

IBM InfoSphere DataStage Server on AWS
Version 1 Release 1

User Guide
(Last updated: 2018-12-14)



Note

Before using this information and the product that it supports, read the information in [Notices and trademarks](#).

Edition Notice

This edition applies to version 1, release 1, modification 0 of IBM InfoSphere DataStage Server on AWS and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2015, 2018.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. KMS (Key Management Services)	1
Creating a new customer master key (CMK).....	1
Chapter 2. Deploying AMI Image	3
Example.....	3
Chapter 3. Architecture	5
DataStage Client.....	5
DataStage Server.....	5
Chapter 4. Connecting to your DataStage Server	7
Advanced.....	7
Modifying allowed IP ranges.....	7
Creating new DataStage Client instances.....	7
Chapter 5. CloudFormation Templates	9
How to Deploy DataStage CloudFormation Templates.....	9
Post-Deployment Steps.....	9
Parameters.....	9
VPC Configuration.....	9
IBM DataStage Instance Configuration.....	10
Advanced.....	10
Templates.....	10
Common.....	10
Premium.....	10
High Availability.....	11
Chapter 6. Route 53 Hosted Zones	13
Proxying additional ports on Network Load Balancer.....	13
Chapter 7. Backing up DataStage	15
S3 Bucket.....	15
Creating S3 Bucket.....	15
WebSphere Application Profiles Backup.....	15
WebSphere Application Configuration Backup.....	15
DataStage Database Backup.....	15
Online full database backups.....	16
Online incremental database backups.....	16
DataStage artifacts backup using ISTOOL.....	16
Files and Directories Backup.....	16
EBS Volume Backup for /opt/IBM disk.....	16
EBS Volume Backup for /Backup disk.....	17
Adding Extra Files or Folders to local directory or S3 bucket as backup.....	17
For directory backup.....	17
For file backup.....	17
Configuring backup process.....	17
Setting backup jobs time and frequency using Crontab.....	18
Chapter 8. Disaster Recovery	21

Scenario 1: Storage is corrupted or not working as expected..... 21
Scenario 2: There is problem with the current availability zone..... 23

Notices.....25

Chapter 1. KMS (Key Management Services)

AWS KMS is a fully managed service that makes it easy to create and control encryption keys on AWS which can then be utilized to encrypt and decrypt data in a safe manner. Compliance and security monitoring are ensured in KMS by using key-based permission policies and integrating with Amazon's Identity and Access Management (IAM) which allows access management to a very granular level.

The Customer Master Key (CMK) is a logical key which represents the primary resource in AWS KMS. Master keys are either customer- or AWS managed and can be used to encrypt up to 4 kilobytes (4096 bytes) of raw data directly. CMKs are unique to your AWS account and the AWS region in which they are used.

Creating a new customer master key (CMK)

To create a new CMK execute the `CMK_Policy.sh` script and provide the inputs for the below prompts:

NOTE: CMKs are unique to your AWS account and the AWS region in which they are used

```
Enter AWS Account Id:
Type your AWS account id.
(Example: 299743145111)

Enter Admin user Id:
Type your user id.
(Example: xyz.corporation.DS)

Enter key User Id
Type the same user id as entered for above prompt.
(Example: xyz.corporation.DS)
```

After providing the inputs, script will display the `CMK KEY ID` value and `CMK ALIAS NAME`. You can use this `KEY ID` or `ALIAS NAME` for encryptions.

Chapter 2. Deploying AMI Image

To copy the snapshots to your region and re-encrypt them using your account's KMS encryption keys in the target region follow the below steps:

1. Login to your ec2 Linux instance as root user through ssh connection.
2. Once logged in to the ec2 instance, switch to the region where snapshots are copied from the source account using "aws configure" command. You can check the region where snapshots are copied using AWS console under "Snapshots" section in the EC2 Dashboard.

Example

Here the snapshots are copied into "Tokyo (ap-northeast-1)" region, you can check the region in the AWS console under "Snapshots" section in the EC2 Dashboard. Change the filter condition to "Private Snapshots".



As our snapshots are in "Tokyo (ap-northeast-1)" region, so changed the region using "aws configure" command as below:



1. Now go to the directory where TargetMoveAMI.sh script is present and execute TargetMoveAMI.sh script in the environment as below:

```
./TargetMoveAMI.sh
```

1. Once above script executes successfully, now switch to the region where you copied these snapshots using "aws configure" command.
2. Once you are in the correct region, execute CreateAMI.sh script as below and provide the input for the prompts:

```
./CreateAMI.sh `` ` Enter Root Volume Snapshot ID: Snapshot id for the ROOT VOLUME
```

```
Enter Second Volume Snapshot ID: Snapshot id for the second volume
```

```
Enter Third Volume Snapshot ID: Snapshot id for the third volume ` ` `
```

1. The CreateAMI.sh shell script will create a new AMI Image which can be used to deploy a new machine.

Chapter 3. Architecture

IBM InfoSphere DataStage is a data integration tool for designing, developing, and running jobs that move and transform data. Using IBM-developed Amazon Machine Images and CloudFormation templates, DataStage can be deployed on your own AWS account. The following guide will use AWS terminology. Please refer to the [AWS Glossary](#) and [AWS Documentation](#) for reference.

DataStage Client

The DataStage Clients are Windows 2012 bastion machines placed in a public subnet. These instances are attached to an Auto Scaling, which confers two benefits: - additional instances can be created (as your license permits) - a problematic instance can simply be terminated and a replacement will be provisioned automatically

Being in a public subnet, each DataStage Client instance has its own public IPv4 address. The Client instances' security groups are set to allow inbound traffic only on port 3389 (RDP) and from the "Allowed Client External Access" IP range, supplied when the templates are provisioned.

From the instances, the DataStage web application can be accessed using Internet Explorer. PuTTY is also available to be able to log in to the DataStage Server instances from the Client.

DataStage Server

The DataStage Server is running on a Red Hat 7.5 instance in a private subnet, with outbound Internet traffic route to a NAT Gateway. This allows Internet communication initiated by the DataStage Server instance (such as for updates or downloading data). The security group of the Server instance only allows traffic from the DataStage Client subnets on select ports. Additionally, traffic over port 9446 is allowed the "Allowed Client External Access" IP range via a Network Load Balancer proxy. This enables Launchpad access from a browser on the Internet.

Access can be modified as described in "Modifying allowed IP ranges" in Chapter 4.

Chapter 4. Connecting to your DataStage Server

The DataStage server can be accessed in the following two ways. Note that only IP addresses in the range specified in the `ClientAppAccessCIDR` parameter are configured to have access.

1. Connecting to the DataStage Client

Upon logging in to the Client instance via RDP, there will be an icon on the Desktop to open the Launchpad.

2. Connecting to DataStage Server Launchpad over HTTPS

The URL for the Launchpad is the Output value `NetworkLoadBalancerDNSName` from the CloudFormation template. Alternatively `NetworkLoadBalancerEndpointURL` will also work if a valid DNS endpoint was provided during creation, and the "Steps to Complete Endpoint Attachment" have been followed.

Advanced

Modifying allowed IP ranges

Access to the Client instances may be modified from the initial value of the `ClientAppAccessCIDR` parameter:

1. On the Instances page, select a Client instance and click on its Security group.
2. In the "Inbound" tab, existing inbound access rules may be edited.

Access to the DataStage Server (Launchpad) via the NLB can similarly be modified by changing/adding inbound access rules in the security group of the DataStage Server.

Creating new DataStage Client instances

Client instances are controlled in an Auto Scaling group. New client instances can be created by modifying Auto Scaling options as follows:

1. Go to `Services > EC2`. Select the `Auto Scaling Groups` page from the left navigation and select the group `<OrderID>-DataStageClientStack- . . .`
2. On the `Actions > Edit` page, change Max and Desired Capacity values.
3. If you have attached a DNS endpoint, add the public IP address of the newly-created Client instance in the Route53 Public Hosted Zone for this deployment.

Note: To scale down, the Desired Capacity can be decreased in the same way, which will delete some Client instances as per the [default termination policy](#). Alternatively, if you wish to control which instance is deleted, select the Instances tab below the Auto Scaling group and `Actions > Detach` the instance that you want to delete. This will also automatically decrement Desired Capacity (only until Min). Then, terminate the instance manually by going to the Instances page from the left nav.

Chapter 5. CloudFormation Templates

There are three or four CloudFormation templates provided for Premium and HA respectively. The set of resources created by a template is referred to as its "stack". Refer to "Templates" below for information about the individual templates.

How to Deploy DataStage CloudFormation Templates

1. Upload all the templates to your private S3 bucket.
2. Update the TemplateURL fields for all the stacks in `ibm-master.template` and `ibm-datastage-master.template` with the new S3 links.
3. Re-upload `ibm-master.template` and `ibm-datastage-master.template`.
4. In AWS, go to the CloudFormation management console and create a stack. Under "Choose a template", select "Specify an Amazon S3 template URL" and enter the URL to `ibm-master.template`. Pressing next will take you to the "Specify Details" page. Populate all the fields and click "Next". Refer to the "Parameters" section below to for reference.



1. Click "Next" on the "Options" page. Nothing needs to be entered here.
2. On the "Review" page, mark any checkboxes in the "Capabilities" section and press "Create".

Post-Deployment Steps



The DataStage Master stack will create additional nested stacks. When the status of the root stack changes from `CREATE_IN_PROGRESS` to `CREATE_COMPLETE`, click on the "Outputs" tab for information. Note that it will take 30 min to 1 hour for the stack to be deployed.

Parameters

Parameters are grouped in three categories: *VPC Configuration*, *IBM DataStage Configuration*, and *Advanced*. Parameters for which no default values provided are marked with an asterisk (*).

VPC Configuration

NB: if you intend to peer the DataStage VPC with another VPC, you must ensure that the CIDR blocks of both VPCs do not overlap. Refer to [sizing a private CIDR block and invalid peering configurations in AWS documentation](#).

- **Availability Zone***: Select any two zones from your region in which to create the subnets.
- **VPC CIDR**: Private IP range assigned to the entire VPC.
- **Private Subnet 1 CIDR**: Private IP range assigned to the private subnet in AZ 1
- **Private Subnet 2 CIDR**: Private IP range assigned to the private subnet in AZ 2
- **Public Subnet 1 CIDR**: Private IP range assigned to the public subnet in AZ 1
- **Public Subnet 2 CIDR**: Private IP range assigned to the public subnet in AZ 2

IBM DataStage Instance Configuration

- **DataStage Server Instance Type:** The underlying EC2 instance for the DataStage server. Please ensure that your choice does not violate licensing restrictions.
- **Key Pair Name*:** Key pair to allow logging in to the instances
- **Allowed Client Application External Access CIDR*:** The public IP block that is allowed to access resources in the VPC. For better security, ensure that this range is as narrow as possible (refer to "Determining my IP access range"). To remove or modify the IP block entered here or to allow additional IP blocks, refer to "Modifying Client Application External Access CIDR" after DataStage has been deployed.
- **DataStage OrderID Details*:** You must use a unique string for each DataStage stack deployment in your AWS account. Resources will be named or tagged with this value to allow identification and avoid conflicts from multiple deployments.
- **Domain Name:** The domain or subdomain endpoint to be attached to the VPC. After the stack is deployed, refer to the section "Route 53 Hosted Zones" for instructions on how to finish endpoint attachment with your DNS management provider.

Advanced

- **Block Device IOPS (/dev/sda1)*:** TODO
- **Block Device Volume Size (/dev/sda1)*:**
- **Block Device IOPS (/dev/sdb)*:**
- **Block Device Volume Size (/dev/sdb)*:**
- **Block Device Volume Size (/dev/sdf)*:**

Templates

Common

- `aws-vpc.template`
AWS-provided template to create a VPC with public and private subnets, included as-is. Further information can be found in [AWS documentation](#).
- `ibm-master.template`
Starting point for deployment into a new VPC. It creates a VPC using `aws-vpc.template` and invokes the `ibm-datastage-master.template` stack, URLs for which must be updated by the user before deployment.
- `ibm-datastage-master.template`
Creates the DataStage Client, DataStage Server and Security stacks.
- `ibm-datastage-security.template`
Creates security groups used by both the DataStage Client and DataStage Server instances.
- `ibm-datastage-client.template`
Creates the Windows Client bastion host in an autoscaling group.

Premium

- `ibm-datastage.template`
Creates the DataStage Server in a private subnet as a Launch Template, an IAM role for it, and creates a target group and listener for port 9443 NLB proxying.

High Availability

- TODO

Chapter 6. Route 53 Hosted Zones

If a valid DNS endpoint was provided during stack creation, a few final steps are required to finish attaching the domain name. Route53 automatically assigns four nameservers to every Hosted Zone. These will be stated in the stack output variable `HostedZoneNameServers` or can be viewed by opening the Hosted Zone records in Route53. Update your registrar's or DNS service's records for the domain with NS records containing these Route53 nameservers. Depending on your provider this may be in the form of a single NS record with delimited values or four different NS records. Further information can be found [in the AWS Developer Guide](#).

Proxying additional ports on Network Load Balancer

By default, the Network Load Balancer allows access from the ClientAppCIDR to the DataStage server over port 9446. If access to another port (such as 22 for SSH) is desired, a target group and listener can be attached to the NLB as follows:

1. Go to Services > EC2
2. Open the "Target Groups" page from the left navigation list and create a target group with the following settings:
 - Target group name: <anything>
 - Target type: Instance
 - Protocol: TCP
 - Port: (port to open)
 - VPC: (VPC where DataStage is installed)
 - (Health check settings) Protocol: TCP
 - (Advanced health check settings) Port: override: (port to open)
 - (Advanced health check settings) Healthy threshold: 3
 - (Advanced health check settings) Interval: 10 seconds
3. Select the newly-created Target Group. Navigate to the Targets tab and click "Edit". Select the DataStage Server instance, "Add to registered" and save. The instance should appear under "Registered targets" with the status "unused".
4. Open the "Load Balancers" page from the left navigation. Select the load balancer for DataStage and from the "Listeners" tab, add a listener with the following settings and save:
 - Protocol: TCP
 - port: (port to open)
 - Default action: Forward to target group created in Step 2.
5. Back on the Target Groups page, the DataStage Server target should show its status as "initial" and then "healthy" within a minute. You can now access the newly-opened port using the Network Load Balancer DNS name.

Chapter 7. Backing up DataStage

IBM DataStage on Cloud Premium and High Availability services provides software and hardware infrastructure for taking backups. Replicating client data from a source server to another server helps to ensure that backed-up data is available for recovery if the source server is damaged. Replication incrementally copies data from the source server to the target server to provide failover and fail-back capability. The data is then moved from local storage to the cloud. In this way, you can improve data backup and archive performance. Backup scripts copies data from local machine to Amazon S3 bucket.

S3 Bucket

A bucket is a logical unit of storage in object storage service, Simple Storage Solution S3. Buckets are used to store objects, which consist of data and metadata that describes the data. A bucket is analogous to a subdirectory, where the object storage in the main directory and the buckets in the object storage can be seen as subdirectories.

Creating S3 Bucket

For copying backup data to S3, you need to create a S3 bucket first. To create a bucket, use the below command:

```
/usr/local/aws/bin/aws s3 mb s3://YOUR-S3-BUCKET --region us-east-2
```

Replace your S3 bucket name with “YOUR-S3-BUCKET” and region “us-east-2” with the region in which you want to create the bucket.

WebSphere Application Profiles Backup

A scheduler script WASProfileBackup_DS.sh is configured to take backups of WAS profile directories. WASProfileBackup_DS.sh executes manageprofiles.sh script to create a zip folder named DS_AppServer_backup.zip. DS_AppServer_backup.zip contains content of WAS profile directory which is located at /opt/IBM/WebSphere/AppServer/profiles/

WebSphere Application Profiles backup will be taken every Sunday At 1:10 hrs.

WebSphere Application Configuration Backup

A scheduler script WASProfileConfigBackup_DS.sh is configured to take backups of WAS profile configuration. WASProfileConfigBackup_DS.sh executes backupConfig.sh script to create a zip folder named DS_AppServer_Configbackup.zip. DS_AppServer_Configbackup.zip contains content of WAS profile directory which is located at /opt/IBM/WebSphere/AppServer/profiles/InfoSphere

WebSphere Application Configuration backup will be taken every Sunday At 1:20 hrs.

DataStage Database Backup

There are two kinds of database backups: full and incremental. It is required that a full backup of the database exists before any incremental backup is taken.

NOTE: Database transaction logs are not part of backups.

Online full database backups

A scheduler script `db2FullBackup.sh` is configured to take backups of online full database.

Online full database backups will be taken every Sunday At 00:00 hrs.

Online incremental database backups

A scheduler script `db2IncrementalBackup.sh` is configured to take backups of online incremental database backups.

Online incremental database backups will be taken from Monday through Saturday At 00:30 hrs.

DataStage artifacts backup using ISTOOL

A scheduler script `istool_assets.sh` is configured to take backups of ISTool export configuration. `istool_assets.sh` executes commands to create an archive file which contains all ISTool export configurations.

- IIS `istool.sh` export backup : Using `istool.sh` export all configurations.

DataStage artifacts backup using ISTOOL will be taken from Monday through Friday At 00:40 hrs.

Files and Directories Backup

Two scheduler scripts `main_FilesDaily1.sh` and `main_FilesDaily2.sh` is configured to take backups of files and folders. `main_FilesDaily1.sh` and `main_FilesDaily2.sh` executes commands to take backup of below files and folders.

- `/opt/IBM/InformationServer/Server/MsgHandlers/*`
- `/opt/IBM/InformationServer/Updates/*`
- `/opt/IBM/InformationServer/Server/Projects/`
- `/opt/IBM/InformationServer/Server/Configurations/*`
- `/opt/IBM/InformationServer/Server/DSODB/*.cfg`
- `/opt/IBM/InformationServer/Server/DSEngine/dsenv`
- `/opt/IBM/InformationServer/Updates/ /Backup/FilesDaily2/$DateDir`
- `/opt/IBM/InformationServer/Version.xml /Backup/FilesDaily2/$DateDir`
- `/etc/sysconfig/iptables*`
- `/etc/services`
- `/etc/inittab`
- `/home/db2inst1/db2keystore.p12`
- `/opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/`
- `/opt/IBM/InformationServer/ASBServer/apps/lib/iis/classes/`

Files and Directories Backup will be taken from from Sunday through Saturday At 00:50 and 1:00 hrs.

EBS Volume Backup for /opt/IBM disk

A scheduler script `EBS_Snap.sh` is configured to create snapshot of the EBS volume containing the `/opt/IBM` directory structure where the Information Server product tiers, DB2, and WebSphere Application Server are all installed.

EBS Volume Backup for `/opt/IBM` disk will be taken from from Sunday through Saturday At 1:30 hrs.

EBS Volume Backup for /Backup disk

To take snapshot of the /Backup EBS volume which contains the local copy of all the backups, run the below command in AWS CLI, replace VOLUME_ID with the volume id of the backup disk and put some values for TAG_KEY and TAG_VALUE.

```
/usr/local/aws/bin/aws ec2 create-snapshot --volume-id VOLUME_ID --description 'EBS Backup for /Backup disk' --tag-specifications 'ResourceType=snapshot,Tags=[{Key='TAG_KEY',Value='TAG_VALUE'}]'
```

Adding Extra Files or Folders to local directory or S3 bucket as backup

If you want to add more folder or files in the existing list of backup objects, you need to follow below steps:

1. Edit main_FilesDaily1.sh script, and add your copy command in the section "Copying files to local dir" as below:

For directory backup

```
sudo cp -R source_directory_location /Backup/FilesDaily1/$DateDir
```

For file backup

```
sudo cp source_file /Backup/FilesDaily1/$DateDir
```

1. Save the main_FilesDaily1.sh script after edit

Configuring backup process

To enable backup process, perform the below steps:

1. Login to your ec2 Linux instance through ssh connection.
2. Change directory to /BackupScripts using cd command and verify for the below 11 shell script files under this directory:

```
backupToS3.sh  
createLocalDir.sh  
CronJobs.sh  
db2FullBackup.sh  
db2IncrementalBackup.sh  
istool_assets.sh  
main_FilesDaily1.sh  
main_FilesDaily2.sh  
WASProfileBackup_DS.sh  
WASProfileConfigBackup_DS.sh  
EBS_Snap.sh
```

1. Change user to super user using below command:

```
sudo su
```

Example: \$ sudo su #

2. To upload the backup files to S3, check for the output of below command:

```
/usr/local/aws/bin/aws ssm get-parameters --names "S3FLAG" --with-decryption --query 'Parameters[0].Value' --output text | xargs
```

If the output is “S3” then the files will be uploaded to S3 bucket. If the output is “Local” then the files will present in the local directory.

Example: # /usr/local/aws/bin/aws ssm get-parameters --names "S3FLAG" --with-decryption --query 'Parameters.Value' --output text | xargs S3

1. To change your preference where the backup needs to be saved “S3” or “Local”, run the below command:

For S3

```
/usr/local/aws/bin/aws ssm put-parameter --name "S3FLAG" --value "S3" --type "SecureString"
```

For Local

```
/usr/local/aws/bin/aws ssm put-parameter --name "S3FLAG" --value "Local" --type "SecureString"
```

1. Verify you are still in the /BackupScripts directory.
2. Now execute backupToS3.sh script as below:

```
./backupToS3.sh
```

Example # ./backupToS3.sh

1. Executing backupToS3.sh will create local directories and buckets in the S3. Also, it will create cron jobs which will schedule the other backup jobs to execute at the specified time and frequency.
2. Execute the below command in the command line to verify whether jobs are created in crontab or not.

```
crontab -l
```

Example:

```
[root@ip-10-0-25-62 BackupScripts]# crontab -l
0 16 * * 0-6 sudo /BackupScripts/db2FullBackup.sh > /tmp/db2FullBackup.txt 2>&1
0 16 * * 0-6 sudo /BackupScripts/db2IncrementalBackup.sh > /tmp/db2IncrementalBackup.txt 2>&1
10 16 * * 0-6 sudo /BackupScripts/istool_assets.sh > /tmp/istool_assets.txt 2>&1
20 16 * * 0-6 sudo /BackupScripts/main_FilesDaily1.sh > /tmp/main_FilesDaily1.txt 2>&1
30 16 * * 0-6 sudo /BackupScripts/main_FilesDaily2.sh > /tmp/main_FilesDaily2.txt 2>&1
50 16 * * 0-6 sudo /BackupScripts/WASProfileBackup_DS.sh > /tmp/WASProfileBackup_DS.txt 2>&1
0 17 * * 0-6 sudo /BackupScripts/WASProfileConfigBackup_DS.sh > /tmp/WASProfileConfigBackup_DS.txt 2>&1
10 17 * * 0-6 sudo /BackupScripts/EBS_Snap.sh > /tmp/EBS_Snap.txt 2>&1
[root@ip-10-0-25-62 BackupScripts]#
```

Setting backup jobs time and frequency using Crontab

To change the backup job time and frequency of execution follow the below steps:

1. Change user to superuser as:

```
sudo su
```

1. Open the CronJobs.sh script in edit mode and make the changes to below commands:

```
crontab <<EOF
0 0 * * 0 sudo /BackupScripts/db2FullBackup.sh > /tmp/db2FullBackup.txt 2>&1
30 0 * * 1-6 sudo /BackupScripts/db2IncrementalBackup.sh > /tmp/db2IncrementalBackup.txt 2>&1
40 0 * * 1-5 sudo /BackupScripts/istool_assets.sh > /tmp/istool_assets.txt 2>&1
50 0 * * 0-6 sudo /BackupScripts/main_FilesDaily1.sh > /tmp/main_FilesDaily1.txt 2>&1
0 1 * * 0-6 sudo /BackupScripts/main_FilesDaily2.sh > /tmp/main_FilesDaily2.txt 2>&1
10 1 * * 0 sudo /BackupScripts/WASProfileBackup_DS.sh > /tmp/WASProfileBackup_DS.txt 2>&1
20 1 * * 1-6 sudo /BackupScripts/WASProfileConfigBackup_DS.sh > /tmp/WASProfileConfigBackup_DS.txt 2>&1
30 1 * * 0-6 sudo /BackupScripts/EBS_Snap.sh > /tmp/EBS_Snap.txt 2>&1
EOF
```

Here are the parameters in order sequence of:

```
* min (0 - 59)
* hour (0 - 23)
* day of month (1 - 31)
* month (1 - 12)
* day of week (0 - 6) (Sunday=0)
* command
```

```
0 0 * * 0 sudo /BackupScripts/db2FullBackup.sh
Execute At 00:00 on Sunday.

30 0 * * 1-6 sudo /BackupScripts/db2IncrementalBackup.sh
Execute At 00:30 on every day-of-week from Monday through Saturday.

40 0 * * 1-5 sudo /BackupScripts/istool_assets.sh
Execute At 00:40 on every day-of-week from Monday through Friday.

50 0 * * 0-6 sudo /BackupScripts/main_FilesDaily1.sh
Execute At 00:50 on every day-of-week from Sunday through Saturday.

00 1 * * 0-6 sudo /BackupScripts/main_FilesDaily2.sh
Execute At 01:00 on every day-of-week from Sunday through Saturday.

10 1 * * 0 sudo /BackupScripts/WASProfileBackup_DS.sh
Execute At 1:10 on Sunday.

20 1 * * 1-6 sudo /BackupScripts/WASProfileConfigBackup_DS.sh
Execute At 01:20 on every day-of-week from Monday through Saturday.

30 1 * * 0-6 sudo /BackupScripts/EBS_Snap.sh
Execute At 01:30 on every day-of-week from Sunday through Saturday.
```

To change the time of execution you need to change the values for first five parameters. For example, if you want /BackupScripts/db2FullBackup.sh to run every day of the week at 2:30 am every day of the week then you need to modify the script as

```
30 2 * * 1-6 sudo /BackupScripts/db2FullBackup.sh
```

1. After the changes are done, to delete the old cron schedules run the below command:

```
crontab -r
```

1. Now to set the new cron schedule, go to /BackupScripts directory and execute the CronJobs.sh script as below:

```
./CronJobs.sh
```

Example # ./CronJobs.sh

1. To verify whether all jobs are scheduled properly or not run the below command to check. It will list all the cron jobs scheduled in crontab file.

```
crontab -l
```

Chapter 8. Disaster Recovery

This document describes the steps to follow when there are issues with the aws account and there is need to recover data in the EBS volume /opt/IBM and /Backup mounted respectively at /dev/sdb and /dev/sdf device.

Following are two scenarios when you need to perform the disaster recovery process:

Scenario 1: Storage is corrupted or not working as expected

In the scenario when there is some problem with the attached EBS storage volumes mounted at /dev/sdb and /dev/sdf devices then you need to perform the below steps:

1. First of all, lookout for the latest snapshot of EBS volume "/opt/IBM".
2. Now select the snapshot and click "Actions" tab and select "Create Volume" option for one or both the snapshots depending on the issue.



1. After you select the "Create Volume" option, In the next window you need to select the "Volume Type" of the new EBS Volume which you can select as the type of previous volume. And select the "Availability Zone" where the new ec2 instance is created. You can add tag as well, click "Create Volume" button at the end of the window to create a new EBS volume.

NOTE: Change the "Volume Type", "Size" and "Availability Zone" values as per your configuration.



1. Click on this "Volume Id" link above, it will take you to the "Volumes" window.



1. Now once volume is created, go to the "Launch Templates" section under "Instances" option in the console.



1. Select the appropriate "Launch Template" and click "Launch Instance from template" from the "Action" tab.



1. In the next page fill all the details like "Source template version", instance type, subnet and other details and then once done click on the "Launch instance from the template" button at the end of the page. It will launch a new ec2 instance.



- 1.
2. First of all, verify the EBS volumes having the issues. Perform the recovery process to one of both of the EBS volume /opt/IBM and /Backup mounted respectively at /dev/sdb and /dev/sdf device depending on the issues.



1. You have to detach the corrupted volume, for that first stop the running instance where this volume is attached.



1. Once the instance is in Stop state, now go the corrupted EBS volume, click the “Action” tab and select “Detach Volume”.



1. Once the volume is detached, check for the latest snapshot of the EBS volume mounted at /dev/sdb and /dev/sdf devices.



1. Now select the snapshot and click “Actions” tab and select “Create Volume” option for one or both the snapshots depending on the issue.



1. After you select the “Create Volume” option, In the next window you need to select the “Volume Type” of the new EBS Volume which you can select as the type of previous volume. And select the “Availability Zone” where the new ec2 instance is created. You can add tag as well, click “Create Volume” button at the end of the window to create a new EBS volume.

NOTE: Change the “Volume Type”, “Size” and “Availability Zone” values as per your configuration.



1. Click on this “Volume Id” link above, it will take you to the “Volumes” window.



1. Select your volume and click the “Actions” tab and select “Attach Volume” option.



1. In the new window it will click on the “Instance” box, it will show the list of all running instances. Select your instance from the dropdown, it will automatically fill the “Device” option for you. Once it is done click on the “Attach” button at the bottom of the window for both the volumes.



Scenario 2: There is problem with the current availability zone

In this scenario you need to create volume from snapshot and attach to the new machine having new IP address.

Perform steps 5 to 13 as mentioned in the Scenario 2, for creating new EBS volumes out of snapshot and attaching new volumes to the new instance.

Notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

Notices

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4

555 Bailey Avenue
San Jose, CA 95141-1003 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session or persistent cookies. If a product or component is not listed, that product or component does not use cookies.

Table 1: Use of cookies by InfoSphere Information Server products and components.

Product module	Component or feature	Type of cookie that is used	Collect this data	Purpose of data	Disabling the cookies
Any (part of InfoSphere Information Server installation)	InfoSphere Information Server web console	<ul style="list-style-type: none"> • Session • Persistent 	User name	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
Any (part of InfoSphere Information Server installation)	InfoSphere Metadata Asset Manager	<ul style="list-style-type: none"> • Session • Persistent 	No personally identifiable information	<ul style="list-style-type: none"> • Session management • Authentication • Enhanced user usability • Single sign-on configuration 	Cannot be disabled
InfoSphere DataStage®	Big Data File stage	<ul style="list-style-type: none"> • Session • Persistent 	<ul style="list-style-type: none"> • User name • Digital signature • Session ID 	<ul style="list-style-type: none"> • Session management • Authentication • Single sign-on configuration 	Cannot be disabled
InfoSphere DataStage	XML stage	Session	Internal identifiers	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere DataStage	IBM InfoSphere DataStage and QualityStage® Operations Console	Session	No personally identifiable information	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere Data Click	InfoSphere Information Server web console	<ul style="list-style-type: none"> • Session • Persistent 	User name	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere QualityStage Standardization Rules Designer	InfoSphere Information Server web console	<ul style="list-style-type: none"> • Session • Persistent 	User name	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere Information Governance Catalog		<ul style="list-style-type: none"> • Session • Persistent 	<ul style="list-style-type: none"> • User name • Internal identifiers • State of the tree 	<ul style="list-style-type: none"> • Session management • Authentication • Single sign-on configuration 	Cannot be disabled

Table 1: Use of cookies by InfoSphere Information Server products and components. (continued)

Product module	Component or feature	Type of cookie that is used	Collect this data	Purpose of data	Disabling the cookies
InfoSphere Information Analyzer	Data Rules stage in the InfoSphere DataStage and QualityStage Designer client	Session	Session ID	Session management	Cannot be disabled

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be

trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The United States Postal Service owns the following trademarks: CASS, CASS Certified, DPV, LACSLink, ZIP, ZIP + 4, ZIP Code, Post Office, Postal Service, USPS and United States Postal Service. IBM Corporation is a non-exclusive DPV and LACSLink licensee of the United States Postal Service.

Other company, product or service names may be trademarks or service marks of others.

