ACLs) can be preserved, but attributes cannot.</p> <p>GuardianOS mode is the default mode of transfer.</p> <p>Windows - Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.</p> <p>Unix - Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX/Linux/ Mac hosts. The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.</p> <p>Off - Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.</p>
Job Options	

Field	Description
Bandwidth Throttle	<p>Specifies the rate at which the source will make the data available to the network, regardless of how fast the machine is sending. Bandwidth Throttle is the rate at which the source reads the data from disk and puts it in the send queue.</p> <p>Click the Calculate Bandwidth icon to specify the bandwidth as a percentage of a specified network value, and click OK. The calculated value appears in the Bandwidth Throttle field. Note that bandwidth limiting is done on each stream connection, so a value specified here is passed to the controlling agent for each template executed, and divided among the streams as follows:</p> <ol style="list-style-type: none"> 1 It is divided by the number of streams the agent will use for each remote agent (typically four). 2 It is further divided by the potential number of concurrent remote agents. This will be the lesser of the maximum allowed number of concurrent agents, and the number of remote agents specified in the template. <p>Note that bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.</p>
Encryption Level	<p>Allows users to specify the encryption level from the following values: High, Medium, Low, or No Encryption - signed. The default value is High. Note that mutual authentication is always used regardless of the encryption level specified.</p> <p>The No encryption - signed option transfers unencrypted (plain text) data, but includes the SSL protocol's message digest calculation and signing to ensure data stream integrity.</p> <p>NOTE: Encryption incompatibilities exist with transfers that involve version 5.2.2 build 1194 and higher (including version 7.2) and version 5.2.2 build 1177 and earlier agents.</p> <p>With Snap EDR version 5.2.2 build 1194 and higher, Snap EDR was enhanced to support AES-256 and AES-128 bit encryption. We recommend upgrading your Management Console and all of your agents. If all of the Snap EDR nodes are upgraded, the system will automatically use the "high" (AES-256) encryption when the replication jobs have been set to "high" (the default).</p> <p>If you choose not to upgrade all of the Snap EDR nodes, you must change all of the Snap EDR jobs to reflect "no" encryption; otherwise, the jobs will fail since the earlier versions of Snap EDR agents do not support the improved encryption.</p> <p>If running Snap EDR version 5.2.2 build 1194 or later, these encryption settings map to these encryption levels:</p> <ul style="list-style-type: none"> • High - AES-256 • Medium - AES-128 • Low - RC4-40
Job Log Detail Level	<p>The type of logging information for this job. Choose from Error, Warn (warning), Info (Information) or Debug. Debug provides the greatest level of detail while Error provides the least.</p>
Scheduling Parameters	
Job Start Date/ Time	The date and time at which you want the job to run.
Frequency	
None	Use none when you have a job that needs to be scheduled, but is run only at irregular, user-defined times.

Field	Description
Hourly, Daily, Weekly, Monthly, Yearly	Run the job once every selected interval.
Monthend	Run the job once every last day of the month.
Once	Run the job only once at start time.
Time Zone	<p>Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time.</p> <p>Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere.</p> <p>Note that the time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.</p>
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	Text that appears in the subject field of the report e-mail message. Note To receive email notifications in Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification.
SNMP Trap Notification	
Send Traps Once	Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither. Note To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.

4. Click **Add Job**.

The new job appears in the list.

How Data to Restore is Resolved

This section describes how the Restore solution is designed to restore data from a central location back to the remote hosts from which it was backed up.

Users are recommended to read the section [File Backup Details](#) before reading this section if the data layout for Remote Backups is not familiar.

Source Host Selection

The Restore solution is designed to restore to specified hosts data that was previously backed up to a central storage location using the Remote Backup solution. During the scheduling of a Restore job, the source host selected will be the target host from a previously run Remote Backup job. This is the host with access to the central storage location containing all of the backup data.

Target Host Selection

The target hosts selected during the scheduling of a Restore job will have been one or more of the source hosts from a previously run Remote Backup job. During the Restore job execution, the specified data to restore will be searched for on a per host basis (as specified in the target hosts for the job), in the corresponding directory based on the host name. This is the default behavior. The data to restore can be redirected to an alternate host (e.g., not the original source of the data), a topic that is covered later in the document.

When a target host is specified and a corresponding directory cannot be found in the backup data set (based on host name), an error is flagged and the restore will not continue for the host in question.

Restore Versions

As described in the section [File Backup Details](#), for each host that was part of a Remote Backup, data is stored in a separate directory. For each backup done where files were modified from the latest version, a backup directory is created and a copy of the files from the latest directory is moved there before the latest version is updated from the source host.

When restoring data, the latest directory (specifying latest in versions of data to restore from) will restore the most recent version of the files from the central storage location to the specified hosts.

Specifying a backup version (1 ... N) from which to restore data will get the version of the file as it existed on the source host prior to when that backup occurred.

When specifying a version to restore for multiple hosts, keep in mind that the backup versions among hosts may not have the same date. If no changes are detected in the latest data set during a backup, a backup directory with the corresponding date for that host will NOT be created. In this case, multiple hosts in the same Remote Backup job may have a different number of backup versions.

If a specified backup version does not exist for a target host, the Restore job will not continue for the host in question and an error will occur. For example, a Restore job is scheduled for two hosts, host 1 and host 2. Host 1 has three backup versions and host 2 has one backup version. Specifying a restore from version 2 will succeed for host 1 and fail for host 2.

Specified vs. Full Restore

There are two types of user data restores that may be configured: specified and full.

- A full restore restores all user data files from the specified version. From latest, all the files will be restored. From a particular backup version, all the files in that backup version will be restored (this is a subset of latest, since only the files that have been modified since the last backup are moved into the backup directories).

- A specified restore works the same way as a full restore, except only the specified directory and subdirectories are restored. The specified directory is searched for in the backup version to restore from. If the directory is located in the specified version, the restore will continue. If the specified directory is not located in the target host's backup data, the restore will fail for the host in question and an error will occur.

Restore Data Location on Target Host

When restoring data to a target host, there are two options. The first option is to restore the data to a specified location. In this case, the end user can then determine where to put the files after the fact. The other option is to have the Restore job put the files back to their original location. In this case, the files will be restored back to the original location from which they were backed up, overwriting any existing files.

Application Data Restore

Application data exported and backed up as part of the Remote Backup job may be restored to a target host as well. The version to restore from has the same behavior as the user data files version to restore from. If the application data specified could not be found for the version requested, the Restore job would fail for the specified target host.

A directory to store the backup application data must be specified for the target host. This directory must be large enough to maintain the exported application data. Once the application data is restored to the target host, the user may import the data back into the required application.

File Transfer Options

The files are transferred between the source host and target hosts using the following file transfer options:

1. Files are always transferred from the source host to the target hosts. Any existing files on the target host are overwritten.
2. File ownership preservation is set to "Windows" by default. Selecting one of the file ownership options when scheduling is recommended:
 - **GuardianOS**—When transferring between two GuardianOS machines, Permissions are extracted and set using the GuardianOS routines. When transferring between Windows, UNIX, Linux, or GuardianOS, user name matching is always done for the owner and group regardless of the source and target OS type, and for Permissions if the source and target are both Windows. If no match is found, a warning is generated. Permissions without matches are dropped. Owners and groups without matches are replaced with the transfer user (root, system or UID 1).

Inherited Permissions are not explicitly copied between Windows machines. For example, a file that inherits its Permissions from a folder on the source is transferred to a folder on the target that has different Permissions. However, the source folder is not transferred. In this case, the file's inherited Permissions on the target will come from the target folder into which it is transferred and not from the source folder. Note that Windows attributes (such as read only, and so on) are not preserved in transfers between Windows and GoS machines. Access Control Lists can be preserved, but attributes cannot.

GuardianOS mode is the default mode of transfer.

- **Windows**—Preserves Windows SIDs in homogenous Windows environments. A security identifier (SID) is a unique value of variable length that is used to identify a security principal (e.g., user or security group) in Windows 2000. Well-known SIDs are a group of SIDs that identify generic users or generic groups - these do not change from system to system. In this mode, the security stream, and all other alternate data streams, are copied; inherited Permissions are explicitly copied. Use when transferring data between Windows hosts only.
 - **Unix**—Tries to match the usernames/groupnames found on the source and target systems. Use when transferring between Windows and UNIX, Windows and Mac, and Unix and Mac hosts. The file ownership transfer capability might not work on some UNIX systems such as Solaris that use Orange Book security standards.
 - **Off**—Do not preserve ownership. Files written to the target are owned by the user ID specified in the target User ID field of this job template. The file will be owned by root on UNIX, NT Authority/System on Windows and admin on GuardianOS.
3. Files are compressed by default when being restored back to the target machine. You can turn off this option when scheduling a job.

Once you create a job, there are a number of tasks you may want to perform on the job, including viewing the job logs or editing and copying the job. For information on these topics, see [Manage Jobs](#).

Once you have created a scheduled job, there are many ways you can manage it, including viewing statistics, forcing the job to run, editing the job's schedule, and so on. Two screens allow users to manage jobs:

- Job List Summary Screen
- Job List Detail Screen

Manage Jobs Using the Job List Summary Screen

The scheduled job list summary screen is accessed by selecting the type of job you want to view (e.g., Aggregate, Distribute, etc.). The table displays summary information about the jobs in the selected Snap solution, and allows users to do the following:

- Update multiple jobs
- Add a job
- Edit, delete, run or cancel a job
- Set the bandwidth limit for a running job

In addition, the summary screen displays information about the job including the name, state, percent complete and status message with one of the following terms:

- **Invalid** - Indicates some type of syntactic/semantic error was found during the evaluation of the job description. An Invalid job cannot run, either automatically or manually.
- **Running** - The job is currently running.
- **Suspended** - Indicates the job's schedule has been suspended (i.e., the job will not run at its next scheduled time unless a user clicks Resume).
- **Starting** - The job is beginning.
- **Error** - Job has run and finished with an error.
- **Completed** - Job has run and finished with an exit code of zero (with no errors).
- **Next Run Scheduled** - The job is scheduled for its next run.
- **Scheduled** - Job has been scheduled and has not run yet.
- **Idle** - Indicates a job that does not have a future automated run time scheduled.
- **Pending** - A multiple push or pull transfer has been configured and there are more agents to push to/pull from than currently can be handled, so those that are in the queue waiting to be serviced will have a status of pending. You will also see a pending record at the startup of a transfer before any data starts to get transferred, but it will appear only briefly. You will see pending records only when you look at the job details for a transfer.

NOTE: If you want to display only jobs of a certain state, choose the state by which you want to filter the list from the Job State drop-down (for example, "Error"). Only jobs that are in that state appear in the list (or a message indicating that no jobs in that state were found). Choose "Remove Filter" to display the full list of jobs.

You can also filter on Job Name by choosing the Job Name drop-down and typing text in the field (note that this field is case-sensitive).

Manage Jobs Using the Job List Detail Screen

The scheduled job list detail screen is accessed by clicking on a scheduled job in the Job List Summary table. It displays detailed information about the selected job, and provides a menu bar of tasks, as well as three distinct areas with information about the job: status, runs and detail. The scheduled job view detail screen allows users to perform the following tasks:

- Refresh the screen
- Edit the job's schedule (by clicking Edit)
- Delete the selected job (by clicking Delete)
- Copy the job (by clicking Copy)
- Force the job to run (by clicking Run Now)
- Suspend a scheduled job (by clicking Suspend)
- View the status of the job and any error messages that may have been generated
- View a graph of the bandwidth used during the course of the run
- View details of the job run (such as duration and the number of files transferred) in the Runs area
- View the statistics for the job (by expanding the Detail area)
- Download the job's log files (by clicking the download icon on the Logs tab in the Detail area)

Check a Job's Running Status

While a job is running, you can check statistics such as what percentage of the job is complete, or the number of unrecoverable errors in the Status area of the Job Views Detail screen. The screen also displays a graph that indicates the amount of bandwidth the job used, at certain intervals.

When the job is not running, a legend that explains the colors on the graph appears in this area.

View All Runs for a Job

The Runs area of the Job List Detail screen displays a table of all the runs of the job, and includes the exit code (a numerical value) indicating how the job completed. A zero indicates that the job ran successfully. Any other number indicates an error. You can click the number to view a list of the exit codes, their meaning and possible resolutions for the problem. For more information, see [Exit Codes](#).

Clicking the **Show empty runs** checkbox displays job run records only for jobs that did not transfer any files.

View Job Statistics

Snap EDR provides extensive job run statistics, including bandwidth throttle, transfer start and end times, the number of directories deleted, and so on.

The Detail area of the Job Views screen displays the statistics as follows:

- categorized into separate tabs (overview, transport, data volume, and so on)
- statistics file (for viewing or downloading)
- interval statistics (for viewing or downloading)
- log file (for viewing or downloading)

View Categorized Statistics

Snap EDR provides job run statistics categorized into different tabs that relate to the following kinds of job run information:

- overview
- settings (transport)
- data volume
- file volume
- performance

In addition to this categorized job run statistical information, Snap EDR provides comprehensive statistics, interval statistics and log files for users to view or download.

View/Download Statistics Files

Snap EDR provides two detailed statistics files for users to view or download: statistics and interval statistics.

- The statistics file provides information including files transferred, files skipped, average transfer rate, RSYNC bytes skipped, and so on - the cumulative statistics for a job run.
- The interval statistics file provides information including the network bytes transferred, the bandwidth throttle, bandwidth ceiling and bandwidth floor, for a job during the default, or user-specified intervals (5 seconds, 10 seconds, and so on).

In the Detail area of the Scheduled Job List Detail screen, click the Logs tab, and click the magnifying glass to view a table of cumulative statistics for the job, or click on the floppy disk icon to download the statistics file to your computer. Click the same icons beside **Interval Statistics** to view statistics taken at user-specified time intervals during the job run.

View/Download a Job Log File

Displays the log information for this file. Whenever you run a scheduled job, the Manager generates a log file. You can either view the log file on screen, or download it to your local computer. Downloading a job log file allows you to save the file to your local machine where you can load the file into different programs (for searching, printing, and so on).

In the Detail area of the Scheduled Job List Detail screen, click the Logs tab, and click the magnifying glass beside **Job Log** to view the job's log file, or click on the floppy disk icon to download the log file to your computer.

Filter Log Information

Log files display a large amount of information. Users may choose to set parameters to display only certain types of information in the log file. Note that filtering the log file affects only the log file displayed in the UI. It does not affect the log file you download. The downloaded log file contains all of the log information regardless of the filtering options users select.

The Log screen displays detailed log information for the selected job run. In addition, the top of the screen provides an area in which users can define filters to change what log information appears on the screen.

The following table describes the fields on which users can filter.

Field	Description
Agent Prefix	The initial letters of the hostname
Category Prefix	The initial letters of the category name
Job Template Name Prefix	The initial letters of the job template name
Message ID	A number associated with a text description in the message catalog. A user could filter on all of the messages with a specific ID number. Users can also exclude a specific message ID by typing the following: !message <ID number> (for example, !30051). None of the messages matching that number will appear in the log file.
Minimum Severity	Filters on the priority level of the messages, showing only messages that have the selected priority and lower. Choose from "Error", "Warn" (warning), "Info" (information) or "Debug" (debugging message). For example, if one chooses Debug (the highest level) one will see all of the messages. If one chooses Warn, one will see errors and warning messages.
Date/Timestamp	When checked, displays a column of dates and times. When unchecked, the dates and times do not appear.
Severity	When checked, displays a column of severity levels (Error, Warning, Information, Debug). When unchecked, the severity levels do not appear.
Category	When checked, displays a column of the type of message generated (such as Application, Network, Configuration, and so on). When unchecked, the categories do not appear.
Message	The text of the log message
Message ID	When checked, displays a column of message identification numbers. Clicking these numbers links to additional help on the message. When unchecked, the message IDs do not appear.
Message Source	When checked, displays a column that identifies the component that generated the message (such as transport manager, command control agent, target execution command, and so on). When unchecked, the message sources do not appear.
Agent	When checked, displays a column of hostnames. When unchecked, the hostnames do not appear.
Job Template	When checked, displays a column of job template names. When unchecked, the job template names do not appear.

Statistics

The Statistics screen displays the following information for each template in the job and each Agent pair between which the template runs, calculated as the cumulative statistics for the job run.

Statistic	Description
Job ID	The number Snap EDR uses to identify this job run.
Template Name	The name of the job template with which the job is associated. Click on the name to open the associated job template.
Source Agent	The name of the host that initiates the transfer.
Target Agent	The name of the host to which the data is being transferred.
Status	Indicates the current state of a transfer (PENDING, RUNNING, COMPLETE, PAUSED).
Template Type	Indicates the kind of template with which the job is associated (FILE_TRF, PROCESS_TRF, REMOTE_CMD).
Transport Type	Indicates the kind of transfer (UDP, TCP).
Effective Data Transferred	Specifies the amount of data transferred, plus what was skipped because it already existed on the target.
Files to Transfer	Specifies the known number of files to transfer.
Directories to Transfer	Specifies the known number of directories (as files) to transfer.
Data to Transfer	Specifies the total number of data bytes in the known files to transfer.
Files Transferred	Specifies the number of files successfully transferred.
Files Skipped	Specifies the number of files skipped.
Files Deleted	Specifies the number of files explicitly deleted by the target agent.
Failed Files	Specifies the number of files that did not transfer.
Controlling Agent Start Time	Specifies the controlling agent process start-up time (system epoch time in microseconds) saved in the database as a timestamp.
Remote Agent Start Time	Specifies the remote agent start-up time (system epoch time in microseconds) saved in the database as a timestamp.
Names Command End Time	Specifies the time the names command has completed for a file transfer.
Transfer Start Time	Specifies the start time for transfers with the remote agent (system epoch time in microseconds) saved in the database as a timestamp.
Transfer End Time	Specifies the end time for transfers with the remote agent (system epoch time in microseconds) saved in the database as a timestamp.
Remote Agent End Time	Specifies the remote agent process end time (system epoch time in microseconds) saved in the database as a timestamp.
Controlling Agent End Time	Specifies the controlling agent process end time (system epoch time in microseconds) saved in the database as a timestamp.
Bandwidth Throttle	Specifies the bandwidth throttle being applied at the end of the reporting interval (in bytes/seconds). Note that a value of 0 indicates that no bandwidth throttling was being applied at the end of the reporting interval.
Average Transfer Rate	Calculated transfer rate. A push equals "ntwk_bytes_sent/ (transfer_end_time - transfer_start_time)/ 1000000". For pulls, "ntwk_bytes_recvd" is used.

Statistic	Description
Maximum Transfer Rate	Highest transfer rate calculated across all reported statistics records for this transfer. By default, statistics are reported every 60 seconds. Transfer rate for a push equals "ntwk_bytes_sent/(transfer_end_time - transfer_start_time)/1000000". For pulls, "ntwk_bytes_recd" is used.
Minimum Transfer Rate	Lowest transfer rate calculated across all reported statistics records for this transfer. By default, statistics are reported every 60 seconds. Transfer rate for a push equals "ntwk_bytes_sent/(transfer_end_time - transfer_start_time)/1000000". For pulls, "ntwk_bytes_recd" is used.
Directories Transferred	Specifies the number of directories (as files) successfully transferred.
Directories Deleted	Specifies the number of directories explicitly deleted by the target agent. Note that it does not include files/directories implicitly deleted during (overlay) transfers.
Directories Skipped	Specifies the number of directories (as files) skipped.
Failed Directories	Specifies the number of directories (as files) that failed to transfer.
Data Transferred	Specifies the actual amount of data transferred, calculated as an aggregate value equal to "file_data_bytes" + "file_bytes_skipped" + "the sum of the data counts of all currently-active streams".
File Bytes Written	Specifies the number of uncompressed file data bytes transferred.
File Attributes	Specifies the number of uncompressed file attribute bytes transferred.
Compressed File Data	Specifies the number of compressed file data bytes transferred.
Compressed File Attributes	Specifies the number of compressed file attribute bytes transferred.
Data Skipped	Specifies the number of file data bytes not transferred because the files were skipped.
Data Deleted	Specifies the number of file data bytes deleted.
RSYNC Bytes Skipped	Specifies the number of file data bytes not transferred because the RSYNC algorithm deemed them unchanged.
Source Non-RSYNC Bytes	Specifies the number of non-data RSYNC overhead bytes flowing from source agent to target agent.
Target Non-RSYNC Bytes	Specifies the number of non-data RSYNC overhead bytes flowing from target agent to source agent.
Source Compressed Bytes	Specifies the number of compressed source-to-target file manifest overhead bytes.
Target Compressed Bytes	Specifies the number of compressed target-to-source file manifest overhead bytes.
Source Protocol Bytes	Specifies the number of source-to-target transfer protocol overhead bytes.
Target Protocol Bytes	Specifies the number of target-to-source transfer protocol overhead bytes.
Source Control Channel Bytes	Control channel overhead source to target.
Target Control Channel Bytes	Control channel overhead target to source.
Security Framework Bytes Received	Specifies the number of bytes received by the controlling agent from the Security Framework (SF) layer.
Security Framework Bytes Sent	Specifies the number of bytes sent by the controlling agent to the Security Framework (SF) layer.
Network Bytes Sent	Specifies the number of bytes sent by the controlling agent on the network.

Statistic	Description
Network Bytes Acknowledged	Specifies the number of bytes acknowledged by the controlling agent from the network.
Network Bytes Received	Specifies the number of bytes received by the controlling agent from the network.
Remote Command Total Units	The total number of units the remote command has processed. The "units" are whatever is defined in the Remote Command template.
Remote Command Units Complete	The number of units the remote command has processed. The "units" are whatever is defined in the Remote Command template.
Unrecovered Errors	Error counter for all unrecoverable exceptions declared by the controlling agent with respect to the remote agent.
Total Errors	General error counter for all exceptions detected by or reported to the controlling agent with respect to the remote agent.
Process/ Streaming Bytes Sent	Specifies either the number of bytes (a) sent by the controlling (source) agent of a process push transfer or (b) received by the controlling (target) agent of a process pull transfer; or the number of bytes sent by the controlling (source) agent across all channels of a streaming transfer.
Process Bytes Unconsumed	Specifies the number of bytes received by the target agent of a process transfer that were not delivered to the target data sink command.
Streaming Transfer Bytes Sent	Specifies the number of bytes the controlling (source) agent sent across all channels of a streaming transfer.
Streaming Transfer Bytes Received	Specifies the number of bytes received by the controlling (source) agent across all channels of a streaming transfer.
Agent Stat Time	Specifies the last time a statistic was reported

Interval Statistics

The Interval Statistics screen displays the following information for the selected job run, calculated at a user-specified interval during the job run (5 seconds 10 seconds, and so on).

Statistic	Description
Job ID	The number Snap EDR uses to identify this job run.
Template Name	The name Snap EDR uses for the template for which the statistics are being displayed (for example, FT_RC.SingleFT.FT).
Source Agent	The name of the host that initiates the transfer.
Target Agent	The name of the host to which the files are being transferred.
Interval Start Time	Specifies the start time of the reporting interval (system epoch time in microseconds).
Interval End Time	Specifies the end time of the reporting interval (system epoch time in microseconds).
Event	Displays event types generated from the agent. The following list indicates what each number represents. 0 - Periodic report 1 - Agent starts 2 - Transfer starts 3 - Transfer paused 4 - Transfer resumed 5 - Transfer complete 6 - Agent connection lost 7 - Agent restarts
Network Bytes Sent	Specifies the number of bytes sent from the source to the target on the network for the interval. Note that for File Transfer or Process Transfer Push: ntwk_bytes_sent, for File Transfer or Process Transfer Pull: ntwk_bytes_rcvd, for Streaming Transfer: (ntwk_bytes_sent + ntwk_bytes_rcvd).

Statistic	Description
Network Bytes Acknowledged	Specifies the number of bytes acknowledged by the controlling agent from the network.
Bandwidth Throttle	Specifies the bandwidth throttle being applied at the end of the reporting interval (in bytes/seconds). NOTE: A value of 0 indicates that no bandwidth throttling was being applied at the end of the reporting interval.
Remote Command Units Complete	Specifies remote command statistics.

Perform a Task on More than One Job

Users can run any of the following tasks against all the jobs associated with a solution:

- Delete
- Cancel running jobs
- Run idle jobs
- Suspend future runs
- Resume suspended jobs
- Set bandwidth limit
- Set bandwidth floor
- Set bandwidth ceiling
- Update variables

NOTE: With the update variables option, all of the selected jobs must be based on the same job template to change the variables.

To update several jobs at once:

1. Navigate to the type of job you want to update (e.g., Aggregate) and click the **Update Multiple Jobs** icon.
2. On the Multi-Action Summary screen, place checks by the jobs you want to update, then click the appropriate action icon.

Update Bandwidth Across Multiple Jobs

The Multi-Action Bandwidth screen allows you to specify a bandwidth limit for more than one job. This is especially useful if you have a certain amount of bandwidth to which you want to limit these jobs. Setting the bandwidth limit on multiple jobs will temporarily override any current bandwidth limit settings specified in the job itself. Once the job has completed, the multi-action bandwidth limit setting will be discarded. If the job is run at a later date and time, any bandwidth setting specified in that job will take effect.

Place a check in the **Spread Bandwidth Evenly Across All <number> Jobs** box, to distribute the specified bandwidth limit equally among the selected running jobs. For example, if you specify 100 mbits and four jobs are running, they would each use 25 mbits. With five jobs, it would be 20 mbits.

Note that bandwidth throttles may also be employed by other network devices and policies (e.g., QoS), therefore, a bandwidth throttle (or target maximum) defined here may not be achievable. If you are having difficulty achieving a particular bandwidth target ensure that other policies are not impacting your ability to reach the desired throughput.

Updating Variables Across Multiple Jobs

The Multi-Action Variables screen allows you to specify changes for multiple variables to more than one job. This is especially useful if you want to make specific changes to a set number of jobs without having to change each job independently. The Multi-Action Variables screen displays the variables associated with the selected job template. Click in the checkbox beside the variable(s) you wish to change and specify the value in the appropriate field. To select all of the variables, click to put a check in the Check All box in the top left of the screen.

The Management Console creates at least one statistics record for each completed job run and stores it in the Rules Database.

-

Report Types

Using the Management Console, you can generate the following types of reports:

Report Type	Description
Stats Summary Report	Provides a summary and totals of the data from all jobs in the selected job group, based on how the user chooses to group the data (by month, day etc.).
Detail Report	Provides a detailed record of the data from the job templates or a specific job template in the selected job group. Note that the number of records the report displays depends on the number of Agents on which the job template runs. The report displays a record for every Agent on which the job template runs.
Custom Query Report	Provides a customized query of detail records from a combination of job name and job templates within a selected job group. Note that the number of records the report displays depends on the number of Agents on which the job runs. The report displays a record for every Agent on which the job runs.

When you create a report, you can assign a template name to it. The template appears in the Report view. Report templates allow you to generate reports without having to change the parameters you initially set up. Templates are especially useful if you plan to generate the same type of report on a regular basis. You can Run, Edit, Remove, or Schedule a report by clicking on the appropriate word beside the report template. Click Schedule to run reports daily, weekly or monthly at specific times.

You can also edit templates if you decide you want to change some of the parameters. For example, if you wish to rename an existing template, select Edit beside the template whose name you want to change, change the name and ensure that a check appears in the overwrite template check box. Click Generate Report. The new name will appear in the template listings.

Stats Summary Report

A Stats Summary Report provides a summary and totals of the data from all jobs in the selected job group, based on how the user chooses to group the data (by month, day and so on).

To create a stats summary report:

1. From the Management Console, click **Snap Solutions > Reports**.
2. Click **New Stats Summary Report**.

The Stats Summary Report screen appears.

3. In the **Select a Job Group** list, click the job group for which you want to generate the report.
4. Choose the roll up value from the drop-down list, the value by which the report will break down the data. Choose from none, day, month or year.
5. In the **Report In** drop-down list, choose the unit in which you want the statistics to appear. Choose from bytes, kilobytes, megabytes or gigabytes.
6. Choose the output format by selecting either HTML or CSV.
CSV stands for Comma Separated Variable. You can load and save this type of file in a spreadsheet.
7. In the **Select Columns to Group Data By** area, click the appropriate check boxes to specify the column types in which you want to group the data.
8. In the **Select Columns to Include** area, click in the check boxes beside the type of information you want to include in the report.
9. In the **Select a Date Range** area, type the start and end date of the time period for which you want to generate the report.
10. In the **Generate Report** field, type a name for the template of this report.
This template will appear in the Reports view, so that you can simply click on it to run a new report.
11. Click **Generate Report**.
The Stats Summary Report appears.

12. Close the report window to return to the report setup window, or use your browser's print command to print the report.

For more information on the fields in the generated report, see [Generated Report Categories](#).

Report Fields

The following table describes the sections in the various report screens (note that not all reports have the same fields in each section):

Section	Description
Select a Job Group/Enter Search Criteria	<p>May include the following fields:</p> <ul style="list-style-type: none"> • a list of job groups • "roll up by" drop-down list (The value by which the report will break down the data. Choose from none, day, month or year.) • "report in" drop-down (The unit in which the statistics will appear. Choose from bytes, kilobytes (KB), megabytes (MB) or gigabytes (GB).) • output format (Choose from HTML or CSV - CSV allows you to load the file in a spreadsheet.) • job name (Associated with custom query reports. Allows users to choose the name of a job.) • job template (Allows users to generate a report from a specific job template.) • search with all/any criteria (Allows users to generate a report that includes all or only part of the criteria specified.)
Select Columns to Group Data By	Allows users to specify the column types in which to group the data.
Select Columns to Include	Allows users to specify the type of information to include in the report.
Select a Date Range	Allows users to specify the start and end date of the time period for which users want to generate the report. Click the calendar icon to display a pop-up calendar from which you can select the date.
Generate Report	Allows users to specify a template name, and create a template of the report. Users can schedule the template to run at regular intervals. Only the user who creates the template has access to it. If this is the second time you have created a report, an overwrite template checkbox appears below the template name. Uncheck the box if you do not want to overwrite the previous template. If you are editing the existing template and wish to overwrite it, make sure a check appears in the overwrite template box.

Detail Summary Report

The Detail Summary Report provides a detailed record of the data from the selected job group. Note that the number of records the report displays depends on the number of Agents on which the job(s) runs. The report displays a record for every Agent on which the job(s) runs. For more information on the fields in the report creation screen, see [Report Fields](#).

To generate a detail report:

1. From the Management Console, click **Snap Solutions > Reports**.

2. Click **New Detail Summary Report**.

The Detail Summary Report screen appears:

1 - Select a Job Group

Job Group: Report In:

Job Template: Output Format:

2 - Select Columns to Include

<input checked="" type="checkbox"/> Job Group	<input checked="" type="checkbox"/> Job Name	<input checked="" type="checkbox"/> Job ID	<input checked="" type="checkbox"/> Job Template Type
<input checked="" type="checkbox"/> Job Template	<input checked="" type="checkbox"/> Source Agent	<input checked="" type="checkbox"/> Target Agent	<input checked="" type="checkbox"/> Start Time
<input checked="" type="checkbox"/> Elapsed Time	<input checked="" type="checkbox"/> Files Transferred	<input checked="" type="checkbox"/> Files Skipped	<input checked="" type="checkbox"/> Files Deleted
<input checked="" type="checkbox"/> Bytes Transferred	<input checked="" type="checkbox"/> Bytes Skipped	<input checked="" type="checkbox"/> Transfer Rate	<input checked="" type="checkbox"/> Errors

3 - Select a Date Range

4 - Generate Report

from to

Save as template:

3. In the **Select a Job Group** list, click the job group for which you want to generate the report.

Use shift+click to select multiple consecutive job groups, or ctrl+click to select multiple non-consecutive job groups.

4. In the **Report In** drop-down, choose the unit in which you want the statistics to appear. (Choose from bytes, kilobytes, megabytes or gigabytes.)

5. In the **Job Template** dropdown list, choose a template to use for the report.

6. Choose the **Output Format** by selecting either HTML or CSV.

CSV stands for Comma Separated Variable. You can load and save this type of file in a spreadsheet.

7. In the **Select Columns to Include** area, click the appropriate check boxes to specify the data you want to include in the report.

8. In the **Select a Date Range** area, type the start and end date of the time period for which you want to generate the report.

9. In the **Generate Report** field, type a name for the template of this report.

This template will appear in the Report view, so that you can simply click on it to run a new report.

10. Click **Generate Report**.

The Detail Report appears.

11. Close the report window to return to the report setup screen, or use your browser's print command to print the report.

For more information on the fields in generated report, see [Generated Report Categories](#).

Custom Query Report

The Custom Query report provides a customized query of detail records from a combination of job names and job templates within a selected job group. Note that the number of records the report displays depends on the number of Agents on which the job runs. The report displays a record for every Agent on which the job runs. For more information about the fields in the Custom Query Report screen, see [Report Fields](#).

To create a Custom Query Report:

1. From the Management Console, click **Snap Solutions > Reports**.
2. Click **New Custom Query Report**.

The Custom Query Report screen appears.

The screenshot shows the 'New Custom Query Report' configuration interface. At the top, there are navigation links: '+ New Stats Summary Report', '+ New Detail Report', and 'Back to Reports'. The main form is organized into four numbered sections:

- 1 - Enter Search Criteria:** Contains dropdowns for 'Job Group' and 'Job Template', a text input for 'Job Name', a 'Report In' dropdown set to 'Bytes', and an 'Output Format' section with 'HTML' and 'CSV' options. A radio button interface allows searching 'with all' or 'any criteria'.
- 2 - Select Columns to Include:** A grid of 12 checkboxes, all of which are checked. The columns are: Job Group, Job Name, Job ID, Job Template Type, Job Template, Source Agent, Target Agent, Start Time, Elapsed Time, Files Transferred, Files Skipped, Files Deleted, Bytes Transferred, Bytes Skipped, Transfer Rate, and Errors.
- 3 - Select a Date Range:** Two date pickers are shown, with the range set from '2008/05/14' to '2008/05/21'.
- 4 - Generate Report:** A text input field for 'Save as template' and a blue 'Generate Report' button.

3. In the **Job Group** drop-down list, click the job group for which you want to generate a report.
4. In the **Job Name** field, enter the name of the job for which you want to generate a report.
5. In the **Job Template** dropdown list, choose a template to use for the report.
6. In the **Report In** drop-down, choose the unit in which you want the statistics to appear. (Choose from bytes, kilobytes, megabytes or gigabytes.)
7. Choose the output format by selecting either HTML or CSV.
CSV stands for Comma Separated Variable. You can load and save this type of file in a spreadsheet.
8. Set the search parameters by choosing one of the following radio buttons:
 - **All** (to generate a report only if all the search criteria are found)
 - **Any criteria** (to generate a report if any of the search criteria are found)
9. In the **Select Columns to Include** area, click the appropriate check boxes to specify the data you want to include in the report.
10. In the **Select a Date Range** area, type the start and end date of the time period for which you want to generate the report.
11. In the **Generate Report** field, type a name for the template of this report.
This template will appear in the Reports view so that you can simply click on it to run a new report.
12. Click **Generate Report**.
The Custom Query Report appears.

13. Close the report window to return to the report setup window.

For more information about the fields in a generated report, see [Generated Report Categories](#).

Generated Report Categories

The following table describes the fields that may appear in the generated report, depending on which columns a user chooses to include in the report. Note that the Source, Target, Start Time and Job ID fields do not appear in the Stats Summary Report. The Date field appears only in the Stats Summary Report.

Generated Report Field	Description
Date	The date on which the job ran. Note that the date may appear as a day, month or year, depending on the roll up value the user chooses.
Job Group	The name of the job group included in the generated report.
Job Name	The name of the job included in the generated report.
Job ID	The Job ID number of the job included in the generated report.
Job Template	The name of the job template associated with the selected job.
Job Name	The name of the job included in the generated report.
Source	The name of the agent on which the transfer originated. Note that no source agent will be listed for a remote command transfer.
Target	The name of the agent to which the transfer was sent.
Start Time	The time at which the job began.
Elapsed Time	The length of time it took the job to complete.
Files Transferred	The number of files transferred.
Files Skipped	The number of files that were not transferred. A file is skipped when it already exists on the target and is the same as the source (i.e., it has not changed).
Files Deleted	The number of files deleted.
MB Transferred	The number of Megabytes transferred. Note that this value could be reported in bytes, kilobytes or gigabytes depending on what the user specified when creating the template.
MB Skipped	The number of Megabytes skipped. Note that this value could be reported in bytes, kilobytes or gigabytes depending on what the user specified when creating the template.
Transfer Rate (MB/sec)	The rate at which the data was transferred (the amount of data divided by the execution time for the transfer). Note that this value could be reported in bytes, kilobytes or gigabytes depending on what the user specified when creating the template.
Average Transfer Rate	The average data transfer speed when multiple jobs have run.
Errors	The number of errors that occurred during the transfer.
Record Count	The total of all the records used to generate the particular entry in the report.

Generate a Report from a Template

When you create a report, you can assign a template name to it. The template appears in the Report view. Each template is associated only with the user who creates it. Templates appear in the Report view only if you create a template. Report templates allow you to generate reports without having to change the parameters you initially set up. Templates are especially useful if you plan to generate the same type of report on a regular basis.

You can also edit templates to change some of the parameters. To generate a report from a report template:

1. From the Management Console, click **Snap Solutions > Reports**.
2. Click **Run** beside the name of the template you want to use to generate a report.
The report you generated appears.
3. Close the window to return to the Reports menu.

NOTE: To edit a report template, click **Edit** beside the template you want to change. Note that if you want to keep the same template name, make sure a check appears in the **overwrite template** check box.

To delete a report template, click **Remove** beside the template you want to delete. You are prompted to confirm that you want to delete the template. Click **OK**.

Schedule a Report from a Template

When you create a report, you can assign the report a template name. The template appears in the Report view. Report templates allow you to generate reports without having to change the parameters you initially set up. Templates are especially useful if you plan to generate the same type of report on a regular basis. You can schedule a report to run on a daily, weekly or monthly basis. The scheduled report is e-mailed to the recipient(s) you specify in the schedule. Note that the report appears within the text of the e-mail and not as an attachment.

To schedule a report from a report template:

1. From the Management Console, click **Snap Solutions > Reports**.
2. Click **Schedule** beside the name of the template you want to schedule.
The Schedule Report screen appears.
3. In the **Frequency** drop-down, choose how often you want to generate the report. Choose from daily, weekly or monthly.
4. In the **Start at** field, type the time at which you want the first report to be generated.
5. In the **E-mail report to** field, type the e-mail address(es) of anyone who you want to receive the report. Separate multiple e-mail addresses with a comma.
6. Click **Add Job**.

NOTE: To edit a report template schedule, click **Edit** beside the report template schedule you want to change. Make changes to the schedule and click **Save**.

By default, Snap EDR creates a number of functions you can schedule to view and trim the Management Console logs. The procedures in this chapter describe how to automate these maintenance tasks, and how to view log files. This chapter also describes how to upgrade you individual Snap Solutions.

-

Trim Manager Database Logs

The Management Console restricts the amount of disk space used for logging by running a daily Log Maintenance job that limits the log directory threshold size to 500 megabytes.

Periodically deleting logs is especially important for the delivery logs written to the transfer log directory on the Management Console (these logs are not in the database). A delivery log is generated every time a data transfer occurs and the accumulated logs can quickly take up a substantial amount of disk space.

A default Log Maintenance job is configured and scheduled to run daily. However, users can change the default values in the existing log maintenance job to specify values suitable to their own environment. It is recommended that users schedule the log maintenance job to run once a day (which is the default value). For information on the values users can specify in the Log Maintenance job and the default values with which the job is configured, see the table in the procedure below.

To schedule the template to trim the log and statistics:

1. Click **Administration > Manager > Log Maintenance**.
2. In the actions area, click **Add a Job**.
3. Set information in the fields. The following table lists the fields in the Log Maintenance screen.

Field	Description
Job Name	A unique name for the job.
Database Maintenance Parameters	
Remove statistics and web audit logs older than	A drop-down list that allows users to specify the upper limit on the number of days statistics and logs generated by the Web UI are kept in the database. Default value is 365 days. Choose Never to disable this feature, and keep all statistics and web audit logs. The age of a job run is determined by the time at which the job run started.

Field	Description
Remove ONCE frequency jobs	Choose Yes to delete jobs with a frequency of ONCE from the database once the statistics from the job have been removed. NOTE: If the Remove statistics and web audit logs older than value is set to Never, ONCE jobs will not be removed at all (since their statistics are never removed).
Remove empty file transfer runs	Choose Yes to delete job run records and associated job log files for job runs where no files are transferred.
Transfer Manager Log Maintenance Parameters	
Keep log directory size below	A drop-down list specifying the maximum size of the log directory, in megabytes. The default value is 5000 MB (5 GB). When the log directory size is greater than the specified value, log files are sorted in reverse chronological order, and the oldest log file is removed until the directory size is below the specified value. When all the logs that can possibly be removed are removed and the directory size is greater than the value, the job will fail and a failure notification will be triggered. Choose Unlimited to disable this feature. Logs will be removed based on directory size.
Keep job logs for at least	A drop-down list specifying the minimum number of days a log file will be kept. A log file whose age is less than the specified minimum number of days will not be deleted. Log files older than the specified minimum are deleted only if the log directory size value is exceeded. The default value is 28.
Keep "Delivery" logs for at least	A drop-down list specifying the length of time (in days) to keep delivery logs. Only delivery logs older than the specified value will be removed when the log directory size exceeds the directory size parameter. The default value is 365. Choose Forever to disable this feature, and keep an unlimited number of delivery logs.
Remove all log files older than	A drop-down list specifying the upper limit (in days) to keep a log file. Log files older than the age limit will be removed, even if the directory size limit has not been reached. The default value is 365 days. Choose Never to disable this feature, and keep all log files.
Job Options	
Log Detail Level	The type of logging information for this job. Choose from Error , Warn (warning), Info (Information), or Debug . Debug provides the greatest level of detail while Error provides the least.
Scheduling Parameters	
Job Start Date/ Time	The time at which you want the job to run (in yyyy/mm/dd hh:mm:ss format). Users can also click the calendar icon to select values.
Frequency	How often you want the job to run. Choose from once, hourly, daily, weekly, monthly, yearly, monthend, or none. In most cases, you will want to run the job once. Choose none to run the job immediately after creating it.

Field	Description
Time Zone	<p>Specifies the time zone in which the displayed times are set in the template. For example, if a solutions developer in an Eastern Time zone specifies a start time of 9:00 am and a time zone of PST, the job runs at 9:00 am Pacific Standard Time or noon Eastern Standard Time.</p> <p>Clicking the globe icon displays the "Choose a Time Zone" selection screen. Click on the area of the map in the time zone you want to set. Geographical locations (such as cities or countries) appear in the area below the map. Click on one of the geographical locations to set it as the new time zone. You can click on the left/right arrows to scroll to time zones before or after the currently-selected time zone. Click the up/down arrow to scroll between the northern hemisphere, the equator and the southern hemisphere.</p> <p>NOTE: The time zone value that appears by default may not precisely reflect the time zone specified by the GuardianOS server. In some cases, the Management Console provides multiple locations for a single time zone (for example, "America Detroit, US Eastern, America Montreal and so on for Eastern Standard Time). The value will default to the first time zone option in the list where multiple options are available.</p>
Email Notification	
Email Notification	<ul style="list-style-type: none"> • Always: Users identified in the notification list receive email notification messages both when a job succeeds or fails. • On Transfer: Users receive an e-mail only when a transfer succeeds. • On Error: Users receive an e-mail only when a job fails. • Never: Users do not receive e-mail notifications of transfers.
Email Subject	<p>Text that appears in the subject field of the report e-mail message.</p> <p>NOTE: To receive email notifications in Snap EDR, you must first configure email notification in GuardianOS at Server > Email Notification.</p>
SNMP Trap Notification	
Send Traps On	<p>Specifies the circumstances under which a Simple Network Management Protocol (SNMP) trap is sent. Choose from Job Success, Job Error, both or neither.</p> <p>NOTE: To send SNMP traps in Snap EDR, you must first configure SNMP in GuardianOS at Network > SNMP.</p>

4. Click **Add job**.

View Transfer Logs

The Transfer Logs menu displays data transfer logs generated by the Management Console during each transfer. Users can click on the logs listed to view them.

Users can specify whether logs are certified or uncertified by selecting the Delivery mode (in the File Transfer Options category) when creating a job. The Log File Name option creates a list of the files transferred, but does not certify them. The Certify File Content option creates certified delivery logs.

If the job runs with certified delivery enabled (either signed or unsigned), the log server creates a certified delivery log file on the Management Console in the "delivery_logs" subdirectory. With certified delivery, the source and target Agents use their private keys to

sign hashes separately for each file they transfer. (A hash is an algorithm that creates a message digest for authentication.) Comparing these hashes determines whether the file transfers complete with no change to the data.

To view transfer logs:

1. From the Management Console, click **Administration > Manager > Transfer Logs**.
2. Click the logs to download and view the details.

Upgrade Applications

Periodically, you may want to upgrade one or more of the Snap Solutions to take advantage of updates or new releases. These applications may exist on CD or another computer.

Upgrade an Application

NOTE: Before you can use the application, you must ensure that it is located on a computer to which you have access.

1. From the Management Console, click **Administration > Manager > Applications**.
2. Use the Browse button to navigate directories and select the file for the solution you want to upgrade. Click **Upgrade Application**.

You can also click on the name of the application to view a summary screen that displays information such as the application's status, version number and so on.

Uninstall Applications

An "x" appears beside the Snap Solutions application entry in the **Administration > Manager > Applications** screen. This allows users to uninstall all of the Snap solutions.

View Application Summary Information

In the **Administration > Manager > Applications** screen, click the application to display summary information about the application, such as its status, version number and so on. Clicking on links that may appear on the screen takes the user to screens associated with those items.

This section describes the recommended configuration options and best practices to deploy the Remote Backup and Remote Restore solutions.

Remote Backup Best Practices

Size the Solution

In most cases, as part of the pre-deployment exercise, it is necessary to calculate the number of components (e.g., receiving Agents) and determine the appropriate job schedule.

You will need to fill in the following table to complete this section.

Characteristic	Value	Notes
Number of sites		Using one transfer Agent per site
Files:		
Average Files per site		Customer provided - user files only
Maximum Files at a site		Customer provided - user files only
Average Size of Files		Customer provided
Rate of Change:		
Average percentage of total files changing per day		Customer provided
Average percentage delta for changed files		Customer provided
Average number of files changing per day		
Maximum number of files changing per day		
Average delta MB		Average percentage delta of (average number of changed files storage)
Maximum delta MB		Average percentage delta of (average number of maximum changed files storage)
Bandwidth and Throughput:		
Available Bandwidth per site		
Normal incremental transfer rate*		There is a wide variation in the incremental transfer rate (based on rate of change and average file size).

Characteristic	Value	Notes
File scanning rate		There is a wide variation in the file scanning rates.
Sizing:		
Target transfer agents		
Desired Completion Time:		
Normal backup		e.g., to be run every evening
Synchronizing backup		e.g., to be run once every weekend

* Incremental transfer rate is the following:

= (Delta MB changes/Time taken to transfer changes when using Transfer File Differences Only option)

The effective throughput is much higher and represents the savings of not having to send the entire file.

= (Total MB changes/Time taken to transfer changes when using Transfer File Differences Only option)

Maximum concurrent Agents can be raised up to a maximum of 32, but the default value is 8 in the Snap EDR environment. Generally, the incremental transfer rate will approach the available bandwidth as the average file size increases.

Regular Backup Volume and Frequency

The average regular backup volume for a single site is calculated by the following:

= avg # of files X avg % of files changing X avg % file delta X avg file size

To determine the time required to complete a regular backup, use the following:

= (File comparison time) + (File transfer time)¹

File Transfer Time = Backup Volume / Incremental Transfer Rate²

File Comparison Time = # of Files / File comparison rate³

Synchronize Backup Volume and Frequency

The average synchronizing backup volume for a single site is calculated by the following:

= avg # of files X avg % of files changing X avg % file delta X avg file size

To determine the time required to complete a synchronizing backup, use the following:

= (File comparison time) + (File transfer time)⁴

File Transfer Time = Backup Volume / Incremental Transfer Rate⁵

File Comparison Time = # of Files / File comparison rate⁶

Required Target Agents

The required number of target Agents varies depending on the desired completion time.

- 1 To simplify calculations, these two values are added together even though the operations run concurrently (i.e., this time prediction is 'pessimistic' using this approach).
- 2 This rate can vary widely—approaching the maximum available bandwidth.
- 3 The file comparison rate is the number of files that can be compared per second between branch and central site locations. This comparison involves the evaluation of the file information—i.e., the target site deciding if it needs an updated version of the file that the branch can send.
- 4 To simplify calculations, these two values are added together even though the operations run concurrently (i.e., this time prediction is 'pessimistic' using this approach).
- 5 This rate can vary widely—approaching the maximum available bandwidth.
- 6 The file comparison rate is the number of files that can be compared per second between branch and central site locations. This comparison involves the evaluation of the file information—i.e., the target site deciding if it needs an updated version of the file that the branch can send.

Simultaneous Running Jobs

The running job limit is minimized by following the best practice of configuring the receiving Agents to manage “many” site servers in a single job. If the “running job limit” is exceeded, the Management Console simply queues the jobs in a ‘delayed’ state (no jobs will be ‘lost’).

Transfer Read/Write Issues

If the Backup job is not using the file ownership transfer option (i.e., Off is selected), the user ID as which the transfer/backup runs must have read/write access to the source directories in order for the transfer/backup to succeed.

On Windows, for file ownership transfers, the Backup job user may have to be in the ‘Backup Operators’ group on the target and source machine.

For file ownership transfers on UNIX/Linux/Mac, the Backup job user must be root.

Changing the SnapServer Name or IP

As long as you have DNS and name resolution is working properly, you can change the name or IP address of your SnapServer (Management Console or Agent) and the change will be resolved across the Snap EDR system. If you do not have a DNS setup or if name resolution is not working correctly, you will need to add a host entry in GuardianOS (Maintenance > Host Editor).

Encrypted Files

If you are using the Backup solution in a workstation environment, users may have encrypted folders/files on their systems. Encrypted folders and files cannot be backed up. You will receive a Permission Denied error and the job will fail.

May 2008 07:41:39	Info	Data transfer in progress - Source host: ottws40 ott.signiant.com
May 2008 07:41:39	Error	Unable to open source file 'c:\testdata\Encrypted_tests\Encrypted.doc'. Permission denied
May 2008 07:41:40	Error	Unable to open source file 'c:\testdata\Encrypted_tests\Encrypted.doc'. Permission denied
May 2008 07:41:40	Error	Retry count exhausted for transmission of file 'c:\testdata\Encrypted_tests\Encrypted.doc' to s

File Ownership Transfer

The Remote Backup solution provides various methods to preserve file ownership attributes during a backup.

When using the File Ownership Transfer option for Windows, the Backup job user must be a member of the “Backup Operators” group on the Windows hosts. As well, ensure that the Backup job user has the following local security policies (the Backup Operators group may already be assigned to these security policies):

- Backup files and directories
- Manage Auditing and security log
- Restore files and directories

Target Directory for Backups

The target directory is required for input into the Remote Backup job. This directory specifies the base directory to store the entire backup data set on a per host basis. The directory should be specified at least one level deep.

Example:

```
c:\remotebackupdata
```

This enables the data to be more easily managed.

Users should create the target directory before running the Backup job. The user assigned to run the Backup job should be given full access to the directory. Although the Backup solution will try to create the target directory if it does not exist, the Backup job user may not have sufficient permissions to create the directory.

Source Directories Specified

The source directories to backup should be consistent across all source hosts of a scheduled Remote Backup job. If the source hosts to backup have different source directories, the source hosts with the common directories should be grouped into the same job.

For example, hosts 1 through 5 must be backed up. Hosts 1 and 3 have common directories to backup, `c:\appdata\userfiles`. Hosts 2, 4, and 5 have common directories to backup, `c:\webservers,e:\documentrepository`.

Create one job for hosts 1 and 3 specifying the source directory `c:\appdata\userfiles`, and another for hosts 2, 4, and 5 specifying the source directories `c:\webservers,e:\documentrepository`.

Target Host High Speed Link to Attached Storage

If the target host is using locally-attached storage to maintain all of the backup data, it is recommended the network connection between the target host and storage device be as fast as possible.

Synchronized Vs. Regular Backup

It is recommended to perform a regular (nonsynchronized) backup on a regular basis. A synchronized transfer can be scheduled to run occasionally to trim the directories on the target host storage. This will avoid the latest directory on the target host from growing far beyond the size of the data on the source hosts, since a synchronizing transfer will remove all directories and files from the target host storage that no longer exist on the source hosts.

The additional scheduling variable on the template “Run on days of week” can be used to achieve this. A regular (nonsynchronizing) backup can be scheduled to run “daily” and the “Days of week to Run” can be set from Monday through Friday. A synchronizing transfer can also be scheduled to run daily, but the “Run on days of week” can be set to Saturday.

As mentioned earlier in this guide, the “Days of Week to Run” specified in the template variables is relative to the time zone of the Snap EDR.

A synchronizing transfer takes longer to run.

Mixed Environment and Preserve File Properties

Preserve file ownership works optimally when the target host and source hosts are the same operating system. Use the following table as a reference

If Your SourceHost Is...	And Your Target Host Is...	Choose...
Off	Off	Off
Windows	Windows	Windows
Windows	GuardianOS	GuardianOS
Windows	UNIX/Linux/Mac	Off
GuardianOS	GuardianOS	GuardianOS
GuardianOS	UNIX/Linux/Mac	UNIX
UNIX/Linux/Mac	UNIX/Linux/Mac	UNIX
UNIX/Linux/Mac	GuardianOS	UNIX
UNIX/Linux/Mac	Windows	Off

Troubleshooting

When performing a backup with Windows machines the following errors may occur:

- Unable to enable the 'BackupAccess' privilege for USER: Not all privileges referenced are assigned to the caller
- Unable to determine characteristics of file 'file': Permission denied
- Unable to open source file 'file': Operation not permitted

These errors occur when the File Ownership Transfer option has been set and the Backup user does not have sufficient privileges to access certain files. To correct the problem, ensure that the Backup job user belongs to the "Backup Operators" group on the machine in question. As well, ensure that the Backup job user has the following local security policies (the Backup Operators group may already be assigned to these security policies):

- Backup files and directories
- Manage Auditing and security log
- Restore files and directories

Remote Restore Best Practices

This section covers scenarios related to restoring data with the Remote Restore solution. It is assumed the data to restore was backed up with the Remote Backup solution. The best practices in this section are illustrated through a set of use cases.

Job Scheduling

A Remote Restore job is not a scheduled job. Most likely, it will run on an ad-hoc basis between the source host (central repository) and a single target host for which data must be restored (note that it is not limited to a single target host). Once a Remote Restore job is created and saved, users must select the "Run Now" link to run the job.

File Ownership Transfer

The selection of the Restore File OwnershipTransfer option must correspond to the File Ownership option used to backup the data.

Use Case Data Set

The following data set will be used for the use cases in this section.

Remote Backup Job Data	
Backup Target Host	tgt.company.com
Backup Target Directory	e:\remotebackups
Directories to Backup	C:\userdocuments
Source Host1	Host1.company.com
Source Host2	Host2.company.com
Source Host3	Host3.company.com
Number of backup versions to maintain	3

The Remote Backup job using the above information runs nightly. The directory `c:\userdocuments` has the following subdirectories:

- `powerpointpresentations`
- `excelspreadsheets`
- `worddocuments`

The Remote Backup job scheduled above has been running for the last three months.

Use Case 1 - Specified Restore

In this case, a PowerPoint presentation must be restored from a version prior to January 05/08. Today's date is January 06/08. The user requesting the file works on host 1. The file is called `presentation1`. The user modified the file yesterday, and needs a version of the file prior to the modification.

Restore job parameters	
Source Host	tgt.company.com
Backup Root Directory	<code>e:\remotebackups</code>
Backup Type	Specified
Directory to Restore	<code>C:\userdocuments\powerpointpresentations</code>
Files to Restore	Presentation1
Target Host	Host1.company.com
Restore From Version	1
Restore to Directory	<code>C:\temp</code>

In this case, the user modified the file on January 05/08. A backup is run every night. The file prior to the version modified on January 05/08 is moved from the latest directory for host 1 and stored in a directory called **080105223030**, the date the backup was run. Since the user discovered the error the day after the last backup, the most recent version of the file prior to the change will be in the most recent backup directory, 1.

The user wants to have the restored version transferred to **c:\temp** to compare differences with the version in **C:\userdocuments\powerpointpresentations**.

Use Case 2 - Full Restore of Changes

In this case, a user wants to restore all files in their previous state that the user modified on January 03/08. Today's date is January 06/08. The user requesting the file works on host 2.

Restore job parameters	
Source host	tgt.company.com
Backup Root Directory	e:\remotebackups
Backup Type	Full
Target Host	Host2.company.com
Restore From Version	3
Restore to Directory	C:\temp

In this case, all the files that had been modified on January 03/08 caused the versions in the latest directory to be moved to the directory **080103223030**. The user would get back all the versions of the files prior to the modifications on January 03/08.

The files are restored to **c:\temp** to compare the differences with the most current versions on host 2.

Use Case 3 - Full Restore from Latest

A full restore must be done for all hosts since a malicious user deleted and/or corrupted the user documents folders on all of the machines.

Restore Job Parameters	
Source Host	tgt.company.com
Backup Root Directory	e:\remotebackups
Backup Type	Full
Target Host	Host1.company.com Host2.company.com Host3.company.com
Restore From Version	Latest
Restore to Directory	

In this case, all the backup data from the last backup, latest, will be restored to all of the hosts specified. Since the root directory to restore to on the target host was specified, the data will be restored back to its original location. Any existing files will be overwritten.

The following tables describe the scheduled job run exit codes, along with possible causes for error codes and steps that can be taken to correct any problems.

Exit Code	Cause	Possible Resolution
0	Job completed with no errors	Not applicable
5	<p>Fault Message Received May be caused by the following:</p> <ul style="list-style-type: none"> • An invalid or unexpected request was sent to the Management Console • A critical server-side memory shortage caused request to be preempted • A critical server-side database connection problem caused the request to be preempted • The job may have been killed by one user, when another is requesting the 'status' of the job 	<ul style="list-style-type: none"> • Check that the database is up and responding. • Clear some system memory by stopping unnecessary running processes.
6	<p>Syntax Error Message Received</p> <p>An invalid protocol message was received by the Management Console. Indicates that the Management Console and client Agents may be out of sync.</p>	Internal error. Requires engineering level resolution.
8	<p>Invalid Command in Current State</p> <p>The scheduled job's state is incompatible with the requested command.</p>	A user has issued a command against a job that is incompatible with the state of the job (e.g., suspend command used against an already suspended job).
9	<p>Validation Error</p> <p>Scheduled job's state was assigned to INVALID during execution of the requested directive due to incorrect scheduling parameters or due to a corrupted internal state. Corrective action and re-evaluation are required to restore the job and de-assert the execution blocking condition.</p>	Internal error. Requires engineering level resolution.
12	<p>Memory Allocation Error Insufficient memory is available to run the Management Console.</p>	Internal error. Requires engineering level resolution.
15-18	<p>Authentication Error/Prompt Error/Missing File Error/Node Definition Error</p>	Internal error. Requires engineering level resolution.
20	Quit Requested	Verify that the operator is not manually stopping the process.
21	Abort Requested	Verify that the operator is not manually stopping the process. Note This message may also appear if a job template link trigger command is incorrectly created.
22	Internal Error	Internal error. Requires engineering level resolution.

Exit Code	Cause	Possible Resolution
23	Process Start Error	The Management Console was unable to start a process that is needed to start as part of running the job. This may be due to an internal error.
24	Prepackage Error	Contact support.
25	Prepackage Error	Contact support.
26	Transfer Error	Contact support.
27	Rollback Cleanup Error	Contact support.
28	Job Template In Use	Contact support.
29	Session In Use	Contact support.
30	Session Dead Error	Contact support.
31	Password Cache Error	Contact support.
32	Job Failure These errors may be caused by, but are not limited to, the following: <ul style="list-style-type: none"> • Misconfigured or invalid grants. • Incorrect source or target parent directory specified (e.g., directory does not exist, or transfer user does not have access to the directory). • One or more hosts are unavailable. • A syntax or logic error in one of the template's commands. 	<ul style="list-style-type: none"> • Check the Manager's log by selecting View beside the most recent job listed in the jobs summary screen, and correct any errors indicated. • Check the Transfer logs on the source and target agents. The logs are written to the log directory. • Schedule the job again, selecting debug in Job Options > Job Log Detail Level. These logs contain information that may be useful in troubleshooting (however, they also require more disk space).
33-48	Authorization errors	Internal error. Requires engineering level resolution.
139	Segmentation Fault	The running process produced a segmentation fault.

Symbols

- .req.pem 17
- > (menu flow indicator) 5

A

- agent groups 17
 - creating 17
 - deleting 17
 - editing 17
- aggregate data management tool 23
- aggregate job
 - creating 24
- alert definitions 5

B

- bandwidth, changing 28

C

- conventions, typographical 5
- creating
 - custom query report 82
 - detail report 80
- custom query report
 - creating 82
- customer support 4

D

- dashboard
 - adding widgets 8
 - changing widget name 8
 - configuring widgets 8
- data layout
 - remote backup 56
- data management tool
 - aggregate 23
 - distribute 30
 - remote backup 48
 - replicate 38
- data resolution 66
 - remote restore 58, 66
- delivery log 87
- detail report
 - creating 80

- directory and file options 24, 31, 39, 49, 50
- distribute
 - creating job 31
- distribute data management tool 30

E

- enable service on start-up 11
- encrypted files 91
- encryption levels 29, 36, 43, 46, 54, 64
- errors
 - installation 15
- exit codes 96

F

- file transfer options
 - remote restore 55
- full restore 66

G

- generating
 - report from template 84

I

- installation
 - agent installation host 13
 - Transfer Agent
 - errors, UNIX 15
 - log file, UNIX 15
 - on LINUX-based systems 15
 - on Mac-based systems 16
 - on Windows-based systems 13
 - unsuccessful 16
- installation errors 15

J

- job list detail 70
- job list summary 69
- job statistics, viewing 71

L

- latest directory
 - remote restore 58

- license keys 18
 - adding 18
 - deleting 18
- log files 72
- log maintenance 85
- logs
 - trimming 87
 - viewing transfer 88

M

- management console 3
 - dashboard 6
 - installing 4
 - overview 6
 - upgrading 5
- managing jobs 69
 - job list detail screen 70
 - job list summary screen 69
 - performaing tasks on more than one job 76
 - updating variables 77
- menu flow indicator 5
- Microsoft SystemState backup 58
- multi-tasking 76

O

- Overland technical support 4

P

- port requirements 21
- preferences 7
- preserve file ownership
 - mixed environments 93
 - remote backup 92

R

- read/write issues 91
- remote backup 48
 - best practices 89
 - creating job 48
 - data layout 56
 - details 67
 - file transfer options 67
 - frequency 90
 - preserve file ownership 91

- required target agents 90
 - sizing 89
 - source directories 92
 - target directory 92
 - troubleshooting 93
- remote restore 58, 66
 - data resolution 58
 - file transfer options 55
 - source host 58, 66
 - target host 58, 66
- remote restore data location 67
- remote restore data management tool 58
- remote restore versions 58
- replicate
 - creating job 38
- replicate data management tool 38
- report
 - custom query 78, 82
 - detail 78, 80
 - generating from template 84
 - scheduling from template 84
 - stats summary 78
 - types 78
- required target agents 90
- restart service 4
- re-sync with management console 21
- revoking certificates 19
- running jobs
 - simultaneous 91

S

- simultaneous running jobs 91
- Snap EDR Agents
 - downloading 12
 - installing LINUX 15
 - installing Mac 16
 - installing Windows 13
 - installing windows 13
 - revoking certificates 19
 - system requirements 10
 - uninstalling 20
 - uninstalling from Linux 20
 - uninstalling from Mac 21
 - uninstalling from SnapServer 20
 - uninstalling from Windows 20
 - unsuccessful installation 16

Snap EDR Express **7, 10, 12, 13, 14, 15, 18**
 source directories
 remote backup **92**
 source host
 remote restore **58, 66**
 specified restore **67**
 stats custom query report
 creating **82**
 stats detail report
 creating **80**
 stats summary report **78**
 creating **79**
 stop service **4**
 synchronized transfer **92**
 synchronizing backup volume **90**
 system requirements, Agents **10**

W

widgets **8**

T

target directory
 remote backup **92**
 target host
 remote restore **58, 66**
 target host selection **66**
 technical support **4**
 template
 generating report **84**
 transfer logs **87**
 troubleshooting
 remote backup **93**
 typographical conventions **5**

U

uninstall service **4**
 uninstalling
 Linux **20**
 Mac **21**
 Snap EDR Agents **20**
 Windows **20**
 unsuccessful installation **16**

V

versioning
 remote restore **58, 66**