



TECHNICAL
WHITE
PAPER

Tintri Cloud Connector

Technology Primer & Deployment Guide

Revision History

Version	Date	Description	Author
1.0	12/15/2017	Initial Release	Bill Roth

Table 1 - Revision history

Contents

- Introduction..... 4
 - Any Schedule, Any Destination, Any Retention Period... 4
 - Cost Reduction 5
 - Transition Point Elimination..... 5
 - Security..... 5
 - Performance Derived from Efficiency..... 5
- Prerequisites..... 6
 - Tintri Global Center 6
 - Resource Downloads 6
 - Cloud Connector Licensing..... 6
 - Amazon Account..... 6
 - IBM COS Account 6
- Considerations..... 6
 - Recovery Point Objective 6
 - Recovery Time Objective 7
 - Passphrase 7
- Deployment 7
 - Cloud Connector License 7
 - Adding Cloud Replication Destinations 8
 - AWS S3 9
 - IBM COS 10
 - Changing the Encryption Passphrase..... 11
 - Route Selection 12
 - Configuring Replication..... 13
 - Tintri Global Center Service Groups..... 13
 - Snapshot Scheduling..... 13
 - Replication 15
- Viewing Snapshots 16
- Cloud Snapshot Download..... 17
 - Recovery Options 19
 - Cloning 19
 - Restoring..... 20
- Conclusion 21

Introduction

The Tintri Cloud Connector, introduced with Tintri OS version 4.4, adds the ability to replicate Tintri snapshots to cloud destinations. Tintri snapshots can now be replicated to AWS S3 (Simple Storage Services) and on-premises IBM COS (Cloud Object Storage).

Tintri systems feature significant built in data protection and disaster recovery features. Tintri snapshots are the basis for creating recovery points. A snapshot is a point in time copy of an individual virtual machine. Tintri asynchronous replication creates copies of snapshots on one or more different Tintri systems. Tintri asynchronous replication has been augmented with the ability to replicate from Tintri systems to supported cloud destinations using the S3 API (Application Programming Interface).

Combining Tintri asynchronous replication destinations with cloud replication destinations within a single protection plan introduces a new level of functionality that enables enhanced data protection and disaster recovery schemas.

Any Schedule, Any Destination, Any Retention Period...

Tintri provides a flexible data protection and disaster recovery solution that includes the ability to deploy multiple snapshot schedules, each with its own local retention value. Each schedule can also be configured to replicate to one, two, three, or up to four different Tintri systems, with a remote retention value. With the Cloud Connector, a snapshot schedule can also be configured to replicate to one or more cloud replication destinations, each with its own distinct cloud specific retention value.

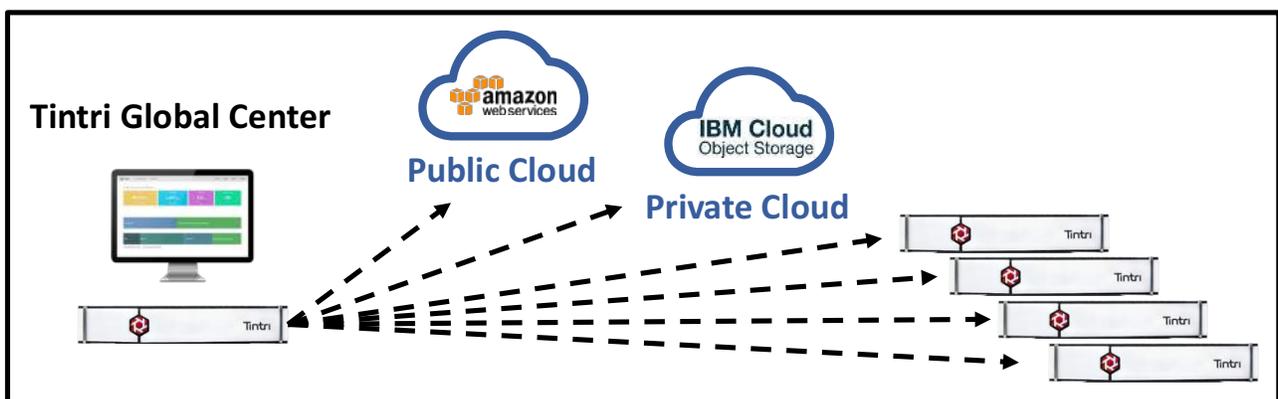


Figure 1 - Snapshot replication destinations

A combination of snapshot schedules, replication destinations, and retention periods are easily configured to create customized protection plans. As an example, a single VM (Virtual Machine) or group of VMs could be protected with:

- Hourly snapshots retained locally for 24 hours.
- Daily snapshots retained locally for 1 week that are also replicated to a different Tintri system where the replica snapshots are retained for 2 weeks.
- Weekly snapshots retained locally for 2 weeks that are also replicated to an on-premises IBM COS destination that are retained for 3 months.
- Monthly snapshots retained locally for 2 months that are also replicated to an AWS S3 region that are retained for 2 years.

Cost Reduction

The use of cloud replication may present a cost reduction opportunity for deployments that:

- Currently use tape media for long term retention.
- Currently use tape media for the creation of offsite backup copies.
- Currently use a purpose built backup appliance to retain backup copies on a different media type.

Transition Point Elimination

The Tintri Cloud Connector replicates scheduled snapshots directly to one or more cloud destinations. There is no requirement to use or pay for a cloud gateway. There is no requirement to use a third party backup application. There is no requirement to stage snapshot data on a local caching device prior to replication.

When a cloud resident snapshot is needed for restoration, it can be easily downloaded to any Tintri system managed by TGC (Tintri Global Center). There is no requirement to download snapshots to the Tintri system on which they originated. Not only does this capability eliminate multiple phase recovery steps that may be required with caching or gateway devices, it also enables the ability to easily and frequently test disaster recovery processes.

Security

Cloud Connector security is based on the Advanced Encryption Standard (AES). AES-128 is used to encrypt snapshot data, at transmission time, to a supported cloud destination. Snapshots remain encrypted throughout their retention period in the cloud. When downloaded to any Tintri system managed by TGC, snapshots are automatically decrypted at arrival time.

When a cloud replication destination is created within TGC, it includes a passphrase parameter. At the time of creation, TGC automatically creates a media encryption key and an authentication key. The media encryption key is a randomly generated 128-bit string that is used to encrypt snapshot data. The authentication key is generated by a one-way hash function that uses the passphrase and random data (commonly referred to as “salt” data) as inputs. The resulting authentication key is 128 bits in length.

There is one media encryption key and one authentication key created for each cloud replication destination. The authentication key is pushed to all Tintri systems managed by TGC. The authentication key is also used to encrypt the media encryption key. The encrypted media encryption key is stored in a file along with the salt data in a file. The file containing the encrypted media encryption key and the salt data are stored in the cloud destination.

When replicating a snapshot to a cloud destination, a Tintri system uses the media encryption key to encrypt snapshot data at the time it is transmitted. When a snapshot is downloaded from a cloud repository to a Tintri system, the encrypted media encryption key is also downloaded. The Tintri system uses the authentication key to decrypt the encrypted media encryption key. The media encryption key is then used to decrypt the snapshot data.

Performance Derived from Efficiency

Tintri compression and deduplication results in the creation of VM snapshots that are space efficient. After an initial VM snapshot is created, new VM snapshots capture only the incremental block level changes that have occurred since the prior snapshot. This further reduces space consumption, and delivers optimized performance in terms of reduced cloud repository space consumption and a reduction in replication payload.

Each snapshot represents a full VM recovery point while consuming only the space required to capture block level changes since the prior snapshot. Downloading a snapshot from a cloud repository to a Tintri system results in the automatic creation of a full VM recovery point image. There is no need to manually synthesize a full recovery point from a collection of individual snapshots. This results in an immediate ability to use a downloaded snapshot for recovery.

Prerequisites

Tintri Global Center

The Tintri Cloud Connector is configured, managed, and monitored with TGC. TGC deployment is required in advance of deploying the Tintri Cloud Connector feature.

Resource Downloads

Documentation is available to assist in deploying TGC. The “Tintri Global Center System Admin Guide” is available for download at <https://support.tintri.com/> and contains detailed information on TGC deployment and use.

The TGC product is available as a download from the Tintri support site at <https://support.tintri.com/>. For Hyper-V environments, TGC is available for download as a “.ZIP” file. For VMware environments, TGC is available for download as an OVA (Open Virtual Appliance). TGC licenses are also available on the Tintri support site.

Cloud Connector Licensing

Each Tintri system that will be replicating scheduled snapshots to a cloud replication destination requires a “Cloud Connector” license. Each Tintri system that will be used as a download target for a cloud repository also requires a “Cloud Connector” license. In the context used within this document, cloud service provider is defined as a supported cloud replication destination. Beginning with Tintri OS release 4.4, supported cloud replication destinations include AWS S3 and on-premises IBM COS deployments.

Amazon Account

Utilizing AWS S3 as a cloud service provider destination requires an AWS account. The AWS account includes an access key and secret key. The use of the access key and secret key are required when creating an AWS S3 cloud replication destination. There may be a need to contact the internal department or team that manages the AWS account(s) in order to obtain this information.

IBM COS Account

Utilizing IBM COS as a cloud service provider destination requires an on-premises IBM COS deployment. When creating an on-premises IBM COS cloud replication destination, an access key, secret key, administrative end point, and data end point are required. There may be a need to contact the internal department or team that administers the on-premises IBM COS deployment in order to obtain this information.

Considerations

Recovery Point Objective

Snapshot schedules are typically configured to achieve a required RPO (Recovery Point Objective). RPO is usually defined as a maximum time period during which data may be lost. For instance, an RPO value of one hour implies that up to one hour of data loss is acceptable. The frequency at which scheduled snapshots are created dictates the recovery point objective that will be achieved. For example, a

snapshot schedule that creates daily snapshots results in a maximum RPO of 24 hours. The retention period associated with a snapshot schedule is also configurable such that the RPO can be achieved over the duration of a timeframe, one week for example.

Dependent on the achievable data transfer rate, it may not be possible to achieve a required RPO when replicating snapshots to a cloud replication destination. Users have the ability to retain snapshots locally on a Tintri system in addition to replicating those snapshots asynchronously to other Tintri systems. Users may choose to architect a protection plan that retains frequent snapshots locally on Tintri system, replicates those snapshots to another Tintri system, and replicates snapshots that occur less frequently to a cloud replication destination.

Unlike snapshots retained locally on a Tintri system, snapshots replicated to a cloud replication destination can be retained without utilizing Tintri file system space. Users may choose to architect a protection plan that achieves a required RPO over the duration of a timeframe that minimizes snapshot utilization of Tintri file system space.

Recovery Time Objective

Defined as the amount of time it takes to recover a VM and the services it provides, RTO (Recovery Time Objective) is an important part of a data protection and disaster recovery strategy. Snapshots residing in a cloud replication repository must be downloaded to a Tintri system before they can be used in a clone operation to create a functional VM. The time it takes to complete a download operation will increase the total recovery time when compared to a snapshot that is being retained locally on a Tintri system. Dependent on the achievable data transfer rate, it may not be possible to achieve a given RTO with snapshots retained in a cloud replication destination. Users may choose to retain a number of snapshots locally on a Tintri system, as well as snapshots replicated to a different Tintri system, in order to achieve a required RTO.

Passphrase

A final important consideration is the use of a passphrase. When creating a cloud replication destination, a passphrase is required. The passphrase is one component that provides encryption based security of the media encryption key.

Simply stated, the importance of passphrase selection and passphrase preservation cannot be over emphasized. There may be a requirement on the part of the user to consult with the local or corporate security team in order to comply with passphrase selection and passphrase preservation requirements.

Deployment

Cloud Connector License

The feature license is formally referred to as “Cloud Connector”. The license key contains the characters “CLDC”. Either terminology, “Cloud Connector” or “CLDC”, refers to the same license entity and entitlement.

The license is available in two different modes, “evaluation” and “permanent”. If activated, the evaluation mode license will remain active for a maximum of 30 days. An applied permanent mode license does not expire.

Cloud Connector licenses must be manually applied to any Tintri system that will be used to replicate snapshots to a cloud replication destination, or that will be used as a download target for a cloud repository.

License management tasks are performed from within the Tintri user interface by clicking:

- Settings > More > Licenses

Evaluation licenses			
Feature	Status	Key Type	Expiration
ReplicateVM	ACTIVE	Evaluation	2017 Nov 16, 5:43PM
SyncVM	ACTIVE	Evaluation	2017 Nov 16, 5:44PM
Cloud Connector	ACTIVE	Evaluation	2017 Nov 16, 5:44PM

Figure 2 - Evaluation licenses

The evaluation licenses depicted above include the Cloud Connector license.

Permanent licenses	
Feature	Key
Cloud Connector	T5040-CLDC- [REDACTED]

Figure 3 - Permanent licenses

An example of a permanent Cloud Connector license. Note that the “Key” field includes the “CLDC” nomenclature as a portion of the license key.

Adding Cloud Replication Destinations

Configuring a cloud replication destination requires that TGC has a network connection that is routable to the cloud service provider replication destination that is being added.

An individual Tintri system can support a maximum of 4 cloud replication destinations, as well as a maximum of 64 Tintri system replication destinations. Users with a collection of VM snapshots that are required to be replicated to more than 4 different cloud replication destinations should consider VM placement such that the collection of VMs residing on a given Tintri system does not require more than 4 different cloud replication destinations.

The user must have access to the required “access” and “secret” keys required for each cloud service provider. The user should also be prepared to use the desired Amazon region, or HTTP (HyperText Transfer Protocol) admin and data end points for use with an on-premises IBM COS deployment. Additionally, the user should also be prepared to use naming conventions that align with organizational naming requirements.

Cloud service provider destinations are administered by means of TGC. To access the Replication Destinations pane, click the “Explore” menu button, followed by the “Settings” “Edit all settings” button, and then the “Replication Destinations” menu item.

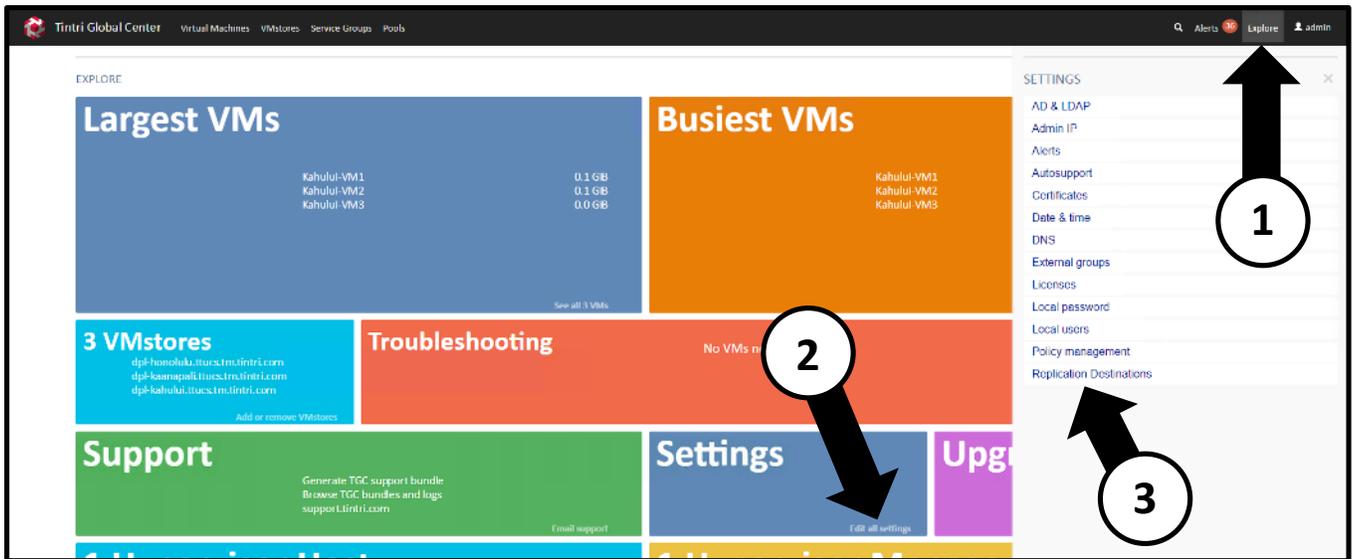


Figure 4 - Navigation: Replication Destinations

The graphic depicted above details the ordered steps required to access the “Replication Destinations” pane.

From within the “Replication Destinations” pane, any existing cloud replication destinations are displayed. New cloud replication destinations are added by first clicking the “Add Destination” button, and then selecting the desired cloud service provider from the pop-up menu.

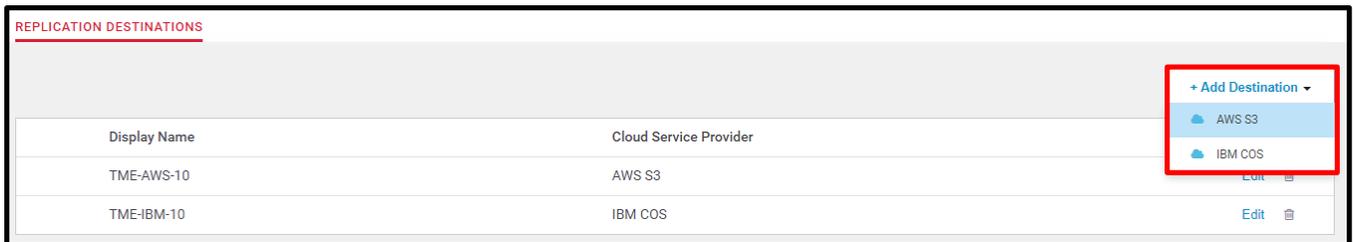


Figure 5 - Replication Destinations - Add Destination

The graphic depicted above details the “Add Destination” button and the available cloud service providers listed in the pop-up menu.

AWS S3

Adding an AWS S3 replication destination is a simple and straightforward task. A total of 6 parameters need to be specified:

1. “Display Name” is a descriptive user defined label that aligns with organizational naming requirements.
2. “Repository ID” is an automatically generated identifier prefixed with “tintri-“. The automatically generated string can be manually edited, if desired, at the time the replication destination is created.
3. “Encryption Passphrase” is a string, and is one component that provides encryption based security of an automatically generated media encryption key. The passphrase must be entered twice, once on the “Encryption Passphrase” line and a second time on the “Confirm Encryption Passphrase” line. The importance of passphrase selection and passphrase preservation cannot be over emphasized.

4. “Region used in Amazon” is a string selected by means of a pull-down menu. The selected string should match the desired Amazon region.
5. “Access Key” is a string provided by Amazon.
6. “Secret Key” is a string provided by Amazon.

Clicking the “Save” button will invoke creation of the AWS S3 replication destination. This action includes the creation of a “bucket” within the AWS S3 infrastructure. The user does not need to access aws.amazon.com to manually create the bucket.

The screenshot shows a form titled "Add Destination - AWS S3" with the Amazon Web Services logo in the top left. The form contains the following fields and controls:

- Display Name ***: Text input with "User Defined Display Name" and an information icon.
- Repository ID ***: Text input with "tintri-userdefineddisplayname-1509489287179" and an information icon.
- Encryption Passphrase ***: Text input with "User Defined Passphrase", a green checkmark, and a "Hide Key" link.
- Confirm Encryption Passphrase ***: Text input with "User Defined Passphrase", a green checkmark, and a "Hide Key" link.
- Region used in Amazon ***: A dropdown menu showing "us-east-2".
- Access Key ***: Text input with "Access Key" and a "Hide Key" link.
- Secret Key ***: Text input with "Secret Key" and a "Hide Key" link.

At the bottom, a yellow warning box contains the text: "The Encryption Passphrase is used to generate an encryption key for access to your data. It is important to keep it in a secure location (eg: password manager). You will need this key when doing key rotations".

Figure 6 - Add Destination - AWS S3

The graphic depicted above details parameters required to add an AWS S3 replication destination.

Warnings

Do not manually create an AWS S3 bucket within aws.amazon.com. TGC automates this process for the user. Manually creating an AWS S3 bucket is not a requirement. Manually creating an AWS S3 bucket is not recommended.

After creating an AWS S3 replication destination within TGC, do not manually add AWS S3 replication paths to a Tintri system. TGC automates this process for the user. When required, TGC will automatically push cloud replication paths to any Tintri systems that require it.

IBM COS

Adding an on-premises IBM COS replication destination is a simple and straightforward task. A total of 7 parameters need to be specified:

1. “Display Name” is a descriptive user defined label that aligns with organizational naming requirements.
2. “Repository ID” is an automatically generated identifier prefixed with “tintri-“. The automatically generated string can be manually edited, if desired, at the time the replication destination is created.
3. “Encryption Passphrase” is a string, and is one component that provides encryption based security of an automatically generated media encryption key. The passphrase must be entered twice, once on the “Encryption Passphrase” line and a second time on the “Confirm Encryption Passphrase” line. The importance of passphrase selection and passphrase preservation cannot be over emphasized.

4. “Admin End Point” is a HTTP IP address or DNS name provided by the IBM COS administrative team for use in establishing an archive.
5. “Data End Point” is a HTTP IP address or DNS name provided by the IBM COS administrative team. This may be the same as the “Admin End Point” in cases where both the admin and data end point paths are the same. In cases where a high performance replication network is utilized, the “Data End Point” path may be different than the “Admin End Point” path.
6. “Access Key” is a string provided by the IBM COS administrative team.
7. “Secret Key” is a string provided by the IBM COS administrative team.

The screenshot shows the 'IBM Cloud Object Storage' configuration interface. It features a grid of input fields for various parameters. On the right side, there are 'Cancel' and 'Save' buttons. A yellow warning box is present at the bottom left of the form area.

Figure 7 - Add Destination - IBM COS

The graphic depicted above details parameters required to add an on-premises IBM COS replication destination.

Warnings

Do not manually create an IBM COS archive. TGC automates this process for the user. Manually creating an IBM COS archive is not a requirement. Manually creating an IBM COS archive is not recommended.

After creating an IBM COS replication destination within TGC, do not manually add IBM COS replication paths to a Tintri system. TGC automates this process for the user. When required, TGC will automatically push cloud replication paths to any Tintri systems that require it.

Changing the Encryption Passphrase

Cloud replication destination passphrases can be changed based on business requirements. Editing an existing cloud replication destination will enable the ability to change an existing passphrase. Clicking the “Change” button initiates the passphrase change workflow.



Figure 8 - Change encryption passphrase

After the “Change” button is clicked, the existing passphrase for the cloud replication destination must be entered. Entering the existing passphrase enables the “Enter New Passphrase” button.

Figure 9 - Enter existing encryption passphrase

After the existing passphrase is entered, and the “Enter New Passphrase” button has been clicked, the user is prompted to enter and confirm the new passphrase. Entering and confirming the new passphrase enables the “Change Passphrase” button.

Figure 10 - Change passphrase

At this point the passphrase change takes effect. Transparent to the user, background processes create a new authentication key based on the new passphrase and salt data. The existing encrypted media encryption key is retrieved from the cloud repository. The media encryption key is decrypted with the original authentication key, encrypted with the new authentication key, and then stored with the salt data in a file. The file is then stored in the cloud replication destination. The new authentication key is also pushed to any Tintri systems managed by TGC.

Route Selection

Each Tintri system that will replicate snapshots to a given cloud replication destination must also have at least one configured network that is routable to the replication destination.

Each Tintri system may have up to 3 networks that can potentially be selected for use when replicating snapshots to a given cloud replication destination. These networks are referred to as the “replication network”, the “data network”, and the “admin network”. A Tintri system will attempt to connect to a cloud replication destination using a prioritized selection order. The selection order is:

1. Replication network
2. Data network
3. Admin network

The “replication network” is defined as the network that is configured for use by an optional dedicated replication NIC (Network Interface Controller). If dedicated replication NICs are not present, the “replication network” is not available and will not be considered for use when selecting a network path. If dedicated replication NICs are present but the network is not routable to the cloud replication destination, the network will not be selected as a network path.

The “data network” is defined as the network that is configured for use by any hypervisors that utilize the Tintri system as a datastore, share, or storage. If the network is not routable to the cloud replication destination, the network will not be selected as a network path.

The “admin network” is defined as the network that is configured for the purpose of administering the Tintri system. If the network is not routable to the cloud replication destination, the network will not be selected as a network path.

Note that with configurations having multiple networks that are routable to a cloud replication destination, the user can override the default network selection. For instance, the “replication network” was automatically selected and configured on the Tintri system and the user wants to use the “admin network” instead of the “replication network”. If the “admin network” is routable to the replication destination, the user can override the default network selection. This is accomplished by manually selecting the desired network over which replication should occur from within the Tintri system user interface.

Configuring Replication

Replication occurs at a VM snapshot level. Replication to a cloud replication destination requires VM snapshots to be created by means of a snapshot schedule. Snapshots that are manually created are not replicated to a cloud replication destination.

Tintri Global Center Service Groups

The creation and use of TGC service groups is not required to replicate scheduled VM snapshots to a cloud replication destination. In this context, the use of one or more TGC service groups is optional.

When using a TGC service group in conjunction with one or more cloud replication destinations, the service group type must support “Cloud Replication”. TGC service group options 1 & 3 support cloud replication.

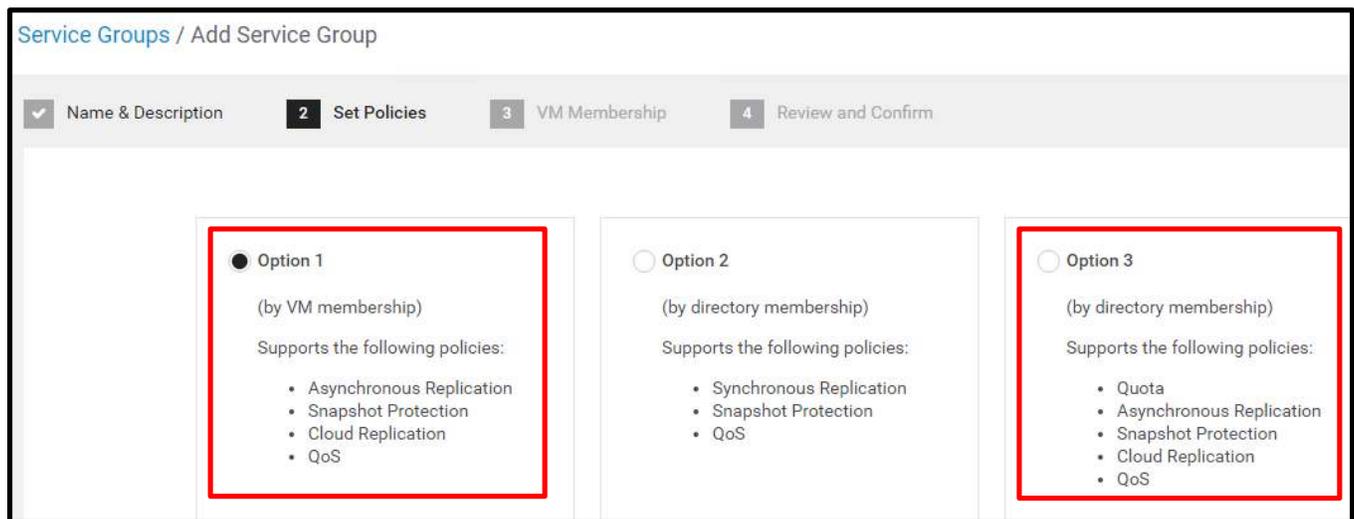


Figure 11 - Tintri Global Center service group options

The graphic depicted above denotes the TGC service group options that support cloud replication.

Snapshot Scheduling

Only snapshots created by means of a snapshot schedule can be replicated to a cloud replication destination. Snapshot schedules can be created for individual VMs as well as for TGC service groups.

The creation and administration of snapshot schedules for individual VMs and TGC service groups is essentially the same. Viewing the “Settings” tab of an individual VM or TGC service group will display any existing snapshot schedules.

Schedule type	Starts at	Specific Day	Retain Local	Retain Async Repl	Consistency
<input checked="" type="checkbox"/> Hourly	5:20,35,50 minutes past the hour	Mon,Tue,Wed,T...	2 Hours	3 Hours	Crash Consistent
<input checked="" type="checkbox"/> Daily	17:25	Everyday	2 Days	2 Days	Crash Consistent
<input checked="" type="checkbox"/> Weekly	7:30	Wed,Sat	14 Days	2 Days	Crash Consistent
<input type="checkbox"/> Monthly	--	--	--	--	--
<input type="checkbox"/> Quarterly	--	--	--	--	--

Figure 12 - Snapshot schedules

The graphic depicted above is a view of snapshot schedules for an individual VM. Note that this VM is using the settings from a TGC service group.

Clicking the “Edit” button enables the ability to alter the snapshot schedules for an individual VM or a TGC service group.

Figure 13 - Edit protection – snapshots

The graphic depicted above is a view of the daily snapshot schedule for a TGC service group. Editing the schedule provides the user with the ability to define the local retention period for snapshots created with the schedule, as well as the ability to define the remote retention period for snapshots replicated asynchronously to one or more Tintri systems. The snapshot consistency parameter can also be selected to create crash consistent or VM consistent snapshots. The days of the week on which the schedule will execute can also be specified. Additionally, the time of day that the schedule executes can also be specified.

New schedules can be enabled after clicking the “Edit” button by clicking the check-box to the left of a given schedule. Likewise, a given active schedule can be disabled by clearing the check-box.



Figure 14 - Enabled schedule

The graphic depicted above is a view of a schedule that has been enabled as the result of clicking the check-box to the left of the schedule.

Replication

Replication destinations can be created for individual VMs as well as for TGC service groups. The creation and administration of replication destinations for individual VMs and TGC service groups is essentially the same. Viewing the “Settings” tab of an individual VM or TGC service group will display all replication destinations that have been configured. Scrolling down below the snapshot schedule section, replications destinations can be viewed.

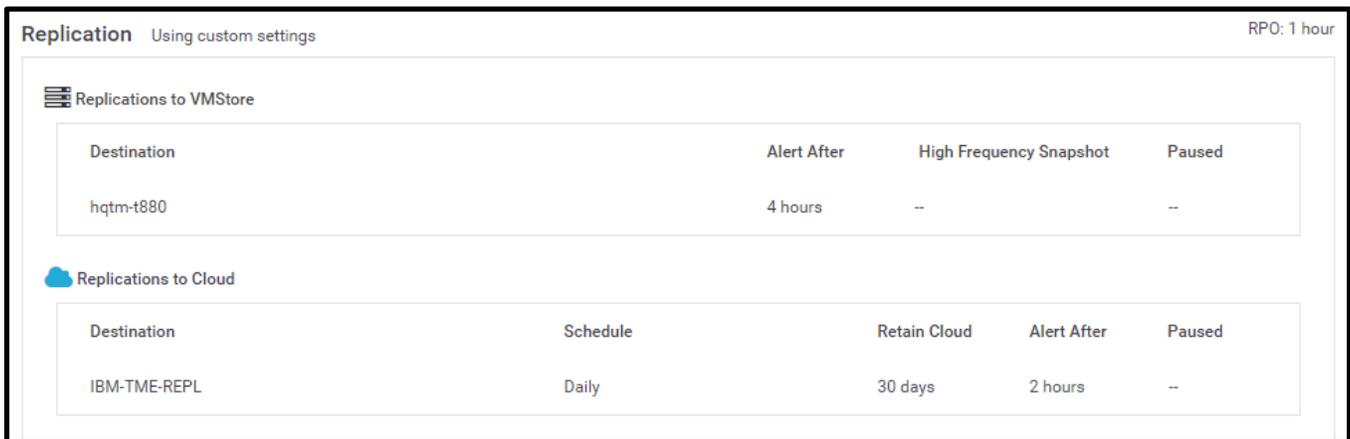


Figure 15 - Replication destinations

The graphic depicted above is a view of replication destinations for a TGC service group. Replication destinations to other Tintri systems as well as cloud destinations are shown.

Editing the protection settings of an individual VM or a TGC service group enables the ability to add, delete, or alter a replication destination. To add a new cloud replication destination, click the “+ Add Cloud Destination” button.



Figure 16 - Add Cloud Destination

After clicking the “+ Add Cloud Destination” button, the “Replicate to Cloud” section of the display allows the user to specify the parameters required to add a previously configured cloud replication destination. The “Destination” pull-down menu allows the user to select a cloud replication destination. The “Replication Schedule” pull-down menu allows the user to select a snapshot schedule to use for the creation of snapshots that will be replicated to the selected cloud replication destination. The “Alert RPO Threshold” field allows the user to specify the numbers of hours after which an alert will be generated if a given snapshot has not been successfully replicated to the selected cloud destination. The “Retain

Cloud” field allows the user to specify the retention period for snapshots replicated to the selected cloud replication destination.

Note that a given cloud replication destination can only be configured once for usage within a service group or individual VM. As a result, a single snapshot schedule is selectable for a given cloud replication destination.

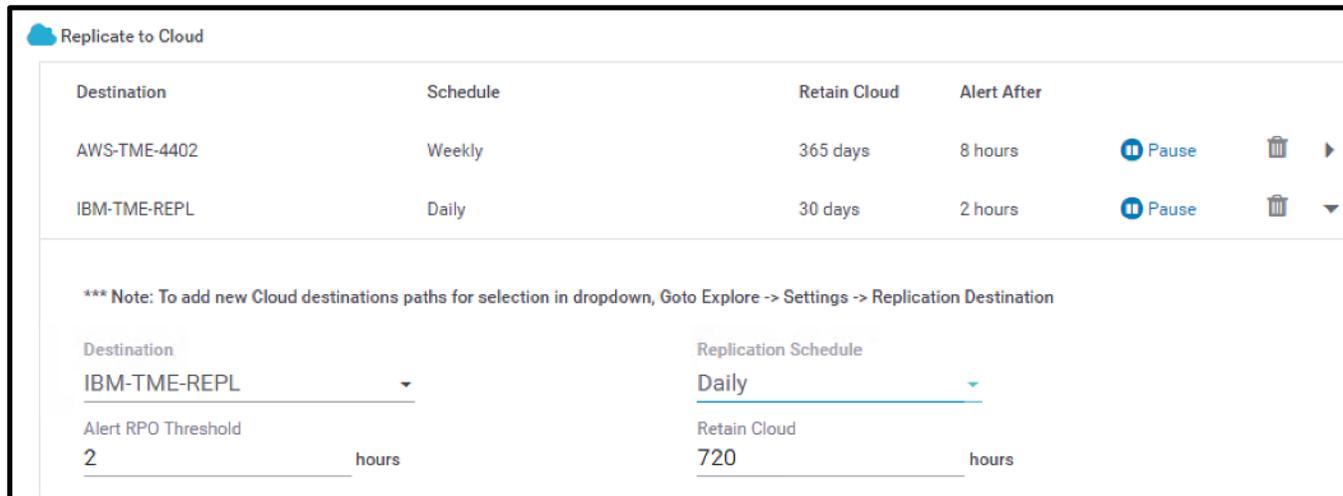


Figure 17 - Cloud destination parameters

The graphic above depicts an on-premises IBM COS replication destination named “IBM-TME-REPL”. Snapshots created by the “Daily” snapshot schedule will be replicated to the selected destination. The “Alert RPO Threshold” parameter has been set to 2 hours. The “Retain Cloud” parameter has been set to 720 hours, which will retain snapshots replicated to the selected destination for 30 days.

Viewing Snapshots

Snapshots can be viewed from within TGC as well as on a Tintri system. From within TGC, snapshots are viewed at a VM level. The user can navigate to given VM, and then click the “Snapshots” tab to view a list of existing snapshots.

Created	Description	Location	Expires
Oct 25 2017, 7:30:06 am	weekly scheduled	AWS-TME-4402	Oct 25 2018, 7:30:06 am
Oct 21 2017, 7:30:06 am	weekly scheduled	hqtm-t5040.ttucs.tm.ti...	Nov 4 2017, 7:30:06 am
Oct 21 2017, 7:30:06 am	weekly scheduled	AWS-TME-4402	Oct 21 2018, 7:30:06 am
Oct 18 2017, 7:30:01 am	weekly scheduled	AWS-TME-4402	Oct 18 2018, 7:30:01 am
Oct 14 2017, 7:30:01 am	weekly scheduled	AWS-TME-4402	Oct 14 2018, 7:30:01 am
Oct 11 2017, 6:30:05 pm	weekly scheduled	AWS-TME-4402	Oct 11 2018, 6:30:05 pm

Figure 18 – Snapshots as viewed from within Tintri Global Center

The graphic depicted above contains a list of snapshots. The displayed columns include the date and time at which a given snapshot was created, the name of the schedule that was used to create the snapshot, the location of the snapshot, and the expiration date of the snapshot.

On a Tintri system, snapshots can also be viewed. After clicking the “Search VM” button, the user can right click on a VM and select “View snapshots” from the pop-up menu.

Source VM	Created	Description	Location
Group2VM	2017 Nov 1, 7:30AM	weekly scheduled	local
Group2VM	2017 Nov 1, 7:30AM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 28, 7:30AM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 28, 7:30AM	weekly scheduled	local
Group2VM	2017 Oct 25, 7:30AM	weekly scheduled	local
Group2VM	2017 Oct 25, 7:30AM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 21, 7:30AM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 21, 7:30AM	weekly scheduled	local
Group2VM	2017 Oct 18, 7:30AM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 14, 7:30AM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 11, 6:30PM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 7, 6:30PM	weekly scheduled	AWS-TME-4402
Group2VM	2017 Oct 4, 6:30PM	weekly scheduled	AWS-TME-4402

Figure 19 – Snapshots as viewed on a Tintri system

Cloud Snapshot Download

Before a snapshot residing in a cloud replication destination can be used for recovery, it must be downloaded from the cloud repository to a Tintri system. Note that a snapshot can be downloaded to the Tintri system from which it originated, or to another Tintri system managed by TGC.

To download a snapshot, select the snapshot to be downloaded by clicking it. Then click the “Download” button.

The screenshot shows a table of snapshots with columns: Created, Description, Location, and Expires. A red box highlights the 'Download' button and a specific snapshot row. An arrow labeled '1' points to the selected snapshot row, and an arrow labeled '2' points to the 'Download' button.

Created	Description	Location	Expires
Oct 21 2017, 7:30:06 am	weekly scheduled	hqtm-t5040.ttucs.tm.ti...	Oct 21 2018, 7:30:06 am
Oct 21 2017, 7:30:06 am	weekly scheduled	AWS-TME-4402	Oct 21 2018, 7:30:06 am
Oct 18 2017, 7:30:01 am	weekly scheduled	AWS-TME-4402	Oct 18 2018, 7:30:01 am

Figure 20 - Select and download a snapshot

After the “Download” button is clicked, the “Download Snapshot” dialog window will appear. The “Download Snapshot” dialog window allows the user to select a destination Tintri system using the “Download To” pull-down menu. The Tintri system selected can be the same Tintri system that originally replicated the snapshot to the cloud destination, or a different Tintri System. The “Local Retention” field needs to be populated by the user with the desired retention period for the snapshot once it has been downloaded.

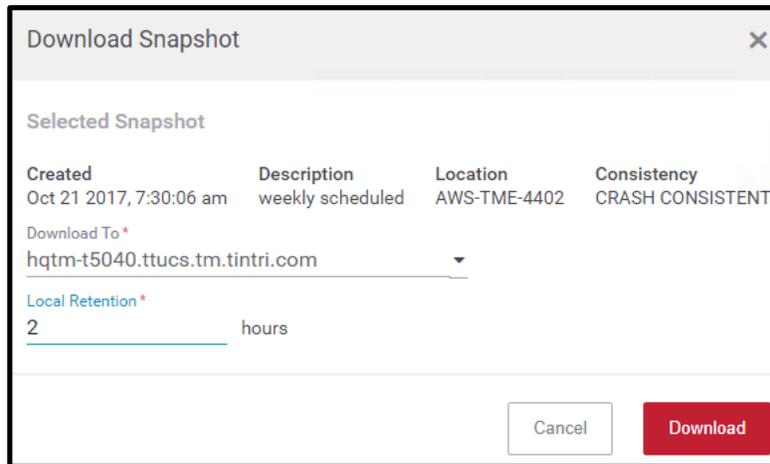


Figure 21 - Download snapshot

After the “Download” button is clicked, the snapshot being downloaded will be displayed with a download progress bar.



Figure 22 - Download snapshot progress

Note that when downloading a snapshot to a different Tintri system, the selected Tintri system must have a Cloud Connector license installed. Also note that the cloud replication destination path will be automatically pushed to the target Tintri system. If pushing the cloud replication destination path would cause the target Tintri system to exceed a total of 4 cloud replication destination paths, the operation will fail.

When the download process completes, the snapshot will be visible on the Tintri system to which it was downloaded. The location of the snapshot will be listed as “local”. Note that the retention period specified at the time the snapshot was downloaded will be reflected in the “Expires” column on the Tintri System. The user may need to enable display of the “Expires” column by selecting it from the pop-up menu after right clicking any of the column header fields.

Note that when downloading a snapshot to a different Tintri system than the one on which the snapshot originated, the VM may appear as a synthetic VM. In this context, the term “synthetic” refers to one or more VM snapshots without having a corresponding VM that has been added to hypervisor inventory.

VM	Replicated to	Replicated from
CriticalDemoApp1	hqtm-t850.ttucs.tm	192.168.200.249
CriticalDemoApp2	hqtm-t850.ttucs.tm	192.168.200.249
CriticalDemo	hqtm-t850.ttucs.tm	192.168.200.249
Group2VM	Oct 07	

Figure 23 - Synthetic virtual machine

The graphic depicted above is a synthetic VM that was created as the result of downloading a snapshot to a different Tintri system than the one on which the snapshot originated. Note that the synthetic VM will appear within the Tintri system user interface subsequent to completion of downloading a snapshot from a cloud repository. The appearance of the synthetic VM may take a short amount of time to occur.

Recovery Options

Two categories of recovery options are available dependent on the use case. The first category is recovery by means of the clone function. The clone function is available for use with a snapshot that has been downloaded to the Tintri system on which it originated or has been downloaded to a different Tintri system.

The second recovery category is recovery by means of the “Synchronize” function. The “Synchronize” function can be performed on a Tintri system where the actual VM that the snapshot was created from resides. This is usually the originating Tintri system, the Tintri system that replicated the snapshot to the cloud replication destination. The “Synchronize” function facilitates recovery of the entire VM, one or more virtual disks, as well as granular folder and file recovery.

Cloning

Cloning is the process by which one or more new virtual machines can be created using the virtual machine retained in a snapshot as the basis or starting point for the new virtual machine(s). Cloned virtual machines operate as independent virtual machines with their own identity.

Cloning can be performed after downloading a snapshot from a cloud repository.

Locate the downloaded snapshot on the Tintri system. This is accomplished by means of the Tintri system user interface. Right clicking the downloaded snapshot will display a pop-up menu and the “Clone” menu item can be selected.

Group2VM	2017 Oct 25, 7:30AM	2017 Nov 8, 6:30AM	weekly scheduled
Group2VM	2017 Oct 25, 7:30AM	, 7:30AM	weekly scheduled
Group2VM	2017 Oct 21, 7:30AM	5:30PM	weekly scheduled
Group2VM	2017 Oct 21, 7:30AM	, 7:30AM	weekly scheduled
Group2VM	2017 Oct 18, 7:30AM	2018 Oct 18, 7:30AM	weekly scheduled

Figure 24 - Clone a downloaded snapshot

After selecting “Clone” from the pop-up menu, the “Create new VMs from” dialog window will appear.

Create new VMs from Group2VM ✕

Clone name:

Datstore:

Host / Cluster:

Customization:

Count:

Figure 25 - Create new VMs dialog window

After clicking the “Clone” button, the cloned VM will be added to Hyper-V or vCenter inventory, dependent on the hypervisor type. The VM will not be powered on automatically.

For comprehensive information about cloning, please reference the “Data Protection Overview and Best Practices with Tintri VMstore and Tintri Global Center” document available on Tintri.com at:

<https://www.tintri.com/resources/white-paper>

Restoring

After a cloud resident snapshot has been downloaded to the Tintri system where the VM from which it was created resides, the snapshot can be used as the recovery point for a variety of recovery scenarios. Tintri restore functionality is trademarked as “SyncVM”, which can also be used to synchronize disks for test and development purposes. This subsection focuses on the use of this technology for the purpose of data recovery. “SyncVM” is a licensed feature that requires enablement through the use of a license key.

The “Synchronize” menu item appears when right-clicking a VM from within the Tintri system user interface. The available menu choices are “Restore VM/files” and “Refresh virtual disks”.

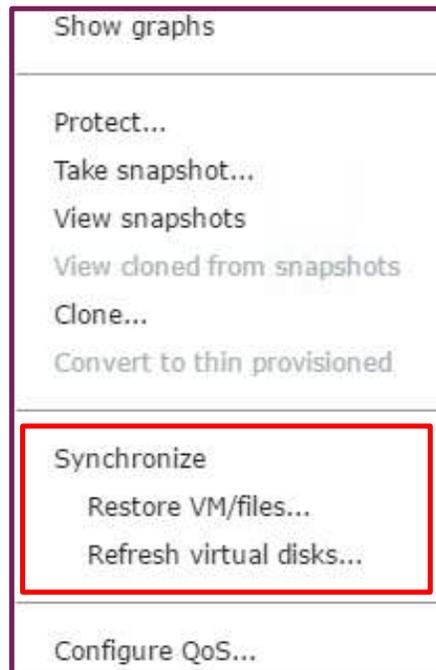


Figure 26 - Synchronize menu

To recover an entire VM to the state it was in at the time the snapshot was created, or to recover at a granular folder or file level, select the “Restore VM/files” menu item. To recover one or more virtual disks to the state they were at the time the snapshot was taken, select the “Refresh virtual disks” menu item.

Full virtual machine recovery is accomplished with a simple process where the snapshot to restore from is selected by means of a pull down menu, and clicking the “Restore” button initiates the restore process.

Virtual disk recovery, where one or more virtual machine disks are restored, is also accomplished by means of a simple process where the snapshot to restore from and specific disk or disks to recover are selected.

The intuitive workflow employed for granular folder and file level recovery adds the appropriate snapshot based virtual disk image as a new drive to the specified virtual machine. Within the guest operating system the new drive is manually brought online and then configured within the operating system. Drag and drop folder and file level recovery, or command line interface recovery is then easily accomplished. By default, any added drives are automatically disconnected after 48 hours.

For comprehensive information about restoring, please reference the “Data Protection Overview and Best Practices with Tintri VMstore and Tintri Global Center” document available on Tintri.com at:

<https://www.tintri.com/resources/white-paper>

Conclusion

The Tintri Cloud Connector adds significant functionality to the existing suite of data protection and disaster recovery features available with Tintri systems. When combined with Tintri asynchronous replication, the Cloud Connector enables the creation of advanced data protection and disaster recovery schemas. Support for AWS S3 and on-premises IBM COS as cloud replication destinations may enable use cases that reduce the overall cost of a data protection and disaster recovery solution. VM snapshots that require longer term retention can be cost-effectively replicated to a cloud destination while being retained for a shorter term duration on a Tintri system. The direct replication of VM snapshots to cloud destinations eliminates any requirements for gateways or caching devices, and may reduce dependency on third party backup applications.

Additional benefits include:

- Snapshot data is encrypted at the time of replication to to a cloud destination, and remains encrypted for the duration of its retention in a cloud repository.
- Tintri VM snapshots are compressed and deduplicated, which reduces the cost of storing them in the cloud.
- After an initial snapshot is created, subsequent snapshots capture only incremental block level changes, increasing efficiency.
- The ability to download a cloud resident snapshot to any Tintri system managed by TGC adds recovery flexibility that may enhance the ability to conduct recovery testing.
- Snapshots downloads result in the creation of full recovery points that are automatically synthesized when they are downloaded from a cloud repository.

© 2017 Tintri, Inc. All rights reserved. Tintri, Tintri VMstore, Tintri Global Center, ReplicateVM, SecureVM, and SyncVM are trademarks of Tintri, Inc., and may be registered in the U.S. Patent and Trademark Office and in other jurisdictions. All other marks appearing in this publication are the property of their respective owners.

Tintri believes the information in this document is accurate as of its publication date. The information in this publication is provided as is and is subject to change without notice. Tintri makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims any implied warranties of merchantability or fitness for a particular purpose.



303 Ravendale Drive
Mountain View CA 94043
+1 650.810.8200
info@tintri.com