DISTRIBUTED DATA COLLECTION FOR REINSURANCE AND RISK ADJUSTMENT:

EDGE SERVER DATA MANAGEMENT

October 2, 2014

Health Insurance Exchange Program Training Series

CMS
CENTERS FOR MEDICARE & MEDICAID SERVICES

HTTPS://WWW.REGTAP.INFO

Session Guidelines

- This is a ninety-minute webinar session.
- For questions regarding content, please submit inquiries to REGTAP at https://www.regtap.info
- For questions regarding logistics and registration, please contact the Registrar at: (800) 257-9520.

Intended Audience

- Amazon and On-Premise EDGE server issuers (Exchange and Non-Exchange)
- Third Party Administrators (TPAs) and Support Vendors
- Issuers deciding which EDGE server option is the best fit for them

Purpose

This Distributed Data Collection (DDC) Session will provide information on the following topics for both Amazon EDGE and On-Premise EDGE servers:

- EDGE Server Data Management including:
  - Backup
  - Recovery/Restore
  - Data Retention
  - Purging

EDGE Server Data Management

Overview of EDGE Server Data Retention

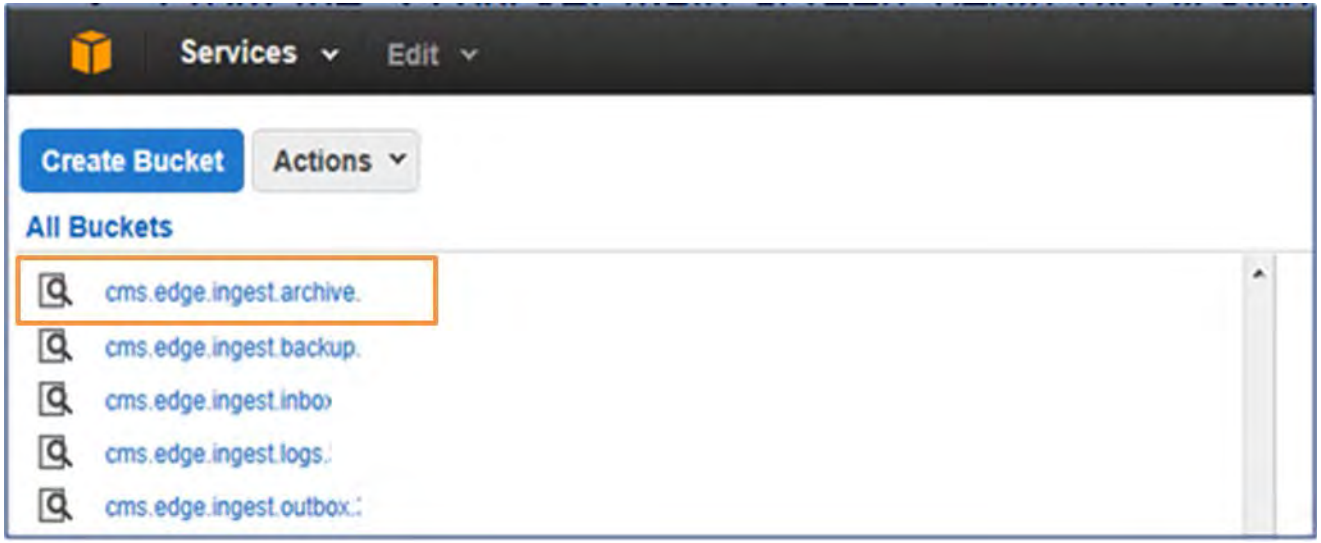
- All Amazon EDGE and On-Premise EDGE issuers are required to retain all inbound and outbound files and all data stored in their MySQL data tables for a period of 10 years.
  - Issuers must maintain their MySQL data in an active, queryable (active) status for a period of 3 years.
  - CMS recommends issuers maintain the inbound and outbound files on their servers for at least the entire reporting period for a coverage year.
  - The MySQL table will store a hash value for every inbound and outbound file.
  - CMS will provide additional guidance on how data should be transferred and retained off of the EDGE server in a job aid.

Amazon EDGE Server: Data Backup

- Amazon EDGE servers will store all inbound, outbound and backup files in their AWS S3 bucket Archive Folder as part of the standard file ingest and outbound report generation process.
- Amazon EDGE issuers have the option of archiving their files to Amazon Glacier.
  - Amazon Glacier offers a lower cost method for archiving files outside the S3 bucket.
  - Issuers may elect to maintain their files in either S3 or Glacier at their own discretion.
  - Amazon EDGE issuers can also export and store their inbound and outbound files off of their EDGE servers.

Storing Files to Amazon Glacier

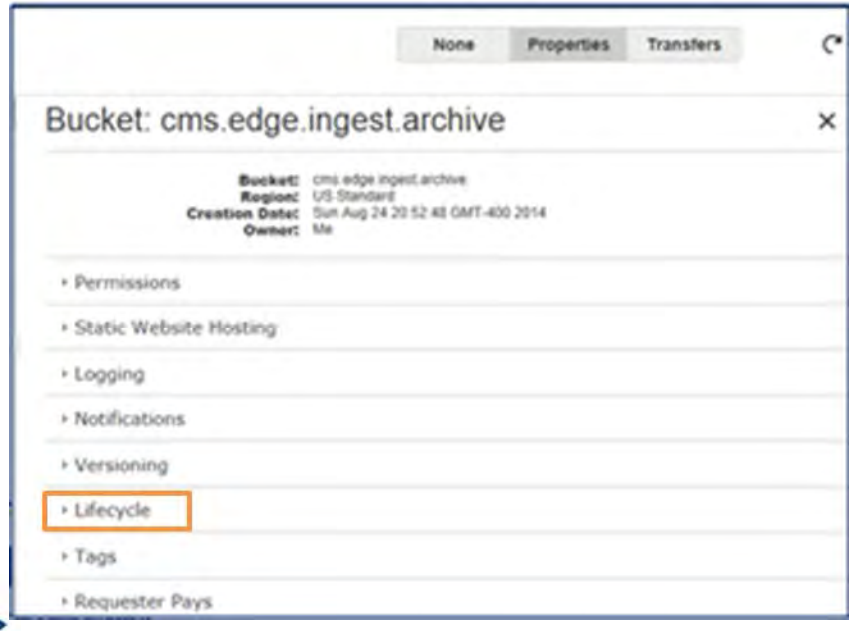
- After uploading files to the S3 bucket and generating outbound reports, Amazon EDGE issuers can archive their files to Amazon Glacier by creating a **Lifecycle Rule** for their S3 bucket.
- From the S3 bucket console, begin by clicking the **Archive Folder** link:



Storing Files to Amazon Glacier

- Once inside the Archive Folder, click the **‘Properties’** button at the top right.

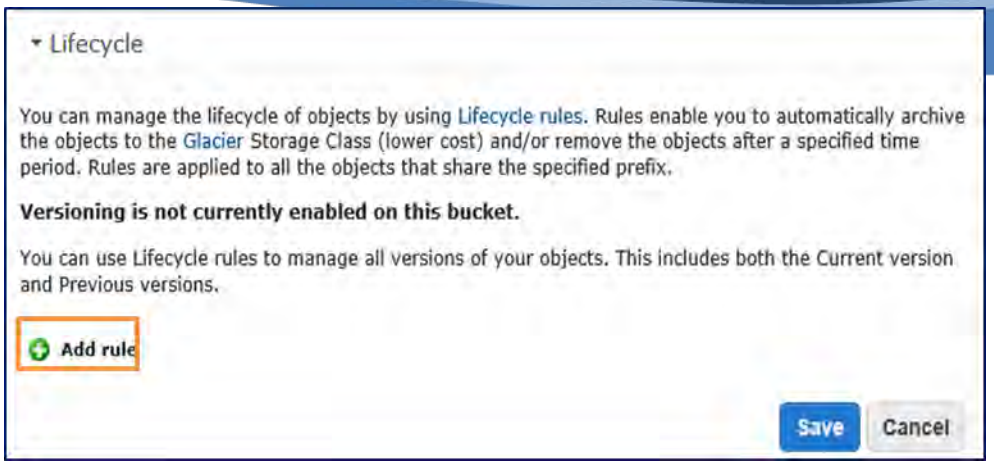




- From the **Properties** tab, click the **‘Lifecycle’** option from the menu.

Storing Files to Amazon Glacier

- In the Lifecycle menu, click **‘Add Rule’**. This will allow you to create a new lifecycle rule that defines conditions for moving S3 files to Amazon Glacier.

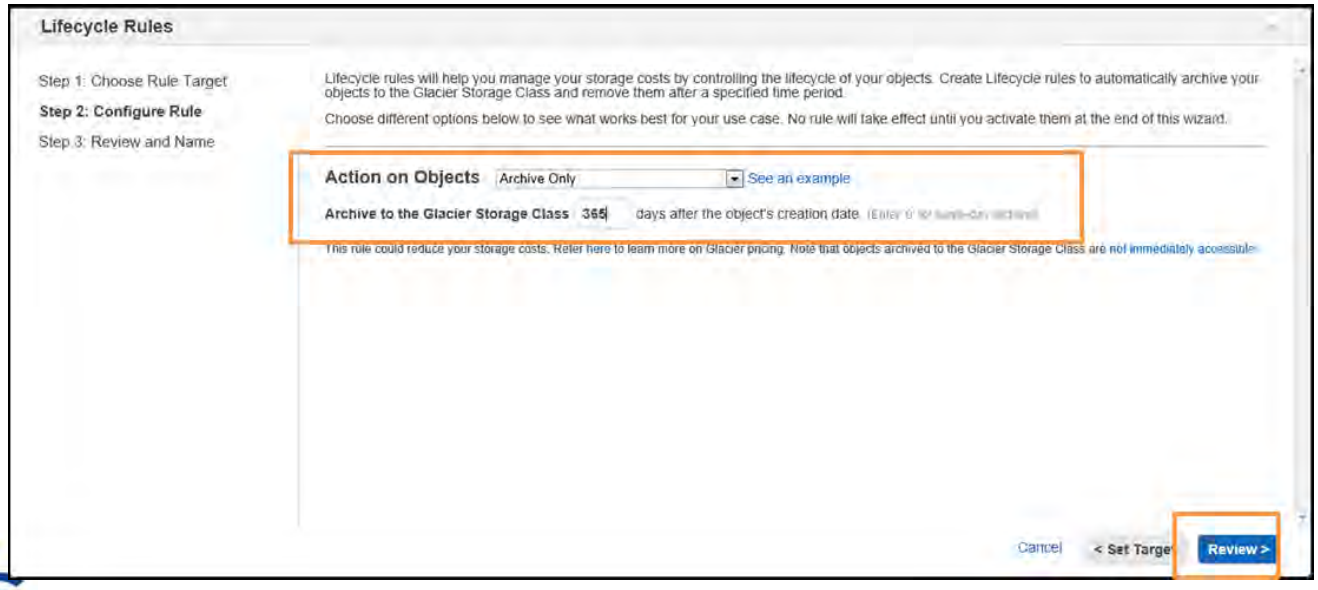




In the Lifecycle Rules window, select to apply the new lifecycle rule to the **Whole Bucket (S3).** Since all files in the S3 bucket will be for the EDGE server, you can select ‘Whole Bucket’ and then click **‘Configure Rule’** button.

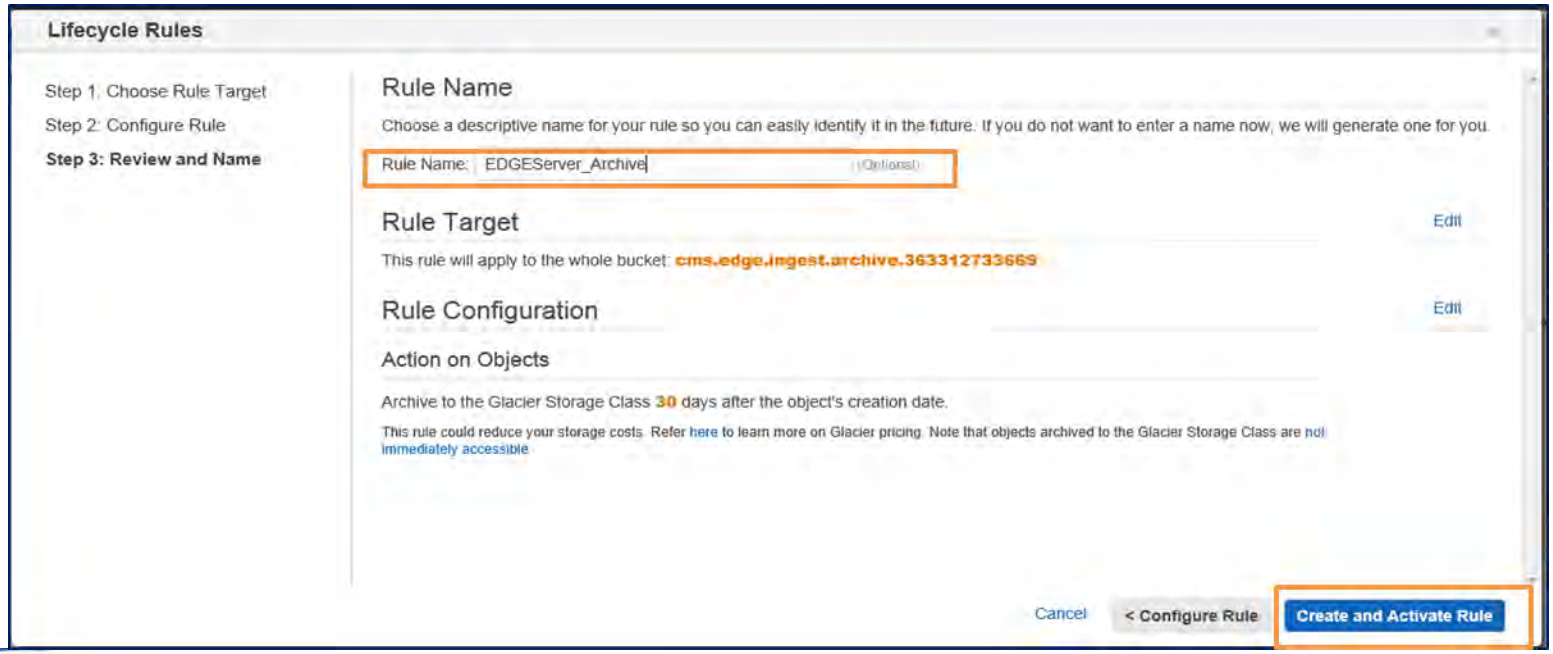
Storing Files to Amazon Glacier

- To configure the new Lifecycle rule, select **'Archive Only'** from the **Action on Objects** dropdown menu. Enter the number of days after the object's creation data to begin archiving.
  - Archived files will move to Amazon Glacier from the S3 bucket after this many days.
- When finished, click the **'Review'** button.



Storing Files to Amazon Glacier

- Finally, review the information and enter a **Rule Name** at the top to be able to identify the rule going forward. When finished, click the **'Create and Activate Rule'** button at the bottom right.



Storing Files to Amazon Glacier

- Once the rule is created, EDGE server files in your S3 bucket will transfer to Amazon Glacier based on the time parameters set in your rule.
- For more information on Amazon Glacier, please access this AWS page: http://aws.amazon.com/blogs/aws/archive-s3-to-glacier/
  - At the bottom of this page, you can find a YouTube video that provides a demonstration on setting Lifecycle rules, and also how to restore archived files from Amazon Glacier.

Amazon EDGE Server: Recovery Process

- Amazon Web Services includes built-in recovery functions for Amazon EDGE servers that demonstrate how to recover from a backup (i.e., EC2 Volume Snapshots).
- Amazon EDGE issuers will need to develop disaster recovery plans based on Amazon's existing recovery functions.
- Issuers should ensure they have recovery points set on at least a quarterly basis.
  - Issuers should also time recovery points to correspond to their data submissions.

EC2 Volume Snapshots

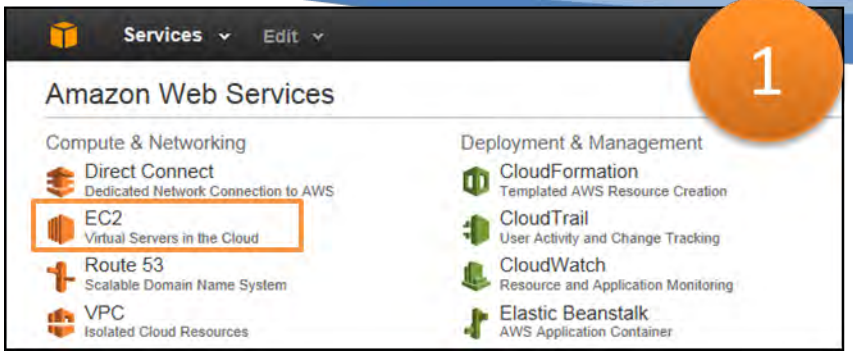
- Amazon EDGE issuers can configure a recovery point for their server instance by taking a snapshot of their EC2 volumes.
- All Amazon EDGE server instances contain two EC2 volumes:
  - OS and EDGE application volume
  - EDGE application data volume
    - This volume contains all the inbound and outbound files and MySQL data volumes.
- Amazon EDGE server issuers will need to take a snapshot of these volumes in order to recover their server instances for a certain point.
- Please remember to shut down any running process, such as the EDGE server application or MySQL database, prior to creating snapshots to avoid any integrity issues.

Create EC2 Volume Snapshot

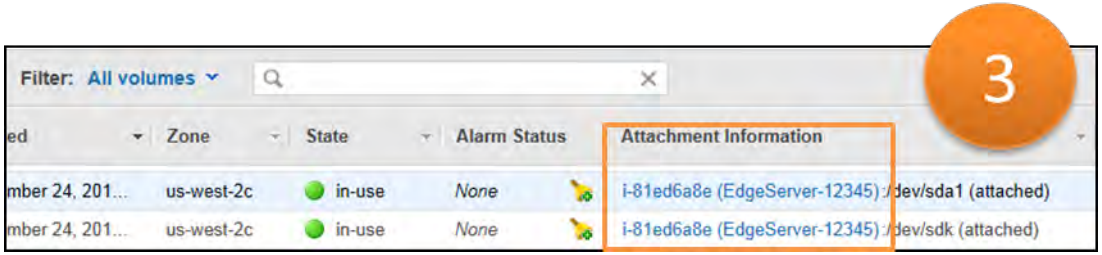
Step 1: From the AWS main portal, select **'EC2'** from the menu to access the available server instances in your account.

Step 2: Select **'Volumes'** from the list of available resources to identify the EC2 volumes attached to your server instance.

Step 3: From the list of available volumes, identify your EC2 volumes by scrolling to the right to find the **Attachment Information** and identify your server names and HIOS ID.





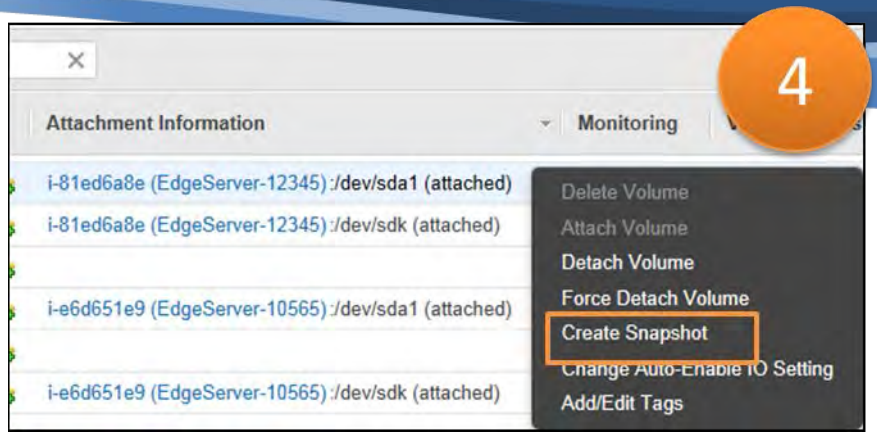


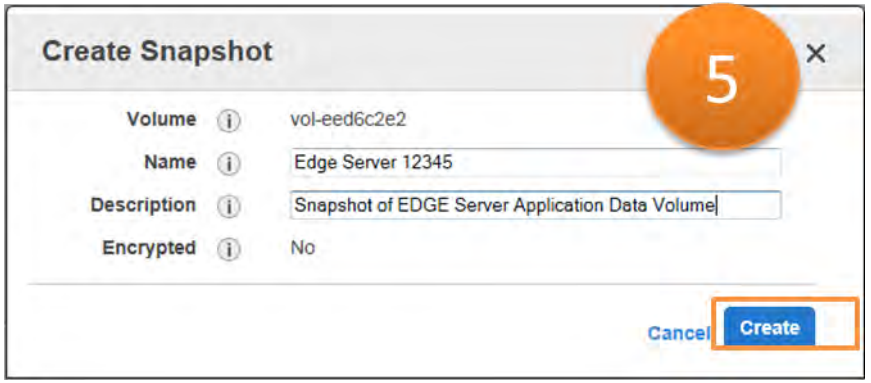
Create EC2 Volume Snapshot

Step 4: Once you have identified your EC2 volumes, right click on a volume and select **'Create Snapshot'** from the menu.

Step 5: In the Create Snapshot menu, enter a name and description for the snapshot being created and then click '**Create**.'

Step 6: When finished, you can select **'Snapshots'** from the EC2 dashboard and then locate the snapshot you created for your EC2 volume.

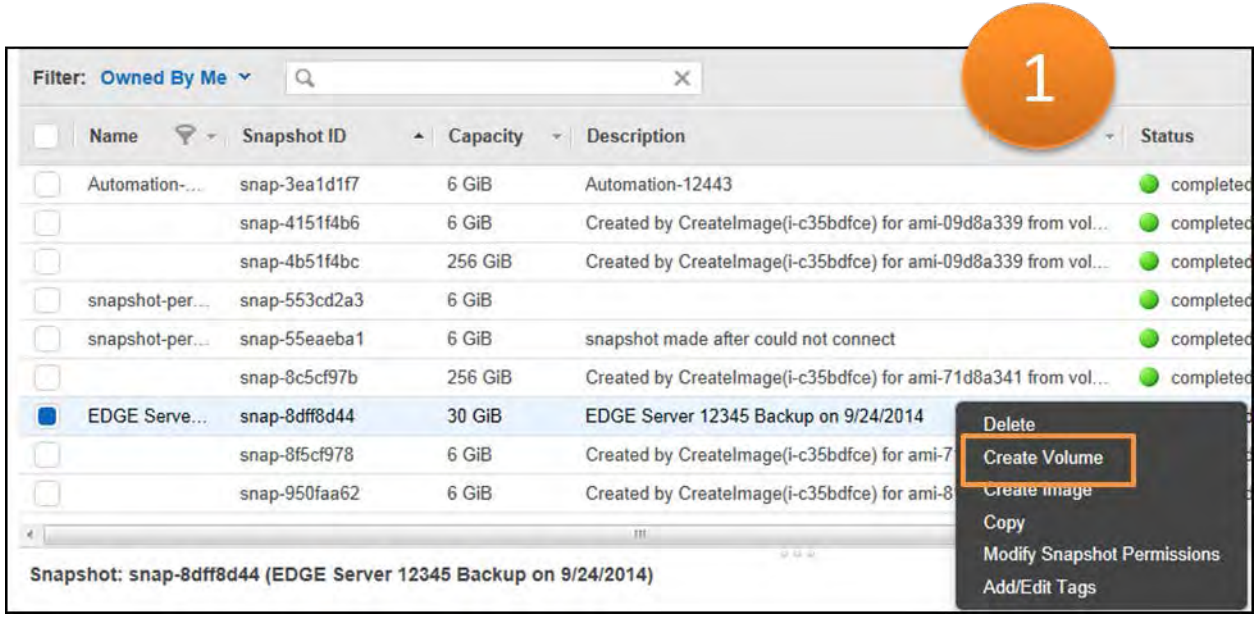






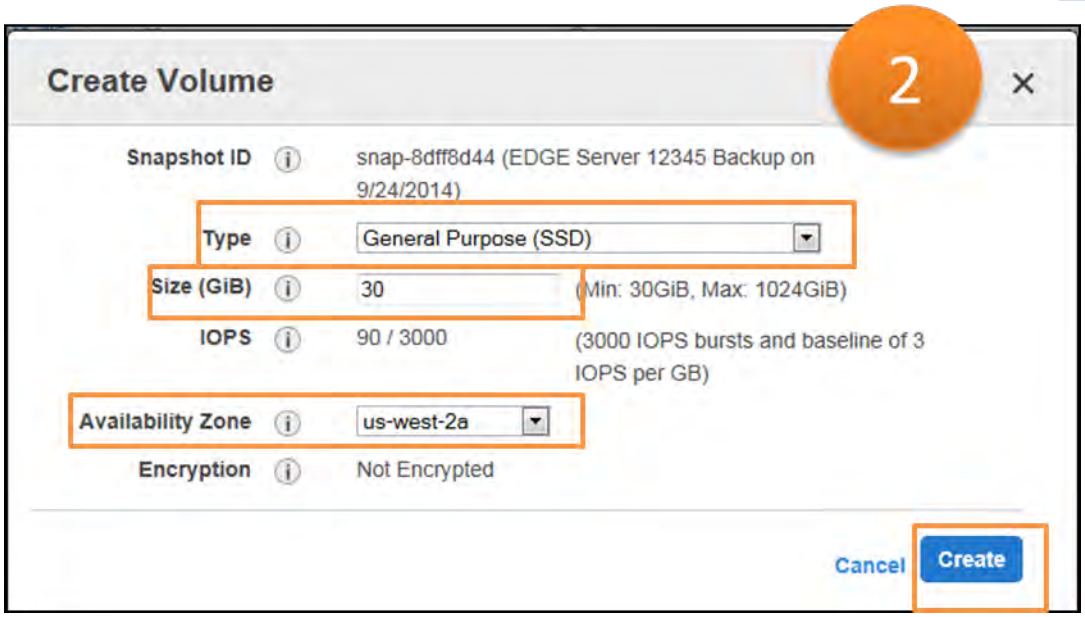
Recover Server Instance from EC2 Snapshot

- In order to recover your Amazon EDGE server to a certain point, you will need to create a new EC2 volume from the snapshot you created and then attach it to your server instance.
- **Please remember to stop your server instance before executing this step.**
- **Step 1:** Right click on your EC2 volume snapshot and select **'Create Volume'** from the menu.



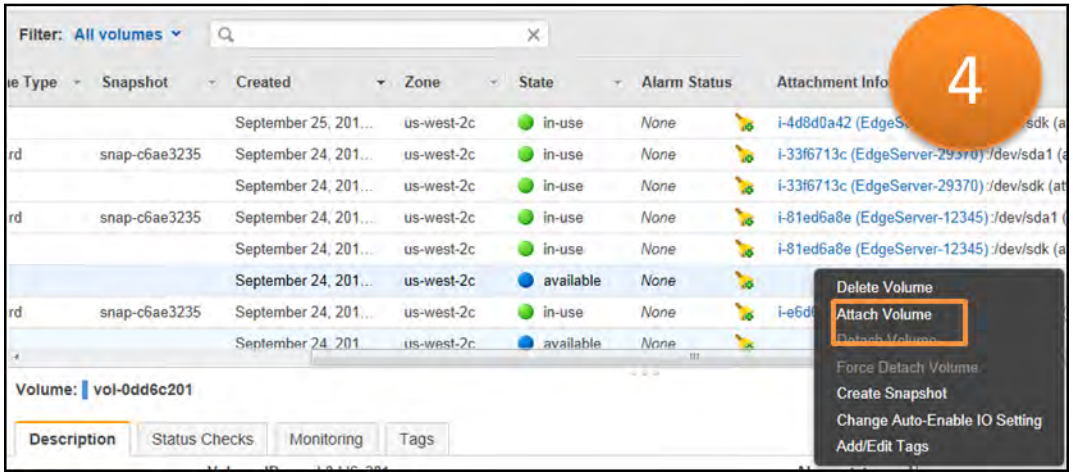
Recover Server Instance from EC2 Snapshot

- **Step 2:** In the Create Volume menu, select **‘General Purpose’** for the volume type, enter **‘30’** for volume size and select the correct availability zone for your volume – this should match the zone of your server instance. Click **‘Create’** when finished.



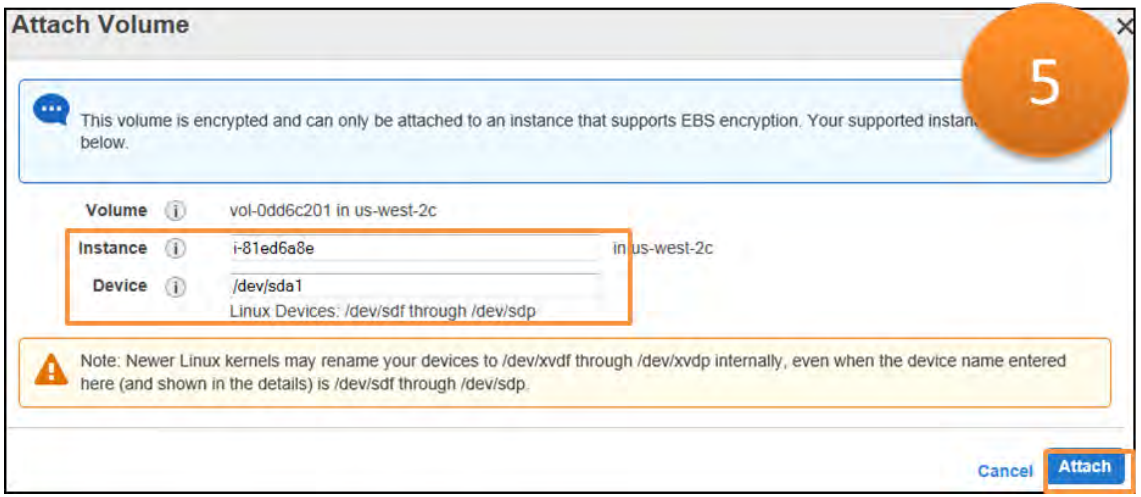
Recover Server Instance from EC2 Snapshot

- **Step 3:** When finished, you can select **‘Volumes’** from the EC2 dashboard and then locate the new volume you created from your EC2 snapshot.
- **Step 4:** Once you locate your volume, right click on it and select **‘Attach Volume’** from the menu.



Recover Server Instance from EC2 Snapshot

- **Step 5:** In the Attach Volume window, enter the name of your server instance and device and click **'Attach.'**
  - **Once attached, the new volume is active.**
- Remember to enter the same device name you use for the new volume as the one you created the snapshot from to provide a seamless restore. If a different name is used, an OS configuration may be required in order to mount the device.



On-Premise EDGE Server Recovery Process

On-Premise EDGE Server: Recovery Process

- On-Premise EDGE issuers are responsible for implementing their own IT recovery solutions.
  - The On-Premise EDGE server application does not have a built-in recovery function.
- On-Premise EDGE issuers should ensure they have recovery points set on at least a quarterly basis.
  - Issuers should also time recovery points to correspond to their data submissions.

Data Table Snapshot

Data Table Snapshots

- The EDGE application will create a flat file containing all of the information in the stored data tables each time the RA/RI calculations are run.
  - The snapshot file is intended to create a record of the data that is used in each calculation run.
  - The file will be assigned a hash value to ensure traceability.
  - Every inbound and outbound file will have a hash value. These values will be stored in the data tables on the servers and will also be sent to CMS.
  - Note: This applies to both On-Premise and Amazon issuers.

EDGE Server Data Purging

Production Zone Data Purging

- In general, all issuers are expected to resolve issues with their stored claims data by submitting void and replacement claims.
- Issuers must resolve enrollment data issues by submitting full, updated replacement enrollment files.
- On a case-by-case basis, CMS may approve requests from issuers to purge all of their production zone data. Issuers will have to demonstrate that they have widespread systemic errors in their stored claims data that require complete purging of all their tables.

Production Zone Data Purging (Continued)

- CMS will provide an execution script and guidance for purging production data.
- Issuers must submit data purging requests to the Help Desk and get approval from CMS prior to purging production data.
- In order to maintain the integrity and traceability of the data collection process, issuers will not be allowed to delete or make modifications to individual stored records in the MySQL data tables.
- CMS will check for changes to the MySQL tables as part of the Reinsurance and Risk Adjustment audit process.

Test Zone Purging

- CMS will allow issuers to purge their data in the test zone of the server in order for their test zone data tables to mirror the production data tables.
- A script will be provided to issuers to run through MySQL command line interface that will purge their test zone data.
- CMS will provide scripts and guidance for purging test data. Issuers do not need CMS approval to purge test data.

CMS Next Steps

- CMS will be providing additional details on both EDGE server options in upcoming webinars and user groups.
- The Distributed Data Collection for Reinsurance and Risk Adjustment Webinar Series II and User Group Series III will be extended through the end of December 2014.
- Future webinar sessions will be held on Thursdays from 3:30 – 5:00 p.m. ET with user group sessions to follow on Tuesdays from 11:30 a.m. – 12:30 p.m. ET.

User Group Topic	Scheduled Event Date and Time
Q&A	October 7, 2014, 11:30 a.m. – 12:30 p.m.
Q&A	October 8, 2014, 11:30 a.m. – 12:30 p.m.
Q&A	October 9, 2014, 3:30 p.m. – 5:00 p.m.

Questions?

To submit questions by phone:

- ❑ dial ‘14’ on your phone’s keypad
- ❑ dial ‘13’ to exit the phone queue

Resources

Resources

Resource	Link/Contact Information
Center for Consumer Information and Insurance Oversight (CCIIO)	http://cms.gov/cciio/
Registration for Technical Assistance Portal (REGTAP) • Registration • Inquiry Tracking and Management System (ITMS) • Resource Library • Frequently Asked Questions (FAQs)	https://www.REGTAP.info/

Inquiry Tracking and Management System (ITMS)

ITMS is available at https://www.regtap.info/

Users can submit questions after the webinar by selecting "Submit an Inquiry" from My Dashboard.



Note: Enter only one (1) question per submission.

Closing Remarks