Introduction to DRS for Cloud

ⓘ **Description**

Provides an overview of the DRS for Cloud feature. To identify if this feature is currently available in a particular location, see How do I Identify Hardware Specifications and Capabilities Available in a Data Center Location.

DRS for Cloud Menu Availability Requires Role

Unlike most CloudControl functions, only Primary Administrators or sub-administrators with the DRS role will even see the DRS for Cloud menu that is needed to view and manage the DRS for Cloud functions. Ensure your users have the proper roles assigned - see How to Manage the Primary Administrator User and other Sub-Administrators as the Primary Administrator

Note that even when the roles are assigned, the menu appears *only where the feature is available.*

What is DRS for Cloud and What is the Use Case?

DRS (Disaster Recovery Services) for Cloud is a feature that enables real-time replication of the Cloud Servers and their associated local disk content between two data center locations in the same Geographic Region. A group of these servers can then be "recovered" with their disks set as of a specific time in the data center to which the disks are replicated. Note the replication only includes local disks - the RAM state of the server and any attached CD-ROM (.iso) or Floppy Fisk (.flp) are not replicated to the target. To deliver this capability, the feature uses EMC's RecoverPoint for Virtual Machines software in conjunction with a series of orchestrations by the CloudControl software to address unique aspects of the Cloud infrastructure.

DRS for Cloud differs from the **Cloud Server Snapshot** feature (see Introduction to Cloud Server Snapshots) in two key aspects:

1. DRS for Cloud continuously replicates all changes to the local disk storage of the Source Server to the Target location, providing much more granular "DRS snapshots" from which a server state can be restored.
   a. The Cloud Server Snapshot takes a snapshot of a single Cloud Server at a specific point in time and allows you to restore only from those specific times when a snapshot is taken, which may be daily or weekly depending on the snapshot plan.
2. DRS for Cloud keeps its local storage snapshot state identical for all servers in the Consistency Group. When you restore from a DRS snapshot, the disk state of all of the servers will represent the disk state of all of the servers at a specific point in time.
   a. Cloud Server Snapshots are server-based. If multiple servers are enabled with Cloud Server Snapshots, the timing of the "snapshots" will be different, even if they have the same snapshot "window". Therefore, if an application involves multiple servers, it will not be possible to restore the servers to their state at the exact same point in time. The timing may be close, but it will not be identical.

There are two main use cases for DRS for Service:

- **MCP Location Migrations**
  - By replicating the server state between locations, the feature can facilitate the migration of the server from one data center to another. This feature is particularly useful in migrating Cloud Servers from MCP 1.0 Data Center locations to MCP 2.0 Data Center locations. For further details, see MCP 1.0 to 2.0 Self-Service Migration
- **A component of a Disaster Recovery plan**
  - The feature does not provide a complete turnkey Disaster Recovery solution in itself as its functionality is limited to replicating the disk state (and some server configuration elements) between two sites. However, its ability to continually replicate and provide "point in time" recovery across multiple Cloud Servers can provide a key component in an overall Disaster Recovery solution.

DRS for Cloud Terminology

The feature uses some unique terminology which should be reviewed before reading through the rest of this article:

- **Source Location** - The data center location of the servers being "protected" using the DRS for Cloud feature. In DR scenarios, this role is normally held by the "production" servers.
  - The servers being protected in the Source location are referred to as **Source Servers.** These servers will run like any other Cloud Server except that their disks will be continually replicated to the Target Servers in the Target Location.

- **Target Location** - The data center location to which you wish to replicate the local disk content of the protected servers. That replication is done to a set of **Target Servers** deployed in the Target Location that have the same disk and network configuration as the Source Servers to which they are paired. In DR scenarios, this role is normally held by the "DR" or "failover" servers.
    - **Target Servers** do not run like normal Cloud Servers. While normal DRS replication is in effect, they are essentially "ghost" servers that run as 1 Standard vCPU, 1 GB RAM servers and accept replication from their corresponding Source Server. They are enabled to act as normal Cloud Servers only when you want to bring the Target Servers live as of a specific point in time. More details on this behavior are described in the sections below.
- **Consistency Group** - A set of up to ten Source Servers that are grouped together and designed to be "failed over" to a corresponding set of Target Servers so that their disks can be configured to represent a specific moment in historical time. The idea is that should there be a problem with the Source location acting as production, users can bring up the Target Servers to match the exact same local disk state as the Source Servers as of a specific point in time in the past.
- **DRS Journal** - In order to perform the DRS replication, the system requires a "journal" to which it replicates all disk writes made to the Source Server. Each disk write to a Source Server is briefly copied to a DRS Journal in the Source location, then transferred across to the corresponding DRS Journal in the Target location across the private MCP interconnect. When a user chooses to bring up the Target Servers with the disks configured as of a point in time, the system uses this journal to reconstruct the Target Server's local disks as of a specific point in time in the life of the Source Servers.
- **DRS Snapshot** - As the system replicates disk writes to the DRS Journals, it will create "DRS snapshots" which represent moments in time to which the Target Servers can be restored. As the system continually replicates writes to the Source Servers, it will produce DRS snapshots that represent points in time in the life of the Source Servers to which the Target Servers can be restored. Once replication is in place, the system makes available 1,014 DRS snapshots across the protection window at any given time.
- **Protection Window** - The Protection Window represents the time of the oldest DRS snapshot to which the Target Servers can currently be restored. This represents (i.e. how far back in time the system can be restored. As the system replicates data from the Source location to the Target location, it begins filling up the space allocated to the DRS Journal. When the system has used 80% of the DRS Journal capacity, it begins dropping older data to make room for newer data. At this time, the Protection Window is established.
    - NOTE: The system reserves 20% of the DRS Journal space for handling changes to the Target Server when a DRS snapshot is applied to a Target which is why replicated data begins aging out at 80% capacity.

A diagram with an overview of the overall solution is shown below. The rest of this article will explain these concepts in more detail.

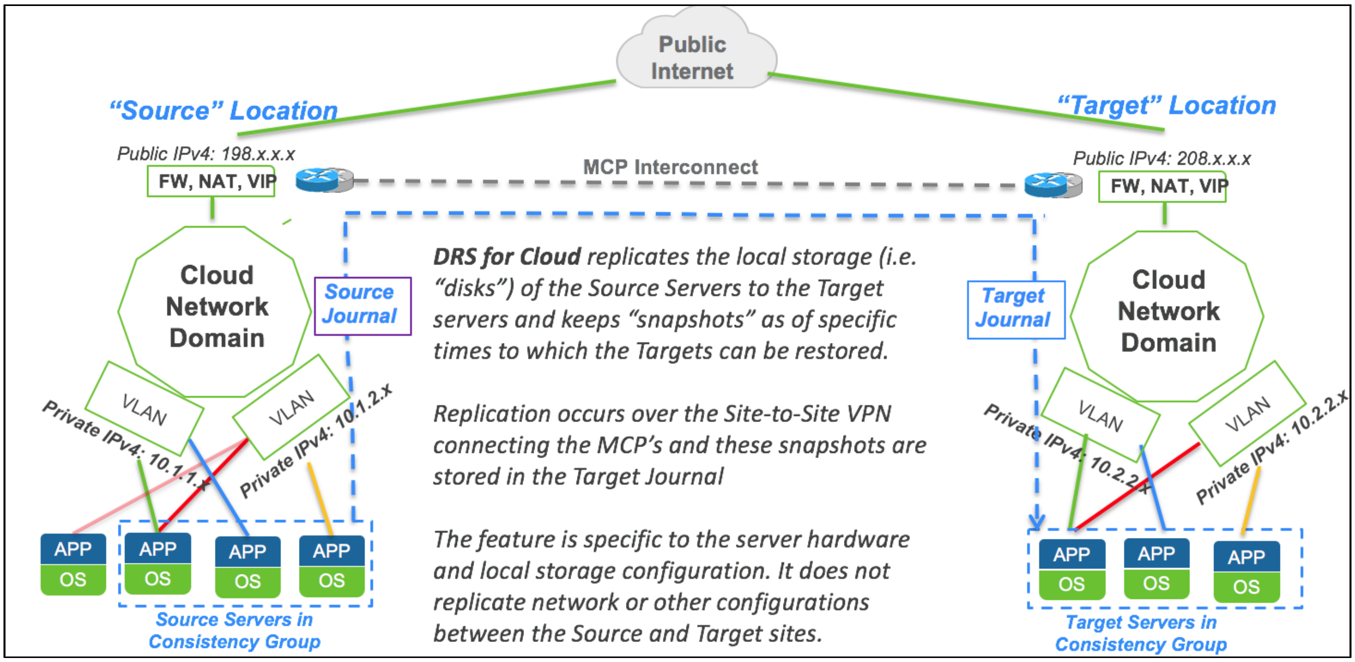

__DRS for Cloud__ replicates the local storage (i.e. "disks") of the Source Servers to the Target servers and keeps "snapshots" as of specific times to which the Targets can be restored.

Replication occurs over the Site-to-Site VPN connecting the MCP's and these snapshots are stored in the Target Journal

The feature is specific to the server hardware and local storage configuration. It does not replicate network or other configurations between the Source and Target sites.

DRS for Cloud Feature Availability

DRS for Cloud is only available in a subset of data center locations. You can identify the availability of the DRS for Cloud service in a Source location and the available Target locations to which a given Consistency Group can be replicated as described in How do I Identify Hardware Specifications and Capabilities Available in a Data Center Location. Note that DRS for Cloud can only replicate between locations in the same Geographic Region. For more information on Geographic Regions, see Introduction to Geographic Regions

User Permissions for DRS for Cloud

DRS for Cloud functions are available only to Primary Administrators and users with a "DRS" role. Sub-Administrators can also be assigned DRS privileges as described in How to Manage the Primary Administrator User and other Sub-Administrators as the Primary Administrator. The exception is DRS-related reports which are also available to users with the Reports role.

Understanding Consistency Groups

Consistency Groups represent the "pairing" of Source Servers and Target Servers for DRS replication purposes. Each Consistency Group consists of 1-10 pairs of Source Servers (i.e. the servers you want to replicate) and corresponding Target Servers (i.e. the servers to which a given Source Server will replicate). Each Source->Target Pair must meet the following criteria:

- All Source Servers in a given Consistency Group must be located in the same data center location (the "Source location")
- All Target Servers in a given Consistency Group must be located in the same data center location (the "Target location")
- The Source Data Center location must have DRS for Cloud enabled with the Target Data Center location listed as a target for that location. Replication is currently available only from a Source location to a single Target location; you cannot replicate to multiple Target locations.
- Each Source Server must match the corresponding Target Server in terms of NIC and disk layout. Specifically:
    - The Target Server must have the same number of NICs as the Source Server to which it is paired
    - The Target Server must have the same SCSI controllers (including SCSI Adapter), IDE controllers, and SATA controllers.
    - The Target Server must have the same size and placement of local storage disks in their appropriate controllers
        - For example, if the Source Server has a 10 GB disk in SCSI 0 of SCSI Controller 0 and another 100 GB disk in SATA Controller 1 in slot 15, the Target Server must also have the same sized disks in the same controllers in the same position.
        - The Target Server's "disk speed" does not need to match the Source Server. However, note that you cannot change the disk speed of a Target Server's disk once the Consistency Group is created, so ensure you have specified the disk speed you wish to have for "Preview" mode as described below.
        - *NOTE: The OS, CPU, CPU Speed, Cores Per Socket and RAM of the Target Server can also be different than the Source Server when they are added to the Consistency Group. However, these settings are overwritten as part of the DRS service when the system replicates the Source Server hardware configuration to the Target Servers.*
        - A matching CD-ROM device and/or Floppy Disk controller is not required. Remember these elements are not replicated - any attached ISO and FLP files are not replicated so the state of the Target Server in terms of these devices and files will remain unchanged after Failover.

At a high level, the process for creating a Consistency Group involves the following steps:

1. **Identify the Servers in the Source Location To Be Added to a Consistency Group** - In choosing which Source Servers to put together in a Consistency Group, users should consider what servers they would want to restore to the same historical point in time in the event of a disaster. If Source Servers are in different Consistency Groups, the available historical DRS snapshots to which they can be restored will likely represent different historical points in time.
    a. *Example: If after a disaster, your solution requires the application server and DB server to be "in sync" in terms of the data they hold (i.e. both be restored to the same point in time), you'd want them in the same Consistency Group.*
2. **Create Target Servers in the Target Location For Each Source Server** - For each Source Server you've identified, you need to create a corresponding Cloud Server that will act as the Target for the Source Server. Each Target Server needs to meet the criteria listed above - the key point being the same number of NICs with the same local storage layout in terms of SCSI ID and disk volume size. This Target Server will have its local disks overwritten as part of the replication process, so there is no need to set up any applications or information on the Target Server. There is a function which does this automatically, see How to Create a DRS Target Server
3. **Mark the Target Servers as "DRS Eligible"** - In order to prevent users from accidentally overwriting a server by placing the "wrong" Target Server in a Consistency Group, the system requires you to explicitly flag the Target Servers you created in the previous step as being "eligible" for DRS replication. For more details, see How to Manage DRS Eligibility of a Cloud Server.
4. **Create the Consistency Group** - The details on this process are available at How to Create a DRS Consistency Group. In addition to providing the pairing of Source/Target Servers you identified above, the system will ask you for the Name, Description, and the DRS Journal Size. For advice on how to size the DRS Journal, see the following section.

Sizing of DRS Journal for a Consistency Group

When you create a Consistency Group, you define the Initial Size of the DRS Journal. The initial size needs to take into account both the Consistency Group's immediate needs and its longer-term size needs. The latter is important because the initial size determines the size to which the Consistency Group can be expanded in the future. Those rules are:

1. **If the Initial DRS Journal Size <= 300 GB**, the maximum size to which the DRS Journal can be expanded is the lesser of:
    a. 1,000 GB
    b. The **Maximum DRS Journal Size** in the Source location
    c. The **Maximum DRS Journal Size** in the Target location
2. **If the Initial DRS Journal Size >= 400 GB**, the maximum size the DRS Journal can be expanded is the lesser of:
    a. Three times the **Initial Journal Size**
    b. The **Maximum DRS Journal Size** in the Source location
    c. The **Maximum DRS Journal Size** in the Target location

For instructions on how to expand the Initial Journal size, see How to Expand Journaling Space for an Existing Consistency Group. Note it is not possible to shrink the size of a DRS Journal. Moving to a lower-sized journal requires deleting and recreating the Consistency Group, thereby temporarily losing the Protection Window and existing historical DRS snapshots.

The DRS Journal size will determine both the Protection Window and DRS Snapshot Granularity for the Consistency Group:

1. **Protection Window** - As disk writes are made to the Source Server local disks, those changes are replicated to the DRS Journal. Once the journal is 80% full, the system begins deleting older changes and therefore the oldest "DRS snapshot" from which you can restore will be removed. The larger your DRS Journal, the more historical disk changes can be stored and the further back your Protection Window will go.
    a. Example: Assuming a constant change rate, a journal that is twice the size will be able to hold twice as many changes, therefore allowing you to go back to a point twice as old. For example, assuming a constant change rate, if a 5,000 GB DRS Journal gives a protection window of 5 days, a 10,000 GB DRS Journal should give you a 10-day protection window.
2. **DRS Snapshot Granularity** - Regardless of the size of the DRS Journal, the system makes available 1,014 DRS snapshots. The system tries to evenly distribute DRS snapshots in terms of approximate size (i.e. the amount of change) across the entire journal. Because the 1,014 DRS snapshot limit is constant, the larger your DRS Journal size, the larger the granularity will be in terms of the size of data between DRS snapshots.
    a. Example: Using the above example, a 5,000 GB DRS journal will give 4,000 GB used for DRS snapshots, and when full, the system should space these DRS snapshots out so each represents about 4 GB of change (4,000 GB / 1,014 = approx. 4 GB). Likewise, a 10,000 GB DRS journal with 8,000 GB DRS snapshot space will end up with 8 GB granularity.

Given this behavior, users should think through both their desired Protection Window and DRS Snapshot Granularity when sizing a Consistency Group. The challenge is that these two axes are related and move in opposite directions from each other. If you want a larger Protection Window (i.e. further back in change size and time), you need a larger journal size - but that also results in a greater amount of time between each DRS snapshot because the system will "space out" DRS snapshots into roughly 1000 equally-sized DRS snapshots in terms of change spread across a larger journal size. Similarly, if you want more granularity in your list of DRS snapshots, you need a smaller journal size - but that reduces the Protection Window.

The third key factor is the **Data Change Rate** on the Source Servers. The more the local disk data changes, the faster a journal will fill up and churn out older data. Therefore, Consistency Groups containing Cloud Servers with high change rates will need larger DRS Journal sizes to maintain the desired protection window than servers with low change rates. Since change rates occur over time, the "time" represented by both Protection Window and DRS Snapshot Granularity is based on the change rate.

- *Example: Using the example above, a 10,000 GB DRS journal will give 8,000 GB of journal space used for DRS snapshots. When full, the system should space these DRS snapshots out so each represents about 8 GB of change (i.e. 8,000 GB / 1014 = approx. 8 GB). If the system experiences a constant change rate of 8 GB per minute, that means the Protection Window will be 1,000 minutes (8,000 GB of journal for DRS snapshots / 8 GB per minute) and the granularity of DRS snapshots available within that Protection window will be approximately one every minute, because the DRS snapshots are approximately 8 GB apart. However, a constant change rate of 16 GB per minute would give you a Protection Window of only 500 minutes (8,000 GB / 16 GB per minute) but granularity of every 30 seconds (since the same 8 GB gap between DRS snapshots = 30 seconds).*

Of course, in most production systems the change rate is not constant. With non-constant change rates, the actual times represented by DRS snapshots of equal sizes will vary. When change rates are high, DRS snapshots representing equal sizes will be closer in points of actual time than when change rates are lower. Therefore, users should evaluate their average change rate for an overall estimate, but realize that the time granularity of DRS snapshots will vary with the change rate. This effect will be particularly noticeable if the Source Servers experience highly variable change rates. For example, if the Source Servers run a monthly batch job that generates a high rate of change, the Protection Window may shrink after the batch job and the 1,000 DRS snapshots available within that window may have more granularity in terms of time for the period in which the batch job ran than during the periods in which the system was comparatively static. Once those batch job changes age out of the journal, the reverse occurs.

One effect of sub-optimal Journal size not well matched to the average change rate will be difficulty in selecting a desired DRS snapshot for Preview due to the snapshot list variability.

Therefore, it is recommended that users pay attention to how the journal behaves after the Consistency Group is created by periodically reviewing the DRS snapshots as described in How to List DRS Snapshots for a Consistency Group.

How Consistency Groups Behave Once Created

Once a Consistency Group is created, the system will go through a series of steps:

1. The system will prevent server actions which change the network or disk configuration from being taken against Source Servers as long as they are members of the Consistency Group. Users should avoid locally shutting down a Source Server for extended amounts of time since such action interferes with replication for the entire Consistency Group. For more details, see What happens to a Source Server when it is added to a Consistency Group?
2. Except for Console access, the system prevents almost all server-level functions against Target Servers as long as they are members of the Consistency Group. This is necessary as the Target Servers are designed to replicate the state of the Source Servers and require a static configuration controlled by the system. For more details, see What happens to a Target Server when it is added to a Consistency Group?
3. The system will create the DRS Journals in both Source and Target locations
4. The system will "shrink" the Target Servers down to a 1 vCPU / 1 GB RAM configuration.
    a. Note these Target Servers remain in a running state and continue to incur usage as 1 vCPU / 1 GB RAM running servers with whatever storage and disk speed is associated with them.
    b. Although the servers will show as being in a running state, accessing the server via remote console will show they are actually in BIOS mode.
5. The system begins an initial replication of the Source Server local disks to the Target Server disks. This will destroy all local disk data on the Target Servers.
6. The system will also remove any Priced Software labels from the Target Server as any such software is wiped out by the disk process.
7. The system will then enter "DRS mode" where it replicates changes to the Source Servers to the DRS Journal and creates historical DRS snapshots as described above.
    a. This replication should normally occur roughly in "near real-time" (i.e. close to synchronous) as long as there is bandwidth available between the Source and Target locations.
    b. If connectivity between the Source and Target locations is interrupted, the system stores changes on the Source location's DRS Journal and will replicate changes to the Target location asynchronously when connectivity is restored. However, this means the replication is effectively asynchronous in this scenario.
8. The system remains in this state until the Consistency Group is deleted or until Preview or Failover actions described below are initiated.

Identifying DRS Snapshots Associated with a Consistency Group

Once replication has begun, the system will begin creating "DRS snapshots" representing the state of all of the Source Servers' local disks as of a specific historical point in time. Users can identify the current Protection Window and the DRS snapshots currently available as described in How to List DRS Snapshots for a Consistency Group. Note this list is dynamic and will change over time.

Preview of a Consistency Group

At any time after the first DRS snapshot was created, users can "Preview" a Consistency Group. Preview brings all of the Target Servers in the Consistency Group into a running state with the same hardware configuration as the Source Servers to which they are paired, but with the local disks set up as of a specific DRS snapshot. Users may wish to use Preview mode for a number of reasons:

- **Retrieve a file or configuration from a previous point in time in the life of the Source Servers** - For example, if someone accidentally deletes a file from a Source Server, Preview mode can be used to restore the Target Server to a DRS snapshot prior to the deletion so you can retrieve it. Users can then stop Preview mode and return to normal DRS replication mode. Note Preview brings up all Target Servers as part of the process.
- **Testing of DRS for Cloud replication as part of an overall Disaster Recovery test** - In this case, a user may want to verify the server replication and restoration as part of an overall Disaster Recovery testing scenario without committing to an actual failover that one would do in an actual disaster event.
- **Actual Disaster Scenario involving the Source Location** – Should the Source location experience a disaster necessitating failover to the Target Servers, Preview mode allows the user to ensure both that they have selected the correct DRS snapshot and that the DRS snapshot represents a valid state prior to committing to the failover. In this case, if the DRS snapshot does not represent the desired state, the user can stop the Preview and try another DRS snapshot. This step is important since Failover is an irrevocable action as described in the Failover section below.

When Preview mode is enabled, the user chooses the DRS snapshot to which the Target Servers should have their local disks synchronized and the following steps occur:

1. The system will "expand" the Target Servers up to the same vCPU, RAM, Cores Per Socket, and Operating System configuration as the current state of the Source Servers (i.e. not the state in which the hardware was in at the time of the DRS snapshot) with the following notes:
    a. CPU Speed for Target Servers is always Standard
    b. Any Priced Software labels on the Source Server are replicated to the Target Server and incur usage while the Target Server is in Preview mode
2. The system sets up the content of all of the Target Server local disks to match those of the Source Server as of the time of the Snapshot
3. Because only the disks are replicated, by default the IP configuration for each NIC within the Guest OS is replicated - see ***"NIC IP Configuration after Preview of a Consistency Group"*** *below for more details*
4. The Target Server is accessible via Console, but changes to the Target Servers are still not allowed in Preview mode since the system is still replicating disk changes from the Source Server
5. Changes to the Target Servers are written to the 20% of the DRS Journal in the Target location that's not used for DRS snapshot replication
    a. Once this space is filled, the system will shut down the Target Servers. Users should stop Preview to restore the Consistency Group to normal replication state if this occurs.
6. The system will remain in this state until Preview mode is stopped or Failover mode is initiated as described below.

For more details on initiating Preview, see:

- How to Start a DRS Preview for a Consistency Group
- How to Stop a DRS Preview for a Consistency Group

Important Note

It is important to note that if you attempt to start DRS Preview with a DRS Snapshot ID that no longer exists, the Start Preview operation will fail and the error will be noted in the Administrator Logs report. The entry in the Administrator Logs report will look like this:
"Snapshot id 163135 no longer exists in the DRS infrastructure. Try again with another Snapshot Id."

If your Journal size is too small, or if you have a high rate of change on a Server, you may encounter this scenario

NIC IP Configuration (Re-IP) after Preview of a Consistency Group

When a DRS Consistency Group Preview is initiated, the system sets the content of all of the Target Server local disks to those of the Source Server as of the time of the Snapshot. However, remember that DRS for Cloud only replicates Cloud Server hardware configuration and local disk contents - not any of the other elements of the Target location's network configuration.

Since the underlying disks in Target Servers are the same as Source Servers, inside the Operating System each Target Server NIC will inherit the same IPv4 and IPv6 addressing as the Source Server in that pair. This is often undesirable because:

- In the case of IPv4 addresses, the target server's NIC may be connected to VLANs with IP address ranges than the source environment, meaning the Target Server NIC's IPv4 addresses will not work as intended.
- In the case of IPv6 addresses, the replicated IPv6 will likely never work on the Target VLAN since IPv6 addresses ranges on VLANs are globally unique.

Given this situation, changes may be required in both the Guest Operating System and/or Network configuration to establish working network connectivity and application functionality.

To help address this limitation, the system offers an optional Re-IP address feature which allows users to define NIC IPv4 configurations to be applied to Target Servers after the replication portion of the Preview is complete. These Re-IP configurations can be optional set up for each Server Pair in a Consistency Group - the function does not require that it be used across all servers. However, when the feature is used, the user must define the Re-IP address approach for all NIC's on the server and must also adjust DNS and a few other guest-level settings along with the IPv4 addresses.

For details on this function, see:

- How to Set Re-IP Address Rules for a Target Server in a DRS Server Pair
- How to Set Re-IP Address Rules for a Target Server in a DRS Server Pair using REST API

Re-IP Supported OS Versions

Re-IP is supported on the following Operating Systems:
Microsoft Windows Server versions: 8, 10, 2008, 2008 R2, 2012, and 2016
Red Hat Linux server versions: 6.5 and 7.1+
Ubuntu Studio: 15.10

In addition, there are some VM Tools limitations:

Centos 7.x must have VMware Tools version 10.1.0.57774 in order for Re-IP to be supported.
Linux SLES12 must have Open VM Tools version 9.4.0.25793 in order for Re-IP to be supported.
Servers running Open VM Tools versions lower than 9.10, Re-IP is not supported unless deployPkg has been manually installed. See VMware KB article 2075048 for detailed information on how to install deployPkg.

Re-IP Address Rules for DRS Target Servers

When a DRS Consistency Group Failover Preview is initiated, the system copies the VMDKs from the Source Server to the Target Server across all server pairs. Since the underlying disks in Target Servers are the same as Source Servers, each Target Server NIC will inherit the same IPv4 and IPv6 addresses as the Source Server in that pair. In the case of IPv4 addresses, this may be undesirable because the target Network Domain and VLAN architecture may not be the same as the source environment, in which case, the Target Server NIC's IPv4 addresses will not work as intended. IPv6 addresses should never work on the Target VLAN since IPv6 addresses are globally unique.

To address this limitation, the system includes a feature which allows users to define NIC configurations to be applied to the Target Servers after Preview is complete. It also allows adjustment of DNS and some other guest-level settings. For each Server Pair in a Consistency Group, the user can view and optionally change the NIC configuration for all Target Server NICs with the new NIC Configuration dialog available from the Server Pair action drop-down menu. The following settings can be changed from that screen for the Target Server:

1. **Hostname** - Text field, for the hostname to be applied to the Target Server. Required for Linux OS.
2. **Domain** - Text field for the Domain to be applied to the Target Server. Required for Linux OS.
3. **DNS Server** - Text field, a series of IPv4 or IPv6 addresses to be applied to the Target Server. Required for all OS types.
   a. On Linux, DNS is a server-level setting. On Windows, there is a DNS entry per NIC. DRS follows this convention, and while DNS Server Addresses have a single field for the entire server, DRS will apply them to each NIC for Windows servers.
4. **DNS Suffix Search Order** - Determines the DNS suffixes to attach to an unqualified single-label name before submission of a DNS query for that name.
   a. This entry can be blank.
5. **NetBIOS** - Selector, which defines whether NetBIOS is enabled or disabled on the Target Server.

The following settings can be changed for each NIC on the Target Server:

1. **Re-IPv4** - new IPv4 address that will be set for that NIC on the Target Server.
2. **Prefix** - new IPv4 Suffix that will be set for that NIC on the Target Server.
3. **Re-IPv4 Gateway** - new IPv4 Gateway address that will be set for that NIC on the Target Server.

To enable Re-IP Address Rules for a Server Pair, refer to How to Set Re-IP Address Rules for a Target Server in a DRS Server Pair.

When setting Re-IP Address rules for a Server pair, you can change the Server and NIC level Re-IP settings. Changing the Server level settings is self-explanatory. Changing NIC level settings requires further explanation. There are four settings for each NIC that can be changed. Or strictly speaking, three settings, but depending on the selected Re-IP Mode, you can change none, two or all three settings as follows:

1. If **Re-IP Mode** selected is "**Use CloudControl**" - the Re-IPv4, Prefix and Re-IPv4 Gateway values for that NIC will be retained as they are currently set in CloudControl for that Server. Those values will be the same as in the IPv4, Gateway and IPv4 Range columns for that NIC (visible just to the left of the Re-IP Mode selector in the same row).
2. If **Re-IP Mode** selected is "**Explicit in Range**" - the Re-IPv4 and Re-IPv4 Gateway values can be changed, with the following limitations - the Re-IPv4 value can be changed only to a value in Range defined by the IPv4 Range column for that NIC just to the left in the Re-IP Mode selector. Also, the Re-IPv4 Gateway must be in the same range as defined by the IPv4 Range column.
3. If **Re-IP Mode** selected is "**Explicit Out Of Range**" - you can change the Re-IPv4, Prefix and Re-IPv4 Gateway values can be changed, with the following limitations - while the Re-IPv4 value can be changed to any valid IPv4 address, and the Prefix can also be set to any value in the drop-down, the Re-IPv4 Gateway has to be in the same Range as defined by the Re-IPv4 and the Prefix value.

For instructions on how to delete Re-IP Address rules for a previously configured Target Server, see How to Delete Re-IP Address Rules for the Target Server in a DRS Server Pair

Failover of a Consistency Group

Once a user has initiated a Preview of a Consistency Group, they can choose to "failover" the Consistency Group. "Failover" will completely commit the current state of the Target Servers (i.e. including any change made while in Preview mode) and dissolve the Consistency Group so the users can manage the Target Servers as normal. It will also shut down the Source Servers. *Failover cannot be reversed so users should ensure the Preview state represents a valid state to which they want to commit the state of the Target Servers*.

When Failover mode is initiated, the following steps will occur:

1. The system will commit the current state of the Target Servers (including any changes to the disks made since Preview was initiated) to the Target Server's local disks
2. The system will briefly "shrink" the Source Servers as if they were to become Target Servers
3. The system will delete the Consistency Group entirely. As a result...
   a. All DRS Journal space will be deleted along with all DRS snapshots and historical change information
   b. The system will "re-inflate" the Source Servers back to their previous state, but will also attempt to shut them down gracefully

4. The system will designate the Source Servers as "DRS Eligible", making it easier for users to reconstruct the Consistency Group with reversed roles if desired, making what were the Target Servers the Source Servers (and vice versa)
5. Because the Consistency Group no longer exists, all server actions against the Source and Target Servers are now permitted

For more details on initiating Failover, see How to Initiate a Failover for a DRS Consistency Group

> **Actual Disaster Involving Source Location May Affect This Behavior**
>
> In the case of an actual disaster scenario that adversely affects the functionality of the Source location's infrastructure, the steps above that need access to the Source location's infrastructure may not be able to be completed. In this case, the Consistency Group may be left in a "Failed Delete" status and require manual intervention by Support personnel. Contact the Support team for assistance if you need to make adjustments to the Target Servers since those actions will still be blocked due to membership in a Consistency Group.

Deleting a Consistency Group

At any point in time except when a Consistency Group is in Preview mode, users can choose to delete the Consistency Group. This will also delete the associated DRS Journal space and all associated DRS snapshots. The specific steps that will occur are as follows:

1. The system will apply a recent DRS snapshot to the Target Server(s) local disks
    a. Due to DRS Snapshot Granularity and Change Rate, it is not always the most recent DRS snapshot, but it will generally represent a recent moment in time.
2. The system will "expand" the Target Servers as follows:
    a. If the Target Servers have previously undergone a Preview, the Target Servers will be "inflated" to the same vCPU, RAM, Cores Per Socket, and Operating System configuration as the state of the Source Servers when the most recent Preview was initiated. In addition, the CPU Speed for the Target Servers will be set to Standard.
    b. If the Target Servers have never undergone a Preview, the Target Servers will be "inflated" to the same vCPU, RAM, Cores Per Socket, and Operating System configuration as they had before they were added to the Consistency Group. In addition, the CPU Speed for the Target Servers will be set to Standard.
3. Any Priced Software labels on the Source Server are replicated to the Target Server and will begin incurring usage
4. The system will delete the Consistency Group entirely. All DRS Journal Space will be deleted along with all DRS snapshots and historical change information.
5. Because the Consistency Group no longer exists, all server actions against the Source and Target Servers are now permitted

For more details on deleting a Consistency Group, see How to Delete a Consistency Group.

DRS for Cloud Usage Reporting

Use of the DRS for Cloud service drives two different usage elements:

1. **DRS Source Server Hours** - Each Source Server in a Consistency Group contributes one hour of usage to this element for each hour it is a member of the Consistency Group. This usage is attributed to the data center location of the Source Server.
    a. ***Example***: *User creates a Consistency Group at Noon with two Source Servers in NA9 paired with two Target Servers in NA12. This configuration will generate 24 DRS Source Server Hours in NA9 on the first day (12 hours x 2 Source Servers = 24 Source Server Hours) and 48 DRS Source Server Hours/day thereafter until the Consistency Group is deleted.*
    b. Details on how this usage is computed can be identified through the DRS Pairs Usage report. This report identifies each Source Server / Target Server pair in a Consistency Group and how usage was computed for the Source Servers in those groups for each day. For details, see How to Generate a DRS Pairs Usage Report.
2. **DRS Storage Hours** - A Consistency Group includes a specific amount of DRS Journal storage in both the Source and Target locations. This usage is attributed to the DRS Storage Hours element separately in each location in the same manner that other disk storage usage is computed - one hour for each GB of DRS Journal storage.
    a. ***Example***: *User creates a Consistency Group at Noon with two Source Servers in NA9 paired with two Target Servers in NA12 and allocates 10,000 GB of DRS Journal Storage to the Consistency Group. On the first day, this configuration will generate 12,000 DRS Storage Hours in NA9 (12 hours x 10,000 GB = 12,000 DRS Storage Hours) and another 12,000 DRS Storage Hours in NA12. Thereafter, both NA9 and NA12 will be attributed 24,000 DRS Storage Hours/day until the Consistency Group is deleted.*
    b. Details on how this usage is computed can be identified through the same Detailed Usage report that tracks storage usage for Cloud Servers, Customer Images, and Cloud Files accounts. Each location's DRS Journal will appear on a separate row in the report with a column identifying the amount of DRS Journal storage and the attributed DRS Storage Hours. For details, see How to Create a Detailed Usage Report.

In summary, the usage impacts of the DRS for Cloud service are as follows:

- **Cloud Server usage for servers involved in a DRS for Cloud Consistency Group**
    - Source Server is attributed usage “normally” for resources used
    - Target Server is attributed usage for resources used:
        - When in "DRS Mode"
            - 1 x “Standard” vCPU, 1 x GB RAM and all storage with the appropriate disk speed
            - No Priced Software charges apply in this state.
        - When In "Preview" mode and after "Failover" is initiated:
            - “Standard” vCPU / RAM usage for the running configuration (that matches Source Server) and all storage with the appropriate disk speed
            - Priced Software charges that match the paired Source Servers.
- **DRS for Cloud usage in addition to the usage above**

    - Storage allocated to the recovery journals for the Consistency Group in the Source and Target locations contribute usage equal to their size under the DRS Storage Hours usage element, computed as hourly usage based on the amount of journal storage. This usage is attributed to each Source/Target location where the journal storage is used.
    - DRS Source Server Hours usage element is incurred hourly for each Source Server in the Consistency Group. This usage is allocated to the Source Data Center location

Recently Updated

- Introduction to Cloud Server NICs in MCP 2.0
- New Features by Release
- Navigating the Virtual Networks and Servers Dashboard
- October 2019 Changes to Cloud Server Snapshot Plans
- Introduction to Usage Reporting