

IBM InfoSphere Information Server on Cloud
Version 1 Release 5

User Guide
(Last updated: 2018-05-06)



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 5, modification 0 of IBM InfoSphere Information Server on Cloud 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. Information Server on Cloud offerings.....	1
Chapter 2. Information Server on Cloud Enterprise Edition.....	5
Overview.....	5
Available configurations.....	7
Layout of IBM Information Server on Cloud Enterprise Edition server and client disks.....	8
Information Server Enterprise Edition on Cloud High Availability.....	10
Technologies and Concepts.....	10
Failure Scenarios.....	12
Backing up IIS Enterprise Edition on Cloud components.....	14
Available configurations.....	14
Available configurations of Information Server on Cloud Enterprise Edition High Availability.....	16
Layout of IBM Information Server on Cloud Enterprise Edition High Availability server and client disks.....	19
Manual Take over by Information Server on Cloud Enterprise Edition Engine Tier on the passive machine.....	22
Information Server on Cloud High Availability Components.....	22
Information roadmap.....	23
Product Overview.....	23
Installing.....	24
Administering.....	24
Managing Metadata.....	24
Getting Started with Consoles.....	24
Connecting to data sources.....	25
New in Information Server 11.7.....	25
Getting started and using IBM Information Server on Cloud Enterprise Edition.....	25
When you connect for the first time, follow these steps:.....	25
After the initial connection, you can do any of the following tasks:.....	27
Backing up IIS Enterprise Edition on Cloud components.....	27
Spectrum Protect setup.....	28
Getting started with IBM Spectrum Protect Operations Center console.....	29
Starting Command Builder.....	31
Retention Policy.....	32
Configuring Object storage.....	34
Limitations and best practices.....	35
IBM Spectrum Protect setup for IIS on Cloud Enterprise Edition Premium service.....	36
IBM Spectrum Protect setup for IIS on Cloud Enterprise Edition High Availability service.....	43
Schedules and Policies.....	57
Protecting the master encryption key.....	58
Spectrum Protect server database backup.....	58
Spectrum Protect server Inventory Expiration.....	59
Restoring Backups.....	60
Chapter 3. Information Server on Cloud Data Quality.....	67
Overview.....	67
Available configurations.....	69
Layout of IBM Information Server on Cloud Data Quality server and client disks.....	70
Information roadmap.....	72
IBM InfoSphere DataStage and QualityStage.....	72
IBM InfoSphere Information Analyzer.....	73

IBM InfoSphere Information Governance Catalog.....	73
New in Information Server 11.7.....	74
Getting started and using IBM Information Server on Cloud Data Quality.....	74
When you connect for the first time, follow these steps:.....	74
After the initial connection, you can do any of the following tasks:.....	75
Chapter 4. DataStage on Cloud.....	77
Overview.....	77
Available configurations.....	78
Layout of IBM DataStage on Cloud server and client disks.....	79
DataStage on Cloud High Availability.....	80
Technologies and Concepts.....	80
Failure Scenarios.....	82
Available configurations.....	84
Available configurations of DataStage on Cloud HA.....	85
Layout of IBM DataStage on Cloud High Availability server and client disks.....	86
Information roadmap.....	88
Product overview.....	88
Getting started.....	88
Using InfoSphere DataStage.....	89
Troubleshooting and support.....	89
Getting started and using IBM DataStage on Cloud.....	89
When you connect for the first time, follow these steps:.....	90
After the initial connection, you can do any of the following tasks:.....	91
Backing up IIS DataStage on Cloud components.....	91
Spectrum Protect setup.....	91
Getting started with IBM Spectrum Protect Operations Center console.....	93
Starting Command Builder.....	94
Retention Policy.....	95
Configuring Object storage.....	97
Limitations and best practices.....	98
IBM Spectrum Protect setup for IIS DataStage on Cloud Premium service.....	99
IIS Engine machine (main).....	99
IIS Compute machine (cmpt).....	103
IBM Spectrum Protect setup for IIS on Cloud Enterprise Edition High Availability service.....	104
IIS Active Engine machine (main).....	105
IIS Passive Engine machine (mainp).....	109
IIS Compute machine (cmpt).....	113
Add new schedulers.....	114
Update existing schedulers.....	115
Create new policies and domain.....	115
Start Spectrum Protect after server restart or reboot.....	115
Protecting the master encryption key.....	115
Spectrum Protect server database backup.....	116
Spectrum Protect server Inventory Expiration.....	116
Restoring Backups.....	117
Restoring DB2 database.....	118
Restoring files and directories.....	119
Restoring WAS artifacts.....	120
Restoring ISTOOL asset.....	121
Viewing and restoring multiple versions of a specific file.....	121
Restoring from Cloud Object Storage.....	122
Chapter 5. DataStage on Cloud Designer Client.....	125
Overview.....	125
Available configuration.....	126
Layout of IBM® DataStage® on Cloud Designer Client disks.....	126

Information roadmap.....	127
Product overview.....	127
Getting started.....	127
Using InfoSphere DataStage.....	127
Troubleshooting and support.....	128
Getting started and using IBM DataStage on Cloud Designer Client.....	128
When you connect for the first time, follow these steps:.....	128
After the initial connection, you can do any of the following tasks:.....	129
Adding and connecting extra DataStage on Cloud Designer Client machines.....	129
Chapter 6. Information Governance Catalog on Cloud.....	131
Overview.....	131
Available configurations.....	133
Layout of IBM Information Governance Catalog on Cloud server and client disks.....	134
Information roadmap.....	135
New in Information Server 11.7.....	136
Getting started and using IBM Information Governance Catalog on Cloud.....	136
When you connect for the first time, follow these steps:.....	137
After the initial connection, you can do any of the following tasks:.....	137
Chapter 7. Connecting to other systems.....	139
Connecting to the IBM InfoSphere DataStage and QualityStage Designer client.....	139
Connecting to an on-premises computer.....	140
Connecting to IBM dashDB.....	140
Example.....	141
Connecting on an on-premises DB2 database instance.....	142
Example.....	143
Chapter 8. Administering cloud offerings on the server.....	145
Open ports on server and client machines.....	145
Open ports for incoming traffic on the server machine.....	145
Open ports for incoming traffic on client machines.....	146
Starting services after the Information Server on Cloud server machine reboots.....	147
Setting up firewall security.....	148
About these tasks.....	148
To change the security level or to manage rules for the Information Server on Cloud server firewall.....	148
To show the Microsoft Windows firewall profiles on the Information Server on Cloud client.....	149
To show a list of all open ports on the Information Server on Cloud client.....	150
To block an open port on the Information Server on Cloud client.....	150
Scripts to change the iptables firewall settings.....	150
Enhancing security of Information Server on Cloud computers.....	151
Setting up SSH keys.....	151
Allowing SSH access to specific IP addresses.....	152
Managing LUKS keys on the IBM Information Server on Cloud server.....	152
Chapter 9. Troubleshooting cloud offerings.....	155
Cannot install a patch on services tier machine.....	155
Symptoms.....	155
Resolving the problem.....	155
Notices.....	157

Chapter 1. Information Server on Cloud offerings

IBM® Information Server on Cloud offerings provide data integration and governance products as a service on the global cloud infrastructure of IBM SoftLayer®. The rich features of on-premises IBM InfoSphere® Information Server components are provided without the cost and complexity of deploying the infrastructure. Optional add-on services can also be added to maintain and manage the infrastructure.

Information Server on Cloud helps to reduce the time that is required to provision and deploy data integration and governance. As a result, your IT resources are free to innovate and develop new solutions.

Information Server on Cloud is based on Information Server version 11.7 and includes the following offerings:

IBM Information Server on Cloud Enterprise Edition

Information Server on Cloud Enterprise Edition provides end-to-end information integration capabilities to help you understand, govern, create, maintain, transform, and deliver quality data.

IBM Information Server on Cloud Data Quality

Information Server on Cloud Data Quality cleanses data and monitors data quality.

IBM DataStage® on Cloud

DataStage on Cloud provides a framework to design and run jobs that transform and cleanse your data.

IBM® DataStage® on Cloud Designer Client

DataStage on Cloud Designer Client provides a framework to create, design, and develop DataStage jobs. This offering includes all IBM InfoSphere Information Server clients. As a result, you can connect to the Information Server on Cloud server and work with the clients.

IBM Information Governance Catalog on Cloud

Information Governance Catalog on Cloud provides a standardized approach to discover and govern your business assets, and to define a common business language.

In the following table, a check mark (✓) indicates that the component is available with the listed offering. An asterisk (*) next to a check mark indicates limits on the product availability. Check the license agreement for details.

Table 1: Available components per offering

InfoSphere Information Server component	Information Server on Cloud Enterprise Edition	Information Server on Cloud Data Quality	DataStage on Cloud	DataStage on Cloud Designer Client	Information Governance Catalog on Cloud
InfoSphere DataStage	✓		✓		✓*
InfoSphere QualityStage®	✓	✓			✓
InfoSphere DataStage and QualityStage Designer	✓	✓	✓	✓	✓*
InfoSphere DataStage and QualityStage Director	✓	✓	✓	✓	✓

Table 1: Available components per offering (continued)

InfoSphere Information Server component	Information Server on Cloud Enterprise Edition	Information Server on Cloud Data Quality	DataStage on Cloud	DataStage on Cloud Designer Client	Information Governance Catalog on Cloud
InfoSphere Data Click	√			InfoSphere Data Click client only	
InfoSphere FastTrack	√			InfoSphere FastTrack client only	
InfoSphere Information Analyzer	InfoSphere Information Analyzer server	InfoSphere Information Analyzer server		InfoSphere Information Analyzer client only	
InfoSphere Information Governance Catalog	√	√*		InfoSphere Information Governance Catalog client only	√
IBM Glossary Anywhere	√	√	√	Glossary Anywhere installer only	√
InfoSphere Information Governance Dashboard	√	√	√	√	√
InfoSphere Information Services Director	√	√	√	InfoSphere Information Services Director client only	
InfoSphere Metadata Integration Bridges and the metadata interchange agent	√	√	√	√	√
InfoSphere Metadata Asset Manager	√	√	√	√	√
InfoSphere Information Server istool command-line utility	√	√	√	√	√

Table 1: Available components per offering (continued)

InfoSphere Information Server component	Information Server on Cloud Enterprise Edition	Information Server on Cloud Data Quality	DataStage on Cloud	DataStage on Cloud Designer Client	Information Governance Catalog on Cloud
InfoSphere Information Server Manager client, Multi-Client Manager	√	√	√	√	√
IBM Business Process Manager Standard	√			IBM Business Process Manager Designer Client	
IBM Cognos® Business Intelligence	√	√			
IBM InfoSphere Data Architect	√			√	√

Chapter 2. Information Server on Cloud Enterprise Edition

IBM® Information Server on Cloud Enterprise Edition provides a hosted environment that you configure and control. You can use Information Server on Cloud Enterprise Edition to extend the reach of your business by leveraging cloud offerings, while you reduce the costs that are associated with providing these services. Different plans are available so that any size business can access the powerful and scalable platform by IBM.

Overview

IBM® Information Server on Cloud Enterprise Edition is a data integration software platform that helps organizations derive more value from the complex, heterogeneous information spread across their systems. Information Server on Cloud Enterprise Edition provides all of the functions of its on-premises counterpart, IBM InfoSphere® Information Server.

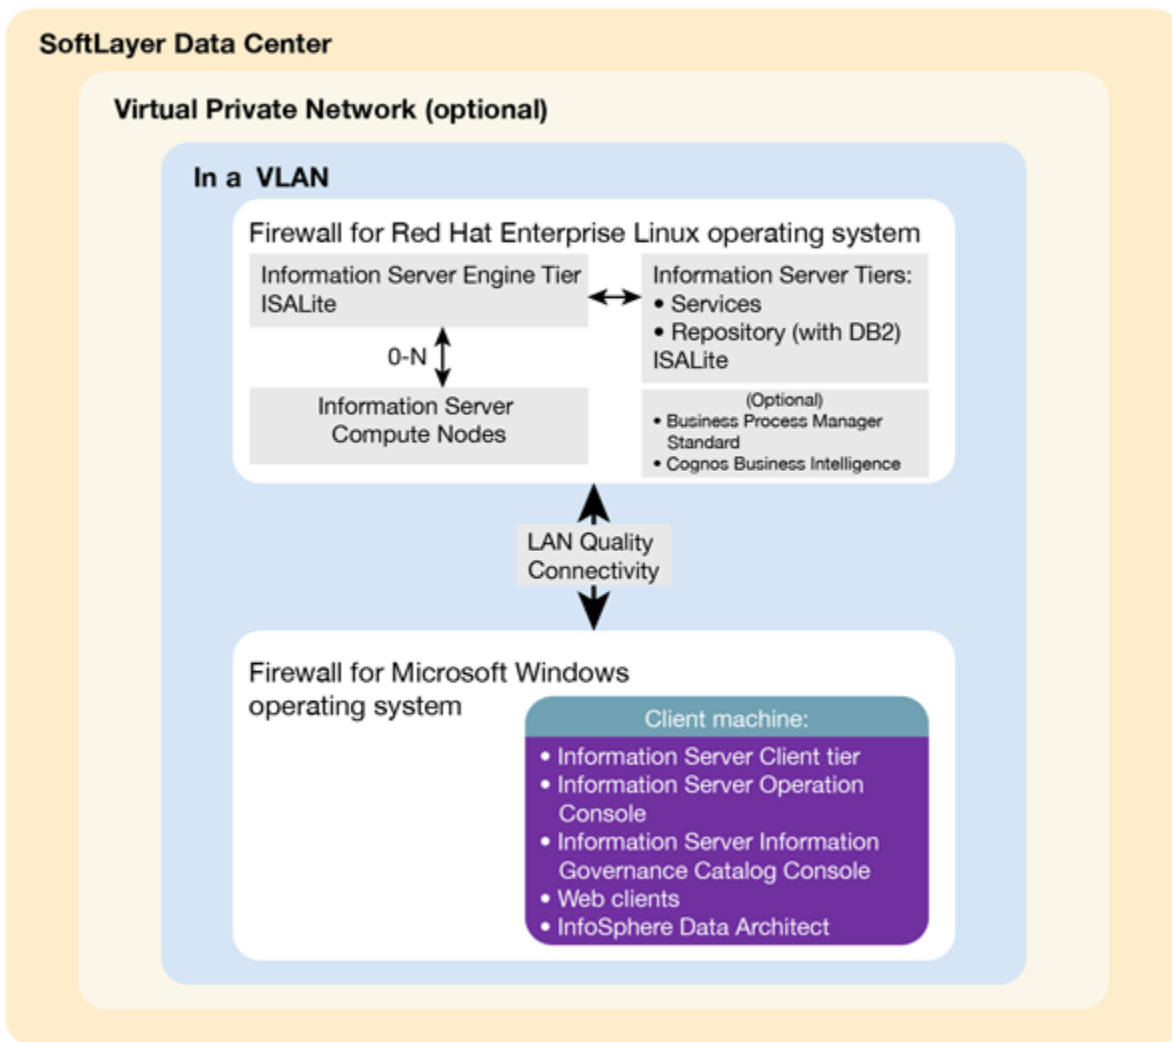
Information Server on Cloud Enterprise Edition includes the following key features:

- Removes the burden of deployment so that you can focus on core development of integration and governance assets
- Integrates and governs products without the need for an on-premises infrastructure
- Reduces time to value with enterprise-grade features
- Optimizes performance and reliability by using a virtual machine or a dedicated bare metal server
- Keeps infrastructure costs in line with the changing needs of the business
- Provides flexibility and agility to meet dynamic business needs
- Works with cloud and on-premises data sources
- Scales according to your business demands

Information Server on Cloud Enterprise Edition uses the characteristics of software-as-a-service (SaaS). You select the plan size based on your needs. IBM provisions the machines and deploys the Information Server on Cloud Enterprise Edition software.

Overview of Information Server on Cloud Enterprise Edition with High Availability

The following figure shows the topology of the server and client machines in a typical deployment without High Availability and backup Configurations.



As a hosted offering, you have the same control over your data in the cloud as in the on-premises system:

- Actively monitor and report any issues that you encounter with IBM Software as a Service (SaaS).
- Maintain the software platform of your cloud offering and the operating system to meet your security standards.
- Maintain software firewalls on servers that face the internet in a manner to provide required protection.
- Develop parallel jobs to transform and cleanse data, and develop server jobs to transform data. Establish connectivity between data sources and applications. Develop your own workload, business rules, monitoring, and scheduling for all jobs. You are responsible for the quality and performance of programs, applications, and jobs that you develop.
- Ensure the continuity, compatibility, and performance of the IBM SaaS platform by installing only permissible software, including any open source packages.
- Regularly upgrade the environment and operating system of your cloud offering.
- Create and maintain regular backups of data.
- Create and maintain high availability configurations.

The following managed add-on services are available to maintain and manage the infrastructure:

Jump start

This setup service provides up to 50 hours of remote consulting time for startup activities.

Accelerator

This setup service provides up to 50 hours of remote consulting time to perform various scoped activities.

Silver

This service provides monthly remote consulting time for operations and maintenance activities.

Gold

This service provides monthly remote consulting time for operations and maintenance activities. The service includes everything that is provided by the Silver service and delivers extra activities.

Note the following limitations and restrictions of Information Server on Cloud Enterprise Edition:

- If your offering is designated as "Non-Production", Information Server on Cloud Enterprise Edition can be deployed only as part of your development and test environments for internal non-production activities. These activities include, but are not limited to: testing, performance tuning, fault diagnosis, internal benchmarking, staging, quality assurance activity, developing internally used additions or extensions to the offering by using published application programming interfaces.
- Users must not modify the configuration file that is needed to run the job for parallel processing.

Available configurations

IBM® Information Server on Cloud Enterprise Edition servers for the small and medium plans are virtual servers with dedicated CPUs. The servers in the large plan are in a bare metal environment.

Select the offering plan that fits your usage and environment needs.

Table 2: Offering sizes: small production and non-production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
Service metadata	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
Enterprise Search	32	8	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
(Optional) BPM, Cognos®	16	4	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN

Table 3: Offering sizes: medium production and non-production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN
Service metadata	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN
Enterprise Search	64	16	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	1TB SAN

Table 3: Offering sizes: medium production and non-production (continued)

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
(Optional) BPM, Cognos	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN

Table 4: Offering sizes: large production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	64	12	1 Gbps with 5000 GB bandwidth	1.7 TB SSD	1.7 TB SSD
Service metadata	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
Enterprise Search	64	12	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
(Optional) BPM, Cognos	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD

Small and medium offerings include the following configuration:

- One virtual server machine with services and repository tiers
- One virtual server machine with the engine tier
- One virtual server machine with the Enterprise Search Node
- Optional: One virtual server machine with IBM Business Process Manager Standard, and IBM Cognos Business Intelligence.
- One client machine for small offerings and three client machines for medium offerings, including IBM InfoSphere® Data Architect.

The number of client machines is based on the number of concurrent users. For small offerings, two concurrent users are allowed. For medium offerings, five concurrent users are allowed. By default, Microsoft Windows operating system allows two concurrent users to access the machine by using Remote Desktop Connection.

Large-size offering includes the following configuration:

- Three bare metal machines with IBM Information Server on Cloud Enterprise Edition. One machine has the services and repository tiers, 2nd machine has the engine tier and the 3rd machine has Enterprise Search Node.
- Optional: One bare metal machine with IBM Business Process Manager Standard, and IBM Cognos Business Intelligence.
- Five client machines based on 10 concurrent users, including IBM InfoSphere Data Architect.

Layout of IBM Information Server on Cloud Enterprise Edition server and client disks

The layout of the Information Server on Cloud Enterprise Edition server and client disks depends on the plan size of your system.

Virtual servers for small and medium plans

Information Server on Cloud Enterprise Edition comes with two virtual servers. One server has the services and repository tiers and the other server has the engine tier. Optionally, you have a third virtual server with IBM® Business Process Manager Standard, and IBM Cognos® Business Intelligence.

The small and medium plans come with two Storage Area Network (SAN) disks. The Red Hat Enterprise Linux operating system is on the first SAN disk in both the small and medium plans. The second SAN disk is encrypted by using Linux Unified Key Setup (LUKS).

The encryption key details are provided in the Welcome letter from the IBM Sales Representative. It is recommended that you add your own key and remove the supplied key before you use the system.

The product is installed on the /opt directory. User data can be stored on /data directory. Both directories are on the partition /dev/xvdc1 that is encrypted.

Table 5: Small and medium disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	LUKS	/dev/xvdc1	500 GB for small plan. 1000 GB for medium plan	Primary disk is /dev/xvdc	Product installation. Directories /data and /opt are created on this disk.
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging

Bare metal server for large plan

The large plan comes with two bare metal machines, and optionally a third one with IBM Business Process Manager Standard and IBM Cognos Business Intelligence. RAID level 1 implementation makes them appear as a single disk.

The disk is divided into four partitions. The Red Hat Enterprise Linux operating system is on a 10 GB partition. The boot data is on a 256 MB partition. The swap space is on a 2 GB partition. The remaining space is on another partition that is encrypted by using LUKS.

Table 6: Bare metal disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	256 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda2	2 GB	/dev/sda	Swap space for paging
/	None	/dev/sda3	10 GB	/dev/sda	Operating system data

Table 6: Bare metal disk layout (continued)

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/disk1	LUKS	/dev/sda5	About 900 GB for the services tier, about 1.6 TB for the engine tier, and optionally about 900 GB for BPM and Cognos	/dev/sda	Product installation. Directories / data, / installables, and /opt are created on this disk.

Client for all plans

The Information Server on Cloud Enterprise Edition client machine configuration is the same for all plans sizes. The client machine has two Storage Area Network (SAN) disks that are 100 GB each. One disk is drive C for the Microsoft Windows operating system. The other disk is drive F, and it is an empty disk.

Information Server Enterprise Edition on Cloud High Availability

IBM® Information Server on Cloud Enterprise Edition is now available with High Availability in small, medium, and large sizes. As part of the High Availability offerings, additional machines are provided in the same data centre and VLAN to act as passive instances that can take over in case of failures to the active instances. The passive instances have pre-installed and configured parts of the application. Clustering is used between the active and passive instances for IBM WebSphere® Application Server. HADR is configured between the active and passive instances IBM DB2® in active-passive mode. And, shared storage is used between them for the Information Server Engine Tier High Availability setup. [The High Availability offering also comes with backup functionality configured.](#)

Technologies and Concepts

HADR

Configuration of two instances of a software or hardware component for High Availability and Disaster Recovery. In the active-passive mode that is used for HADR configuration between the two instances of DB2, the passive instance keeps getting logs from the active one and apply the transactions on it's local database. So, the passive instance would be ready to take over whenever the active instance can't serve it's purpose.

Automatic Client Reroute

This term refers to the configuration of data sources to connect to the available database server in case the primary one goes down.

Portable IP

Portable IP is an IP that can be assigned to any of the active and passive instances and is initially assigned to the active instance or server.

Pacemaker® (PM)

A High availability software that helps detect a failure to an instance of software component and automatic assignment of a portable IP to another machine.

Information Server Enterprise Edition on Cloud High Availability comes with four server machines two active and two passive, and a number of client machines which is dependent on the [size of offering](#) chosen.

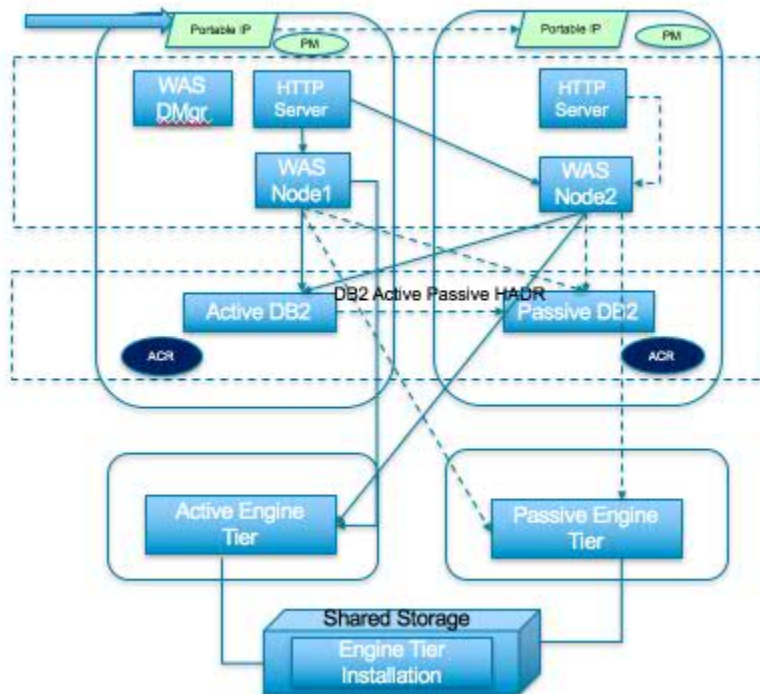
Information Server Engine Tier

Two servers act as active and passive Engine Tier machines, having the Engine Tier installed on a storage volume shared by both of them. A portable IP, also called virtual IP is initially assigned to active machine. A hostname called virtual Engine Tier hostname is associated to this virtual IP via /etc/hosts. The Information Server Engine Tier installation is done with this virtual hostname that helps us be able to bring Engine Tier on the passive machine when the active machine goes down.

Information Server Repository and Services Tiers

The other two servers have IBM DB2® and IBM WebSphere® Application Server installed on each of them, with one of them having the Repository and Services Tiers installed in the non-shared storage. Another virtual IP is assigned to the first of these machines called active machine, and another virtual host name associated with this virtual IP is used for the Information Server installation Repository and Services Tier installation. The other server that may be assigned the virtual IP when the active machine goes down is called the passive machine.

A WebSphere Application Server cluster is built using the instances of WebSphere installed on both the machines, and DB2® HADR in active passive mode setup between the DB2® instances on them. All the WebSphere Application Server data sources are configured for Automatic Client Rerouting so that when an active instance fails or is not reachable, the transactions are automatically redirected to passive instance of DB2. All the default databases that come with the product, DSODB, ESDBDB2, IADB, and XMETA are configured for HADR. Both of the server machines have HTTP Server installed and configured on them. HTTP Server takes care of distributing the load between the WebSphere Application Server nodes in a round robin fashion. When the active machine goes down, the passive machine gets assigned the virtual IP and all the HTTP requests using the virtual IP or the associated virtual host name get diverted to the passive machine automatically.



WebSphere Application Server on the passive server machine has a WebSphere node and HTTP Server installed while active server machine also has WebSphere Deployment Manager installed along with these. The HTTP Server takes care of distributing the load between the WebSphere nodes in a round robin fashion.

All the WebSphere data sources are configured for Automatic Client Rerouting so that when an active instance fails or is not reachable, the transactions are automatically redirected to passive instance of DB2. All the default databases that come with the product, DSODB, ESDBDB2, IADB and XMETA are configured for HADR.

The passive Engine Tier on the passive server machine can be brought up with a few manual steps executed on it if the active Engine Tier fails as it is installed on the shared storage.

The dotted lines in the above diagram show the connections to the passive instances of different components while the hard lines show connections to the active instances.

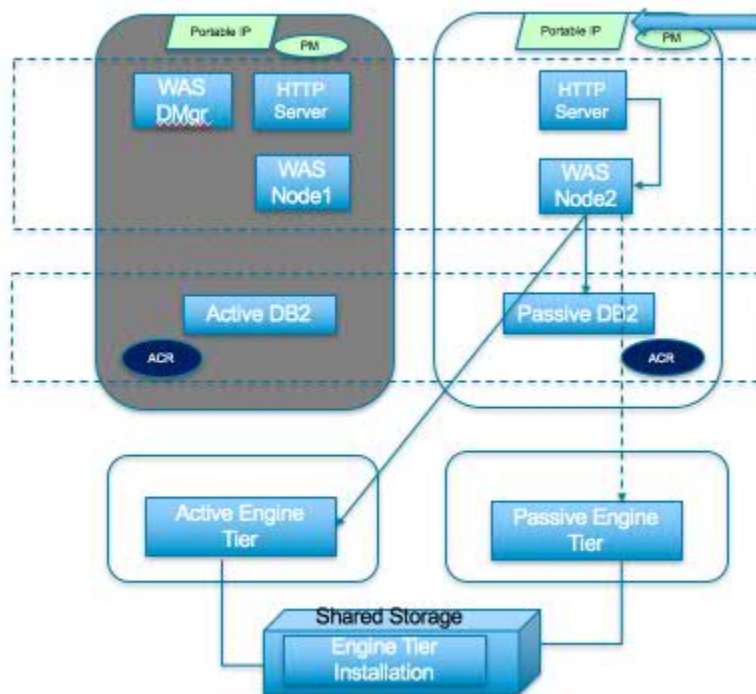
The high availability software Pacemaker installed on both the servers helps switch the IP to the passive instances whenever the active Services Tier machine goes down. Both public and private virtual IPs are used. These IPs should be used when integrating with other applications for making sure that High Availability implementation works with the integration.

Communication between different components or resources in High Availability implementation is explained here.

While on the failures of the some of the components the passive instances take over automatically, some would need manual steps for the take over by the corresponding passive instance to happen. The following scenarios discuss some of the failure cases and the take over procedures in such cases.

Failure Scenarios

Active Services Tier machine fails



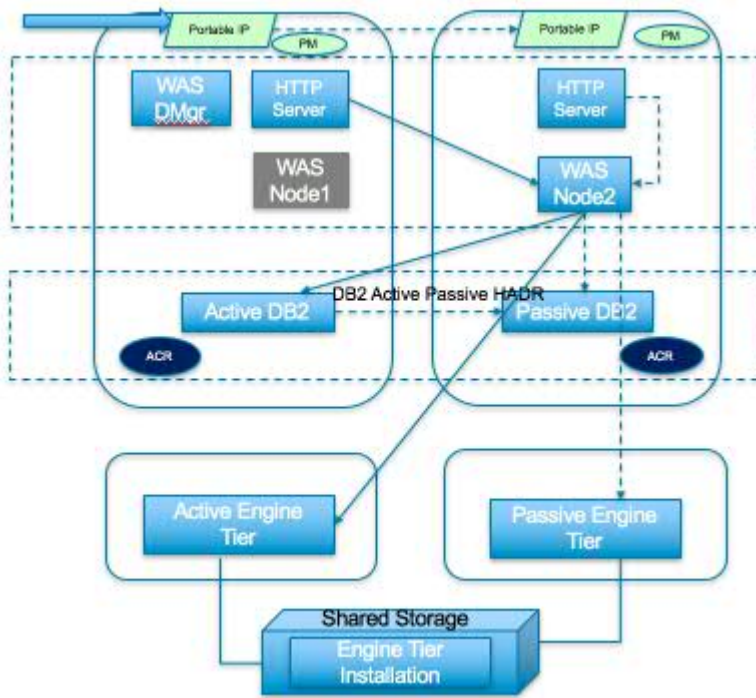
Failure: The machine in dark grey colour has failed all the communication to that machine is stopped.

Take over procedure: Pacemaker detects that failure and assigns the portable IP to the passive machine. HTTP Server on the passive instance receives all the HTTP requests, and sends to the only available WebSphere node, Node2.

DB2 take over steps should be performed on the passive instance of DB2 running on passive Repository and Services Tier machine. Take over procedure is described in the failure case **Active DB2 fails** below.

db2 takeover hadr on database <database_name> by force

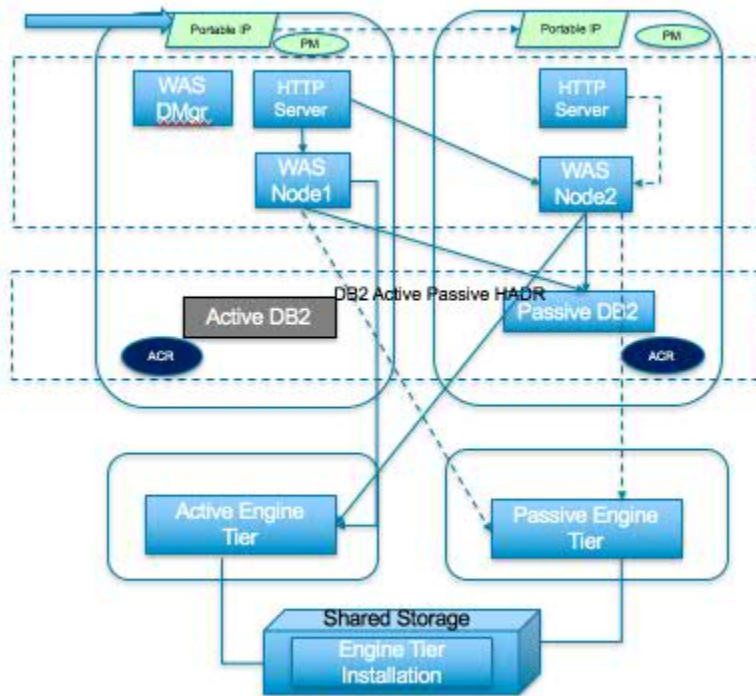
WebSphere Node1 fails



Failure: The WebSphere node, Node 1 shown in dark grey colour goes away, and all the connections between this and other processes are broken.

Take over procedure: HTTP Server detects that one of the WebSphere nodes is not available and directs all requests to the available WebSphere node, Node2 automatically.

Active DB2 fails



Failure: Active DB2 shown in the dark grey colour goes away, and all database transactions fail.

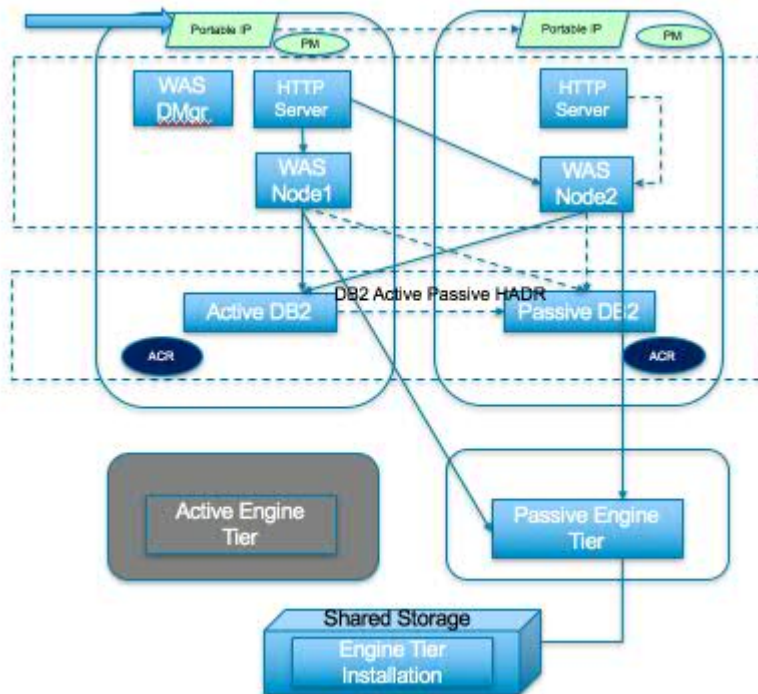
Take over procedure: Administrator should run take over commands on the passive DB2. Automatic client re-route feature makes redirect all transactions to passive DB2 automatically.

Steps for DB2 take over on the passive machine:

- Login to the Services Tier passive machine as the user db2inst1
- Change to the home directory of the user, db2inst1 by running the following command: `cd /home_/db2inst1`
- Execute DB2 profile on the current shell: `. sqllib/db2profile`
- Run the DB2 take over command for taking over on the passive/standby machine: `db2 takeover hadr on database <db_name>` . For example to take over XMETA on the passive machine, the command is "db2 takeover hadr on database xmeta".

Update the database URL in the file `/opt/IBM/InformationServer/Server/DSODB/DSODDBConnect.cfg` to point to the host name of the new machine that is running DB2 Primary.

Active Engine Tier fails



Failure: DataStage jobs failing due to active Engine Tier not being available.

Take over procedure: The passive Engine Tier machine is also mounted with the same storage containing the Engine Tier and is ready to take over. Run the scripts and steps provided along with the machines for take over on the passive Engine Tier machine.

Backing up IIS Enterprise Edition on Cloud components

IBM Information Server on Cloud Premium and High Availability services provides software and hardware infrastructure for taking backups. Backup capability is based on IBM Spectrum Protect version 8.1.3 product (formerly known as Tivoli Storage Manager). One dedicated server machine with the installation of IBM Spectrum Protect server version 8.1 is provided with each deployment of Information Server on Cloud Premium and High Availability service.

More information on the backup process is given in the [Backup Section](#)

Available configurations

IBM® Information Server on Cloud Enterprise Edition servers for the small and medium plans are virtual servers with dedicated CPUs. The servers in the large plan are in a bare metal environment.

Select the offering plan that fits your usage and environment needs.

Table 7: Offering sizes: small production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk	Performance disk
Engine	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN	
Service metadata	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN	
Enterprise Search	32	8	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN	1TB 4000 IOPS
Backup Server	64	8	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	4 TB SAN	
(Optional) BPM, Cognos®	16	4	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN	

Table 8: Offering sizes: medium production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk	Performance disk
Engine	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN	
Service metadata	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN	
Enterprise Search	64	16	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	1TB SAN	1TB 4000 IOPS
Backup Server	64	8	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	4 TB SAN	
(Optional) BPM, Cognos	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN	

Table 9: Offering sizes: large production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	64	12	1 Gbps with 5000 GB bandwidth	1.7 TB SSD	1.7 TB SSD
Service metadata	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
Enterprise Search	64	12	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
Backup Server	128	12	1 Gbps with 5000 GB bandwidth	6.8TB SSD	8 TB SATA
(Optional) BPM, Cognos	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD

Small and medium offerings include the following configuration:

- One virtual server machine with services and repository tiers
- One virtual server machine with the engine tier
- One virtual server machine with the Enterprise Search Node
- One virtual server machine used for storing backup artifacts
- Optional: One virtual server machine with IBM Business Process Manager Standard, and IBM Cognos Business Intelligence.
- One client machine for small offerings and three client machines for medium offerings, including IBM InfoSphere® Data Architect.

The number of client machines is based on the number of concurrent users. For small offerings, two concurrent users are allowed. For medium offerings, five concurrent users are allowed. By default, Microsoft Windows operating system allows two concurrent users to access the machine by using Remote Desktop Connection.

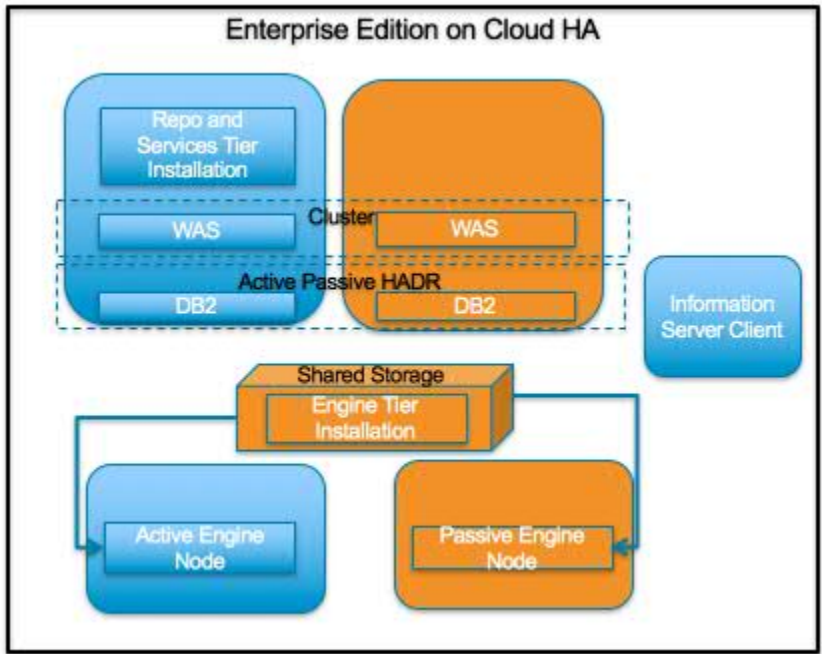
Large-size offering includes the following configuration:

- Two bare metal machines with IBM Information Server on Cloud Enterprise Edition. One machine has the services and repository tiers while the other machine has the engine tier.
- One bare metal machine used for storing backup artifacts
- One bare metal machine with the Enterprise Search Node.
- Optional: One bare metal machine with IBM Business Process Manager Standard, and IBM Cognos Business Intelligence.
- Five client machines based on 10 concurrent users, including IBM InfoSphere Data Architect.

Available configurations of Information Server on Cloud Enterprise Edition High Availability

High Level Model of Information Server Enterprise Edition High Availability

The High Availability offering of IBM® Information Server on Cloud Enterprise Edition will have four server machines as shown in the below diagram, along with one or more client machines depending on the size chosen.



Virtual servers with dedicated CPUs are provided for small and medium plans of IBM Information Server on Cloud Enterprise Edition High Availability. The servers in the large plan are in a bare metal environment. Each of the plans below are given with four server machines.

Select the offering plan that fits your usage and environment needs.

Small Production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine Tier Active and Passive machines	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
Service metadata Active and Passive machines	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
Enterprise Search	32	8	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
(Optional) BPM, Cognos®	16	4	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN

- Engine Tier machines are also mounted with 500 GB Shared File Storage.

Medium production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine Active and Passive machines	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN
Service metadata Active and Passive machines	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN
Enterprise Search	64	16	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	1TB SAN
(Optional) BPM, Cognos	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN

- Engine Tier machines are also mounted with 1000 GB Shared File Storage.

Large production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine Active and Passive machines	64	12	1 Gbps with 5000 GB bandwidth	1.7 TB SSD	1.7 TB SSD
Service metadata Active and Passive machines	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
Enterprise Search	64	12	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
(Optional) BPM, Cognos	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD

- Engine Tier machines are also mounted with 1000 GB Shared File Storage.

Small and medium offerings include the following configuration:

- Four virtual server machines are provided IBM Information Server on Cloud Enterprise Edition High Availability.
- One virtual server machine with services and repository tiers along with IBM WebSphere® Application Server and IBM® DB2®, another virtual server machine with WebSphere Application Server and DB2 alone.
- Two virtual servers one active and one passive, that shared the same storage where Information Server engine tier is installed.
- One virtual server with the Enterprise Search Node.
- Optional: One virtual server machine with IBM Business Process Manager Standard, and IBM Cognos Business Intelligence.

- One client machine for small offerings and three client machines for medium offerings, including IBM InfoSphere® Data Architect.

The number of client machines is based on the number of concurrent users. For small offerings, two concurrent users are allowed. For medium offerings, five concurrent users are allowed. By default, Microsoft Windows operating system allows two concurrent users to access the machine by using Remote Desktop Connection.

Large-size offering includes the following configuration:

- Five bare metal machines are provided with IBM Information Server on Cloud Enterprise Edition High Availability.
- One machine has the services and repository tiers along with WebSphere Application Server and DB2, and another machine with WebSphere Application Server and DB2 alone.
- One machine with the engine tier, another machine having the same shared storage mounted where the former installed engine tier so that it can take over when needed..
- One machine with the Enterprise Search Node.
- Optional: One bare metal machine with IBM Business Process Manager Standard, and IBM Cognos Business Intelligence.
- Five client machines based on 10 concurrent users, including IBM InfoSphere Data Architect.

Layout of IBM Information Server on Cloud Enterprise Edition High Availability server and client disks

The layout of the Information Server on Cloud Enterprise Edition server and client disks depends on the plan size of your system.

Virtual servers for small and medium plans

Information Server on Cloud Enterprise Edition comes with five virtual servers. Two servers implement the active and passive instances of WebSphere Application Server and DB2. The active one has services and repository tiers. The other couple of servers have the engine tier installed on storage shared among them. The other server hosts Information Server Enterprise Search node. Optionally, you have another virtual server with IBM® Business Process Manager Standard, and IBM Cognos® Business Intelligence.

The small and medium plans come with two Storage Area Network (SAN) disks on all the servers while the engine tier machines also has File Endurance storage used as shared storage. The Red Hat Enterprise Linux operating system is on the first SAN disk in both the small and medium plans. The second SAN disk is encrypted by using Linux Unified Key Setup (LUKS).

The encryption key details are provided in the Welcome letter from the IBM Sales Representative. It is recommended that you add your own key and remove the supplied key before you use the system.

The Information Server product tiers, DB2, and WebSphere Application Server are all installed on the /opt directory. User data can be stored on /data directory. /opt is on the shared storage on the engine tier machines, and on encrypted second SAN disk in the services tier machines.

Small and medium services tier and enterprise search node's disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	LUKS	/dev/xvdc1	500 GB for small plan. 1000 GB for medium plan	Primary disk is /dev/xvdc	Directories / data and /opt are created on this disk. Products DB2 and WebSphere Application Server are installed under /opt, along with the Information Server repository and service tiers.

Small and medium engine tier machines disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	Provided by Softlayer	Dynamic Name	500 GB for small plan. 1000 GB for medium plan	Dynamic Name	Directories / data and /opt are created on this disk. Information Server engine tier is installed under /opt.
/disk2	LUKS	/dev/xvdc1	500 GB for small plan. 1000 GB for medium plan	/dev/xvdc	

Bare metal server for large plan

The large plan comes with two bare metal machines, and optionally a third one with IBM Business Process Manager Standard and IBM Cognos Business Intelligence. RAID level 1 implementation makes them appear as a single disk.

All the bare metal machines have a non-shared disk that is divided into four partitions. The Red Hat Enterprise Linux operating system is on a 10 GB partition. The boot data is on a 256 MB partition. The swap space is on a 2 GB partition. The remaining space is on another partition that is encrypted by using LUKS.

Bare metal services tier and enterprise search node's disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	256 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda2	2 GB	/dev/sda	Swap space for paging
/	None	/dev/sda3	10 GB	/dev/sda	Operating system data
/disk1	LUKS	/dev/sda5	About 900	Primary disk is /dev/xvdc	Directories / data and /opt are created on this disk. Products DB2 and WebSphere Application Server are installed under /opt, along with the Information Server repository and service tiers.

Bare metal engine tier disk layout

Along with the non-shared disk, the engine tier machines also have a shared storage for installing engine tier.

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	256 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda2	2 GB	/dev/sda	Swap space for paging
/	None	/dev/sda3	10 GB	/dev/sda	Operating system data
/disk1	Provided by Softlayer	Dynamic Name	About 1000 GB	Dynamic name	Directories / data and /opt are created on this disk. Information Server engine tier is installed under /opt.
/disk2	LUKS	/dev/sda5	About 1.6 TB	/dev/sda	

Client for all plans

The Information Server on Cloud Enterprise Edition client machine configuration is the same for all plans sizes. The client machine has two Storage Area Network (SAN) disks that are 100 GB each. One disk is drive C for the Microsoft Windows operating system. The other disk is drive F, and it is an empty disk.

Manual Take over by Information Server on Cloud Enterprise Edition Engine Tier on the passive machine

Perform the following steps on different machines in the below order. Required credentials for logging into the machines provided with your order are available in the welcome letter.

Login to the Information Server Engine Tier passive machine which is provided with your order as root and perform the following steps.

1. Make sure that the database URL in the file `/opt/IBM/InformationServer/Server/DSODB/DSODBCConnect.cfg` points to the hostname of the machine running DB2 in PRIMARY mode. It initially does so. If there was a take over of DB2 to the other machine, it has to be changed to that machine's host name.
2. Run the command `/opt/IBM/ha_scripts/runSignerCerts.sh`
3. Start the engine processes on the passive machine by running `/opt/IBM/ha_scripts/startEngine.sh`
4. Run the below commands to first check the cluster status, switch the portable IP to the passive engine tier's machine, and check the status again. When the active engine tier machine goes down, Pacemaker automatically switches the portable IP to the passive machine, in such a case the script `switchPortableIP.sh` need not be run.
 - `pcs status`
 - `/opt/IBM/ha_scripts/switchPortableIP.sh`
 - `pcs status`
5. Check the status of engine on the passive machine, by running `/opt/IBM/ha_scripts/statusEngine.sh` . The output of this command should be showing as RUNNING for all the processes.


Reverse take over steps by the active engine tier

After fixing the issues with active engine tier, take over can be done on this machine to continue using it as active engine tier.

1. Login to the Information Server Engine Tier passive machine as root. Stop the engine tier processes on this machine if they are already not stopped. For stopping, run the command `/opt/IBM/ha_scripts/stopEngine.sh`
2. Login to the Information Server Engine Tier active machine as root. Repeat the steps mentioned above for take over on the passive machine.

Information Server on Cloud High Availability Components

IBM® Information Server on Cloud High Availability offerings come with different fail over components that can start serving the purpose of active instances, in case those active instances fail. This section describes how those components are linked so that they can be used with minimal or no interruptions to the applications.



The two dotted lined rectangular boxes show one set of components each, of the active and passive instances of them.

The Information Server on Cloud Client machine, named *IS Client* can communicate

- with the engine tier machine using the virtual hostname of the engine tier machine
- with HTTP Server using the virtual host name of services tier machine

while the virtual host names are associated to the corresponding portable IP in `/etc/hosts`.

HTTP Server communicates with the WebSphere® Application Server processes using the actual host names of the machine.

Engine Tier machine is linked to the WebSphere Application Server processes with the virtual host name while DB2 is linked using actual hostname of the machine it is running on.

HTTP Server communicates with the WebSphere Application Server using the actual host name.

External client machines access HTTP services using the portable IP.

Information roadmap

This roadmap lists information resources that are available for users who are new to the IBM® Information Server on Cloud Enterprise Edition products. These resources provide information about various subject areas.

Product Overview

- **Introduction to InfoSphere® Information Server** InfoSphere Information Server features an integrated set of product modules, or suite components, including the following: IBM InfoSphere Information Governance Catalog, IBM InfoSphere DataStage®, IBM InfoSphere FastTrack, IBM InfoSphere

Information Analyzer, IBM InfoSphere Information Services Director, and IBM InfoSphere QualityStage®.

- **Integration scenarios** Information integration is a complex activity that affects every part of an organization. To address the most common integration business problems, these integration scenarios show how you can deploy and use InfoSphere Information Server and the InfoSphere Foundation Tools components together in an integrated fashion. The integration scenarios focus on data quality within a data warehouse implementation.
- **Release notes** The release notes contain information that is important for the successful installation and use of the product.

Installing

- **Installing InfoSphere Information Serversoftware.** This section includes information about planning, preparing the target computers, installing, and configuring the software. It also explains how to create highly available, scalable configurations.
- **Troubleshooting installations.** This section provide descriptions of possible installation problems and the steps to correct them.
- **Migrating to InfoSphere Information Server, Version 11.5** This section describes how to migrate from existing installations to new InfoSphere Information Server, Version 11.5, installations.

Administering

- **Managing Logs** You can configure the log messages that are generated when activities run in the suite.
- **Managing Schedules** In the IBM InfoSphere Information Server Web console, you can query all of the schedules that are defined across all of the suite components, check their status, history, and forecast, perform maintenance tasks such as purging the schedule execution history, and stop or start existing schedules to prevent system overload.
- **Temporary file cleanup and database optimization** If you are concerned about disk space, you can occasionally clear some files and directories to lower your disk usage. You can also run database tools to ensure good repository database performance.

Managing Metadata

- **Importing and managing assets by using IBM InfoSphere Metadata Asset Manager** You can use bridges and connectors to import metadata into a staging area, where you can analyze and preview the contents of the import before you share it to the metadata repository. You can browse, search, and manage assets that are in the metadata repository. You can export database assets.
- **Exchanging metadata by using IBM InfoSphere Metadata Interchange Bridges** By using bridges you can import metadata from a wide range of sources, including IBM Cognos, SAP BusinessObjects, CA ERwin, and other tools. The prerequisites and import parameters for bridges are documented in detail in this section.
- **Common metadata assets** Common metadata assets are stored in the metadata repository and shared between tools in the InfoSphere™ Information Server suite. Common metadata assets include logical and physical data model assets, implemented data resources, and business intelligence (BI) assets.
- **Managing assets by using the command line** You can use the command line to move assets between different metadata repositories for environments such as development, test, and production. You can also query and delete common metadata assets and generate glossary assets from BI assets and logical data models.

Getting Started with Consoles

- **Opening consoles and clients by using the Launchpad** (https://www.ibm.com/support/knowledgecenter/SSZJPZ_11.7.0/com.ibm.swg.im.iis.found.admin.common.doc/topics/t_open_with_launchpad.html) The Launchpad is a standard, single web interface for opening the various clients or consoles for IBM InfoSphere Information Server.

- [****Opening consoles and clients without using the Launchpad in IIS ****](https://www.ibm.com/support/knowledgecenter/en/SSZJPZ_11.7.0/com.ibm.swg.im.iis.found.admin.common.doc/topics/t_open_without_launchpad.html) You can open the various clients or consoles for InfoSphere Information Server by using their URL.
- **Automatic login from web-based applications** The login page of the IBM InfoSphere Information Server console and web-based applications contain an option that enables you to log in automatically in your subsequent attempts.
- **IBM InfoSphere Information Server console overview** The IBM InfoSphere Information Server console is a rich-client-based interface for activities such as creating and managing projects, setting project-level security, analyzing data with IBM InfoSphere Information Analyzer, enabling information services with IBM InfoSphere Information Services Director, and running reports.

Connecting to data sources

- **Connecting to data sources** IBM InfoSphere Information Server connectivity options enable jobs to transfer data between InfoSphere Information Server and data sources.

New in Information Server 11.7

- Various new features introduced in Information Server 11.7 can be found [here](#)
- A key addition to Information Server 11.7 is [Information Server Enterprise Search](#)

Getting started and using IBM Information Server on Cloud Enterprise Edition

You must set up your connection to Information Server on Cloud Enterprise Edition. Information Server on Cloud Enterprise Edition provides all of the functions of its on-premises counterpart, IBM® InfoSphere® Information Server. It is in an IBM SoftLayer® hosted environment.

Prerequisite: You must know the IP address and the credentials of an account on the Information Server on Cloud Enterprise Edition server and client computers. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM Sales Representative.

The Information Server on Cloud Enterprise Edition client is on a Microsoft Windows virtual machine that is hosted on SoftLayer. When you connect to the client, you can access the client user interfaces. McAfee anti-virus software is installed on the client machine.

The Information Server on Cloud Enterprise Edition servers are on Red Hat Enterprise Linux virtual or bare metal computers that are hosted on SoftLayer. The services and repository tiers are on one computer. The engine tier is on the other computer. Optionally, IBM Business Process Manager Standard and IBM Cognos® Business Intelligence are on the separate Red Hat Enterprise Linux virtual or bare metal computers. When you connect to the servers, you can access the IBM InfoSphere Information Server engine, services, and repository tiers. You can restart Information Server on Cloud Enterprise Edition and do administrative tasks.

The default firewall configuration of server machines allows SSH connections only from client machines. You must first connect to a client machine by using a remote desktop connection, and then from the client machine you can connect to server machines by using SSH. After you log in to a server machine, you can change the firewall configurations to allow SSH connections from other machines. Communication between the server and client systems happens through a private IP. If you want to access the server from an on-premises client machine, you must modify the iptable rules.

Note: The Add Subscription window in Subscription Manager is disabled when you use a host name in the URL. To enable the window, use an IP address in the Subscription Manager URL.

When you connect for the first time, follow these steps:

1. SSH into the any server machine using unique user (order id) provided in the welcome letter using port 4362.

2. Change the password and su to root user.
3. On the Information Server on Cloud Enterprise Edition servers, run the ISALite tool on all tiers. The *IS_install_path* for the server computers is `/opt/IBM/InformationServer`.
4. Connect the Information Server on Cloud Enterprise Edition client to the Information Server on Cloud Enterprise Edition servers by following these steps:
 - **If your local computer is in a Microsoft Windows environment**
 - a. On your local computer, go to the **Start** menu. Click **Accessories > Remote Desktop Connection**.
 - b. Enter the IP address of the Microsoft Windows computer that hosts your Information Server on Cloud Enterprise Edition client. Click **Connect**.
 - c. In the Windows Security window, enter the user name and password for the Information Server on Cloud Enterprise Edition client. The user ID, password, and IP address of the client are in your Welcome letter. **Important:** Do not include the domain name with the user name.
 - d. In the Information Server on Cloud Enterprise Edition client, open the file `C:\Windows\System32\drivers\etc\hosts`. Make sure that entries with the private IP exist for the Information Server on Cloud Enterprise Edition servers that you are connecting to. **Important:** The server IP must be a private IP. You cannot connect to Information Server on Cloud Enterprise Edition servers when you use a public IP.
 - **If your local computer is in an Apple Mac environment**
 - a. On your local computer, install Microsoft Remote Desktop from the Apple App Store.
 - b. Click the Microsoft Remote Desktop icon, and then click **Open**.
 - c. In the Microsoft Remote Desktop window, click **New**.
 - d. In the Edit Remote Desktops window, supply the following information:
 - In the PC name field, type in the IP address of the cloud client machine.
 - In the User name and Password fields, type in the Windows user name and password that are in the Welcome letter.
5. Verify the connection and installation on the Information Server on Cloud Enterprise Edition client by following these steps:
 - Run the ISALite tool. The *IS_install_path* for the client computer is `C:\IBM\InformationServer`.
 - Test the installation of the Information Server on Cloud Enterprise Edition client.
6. Optional: Enable multiple users to open remote sessions to the Information Server on Cloud Enterprise Edition client by following these steps on the client computer:
 - a. Create user accounts.
 - b. Give users permission to do a remote desktop connection. Each Windows machine allows two concurrent user sessions. The number of Windows client machines that come with different sizes of your Information Server on Cloud Enterprise Edition offering are based on the number of concurrent users that are allowed in that offering. The small offering size has a maximum of two concurrent sessions. The medium offering size has a maximum of five concurrent sessions. The large offering size has a maximum of 10 concurrent sessions.
7. Reset the password for users and administrators on the Information Server on Cloud Enterprise Edition servers.
8. Optional: If you choose to include IBM Business Process Manager Standard, or IBM Cognos Business Intelligence in your environment, you must configure them. For details, see:
 - IBM Business Process Manager Standard: Installing and configuring IBM BPM Standard
 - IBM Cognos Business Intelligence:
 - IBM InfoSphere Information Governance Dashboard
 - Configure the IBM InfoSphere Business Glossary URI

Note: IBM Business Process Manager Standard Process Center and Process Server are installed and running. If you do not plan to use both of them, stop the service.

After the initial connection, you can do any of the following tasks:

- [Open the InfoSphere DataStage® and QualityStage® Designer client](#)
- [Connect to an on-premises computer](#)
- [Connect to an IBM dashDB™ database](#)
- [Connect to an on-premises DB2® database instance](#)
- [Perform general administration and security tasks](#)
- Open Information Server on Cloud Enterprise Edition clients from the client machine by following either method.
 - For browser clients, open Microsoft Internet Explorer. In the Favorites Bar,
 1. click the **IIS_Launchpad** bookmark. From the Launchpad window, click the icon of the client of InfoSphere Information Server.
 2. click the **Kubernetes_Dashboard** bookmark to login to the Kubernetes Dashboard sitting on your Enterprise Search node.
 - For thick clients, on your desktop click **Start > Programs > IBM Infosphere Information Server > client-name**.

You can connect to any of the following InfoSphere Information Server clients that were provisioned with their server component:

- IBM InfoSphere Information Governance Catalog
- IBM InfoSphere Information Analyzer
- IBM InfoSphere Metadata Integration Bridges and the metadata interchange agent
- IBM InfoSphere Metadata Asset Manager
- IBM InfoSphere Information Server istool command-line utility
- IBM InfoSphere Information Server Manager client
- Multi-Client Manager
- IBM InfoSphere DataStage and QualityStage Administrator
- IBM InfoSphere DataStage and QualityStage Designer
- IBM InfoSphere DataStage and QualityStage Director.

Related information

- [Enhancing security of Information Server on Cloud computers](#)


Backing up IIS Enterprise Edition on Cloud components

IBM Information Server on Cloud Premium and High Availability services provides software and hardware infrastructure for taking backups. Backup capability is based on IBM Spectrum Protect version 8.1.3 product (formerly known as Tivoli Storage Manager). One dedicated server machine with the installation of IBM Spectrum Protect server version 8.1.3 is provided with each deployment of Information Server on Cloud Premium and High Availability service. Spectrum Protect Server version used is 8.1.3 and Spectrum Protect Client version used is 8.1.2.

Sample configuration templates are also provided. You can create new configurations or customize the sample templates to take regular backups.

For more information on IBM Spectrum Protect, you can check [IBM Spectrum Protect Knowledge Center](#).

Spectrum Protect setup



IBM Spectrum Protect Server installation

Two instances of IBM Spectrum Protect server are configured on a dedicated machine for each deployment of IIS on Cloud Premium and High Availability service. To store backup data on the storage mounted on the server machine, the first instance is configured with Directory Container Storage Pool. The second instance is configured with Cloud Container Storage Pool to store data in Object Store. For the description of different data pools types, check [Storage pool types](#).

By default, operating system IP table rules allows communication to the Spectrum Protect backup server only from those machines where IBM Spectrum Protect Client is installed. In order to use Spectrum Protect Operation Console (web application), you need to open 11090 port for specific IPs in the IP table firewall rules. The port 11090 will be open for the IIS Windows Client to communicate with the Spectrum Protect backup server.

For small & medium plans, Directory Container Storage Pool is mapped to SAN storage and DB2 (which is used by IBM Spectrum Protect Server) is installed on Performance Storage. SAN storage is encrypted using operating system's LUKS encryption. IBM SoftLayer provides default encryption for Performance Storage in most of the data centers, for remaining data centers encryption is done at operating system level using LUKS. For data center list you can check [IBM SoftLayer documentation](#).

Directory Container Storage Pool (Instance 1)

A Directory Container Storage Pool is configured in Instance 1, this pool is used to store backup data locally. To know more about Directory Container Storage Pool, check [Directory-container storage pools FAQs](#).

Domain configurations are created for all client machines; these domains are linked with Directory Container Storage pool. For details about policy domain configuration check [Creating a policy domain](#).

Cloud Container Storage Pool (Instance 2)

The Cloud Container Storage Pool configured on Instance 2 is used to store data in cloud storage. The cloud-container storage pools that are provided by IBM Spectrum Protect can store data to cloud storage that is object-based. By storing data in cloud-container storage pools, you can exploit the cost per unit advantages that clouds offer along with the scaling capabilities that cloud storage provides. IBM Spectrum Protect manages the credentials, security, read and write I/Os, and the lifecycle for data that is stored to the cloud. You can back up and restore data or archive and retrieve data directly from the cloud-container storage pool. To understand more, check [Cloud-container storage pools FAQs](#) and [Configuring a cloud-container storage pool pages](#).

Before sending data to Object Store, Spectrum Protect server encrypts the data using encryption key. For encryption configurations details, check (https://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.0/srv.admin/t_cloud_encryption.html).

Like Directory Container Storage Pool, policy domain configurations are created for all Client machines; these domains are linked with Cloud Container Storage pool.

Node replication

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 failback capability.

In the setup provided to you, replication is enabled for all clients and backed-up data is replicated from Spectrum Protect Server Instance 1 to Instance 2. If required you can change replication settings by following instructions available at [Replicating client data to another server](#).

Data is replicated from Spectrum Protect instance 1 to Spectrum Protect instance 2 using node replication. The administrative schedule is configured for this purpose. There are two schedules replicate_nodes_weekend and replicate_nodes_weekday.

schedule replicate_nodes_weekday replicates data from Spectrum Protect instance 1 to Spectrum Protect instance 2 for every 3 hours.

schedule replicate_nodes_weekend replicates data from Spectrum Protect instance 1 to Spectrum Protect instance 2 for every 12 hours.

By default this schedule is stopped, run the following command to start it.

```
update schedule replicate_nodes_weekday type=administrative expiration=never
update schedule replicate_nodes_weekend type=administrative expiration=never
```

Run above commands from the "Command Builder" of Spectrum protect operation center.

IBM Spectrum Protect client installation

IBM Spectrum Protect Client is installed on all the machines except the one on which IBM Spectrum Protect Server is installed and the IIS Windows Client machine in which the IIS Designer Client and the other thick clients are installed. IBM Spectrum Protect Client is configured to send backup data/metadata/configuration files to Spectrum Protect Server over SSL. Client communicates with IBM Spectrum Protect server using server's private IP.

To learn about IBM Spectrum Protect Client, check [IBM Spectrum Protect Knowledge Center](#)

To enable access to IBM Spectrum Protect Client user interfaces, VNC server is installed on Client machines. By default, IP table firewall rules does not allow communication over 5901 port which is used by VNC server. To allow communication from the machine on which you want to access the user interfaces, you need to update IP Table firewall rules for port 5901. VNC server communication is not encrypted, if your organization mandates this communication to be secure, then you can use some other tool which support encryption.

Ports exposed

The following ports are opened to and from both the Spectrum Protect server and Spectrum Protect client machines:

- 1550
- 1552
- 1553
- 1650
- 1652
- 1653

Port 11090 is exposed from Spectrum Protect Server machine to the IIS Windows Client machine to access the Operations Center.

Port 4362 is also opened to access Spectrum Protect Server from the IIS Windows Client machines.

Apart from these ports, all the other ports are blocked for communication in the Spectrum Protect Server.

Getting started with IBM Spectrum Protect Operations Center console

IBM Spectrum Protect provides a web application called Operations Center for managing IBM Spectrum Protect environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

More details about Operations Center are available at [Managing the Operations Center](#).

In the setup provided to you, Operations Center is accessible using port 11090. Default IP table firewall rules on the IBM Spectrum Protect server machine does not allow communication with port 11090 from external machines. By default, the only machine which can access Spectrum Prtotect server using port 11090 is the IIS Windows Client machine. This is done for security purposes.

You can follow below steps to enable Operations Center access from an external machine.

1. Connect to Spectrum Protect server machine using putty or terminal. (Not required if accessing using IIS on Cloud Windows Client machine)

2. Go to scripts directory. (Not required if accessing using IIS on Cloud Windows Client machine)

```
cd /bckp/opt/IBM/scripts
```

3. Execute below command after replacing <IP_ADDRESS> with IP address of the machine from where you want to access Operations Console. (Not required if accessing using IIS on Cloud Windows Client machine)

```
./allow_port_from_ip.sh <IP_ADDRESS> 11090
```

When Operations Center is opened for the first time, it asks for some inputs. You must follow below steps to provide inputs when you open Operations Center for the first time

1. Open following URL in browser after replacing <Spectrum_Protct_Server_IP_Address> with your Spectrum Protect server machine IP.

```
https://<Spectrum_Protct_Server_IP_Address>:11090/oc
```

2. When you open Operations Console for the first time, it will ask for credentials.



3. Replace default details with correct values.

```

Localhost:1500      -- <Spectrum Protect server PUBLIC or PRIVATE IP>:1550
Administrator      -- tsminst
Password           -- Password for tsminst user is provided in welcome letter.


```

4. In the next page you will be asked to provide password (two times) for "Administrator ID". Provide password.
5. After providing password details, in the next page you need to specify how frequently you want to collect data. Depending on your requirement you can select 1 minute to 1 hour.
6. Follow instructions on the user interface to finish the wizard.

Note that "Instance 2" may be down and may take few minutes to start. You can check status of both instances under Overview tab of Operations Center console.


Starting Command Builder

Open following URL in browser https://<Spectrum_Protct_Server_IP_Address>:11090/oc to access Operation Center. To open the command-line interface, hover over the globe icon in the Operations Center menu bar, and click Command Builder.



IBM Spectrum Protect processes administrator commands either in the foreground or in the background. Commands that process in the foreground must complete before you can issue another command. When commands are processing in the background, you can issue additional commands at any time. Most IBM Spectrum Protect commands process in the foreground.

Spectrum Protect server contains two instances of servers which are connected using node replication feature for the fail-over scenario. In Command builder left side down you can see both Spectrum Protect instance. By default first instance of Spectrum Protect server is selected. You need to select the second instance of Spectrum protect server if you need to execute any commands against the second instance of Spectrum protect server. In "Command Builder" you can select second instance which has name like <ORDER_ID>/_s/m/l_bckp1. Select drop down menu from left downside corner, as shown in the image.



Retention Policy

Life cycle of backup data objects

A backup object exists in three states, active, inactive, and expired, before being purged from the Spectrum Protect server. The four steps involved in the life cycle of a backup data object are listed here.

1. A copy of the client data is sent to the Spectrum Protect server as a backup object. When a backup object is sent to the Spectrum Protect server, it becomes the active version.
2. It remains in an active state until the Spectrum Protect client program deletes the backup object manually, or a newer version of the backup object is sent. The backup object changes state from active to inactive.
3. The backup object remains inactive until it exceeds its retention settings. A backup object can exceed retention settings by either time or number of versions. The backup object changes state from inactive to expired.
4. The backup object remains in the expired state until expiration processing runs on the Spectrum Protect server. This process is invoked by a Spectrum Protect administrator with the expire inventory command. When expiration processing encounters a backup object in the expired state, it purges that object from the Spectrum Protect database and frees up the storage space where the backup object resided.

Spectrum Protect server sample domains for directory container storage pool (Spectrum Protect server local storage) are configured with backretention=30 archretention=30 Spectrum Protect server sample domains for cloud container storage pool (Object storage) are configured with backretention=365 archretention=365

BACKREtention :

Specifies the number of days (from the date the backup versions became inactive) to retain backup versions of files that are no longer on the client file system. This parameter is optional. You can specify an integer from 0 to 9999. The default value is 30. The server uses the backup retention value to manage

inactive versions of files when any of the following conditions occur: - A file is rebound to a new management class, but the new management class and the default management class do not contain a backup copy group. - The management class to which a file is bound no longer exists. The default management class does not contain a backup copy group. - The backup copy group is deleted from the management class to which a file is bound. The default management class does not contain a backup copy group.

ARCHREtention :

Specifies the number of days (from the date of archive) to retain archive copies. This parameter is optional. You can specify an integer from 0 to 30000. The default value is 365. The server uses the archive retention value to manage archive copies of files when either of the following conditions occur: - The management class to which a file is bound no longer exists. The default management class does not contain an archive copy group. - The archive copy group is deleted from the management class to which a file is bound. The default management class does not contain an archive copy group.

More details about domain configuration details are [here](#)

Below copygroup is defined for directory container storage pool (Spectrum Protect server local storage)
Spectrum Protect server sample copygroup defined with domain for backup is configured with
VEREXISTS=NOLimit VERDEL=NOLimit RETEXTRA=30 RETONLY=30

Below copygroup is defined for cloud container storage pool (Object storage) Spectrum Protect server
sample copygroup defined with domain for backup is configured with VEREXISTS=NOLimit
VERDEL=NOLimit RETEXTRA=365 RETONLY=365

Domain and copygroup created for each Spectrum Protect client machine have same settings.

VERExists :

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional. The default value is 2.

VEREXISTS=NOLimit Specifies that you want the server to retain all backup versions. The number of backup versions to retain is controlled by this parameter until versions exceed the retention time specified by the RETEXTRA parameter.

VERDeleted :

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using IBM Spectrum Protect. This parameter is optional. The default value is 1. If a user deletes a file from the client file system, the next incremental backup causes the server to expire the oldest versions of the file in excess of this number. The expiration date for the remaining versions is determined by the retention time specified by the RETEXTRA or RETONLY parameter.

VERDEL=NOLimit Specifies that you want the server to retain all backup versions for files that are deleted from the client file system after being backed up.

RETEtra :

Specifies the number of days to retain a backup version after that version becomes inactive. A version of a file becomes inactive when the client stores a more recent backup version, or when the client deletes the file from the workstation and then runs a full incremental backup. The server deletes inactive versions based on retention time even if the number of inactive versions does not exceed the number allowed by the VEREXISTS or VERDELETED parameters. This parameter is optional. The default value is 30 days.

RETOly :

Specifies the number of days to retain the last backup version of a file that has been deleted from the client file system. This parameter is optional. The default value is 60.

You can change sample retention policy values according to your requirement, keeping Spectrum Protect server storage space in mind.

More details about copygroup configuration details are [here](#)

Configuring Object storage

You can store deduplicated data and non-deduplicated data in a cloud-container storage pool and restore the data as required. You can configure IBM Spectrum Protect to temporarily store data in one or more local storage pool directories during data ingestion. The data is then moved from local storage to the cloud. In this way, you can improve data backup and archive performance.

After you define a storage pool directory, the IBM Spectrum Protect server uses that directory as a temporary landing spot for the data that you are transferring to cloud object storage. The server uses an automated background process to transfer data from local storage in the directory to cloud object storage. You do not need to take any additional steps to start or manage this transfer process. After the server successfully moves the data from local storage to cloud object storage, the server deletes the data from the directory and releases space for more incoming data.

If storage pool directories contain no more free space, backup operations stop prematurely. To avoid this situation, you can allocate more storage pool directories. You can also wait for the data to be automatically removed from the local directories after the data moves to the cloud.

Spectrum Protect server supports these cloud service providers.

- Amazon S3
- IBM Cloud Object Storage
- IBM SoftLayer
- OpenStack Swift

Amazon S3 API object storage has been used for the sample domains, policies and schedules. Object store is configured with dummy credentials and URL. Once the user has created an object storage with S3 API of their own, they can input the appropriate values for credentials and URLs and other necessities required to configure an object storage to Spectrum Protect server.

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. In the sample policy provided, each Spectrum Protect client has its own bucket, ie. its own subdirectory in the object storage. Bucket names and unique IDs (called "keys" in S3) are used to access data from object storage in Spectrum Protect.

Use Operation center in updating the details of Object store details or use "Command Center" in Operation center to update values.

In the sample policy provided, data retention policy for Object store is set to 365 days, which means data stored in Object store is available for 365 days. These are specified in the backretention and archretention parameters of a domain, as mentioned in the [Retention Policy](#) section.

Each machine contains 2 cloud pools, one used for backup and another for archive , where dummy credentials are provided.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

More details about SoftLayer Object store is [here](#)

More details about cloud object storage details are [here](#)

More details about configuring cloud-container storage pools for IBM SoftLayer is [here](#)

More details about encrypting data for cloud-container storage pools is [here](#)

Limitations and best practices

Spectrum Protect server setup and client installations are provided only with IIS on Cloud Premium and High Availability services. If you have also subscribed to Non-Production offerings then you need to develop your own backup artifacts for this offering.

Spectrum Protect server installation which is provided with IIS on Cloud Premium and High Availability offerings can only be used for taking backups of applications and files which are part of IIS on Cloud deployment.

IIS on Cloud setup consists of many applications and components. When you configure different policies & schedulers to take backup of different applications and components, backups are created at different timestamps. Hence after restoration of specific application backup, its data may not be in complete sync with related data elements in other applications or components.

Spectrum Protect server and Operation Center are installed in same system, so communication between Spectrum Protect server and Operation Center is through non-SSL. Spectrum Protect server has 2 instances which are connected through node replication. Both the instances are in same system, so communication between these 2 instances is through non-SSL.

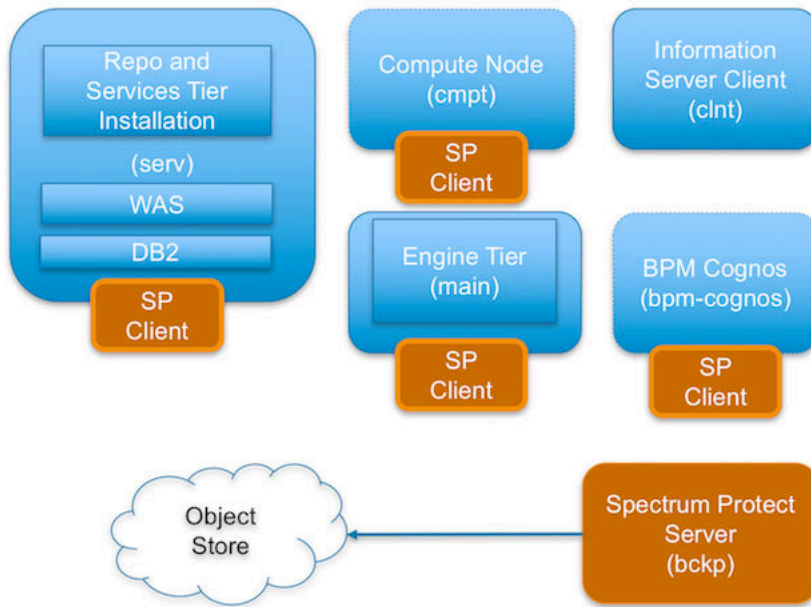
Cloud storage is connected to 2nd instance of Spectrum Protect server, which acts as a fail-over server. There are some limitations in connecting to cloud storage from Spectrum Protect client. You can't connect to Cloud storage and retrieve data when 1st instance of Spectrum Protect server is up and running. If 1st instance of Spectrum Protect server is down, then you can connect to 2nd instance of Spectrum Protect server. You'll have only read only access, which means it's used only to retrieve data, you can't take backup and archive using this.

You should follow product or application specific documentation and best practices while developing artifacts to take backups, below are some examples: - IBM Information Server does not support hot backups. For details, check [Backing up IBM InfoSphere Information Server components](#). - WebSphere Application Server documentation suggests that servers are stopped while taking backup of node configurations. For details, check [backupConfig command](#)

IMPORTANT

- Spectrum Protect server's database backup files are not backed-up automatically. These files are mandatorily required to restore a SPectrum Protect Server in case of a failure scenario. Hence, it is recommended that the user backs the artifacts related to this to a secure location. Details about Spectrum Protect server database backup is available at "Spectrum Protect server database backup" section. [Spectrum Protect server database backup](#)
- Spectrum Protect server master encryption key is stored in the server password file, dsmserv.pwd. This file takes care of encryption and decryption of data being transferred for backup and restore. Hence, it is recommended that the user backs the artifacts related to this to a secure location. Details about master key is available at "Protecting the master encryption key" section. [Protecting the master encryption key](#)

IBM Spectrum Protect setup for IIS on Cloud Enterprise Edition Premium service



IIS Enterprise Edition on Cloud has the following machines:

- IIS Services machine
- IIS Engine machine
- Spectrum Protect Backup server
- IIS Windows Client machine
- IIS Compute machines (Optional)
- BPM/Cognos machine (Optional)

Spectrum Protect server software is installed in the Spectrum Protect server machine and the Spectrum Protect Client software is installed in the rest except the IIS Windows Client machine.

By default all the sample schedulers are disabled by setting expiration value to -1. In order to use sample scheduler, you need to enable them by setting appropriate expiry date. This is discussed in detail for each of the machines in these offerings in the following sections.

Backup files will be available only for 30 days, later they are removed from Spectrum Protect server directory storage pool. If cloud container is configured backup files will be available for 365 days, later they are removed from Spectrum Protect server cloud storage pool. You can change these settings according to your requirement, more details on modifying these settings is available in "Retention Policy" section. Make sure storage in Spectrum Protect server is limited.

IIS Services machine (serv)

To take backups of the artifacts present in the IIS Services machine, a sample policy domain configuration named as SERVICE-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of WAS profile directory, WAS profile configuration, database full, database incremental online backups and other files which are located in IIS Services machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine, execute commands to create backup archive files, copy files to specific location, executes commands to take db2 full and incremental backups and invoke Spectrum Protect commands to copy files to Spectrum Protect Server machine.

WebSphere Application Profiles Backup

A sample schedule named as SERVICE-WAS-WEEKLY is configured to take backups of WAS profiles. This schedule invokes WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder

inside IIS Services Machine. Before enabling SERVICE-WAS-WEEKLY schedule, you must update WASProfileBackup_IIS.sh and provide value for PASSWORD field.

SERVICE-WAS-WEEKLY schedule is configured to execute weekly once, on Sundays at midnight.

WASProfileBackup_IIS.sh executes manageprofiles.sh utility provided by WAS to create profile backup. A zip file named IIS_AppServer_backup.zip is created in /home_/WAS_Backup folder. This zip file contains the backup of profile InfoSphere.

- WAS InfoSphere profile folder : /opt/IBM/WebSphere/AppServer/profiles/InfoSphere

IIS_AppServer_backup.zip : Contains backup of InfoSphere profile created using manageprofiles.sh utility where profile is located at /opt/IBM/WebSphere/AppServer/profiles/InfoSphere

WASProfileBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServer_backup.zip file.

IMPORTANT

WAS guidelines for backup of WAS profile mandates that the WAS services be stopped before the profile is being backed up. This schedule forces to stop servers related to that WAS profile.

As this schedule stop server related to WAS profile, it's up to you to take a decision on whether to run this schedule on weekly basis or run this when it's needed. It is recommended to take backup of profiles when there are changes to WAS like adding a new CBA or deleting a CBA etc..

If you want to run this schedule manually whenever it's needed instead of running weekly, a shell script named WASProfileBackup_IIS.sh is available at /opt/tivoli/tsm/client/ba/bin folder. Execute this script as a root user, it'll stop the profile and related servers, take the backup of profile, send it to Spectrum Protect server and start the profile and servers.

More information on the backup of WAS profiles is mentioned in the following links.

https://www.ibm.com/support/knowledgecenter/en/SSAW57_8.5.5/com.ibm.websphere.nd.doc/ae/rxml_manageprofiles.html

Sample template used for taking backup of WAS profile :

```
/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin/stopServer.sh server1 -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/bin/manageprofiles.sh -backupProfile -profileName InfoSphere -backupFile /home_/WAS_Backup/IIS_AppServer_backup.zip
/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin/startServer.sh server1
dsmc sel "/home_/WAS_Backup/*" -subdir=yes
```

The server is stopped first, then the manageprofiles.sh command is used to take backup of the WAS profile and then the server is started again. *dsmc* command is used to transfer the backup files from IIS Services machine to the Spectrum Protect Server.

In order to enable SERVICE-WAS-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at [Starting Command Builder](#) section.

If you want to run this schedule weekly, execute below command.

```
update schedule SERVICE-DOMAIN SERVICE-WAS-WEEKLY expiration=never
```

Above command enable SERVICE-WAS-WEEKLY schedule without expiry date. After enabling SERVICE-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Services machine. dsmcad provides a light-weight timer which automatically starts and stops the schedule process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

After running the SERVICE-WAS-WEEKLY schedule or running WASProfileBackup_IIS.sh manually, make sure IIS_AppServer_backup.zip is created at /home_/WAS_Backup folder. If you didn't find these files, there might be an issue while running WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder. Run this script manually and fix the issues.

WebSphere Application Configuration Backup

A sample scheduler named as SERVICE-WAS-DAILY is configured to take backups of WAS profile configuration. This schedule invokes WASProfileConfigBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Services Machine. Before enabling SERVICE-WAS-DAILY schedule, you must update WASProfileConfigBackup_IIS.sh and provide value for PASSWORD in the shell script.

SERVICE-WAS-DAILY scheduler is configured to execute daily at midnight.

WASProfileConfigBackup_IIS.sh executes commands to create an archive file named IIS_AppServerConfig.zip, which contains all profile configurations.

- WAS InfoSphere profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/InfoSphere

IIS_AppServerConfig.zip contains InfoSphere profile configuration which is generated by using backupConfig command.

Archive files are stored in /home_/WAS_Conf/ folder. WASProfileConfigBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServerConfig.zip file.

In order to enable SERVICE-WAS-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section.
[Starting Command Builder](#)

```
update schedule SERVICE-DOMAIN SERVICE-WAS-DAILY expiration=never
```

Above command enable SERVICE-WAS-WEEKLY scheduler without expiry date.

After enabling SERVICE-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services and run `service dsmcad restart` using root user.

IIS Database Backup

Database is configured with linear logging, which means all the transaction (archive) logs of database are stored in the IIS Services machine. It is recommended that these logs should be stored in Spectrum Protect server itself. In case of a failover scenario, the logs are safe and secure in the Spectrum Protect server.

As transaction (archive) logs of database are stored in the IIS Services machine, it is the user's responsibility to clean up unused logs, as these logs are not cleared by default. Once you move transaction (archive) logs from local disk to Spectrum Protect server, it's the responsibility of Spectrum Protect server to clean unused logs. Use below commands in moving transaction (archive) logs from local disk to Spectrum Protect server:

Open Putty or terminal in the IIS Services machine, switch to db2inst1 user.

```
db2 update database configuration for xmeta using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for iadb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for dsosb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for esbdb2 using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager
```

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. Hence before the sample schedules are started, you have to make sure that one full backup of the respective database is available. Verify available backups for a database by using 'db2adutl query db <DATABASE_NAME>'. Open Putty or terminal in the IIS Services machine, switch to db2inst1 user to run this command. You have to execute the shell script

db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in the IIS Services Machine, incase there are no full backups available.

Before running db2FullBackup.sh you must consider, where to store database archive logs.

Online full database backups

A sample scheduler named as SERVICE-DB-FULL-WEEKLY is configured to take backups of online full database. This schedule invokes db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Services Machine.

SERVICE-DB-FULL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

In order to enable SERVICE-DB-FULL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DB-DOMAIN SERVICE-DB-FULL-WEEKLY expiration=never
```

Above command enable SERVICE-DB-FULL-WEEKLY scheduler without expiry date. After enabling SERVICE-DB-FULL-WEEKLY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services and run `service dsmcad restart` using root user.

Online incremental database backups

A sample scheduler named as SERVICE-DB-INCREMENT-DAILY is configured to take online incremental database backups. This schedule invokes db2IncrementalBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Services Machine.

SERVICE-DB-INCREMENT-DAILY scheduler is configured to execute from Monday to Saturday at midnights, except Sunday.

In order to enable SERVICE-DB-INCREMENT-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DB-DOMAIN SERVICE-DB-INCREMENT-DAILY expiration=never
```

Above command enable SERVICE-DB-INCREMENT-DAILY scheduler without expiry date. After enabling SERVICE-DB-INCREMENT-DAILY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

Files and Directories Backup

A sample scheduler named as SERVICE-FILES-DAILY is configured to take backups of files and folders. This schedule invokes service_FilesDaily.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Services Machine.

SERVICE-WAS-DAILY scheduler is configured to execute everyday at midnight.

service_FilesDaily.sh executes commands to take backup of below files and folders.

- /root/keyfile
- /home_/db2inst1/db2keystore.p12
- /etc/sysconfig/iptables
- /opt/IBM/InformationServer/Updates/*
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/*
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/classes/*
- /opt/IBM/InformationServer/Version.xml
- /etc/services

- /etc/inittab

service_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable SERVICE-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DOMAIN SERVICE-FILES-DAILY expiration=never
```

Above command enable SERVICE-FILES-DAILY scheduler without expiry date.

After enabling SERVICE-FILES-DAILY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script /opt/tivoli/tsm/client/ba/bin/service_FilesDaily.sh available in IIS Services machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Services machine is *service*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

```
update stgpool serv-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool serv-cloud-pool identity=<USERNAME>
update stgpool serv-cloud-pool password=<PASSWORD>

update stgpool serv-arc-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool serv-arc-cloud-pool identity=<USERNAME>
update stgpool serv-arc-cloud-pool password=<PASSWORD>
```

IIS Engine machine (main)

To take backups of the artifacts present in the IIS Engine machine, a sample policy domain configuration named as MAIN-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Engine machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

IIS artifacts backup using ISTOOL

A sample scheduler named as MAIN-ISTOOL-WEEKLY is configured to take backups of ISTool export configuration. This schedule invokes `istool_assets.sh` which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Engine Machine. Before enabling MAIN-ISTOOL-WEEKLY schedule, you must update `istool_assets.sh` and provide value for PASSWORD field.

MAIN-ISTOOL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

`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 configuration

Archive file generated by istool.sh is stored in /home_/istool/. istool_assets.sh executes Spectrum Protect server selective backup command to take backup of istool generated file.

In order to enable MAIN-ISTOOL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section.
[Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-ISTOOL-WEEKLY expiration=never
```

Above command enable MAIN-ISTOOL-WEEKLY scheduler without expiry date.

After enabling MAIN-ISTOOL-WEEKLY schedule, start 'dsmcad' service from IIS machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

Files and Directories Backup

A sample scheduler named as MAIN-FILES-DAILY is configured to take backups of files and folders. This schedule invokes main_FilesDaily.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Engine Machine.

MAIN-FILES-DAILY scheduler is configured to execute everyday at midnight.

main_FilesDaily.sh executes commands to take backup of below files and folders.

- /root/keyfile
- /opt/IBM/InformationServer/Server/Projects/*
- /etc/sysconfig/iptables
- /opt/IBM/InformationServer/Server/MsgHandlers/*
- /opt/IBM/InformationServer/Server/Configurations/*
- /opt/IBM/InformationServer/Updates/*
- /opt/IBM/InformationServer/Server/DSODB/*.cfg
- /opt/IBM/InformationServer/Server/DSEngine/dsenv
- /etc/services
- /etc/inittab

main_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable MAIN-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section.
[Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-FILES-DAILY expiration=never
```

Above command enable MAIN-FILES-DAILY scheduler without expiry date.

After enabling MAIN-FILES-DAILY schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script /opt/tivoli/tsm/client/ba/bin/main_FilesDaily.sh available in IIS Engine machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Engine machine is *main*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

```
update stgpool main-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool main-cloud-pool identity=<USERNAME>
update stgpool main-cloud-pool password=<PASSWORD>

update stgpool main-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool main-arc-cloud-pool identity=<USERNAME>
update stgpool main-arc-cloud-pool password=<PASSWORD>
```

IIS Compute machine (cmpt)

To take backups of the artifacts present in the IIS Compute machine, a sample policy domain configuration named as COMPUTE-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Compute machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

Files and Directories Backup

A sample scheduler named as COMPUTE-FILES-DAILY is configured to take backups of files and folders. This schedule invokes compute_FilesDaily.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Compute Machine.

COMPUTE-FILES-DAILY scheduler is configured to execute everyday at midnight.

compute_FilesDaily.sh executes commands to take backup of below files and folders.

- /root/keyfile
- /etc/sysconfig/iptables
- /etc/services
- /etc/inittab

compute_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable COMPUTE-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule COMPUTE-DOMAIN COMPUTE-FILES-DAILY expiration=never
```

Above command enable COMPUTE-FILES-DAILY scheduler without expiry date.

After enabling COMPUTE-FILES-DAILY schedule, start 'dsmcad' service from IIS Compute machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Compute machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script /opt/tivoli/tsm/client/ba/bin/main_FilesDaily.sh available in IIS Compute machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS compute machine is *compute*.

Use Operation center in updating the details of Object store details.

IIS Compute machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

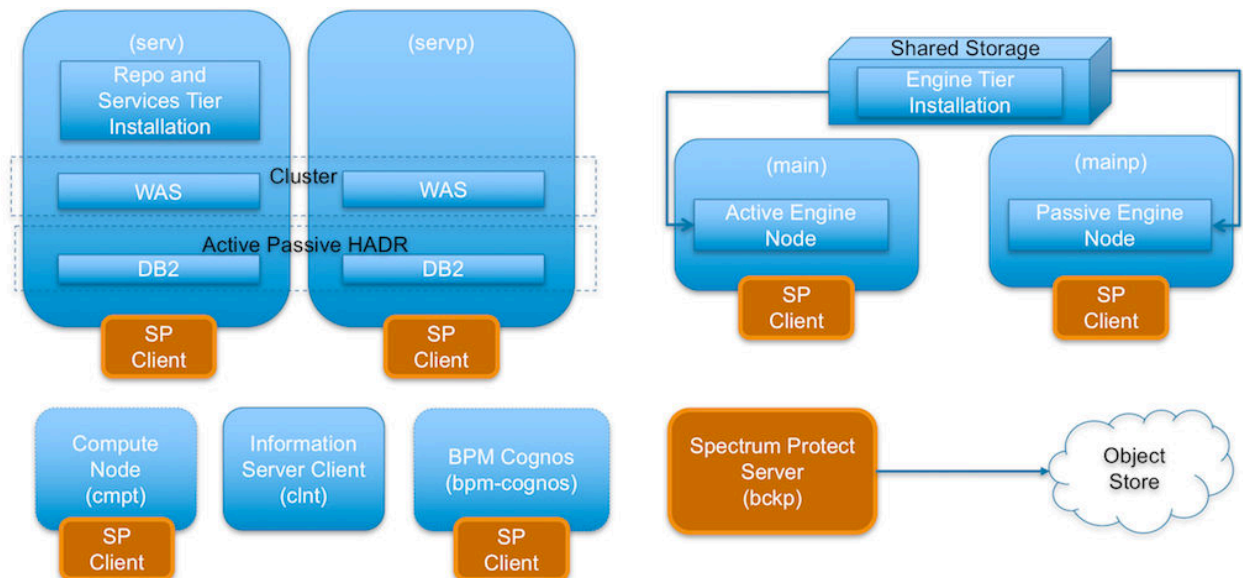
```
update stgpool compute-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool compute-cloud-pool identity=<USERNAME>
update stgpool compute-cloud-pool password=<PASSWORD>

update stgpool compute-arc-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool compute-arc-cloud-pool identity=<USERNAME>
update stgpool compute-arc-cloud-pool password=<PASSWORD>
```

IBM Spectrum Protect setup for IIS on Cloud Enterprise Edition High Availability service

IIS Enterprise Edition on Cloud HA offering has the following machines:

- IIS Services machine (Active)
- IIS Services machine (Passive)
- IIS Engine machine (Active)
- IIS Engine machine (Passive)
- Spectrum Protect Backup server
- IIS Windows Client machine
- IIS Compute machines (Optional)
- BPM/Cognos machine (Optional)



Spectrum Protect server software is installed in the Spectrum Protect server machine and the Spectrum Protect Client software is installed in the rest except the IIS Windows Client machine.

IIS Active Services machine (serv)

To take backups of the artifacts present in the IIS Services machine, a sample policy domain configuration named as SERVICE-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of WAS profile directory, WAS profile configuration, database full, database incremental online backups and other files which are located in IIS Active Services machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine, execute commands to create backup archive files, copy files to specific location, executes commands to take db2 full and incremental backups and invoke Spectrum Protect commands to copy files to Spectrum Protect Server machine.

WebSphere Application Profiles Backup

A sample schedule named as SERVICE-WAS-WEEKLY is configured to take backups of WAS profiles. This schedule invokes WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Active Services machine. Before enabling SERVICE-WAS-WEEKLY schedule, you must update WASProfileBackup_IIS.sh and provide value for PASSWORD field.

SERVICE-WAS-WEEKLY schedule is configured to execute weekly once, on Sundays at midnight.

WASProfileBackup_IIS.sh executes manageprofiles.sh utility provided by WAS to create profile backup. The profiles are being backed up into the following files: IIS_Dmgr01_backup.zip and IIS_Custom01_backup.zip, which contains Dmgr01 and Custom01 profile backup.

- WAS Dmgr01 profile folder : /opt/IBM/WebSphere/AppServer/profiles/Dmgr01
- WAS Custom01 profile folder : /opt/IBM/WebSphere/AppServer/profiles/Custom01

Generated files are stored in /home_/WAS_Backup folder.

WASProfileBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_Dmgr01_backup.zip and IIS_Custom01_backup.zip file.

IMPORTANT

WAS guidelines for backup of WAS profile mandates that the WAS services be stopped before the profile is being backed up. This schedule forces to stop servers related to that WAS profile.

As this schedule stop server related to WAS profile, it's up to you to take a decision on whether to run this schedule on weekly basis or run this when it's needed. It is recommended to take backup of profiles when there are changes to WAS like adding a new CBA or deleting a CBA etc.

If you want to run this schedule manually whenever it's needed instead of running weekly, a shell script named WASProfileBackup_IIS.sh is available at /opt/tivoli/tsm/client/ba/bin folder. Execute this script as a root user, it'll stop the profile and related servers, take the backup of profile, send it to Spectrum Protect server and start the profiles and servers.

More information on the backup of WAS profiles is mentioned in the following links.

https://www.ibm.com/support/knowledgecenter/en/SSAW57_8.5.5/com.ibm.websphere.nd.doc/ae/rxml_manageprofiles.html

Shell script used for taking backup of WAS profile :

```
/opt/IBM/WebSphere/AppServer/profiles/Custom01/bin/stopServer.sh webserver1 -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/profiles/Custom01/bin/stopServer.sh server1 -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/profiles/Custom01/bin/stopNode.sh -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/bin/manageprofiles.sh -backupProfile -profileName Custom01 -backupFile ${ WAS_BACKUP_PATH }/IIS_Custom01_backup.zip

/opt/IBM/WebSphere/AppServer/profiles/Dmgr01/bin/stopManager.sh -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/bin/manageprofiles.sh -backupProfile -profileName Dmgr01 -backupFile ${ WAS_BACKUP_PATH }/IIS_Dmgr01_backup.zip

/opt/IBM/WebSphere/AppServer/profiles/Dmgr01/bin/startManager.sh
```

```
/opt/IBM/WebSphere/AppServer/profiles/Custom01/bin/startNode.sh
/opt/IBM/WebSphere/AppServer/profiles/Custom01/bin/startServer.sh webserver1
/opt/IBM/WebSphere/AppServer/profiles/Custom01/bin/startServer.sh server1
```

The server and nodeagent are stopped first, then the manageprofiles.sh command is used to take backup of the WAS profile and then the nodeagent, the server is started again. *dsmc* command is used to transfer the backup files from IIS Active Services machine to the Spectrum Protect Server. In the case of Dmgr (deployment manager) profile, it'll stop deployment manager and start it back. There are no servers or node agents attached to it.

In order to enable SERVICE-WAS-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at [Starting Command Builder](#) section.

If you want to run this schedule weekly, execute below command.

```
update schedule SERVICE-DOMAIN SERVICE-WAS-WEEKLY expiration=never
```

Above command enable SERVICE-WAS-WEEKLY schedule without expiry date. After enabling SERVICE-WAS-WEEKLY schedule , start 'dsmcad' service from IIS Active Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the schedule process as needed.

Open Putty or terminal in IIS Active Services machine and run `service dsmcad restart` using root user.

After running the SERVICE-WAS-WEEKLY schedule or running WASProfileBackup_IIS.sh manually, make sure IIS_Dmgr01_backup.zip and IIS_Custom01_backup.zip are created at /home_/WAS_Backup folder. If you didn't find these files, there might be an issue while running WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder. Run this script manually and fix the issues.

WebSphere Application Configuration Backup

A sample scheduler named as SERVICE-WAS-DAILY is configured to take backups of WAS profile configuration. This schedule invokes WASProfileConfigBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Services Machine. Before enabling SERVICE-WAS-DAILY schedule, you must update WASProfileConfigBackup_IIS.sh and provide value for PASSWORD in the shell script.

SERVICE-WAS-DAILY scheduler is configured to execute daily at midnight.

WASProfileConfigBackup_IIS.sh executes commands to create an archive file named IIS_AppServerConfig.zip, which contains all profile configurations.

- WAS Dmgr01 profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/Dmgr01
- WAS Custom01 profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/Custom01

IIS_AppServerConfig.zip contains InfoSphere profile configuration which is generated by using backupConfig command.

Archive files are stored in /home_/WAS_Conf/ folder. WASProfileConfigBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServerConfig.zip file.

In order to enable SERVICE-WAS-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DOMAIN SERVICE-WAS-DAILY expiration=never
```

Above command enable SERVICE-WAS-WEEKLY scheduler without expiry date.

After enabling SERVICE-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

IIS Database Backup

Database is configured with linear logging, which means all the transaction (archive) logs of database are stored in the IIS Services machine. It is recommended that these logs should be stored in Spectrum Protect server itself. In case of a failover scenario, the logs are safe and secure in the Spectrum Protect server.

As transaction (archive) logs of database are stored in the IIS Services machine, it is the user's responsibility to clean up unused logs, as these logs are not cleared by default. Once you move transaction (archive) logs from local disk to Spectrum Protect server, it's the responsibility of Spectrum Protect server to clean unused logs.

Database is configured with HADR scenario, verify HADR status before applying any database configuration changes.

Open Putty or terminal in the IIS Services machine, switch to db2inst1 user.

```
db2pd -db xmeta -hadr db2pd -db iadb -hadr db2pd -db dsodb -hadr db2pd -db esbdb2 -hadr
```

If it's proper then follow these steps.

Use below commands in moving transaction (archive) logs from local disk to Spectrum Protect server:

Open Putty or terminal in the IIS Services machine, switch to db2inst1 user.

```
db2 update database configuration for xmeta using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for iadb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for dsosb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for esbdb2 using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager
```

Wait for 2 or 3 minutes and check HADR status.

```
db2pd -db xmeta -hadr
db2pd -db iadb -hadr
db2pd -db dsodb -hadr
db2pd -db esbdb2 -hadr
```

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. Hence before the sample schedules are started, you have to make sure that one full backup of the respective database is available. Verify available backups for a database by using 'db2adutl query db <DATABASE_NAME>'. Open Putty or terminal in the IIS Services machine, switch to db2inst1 user to run this command. You have to execute the shell script db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in the IIS Services Machine, incase there are no full backups available.

Before running db2FullBackup.sh you must consider, where to store database archive logs.

Online full database backups

A sample scheduler named as SERVICE-DB-FULL-WEEKLY is configured to take backups of online full database. This schedule invokes db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Services Machine.

SERVICE-DB-FULL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

In order to enable SERVICE-DB-FULL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DB-DOMAIN SERVICE-DB-FULL-WEEKLY expiration=never
```

Above command enable SERVICE-DB-FULL-WEEKLY scheduler without expiry date. After enabling SERVICE-DB-FULL-WEEKLY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

Online incremental database backups

A sample scheduler named as SERVICE-DB-INCREMENT-DAILY is configured to take online incremental database backups. This schedule invokes `db2IncrementalBackup.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Services Machine.

SERVICE-DB-INCREMENT-DAILY scheduler is configured to execute from Monday to Saturday at midnights, except Sunday.

In order to enable SERVICE-DB-INCREMENT-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DB-DOMAIN SERVICE-DB-INCREMENT-DAILY expiration=never
```

Above command enable SERVICE-DB-INCREMENT-DAILY scheduler without expiry date. After enabling SERVICE-DB-INCREMENT-DAILY schedule, start 'dsmcad' service from IIS Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

Files and Directories Backup

Two sample scheduler named as SERVICE-FILES-DAILY-1 and SERVICE-FILES-DAILY-2 is configured to take backups of files and folders. This schedule invokes `service_FilesDaily1.sh` and `service_FilesDaily2.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Services Machine.

SERVICE-FILES-DAILY scheduler is configured to execute everyday at midnight.

`service_FilesDaily1.sh` and `service_FilesDaily2.sh` executes commands to take backup of below files and folders.

- /root/keyfile
- /home_/db2inst1/db2keystore.p12
- /etc/sysconfig/iptables
- /opt/IBM/InformationServer/Updates/*
- /etc/services
- /etc/inittab
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/*
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/classes/*
- /opt/IBM/InformationServer/Version.xml
- /opt/IBM/HTTPServer/conf/*
- /opt/IBM/WebSphere/AppServer/bin/configurewebserver1.sh
- /opt/IBM/WebSphere/AppServer/bin/configurewebserver2.sh
- /opt/IBM/Plugins/config/*
- /opt/IBM/sslkey.kdb
- /opt/IBM/sslkey.rdb
- /opt/IBM/sslkey.sth

`service_FilesDaily1.sh` and `service_FilesDaily2.sh` executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable SERVICE-FILES-DAILY-1 and SERVICE-FILES-DAILY-2 schedules you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICE-DOMAIN SERVICE-FILES-DAILY-1 expiration=never
update schedule SERVICE-DOMAIN SERVICE-FILES-DAILY-2 expiration=never
```

Above command enables the schedulers without expiry date.

After enabling the schedules, start 'dsmcad' service from IIS Active Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Services machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script `/opt/tivoli/tsm/client/ba/bin/service_FilesDaily1.sh` or `/opt/tivoli/tsm/client/ba/bin/service_FilesDaily2.sh` available in IIS Services machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Services is *service*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like `<ORDER_ID>-/s/m/l-bckp1`.

```
update stgpool service-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool service-cloud-pool identity=<USERNAME>
update stgpool service-cloud-pool password=<PASSWORD>

update stgpool service-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool service-arc-cloud-pool identity=<USERNAME>
update stgpool service-arc-cloud-pool password=<PASSWORD>
```

IIS Passive Services machine (servp)

To take backups of the artifacts present in the IIS Passive Services machine, a sample policy domain configuration named as SERVICEP-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of WAS profile directory, WAS profile configuration, database full, database incremental online backups and other files which are located in IIS Services machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine, execute commands to create backup archive files, copy files to specific location, executes commands to take db2 full and incremental backups and invoke Spectrum Protect commands to copy files to Spectrum Protect Server machine.

WebSphere Application Profiles Backup

A sample schedule named as SERVICEP-WAS-WEEKLY is configured to take backups of WAS profiles. This schedule invokes WASProfileBackup_IIS.sh which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Passive Services machine. Before enabling SERVICEP-WAS-WEEKLY schedule, you must update WASProfileBackup_IIS.sh and provide value for PASSWORD field.

SERVICEP-WAS-WEEKLY schedule is configured to execute weekly once, on Sundays at midnight.

WASProfileBackup_IIS.sh executes manageprofiles.sh utility provided by WAS to create profile backup. The profile is being backed up into the following file: IIS_node2_backup.zip. The zip file contains backup of node2 profile. This zip file is stored in /home_/WAS_Backup folder.

- WAS node2 profile folder : /opt/IBM/WebSphere/AppServer/profiles/node2

Generated files are stored in /home_/WAS_Backup folder. WASProfileBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_node2_backup.zip.

IMPORTANT

WAS guidelines for backup of WAS profile mandates that the WAS services be stopped before the profile is being backed up. This schedule forces to stop servers related to that WAS profile.

As this schedule stop server related to WAS profile, it's up to you to take a decision on whether to run this schedule on weekly basis or run this when it's needed. It is recommended to take backup of profiles when there are changes to WAS like adding a new CBA or deleting a CBA etc..

If you want to run this schedule manually whenever it's needed instead of running weekly, a shell script named WASProfileBackup_IIS.sh is available at /opt/tivoli/tsm/client/ba/bin folder. Execute this script as a root user, it'll stop the profile and related servers, take the backup of profile and send it to Spectrum Protect server.

More information on the backup of WAS profiles is mentioned in the following links.

https://www.ibm.com/support/knowledgecenter/en/SSAW57_8.5.5/com.ibm.websphere.nd.doc/ae/rxml_manageprofiles.html

Shell script used for taking backup of WAS profile :

```
/opt/IBM/WebSphere/AppServer/profiles/node2/bin/stopServer.sh webserver2 -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/profiles/node2/bin/stopServer.sh server2 -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/profiles/node2/bin/stopNode.sh -username wasadmin -password PASSWORD
/opt/IBM/WebSphere/AppServer/bin/manageprofiles.sh -backupProfile -profileName node2 -backupFile ${ WAS_BACKUP_PATH }/IIS_node2_backup.zip

/opt/IBM/WebSphere/AppServer/profiles/node2/bin/startNode.sh
/opt/IBM/WebSphere/AppServer/profiles/node2/bin/startServer.sh webserver2
/opt/IBM/WebSphere/AppServer/profiles/node2/bin/startServer.sh server2
```

The server and nodeagent are stopped first, then the manageprofiles.sh command is used to take backup of the WAS profile and then the nodeagent, the server is started again. *dsmc* command is used to transfer the backup files from IIS Passive Services machine to the Spectrum Protect Server.

In order to enable SERVICEP-WAS-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at [Starting Command Builder](#) section.

If you want to run this schedule weekly, execute below command.

```
update schedule SERVICEP-DOMAIN SERVICEP-WAS-WEEKLY expiration=never
```

Above command enable SERVICEP-WAS-WEEKLY schedule without expiry date. After enabling SERVICEP-WAS-WEEKLY schedule , start 'dsmcad' service from IIS Passive Services machine. dsmcad provides a light-weight timer which automatically starts and stops the schedule process as needed.

Open Putty or terminal in IIS Passive Services machine and run `service dsmcad restart` using root user.

After running the SERVICEP-WAS-WEEKLY schedule or running WASProfileBackup_IIS.sh manually, make sure IIS_node2_backup.zip is created at /home_/WAS_Backup folder. If you didn't find these files, there might be an issue while running WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder. Run this script manually and fix the issues.

WebSphere Application Configuration Backup

A sample scheduler named as SERVICEP-WAS-DAILY is configured to take backups of WAS profile configuration. This schedule invokes WASProfileConfigBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Passive Services Machine. Before enabling SERVICEP-WAS-DAILY schedule, you must update WASProfileConfigBackup_IIS.sh and provide value for PASSWORD in the shell script.

SERVICEP-WAS-DAILY scheduler is configured to execute daily at midnight.

WASProfileConfigBackup_IIS.sh executes commands to create an archive file named IIS_AppServerConfig.zip, which contains all profile configurations.

- WAS "node01" profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/node01

IIS_AppServerConfig.zip contains InfoSphere profile configuration which is generated by using backupConfig command.

Archive files are stored in /home_/WAS_Conf/ folder. WASProfileConfigBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServerConfig.zip file.

In order to enable SERVICEP-WAS-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section.
[Starting Command Builder](#)

```
update schedule SERVICEP-DOMAIN SERVICEP-WAS-DAILY expiration=never
```

Above command enable SERVICEP-WAS-WEEKLY scheduler without expiry date.

After enabling SERVICEP-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Passive Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Services and run `service dsmcad restart` using root user.

IIS Database Backup

As this machine acts as passive, the user must not run any database schedulers or configuration changes provided for this machine until the machine behaves as active one.

Don't run below schedulers or configuration changes at the starting. Start these database schedulers or configuration changes only if this machine is taken over as an Active Services machine

Database is configured with linear logging, which means all the transaction (archive) logs of database are stored in the IIS Passive Services machine. It is recommended that these logs should be stored in Spectrum Protect server itself. In case of a failover scenario, the logs are safe and secure in the Spectrum Protect server.

As transaction (archive) logs of database are stored in the IIS Passive Services machine, it is the user's responsibility to clean up unused logs, as these logs are not cleared by default. Once you move transaction (archive) logs from local disk to Spectrum Protect server, it's the responsibility of Spectrum Protect server to clean unused logs.

Open Putty or terminal in the IIS Passive Services machine, switch to db2inst1 user.

```
db2pd -db xmeta -hadr
db2pd -db iadb -hadr
db2pd -db dsodb -hadr
db2pd -db esbdb2 -hadr
```

If it's proper then follow these steps.

Use below commands in moving transaction (archive) logs from local disk to Spectrum Protect server:

Open Putty or terminal in the IIS Passive Services machine, switch to db2inst1 user.

```
db2 update database configuration for xmeta using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for iadb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
```



```

db2 start db manager

db2 update database configuration for dsosb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

db2 update database configuration for esdbdb2 using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager

```

Wait for 2 or 3 minutes and check HADR status.

```
db2pd -db xmeta -hadr db2pd -db iadb -hadr db2pd -db dsodb -hadr db2pd -db esdbdb2 -hadr
```

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. Hence before the sample schedules are started, you have to make sure that one full backup of the respective database is available. Verify available backups for a database by using 'db2adutl query db <DATABASE_NAME>'. Open Putty or terminal in the IIS Passive Services machine, switch to db2inst1 user to run this command. You have to execute the shell script db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in the IIS Passive Services Machine, incase there are no full backups available.

Before running db2FullBackup.sh you must consider, where to store database archive logs.

Online full database backups

A sample scheduler named as SERVICEP-DB-FULL-WEEKLY is configured to take backups of online full database. This schedule invokes db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Passive Services Machine.

SERVICEP-DB-FULL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

In order to enable SERVICEP-DB-FULL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICEP-DB-DOMAIN SERVICEP-DB-FULL-WEEKLY expiration=never
```

Above command enable SERVICEP-DB-FULL-WEEKLY scheduler without expiry date. After enabling SERVICEP-DB-FULL-WEEKLY schedule, start 'dsmcad' service from IIS Passive Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Services and run `service dsmcad restart` using root user.

Online incremental database backups

A sample scheduler named as SERVICEP-DB-INCREMENT-DAILY is configured to take online incremental database backups. This schedule invokes db2IncrementalBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Passive Services Machine.

SERVICEP-DB-INCREMENT-DAILY scheduler is configured to execute from Monday to Saturday at midnights, except Sunday.

In order to enable SERVICEP-DB-INCREMENT-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICEP-DB-DOMAIN SERVICEP-DB-INCREMENT-DAILY expiration=never
```

Above command enable SERVICEP-DB-INCREMENT-DAILY scheduler without expiry date. After enabling SERVICEP-DB-INCREMENT-DAILY schedule, start 'dsmcad' service from IIS Passive Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Services machine and run `service dsmcad restart` using root user.

Files and Directories Backup

A sample scheduler named as SERVICEP-FILES-DAILY is configured to take backups of files and folders. This schedule invokes service_FilesDaily.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Passive Services Machine.

SERVICEP-FILES-DAILY scheduler is configured to execute everyday at midnight.

service_FilesDaily.sh executes commands to take backup of below files and folders.

- /root/keyfile
- /home_/db2inst1/db2keystore.p12
- /etc/sysconfig/iptables
- /etc/services
- /etc/inittab
- /opt/IBM/sslkey.sth
- /opt/IBM/sslkey.kdb
- /opt/IBM/sslkey.rdb
- /opt/IBM/Plugins/config/*
- /opt/IBM/HTTPServer/conf/*

service_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable SERVICEP-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule SERVICEP-DOMAIN SERVICEP-FILES-DAILY expiration=never
```

Above command enable SERVICEP-FILES-DAILY scheduler without expiry date.

After enabling SERVICEP-FILES-DAILY schedule, start 'dsmcad' service from IIS Passive Services machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Services machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script /opt/tivoli/tsm/client/ba/bin/service_FilesDaily.sh available in IIS Passive Services machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Passive Services machine is *servicep*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

```
update stgpool servicep-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool servicep-cloud-pool identity=<USERNAME>
update stgpool servicep-cloud-pool password=<PASSWORD>

update stgpool servicep-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
```

```
update stgpool servicep-arc-cloud-pool identity=<USERNAME>
update stgpool servicep-arc-cloud-pool password=<PASSWORD>
```

IIS Active Engine machine (main)

To take backups of the artifacts present in the IIS Engine machine, a sample policy domain configuration named as MAIN-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Engine machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

IIS artifacts backup using ISTOOL

A sample scheduler named as MAIN-ISTOOL-WEEKLY is configured to take backups of ISTool export configuration. This schedule invokes `istool_assets.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder in IIS Engine Machine. Before enabling MAIN-ISTOOL-WEEKLY schedule, you must update `istool_assets.sh` and provide value for PASSWORD field.

MAIN-ISTOOL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

`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 configuration

Archive file generated by `istool.sh` is stored in `/home_/istool/`. `istool_assets.sh` executes Spectrum Protect server selective backup command to take backup of `istool` generated file.

In order to enable MAIN-ISTOOL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-ISTOOL-WEEKLY expiration=never
```

Above command enable MAIN-ISTOOL-WEEKLY scheduler without expiry date.

After enabling MAIN-ISTOOL-WEEKLY schedule, start 'dsmcad' service from IIS machine. `dsmcad` provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

Files and Directories Backup

A sample scheduler named as MAIN-FILES-DAILY is configured to take backups of files and folders. This schedule invokes `main_FilesDaily.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Engine Machine.

MAIN-FILES-DAILY scheduler is configured to execute everyday at midnight.

`main_FilesDaily.sh` executes commands to take backup of below files and folders.

- `/root/keyfile`
- `/opt/IBM/InformationServer/Server/Projects/*`
- `/etc/sysconfig/iptables`
- `/opt/IBM/InformationServer/Server/MsgHandlers/*`
- `/opt/IBM/InformationServer/Server/Configurations/*`
- `/opt/IBM/InformationServer/Updates/*`
- `/opt/IBM/InformationServer/Server/DSODB/*.cfg`
- `/opt/IBM/InformationServer/Server/DSEngine/dsenv`
- `/etc/services`
- `/etc/inittab`

main_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable MAIN-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-FILES-DAILY expiration=never
```

Above command enable MAIN-FILES-DAILY scheduler without expiry date.

After enabling MAIN-FILES-DAILY schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script `/opt/tivoli/tsm/client/ba/bin/main_FilesDaily.sh` available in IIS Engine machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Engine machine is *main*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like `<ORDER_ID>-/s/m/l-bckp1`.

```
update stgpool main-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool main-cloud-pool identity=<USERNAME>
update stgpool main-cloud-pool password=<PASSWORD>

update stgpool main-arc-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool main-arc-cloud-pool identity=<USERNAME>
update stgpool main-arc-cloud-pool password=<PASSWORD>
```

IIS Passive Engine machine (mainp)

To take backups of the artifacts present in the IIS Passive Engine machine, a sample policy domain configuration named as MAINP-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Passive Engine machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

As this machine acts as passive, the user must not run any schedulers or configuration changes provided for this machine until the machine behaves as active one.

Don't run below schedulers or configuration changes at the starting. Start these schedulers and make the configuration changes only if this machine is taken over as an Active Engine machine

IIS artifacts backup using ISTOOL

A sample scheduler named as MAINP-ISTOOL-WEEKLY is configured to take backups of IStool export configuration. This schedule invokes `istool_assets.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder in IIS Passive Engine Machine. Before enabling MAINP-ISTOOL-WEEKLY schedule, you must update `istool_assets.sh` and provide value for PASSWORD field.

MAINP-ISTOOL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

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 configuration

Archive file generated by istool.sh is stored in /home_/istool/. istool_assets.sh executes Spectrum Protect server selective backup command to take backup of istool generated file.

In order to enable MAINP-ISTOOL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section.
[Starting Command Builder](#)

```
update schedule MAINP-DOMAIN MAINP-ISTOOL-WEEKLY expiration=never
```

Above command enable MAINP-ISTOOL-WEEKLY scheduler without expiry date.

After enabling MAINP-ISTOOL-WEEKLY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

Files and Directories Backup

A sample scheduler named as MAINP-FILES-DAILY is configured to take backups of files and folders. This schedule invokes main_FilesDaily.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Passive Engine Machine.

MAINP-FILES-DAILY scheduler is configured to execute everyday at midnight.

main_FilesDaily.sh executes commands to take backup of below files and folders.

- /root/keyfile
- /opt/IBM/InformationServer/Server/Projects/*
- /etc/sysconfig/iptables
- /opt/IBM/InformationServer/Server/MsgHandlers/*
- /opt/IBM/InformationServer/Server/Configurations/*
- /opt/IBM/InformationServer/Updates/*
- /opt/IBM/InformationServer/Server/DSODB/*.cfg
- /opt/IBM/InformationServer/Server/DSEngine/dsenv
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/*
- /etc/services
- /etc/inittab

main_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable MAINP-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section.
[Starting Command Builder](#)

```
update schedule MAINP-DOMAIN MAINP-FILES-DAILY expiration=never
```

Above command enable MAINP-FILES-DAILY scheduler without expiry date.

After enabling MAINP-FILES-DAILY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script `/opt/tivoli/tsm/client/ba/bin/main_FilesDaily.sh` available in IIS Passive Engine machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Engine machine is *mainp*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like `<ORDER_ID>-/s/m/l-bckp1`.

```
update stgpool mainp-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool mainp-cloud-pool identity=<USERNAME>
update stgpool mainp-cloud-pool password=<PASSWORD>

update stgpool mainp-arc-cloud-pool cloudurl=https://<PRIVATE_URL>
update stgpool mainp-arc-cloud-pool identity=<USERNAME>
update stgpool mainp-arc-cloud-pool password=<PASSWORD>
```

IIS Compute machine (cmpt)

To take backups of the artifacts present in the IIS Compute machine, a sample policy domain configuration named as COMPUTE-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Compute machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

Files and Directories Backup

A sample scheduler named as COMPUTE-FILES-DAILY is configured to take backups of files and folders. This schedule invokes `compute_FilesDaily.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Compute Machine.

COMPUTE-FILES-DAILY scheduler is configured to execute everyday at midnight.

`compute_FilesDaily.sh` executes commands to take backup of below files and folders.

- `/root/keyfile`
- `/etc/sysconfig/iptables`
- `/etc/services`
- `/etc/inittab`

`compute_FilesDaily.sh` executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable COMPUTE-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule COMPUTE-DOMAIN COMPUTE-FILES-DAILY expiration=never
```

Above command enable COMPUTE-FILES-DAILY scheduler without expiry date.

After enabling COMPUTE-FILES-DAILY schedule, start 'dsmcad' service from IIS Compute machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Compute machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script `/opt/tivoli/tsm/client/ba/bin/compute_FilesDaily.sh` available in IIS Compute machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS compute machine is *compute*.

Use Operation center in updating the details of Object store details.

IIS Compute machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like `<ORDER_ID>-/s/m/l-bckp1`.

```
update stgpool compute-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool compute-cloud-pool identity=<USERNAME>
update stgpool compute-cloud-pool password=<PASSWORD>

update stgpool compute-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool compute-arc-cloud-pool identity=<USERNAME>
update stgpool compute-arc-cloud-pool password=<PASSWORD>
```

Schedules and Policies

Add new schedulers

Spectrum protect backup server comes with sample schedulers and the user can create new schedulers based on their requirement.

More information on creating new schedulers is available [here](#)

Update existing schedulers

Spectrum protect backup server comes with sample schedulers and the user can change existing schedulers according to their requirement or they can create a new scheduler.

The user can update the starting time of the scheduler, when to execute the scheduler or disable scheduler based on their requirements.

More information on updating schedulers is available [here](#)

Create new policies and domain

Spectrum protect backup server comes with default policies and domain for each machine. The user can new policies and domains based on their requirement.

More information about adding new policies and domain is available [here](#)

The user can specify their own rules on when to take backup, archive and data retention requirements. More information about these are available [here](#)

Start Spectrum Protect after server restart or reboot

If Spectrum Protect server is restarted or rebooted, follow these steps to start both Spectrum Protect instances and Operation center.

Open Spectrum Protect server using putty or terminal from the IIS Windows Client machine. Execute the following commands.

```
cd /bckp/opt/tivoli/tsm/server/bin/  
./bckp/tsminst1/sql1lib/db2profile  
./dsmsevr -u tsminst1 -i /bckp/tsminst1 -q &
```

This will start Spectrum Protect server 1.

```
./bckp/tsminst2/sql1lib/db2profile  
./dsmsevr -u tsminst2 -i /bckp/tsminst2 -q &
```

This will start Spectrum Protect server 2.

Starting Operation center

```
cd /bckp/opt/tivoli/tsm/ui/Liberty/bin  
service opscenter.rc status  
service opscenter.rc start
```

Initially Spectrum Protect server second instance will be down, it'll take 5 to 10 minutes to come up. You can check if both instances are up, under overview tab in operation center console.

https://<PUBLIC_IP>:11090/oc or https://<PRIVATE_IP>:11090/oc

If any of Spectrum protect client (like Services or Engine machine) is restarted run `dsmcad service Open` Putty or terminal in Spectrum protect client machine which is restarted and run `service dsmcad start` using root user.

Protecting the master encryption key

Data encryption and decryption is handled automatically by the Spectrum Protect server and does not require any user action apart from some initial configuration. To encrypt data for cloud-container storage pools, the server uses a master encryption key, which is created when the server password is set. The master encryption key is itself encrypted, and is stored as part of the server password file.

The master encryption key is stored in the server password file, `/bckp/tsminst1/dsmsevr.pwd` for the first server instance and `/bckp/tsminst2/dsmsevr.pwd` for the second server instance in the Spectrum Protect server machine. The master encryption key is encrypted by a different key, so the master encryption key is itself protected. The master encryption key is re-encrypted whenever the server password is set by the `SET SERVERPASSWORD` command, so the user can issue this command periodically to further protect the key.

To decrypt data that was sent to encrypted cloud-container storage pools, the master encryption key is required. For this reason, it is important that the server password file is protected. If the server password file is lost or corrupted, the server cannot decrypt the data.

It is recommended that the user copy these files to Object Store or some secure location.

More details about master key is [here](#)

Spectrum Protect server database backup

In case of a scenario, where the Spectrum Protect server is no longer accessible, the Spectrum Protect server can be restored back to its latest state. To recover or restore any Spectrum Protect server, the following artifacts pertaining to that are required:

- Database used by Spectrum Protect for its functionality (.dbv)
- Metadata volume history file (volhist.out)
- Device configuration file (devconfig.out)

Since there are two Spectrum Protect server instances used in this offering, we need to backup two sets of the above mentioned artifacts. Scripts have been created to store these artifacts in a specific location locally. These artifacts are mandatorily required in the scenario where the Spectrum Protect server has to be restored. Hence, **it is recommended that the user should back transfer these files to a secure place.**

Scripts have been created within the Spectrum Protect server to take backup of these artifacts. Administrative schedulers are used to run these scripts. The details of these scripts and their relation to the two server instances can be found below.

First instance

An administrative schedule named *tminst1db*, created in the first instance, is used to call this script named *backuptsms*. The schedule calls this script everyday at 18:00:00. The script, *backuptsms*, creates the artifacts required to restore the Spectrum Protect Server in the */bckp/tsmdbbckps/* location in the Spectrum Protect server. The script also deletes the old backup artifacts and maintains only the latest two versions at any point of time. The script can be queried by using the following command in the command builder.

```
query script backuptsms format=lines
```

Second instance

An administrative schedule named *tminst1db*, created in the second instance, is used to call this script named *backuptsmt*. The schedule calls this script everyday at 19:00:00. The script, *backuptsmt*, creates the artifacts required to restore the Spectrum Protect Server in the */bckp/tsmdbbckpt/* location in the Spectrum Protect server. The script also deletes the old backup artifacts and maintains only the latest two versions at any point of time. The script can be queried by using the following command in the command builder.

```
query script backuptsmt format=lines
```

These scripts also clears the Spectrum Protect database archive logs.

Spectrum Protect server Inventory Expiration

As mentioned in the [Retention Policy](#) section, a file can be present in 3 states: active, inactive and expired. Once the file has reached the "expired" mode, it has to be manually deleted from the Spectrum Protect server to free up the space in order to take continuous backup. Inventory Expiration enables us to delete these expired artifacts.

Expire Inventory command might take several minutes to hours some times which results in slowness of the Spectrum Protect Server. This command should be scheduled to run only when the Spectrum Protect server is not busy.

Schedules have been created for both the Spectrum Protect server instances, which call scripts to execute the Inventory Expiration command. Information on these scripts and schedules have been given below.

First instance

An administrative schedule named *EXPIRES* is used to call the script, *EXPIRES*, every day at 07:00:00. It can be queried by using the following command in the command builder.

```
query script EXPIRES format=lines
```

Second instance

An administrative schedule named *EXPIRET* is used to call the script, *EXPIRET*, every day at 07:00:00. It can be queried by using the following command in the command builder.

```
query script EXPIRET format=lines
```

More details about expire inventory details are [here](#)

Restoring Backups

This section covers the steps involved to restore various IIS artifacts to IIS Services and Engine machine. Before any steps to restore are executed, the services of IIS must be stopped. Follow the following steps to stop the services and then restore the artifacts. After restoring the artifacts, start the services again.

Steps to stop IIS services.

1. Login to IIS Engine machine as root and execute the following in Putty.

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSEngine
./dsenv; bin/uv -admin -stop;
```

2. Stop Node Agents

```
su - root
cd /opt/IBM/InformationServer/ASBNode/bin
./opt/IBM/InformationServer/Server/DSEngine/dsenv
./NodeAgents.sh stop
```

3. Stop Appwatcher process

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSODB/bin
./DSAppWatcher.sh -stop
```

4. Stop WAS Login to IIS Services machine

```
su - root
cd /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin
./stopServer.sh server1 -user wasadmin -password <PASSWORD>
```

Steps to start IIS services

1. Start WAS Login to IIS Services machine

```
su - root
cd /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin
./startServer.sh server1 -user wasadmin -password <PASSWORD>
```

2. Start DS Engine services Login to Engine tier machine

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSEngine
./dsenv; bin/uv -admin -start;
```

3. Start node Agents

```
su - root
cd /opt/IBM/InformationServer/ASBNode/bin
./opt/IBM/InformationServer/Server/DSEngine/dsenv
./NodeAgents.sh start
```

4. Start Appwatcher process

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSODB/bin
./DSAppWatcher.sh -start
```

Restoring DB2 database

The following db2 databases are available in IIS on Cloud Enterprise Edition in IIS Services machine:

- xmeta
- iadb

- dsodb
- esbdb2

The databases are referred to as <DATABASE_NAME> in the upcoming sections.

Follow the mentioned steps in order to restore a specific database back to a specific timestamps

1. Drop existing database.
2. Restore database.
3. Verify if the database is restored.

Drop existing database

1. Open terminal for IIS Services machine, switch to db2inst1 user using command `su - db2inst1`
2. `db2 LIST APPLICATIONS` , to check if any applications are connected to this database.
3. Using `db2adutl` command , check available full or incremental backups. Note down timestamp which you are going to restore. Once database is dropped you can't see available full or incremental backups.
4. Drop the database using `db2 drop db <DATABASE_NAME>` , if you face any issues that means database is connected to applications , we need to close all connections to database before dropping it.

```
db2 connect to <DATABASE_NAME> user <DATABASE_USER> using <PASSWORD>
db2 quiesce db immediate force connections
db2 connect reset;
db2 LIST APPLICATIONS
db2 terminate
db2 force application all
db2 drop database <DATABASE_NAME>
```

5. In order to make sure , there is no database named <DATABASE_NAME> , execute list command. `db2 list db directory`

Restore Database

1. Using `db2adutl` command , check available full or incremental backups.
2. Select full or incremental backup timestamp which needs to be restored.
3. `db2 restore db <DATABASE_NAME> use tsm taken at 20170526063832 ENCRYPT`
4. In above statement '20170526063832' is timestamp when the DB backup was taken.
5. "use TSM" is used which means , we are restoring a database which is stored in Spectrum Protect server.
6. "ENCRYPT" is used as existing database is db2 native encrypted.
7. While taking backup "include logs" is used , so if needed you can use "include logs" option, when there is a need to extract transaction logs needed in restore scenario.
8. Select incremental backup timestamp if you want to restore till that date.


```
db2 restore db <DATABASE_NAME> incremental automatic use tsm taken at 20170605134603 1. On top of full backup , we are restoring incremental backup.
```
9. `db2 rollforward db <DATABASE_NAME> to end of logs and complete.`
 - a. Used to rollforward database till end of logs.

Verify database

1. Connect to <DATABASE_NAME> database using following command. `db2 connect to <DATABASE_NAME> user <DATABASE_USER> using <PASSWORD>`
2. Verify any table to check restore process is done and expected data is available.
3. Terminate database using command, `db2 terminate`
4. Verify if db2 native encryption is available.


```
db2 connect to <DATABASE_NAME> user db2inst1 using <PASSWORD>
```

```
db2 "SELECT * FROM TABLE (SYSPROC.ADMIN_GET_ENCRYPTION_INFO())"
```

5. Verify if AES is available in the output.

Restoring files and directories

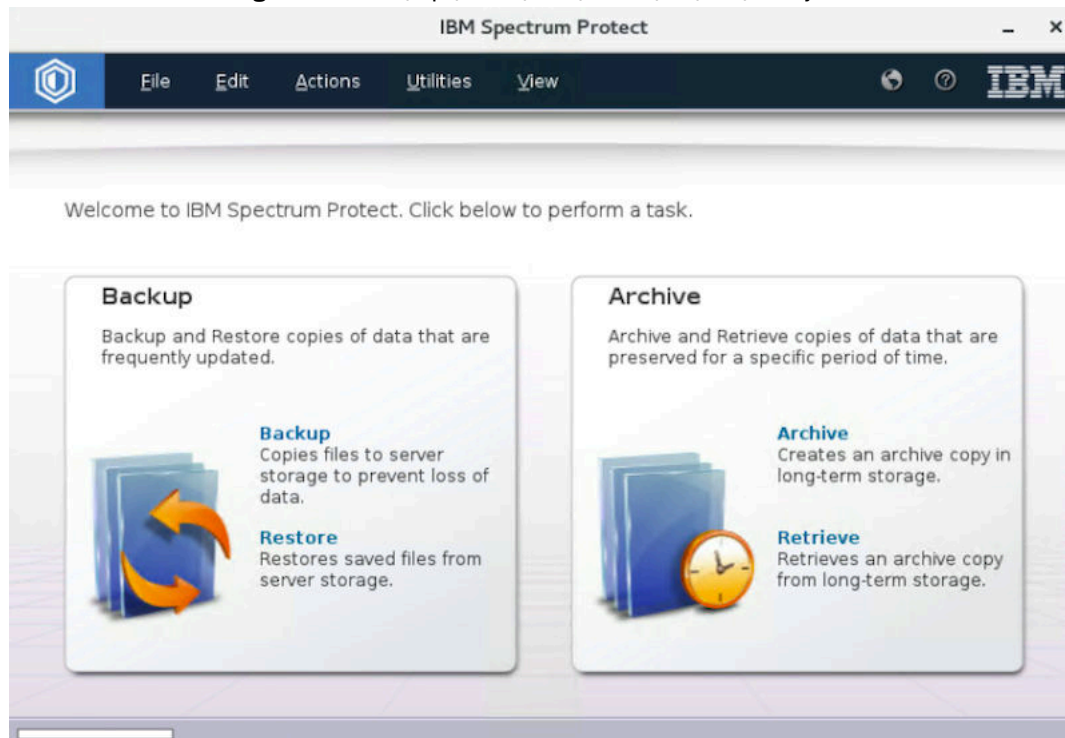
As has been discussed in the Backup Section of IIS on Cloud Enterprise Edition, sample policies and schedules are created to backup specific files and directories to the Spectrum Protect Server from individual machines like the IIS Services machine, IIS Engine machine and IIS Compute machine.

There are two things to be particular about when a file/directory is being restored

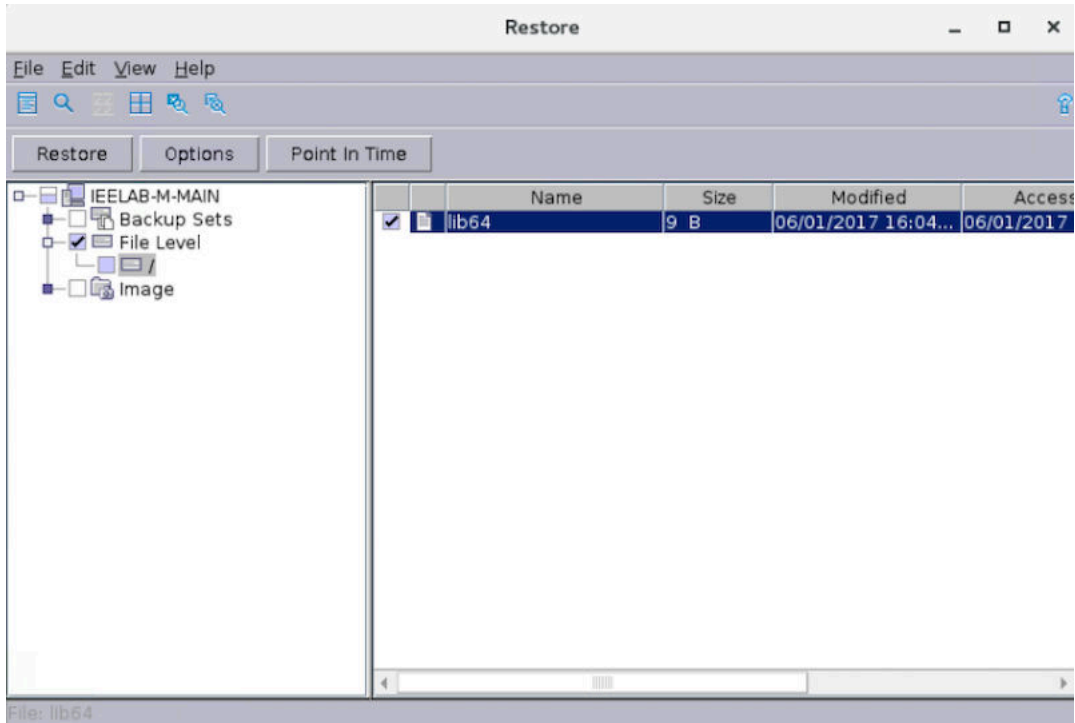
- Whichever file/directory is being restored, make sure that a copy of the same is created as a temporary file.
- The user privileges of the file/directory must be noted before restore. After restore, if there are any discrepancies, the user privileges must be set by the user for the restored files/directories so that they match with the ones of the current file/directory.

The method to restored has been explained for the IIS Services machine. The procedure is the same for all other machines as well where Spectrum Protect Client is installed. The files/directories can be restored in the following way:

1. Log into the IIS Services machine as root in a GUI session.
2. Execute the following command: `/opt/tivoli/tsm/client/ba/bin/dsmj`



- Under the "Restore" section, all the files and directories which have been backed up will be available to be selected for restore.



- Select the files to be restored and click on the restore button.

Restoring WAS artifacts

The following WAS artifacts are backed up to the Spectrum Protect Server. WAS artifacts are to be restored in the IIS Services machine.

- InfoSphere profile
- InfoSphere profile configurations

InfoSphere profile

The backup file of InfoSphere profile is stored in the following location: /home_/WAS_Backup/IIS_AppServer_backup.zip

The user can use the file current in this location, which is the backup taken on the previous Sunday. Or the user can choose different backup versions of this file. Different versions of backup file can be restored using the /opt/tivoli/tsm/client/ba/bin/dsmj. Once the required backup file has been made available, execute the following commands as "root" user to restore the InfoSphere profile.

```
/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin/stopServer.sh server1 -username wasadmin -password PASSWORD
```

Rename the "InfoSphere" profiles directory in /opt/IBM/WebSphere/AppServer/profiles to "InfoSphere1". Execute the following commands.

```
/opt/IBM/WebSphere/AppServer/bin/manageprofiles.sh -restoreProfile -backupFile <location_of_backup_file>
/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin/startServer.sh server1
```

InfoSphere profile configurations

The backup file of InfoSphere profile is stored in the following location: /home_/WAS_Conf/IIS_AppServerConfig.zip

The user can use the file current in this location, which is the backup taken on the previous Sunday. Or the user can choose different backup versions of this file. Different versions of backup file can be restored

using the /opt/tivoli/tsm/client/ba/bin/dsmj. Once the required backup file has been made available, follow the procedure to restore the InfoSphere profile configurations.

1. Open terminal for IIS Services machine using root user.
2. Execute the following command: /opt/IBM/WebSphere/AppServer/bin/restoreConfig.sh
<location_of_backup_file> -nostop -username wasadmin -password <PASSWORD> -profileName InfoSphere

Restoring ISTOOL asset

InfoSphere Information Server provides a backup utility called "ISTOOL" which enable us to backup assets related to Information Governance Catalog, Information Analyzer etc. If these components are to be restored, they are backed up in the following location in the IIS Engine machine.

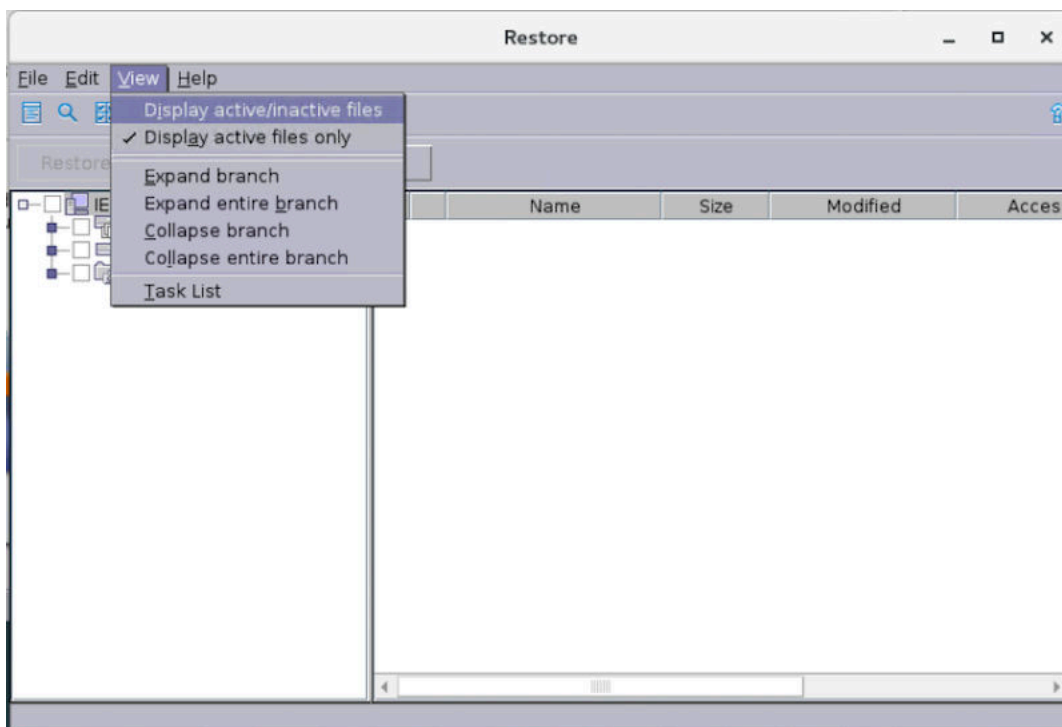
```
/home_/istool/istool.isx
```

Follow the procedure to restore these assets from the .isx archive file

```
cd /opt/IBM/InformationServer/Client/istools/cli  
import -dom <URL_IIS_ENGINE>:9443 -u isadmin -p <PASSWORD> -archive "/home_/istool/istool.isx" -  
all
```

Viewing and restoring multiple versions of a specific file

Spectrum Protect enables the user to restore not just the latest version of the file backup but older versions as well, if they are available in the Spectrum Protect server. In order to view the available versions, you can toggle the "Display active/inactive files" option in the "View" tab of Restore window.



The versions available in this option will not be more that 30 days old. If an older version is required, the user can use the cloud object storage to restore backup versions which are older that 30 days but not more than 365 days.

Restoring from Cloud Object Storage

As mentioned in the Spectrum Protect Setup, there are two Spectrum Protect Server instances. First instance is connected to local storage with the Spectrum Protect server and the second instance uses the Object Storage for storing backup.

Cloud Object storage stores backup of artifacts which are up to 365 days old. In order to restore the from the cloud object storage, follow the procedure.

Halt the first instance of the Spectrum Protect Server

Shut down the first instance of Spectrum Protect Server (tsminst1), so that the Spectrum Protect client can connect to the second instance of the Spectrum Protect server. The second instance of Spectrum Protect server is connected to the Cloud Object Storage.

1. Connect to the Operations Center. Start the command builder.
2. Execute the following commands.

```
DISABLE SESSIONS
QUERY SESSIONS
CANCEL SESSIONS
HALT
```

This stops the first server instance (tsminst1) and connects the Spectrum Protect Clients to the second server instance (tsminst2).

Change the SSL Certificate to connect to second instance of Spectrum Protect server

Certificate files of first instance of spectrum protect must be deleted so that the client machine can connect to the second instance. These files are located at the /opt/tivoli/tsm/client/ba/bin/ location of the Spectrum Protect client machine.

```
dsmcert.crl
dsmcert.kdb
dsmcert.rdb
dsmcert.sth
```

Execute the following commands to create the ssl certificate key for the second server instance.

```
cd /opt/tivoli/tsm/client/ba/bin/
gsk8capicmd_64 -keydb -create -populate -db dsmcert.kdb -pw <Spectrum_Protect_Server_Password> -
stash
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed -label "<description>" -file /opt/tivoli/tsm/
client/ba/bin/tsminst2/cert256.arm -format ascii
chmod 777 dsmcert.*
```

Use the /opt/tivoli/tsm/client/ba/bin/dsmj in the client machine to restore the required files/directories. The GUI will be opened in read only format. This implies we can only restore the files that have already been backed up.



After the restoration is completed, to revert back to the original configuration (ie. connecting back to first instance of Spectrum Protect Server), follow the procedure in the next sub-section.

Connecting back to first instance of Spectrum Protect server

The first server instance (tsminst1) has to be started to revert back to the original configuration of backup and its schedules. Open the terminal and execute the following commands in Spectrum Protect Server machine.

```
./bckp/tsminst1/sqllib/db2profile
/bckp//opt/tivoli/tsm/server/bin/dsmserve -u tsminst1 -i /bckp/tsminst1
```

The above mentioned commands starts the first server instance. Next, the SSL certificates in the Spectrum Protect client machine has to be made compliant to the first server instance.

In the Spectrum Protect client machines, execute the following steps.

1. Delete the following files in /opt/tivoli/tsm/client/ba/bin/

```
dsmcert.crl  
dsmcert.kdb  
dsmcert.rdb  
dsmcert.sth
```

2. Execute the following commands to change the ssl certificate key.

```
cd /opt/tivoli/tsm/client/ba/bin/  
gsk8capicmd_64 -keydb -create -populate -db dsmcert.kdb -pw  
<Spectrum_Protect_Server_Password> -stash  
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed -label "<description>" -file /opt/  
tivoli/tsm/client/ba/bin/tsminst1/cert256.arm -format ascii  
chmod 777 dsmcert.*
```

Chapter 3. Information Server on Cloud Data Quality

IBM® Information Server on Cloud Data Quality is included in Information Server on Cloud Enterprise Edition, but it can be used as an independent service. Use Information Server on Cloud Data Quality to cleanse data and monitor data quality.

Overview

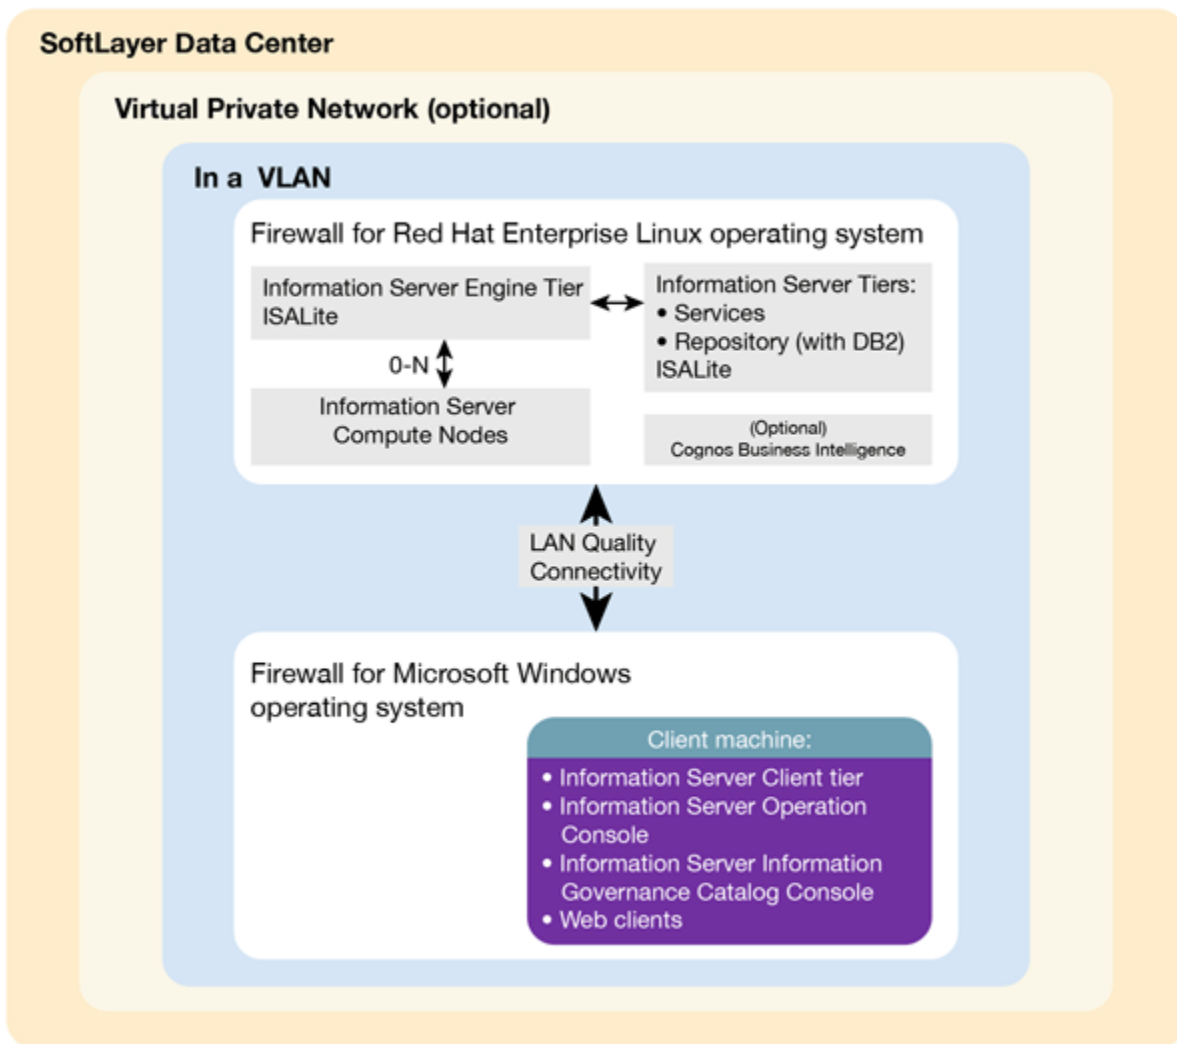
IBM® Information Server on Cloud Data Quality is a data integration software platform that helps organizations derive more value from the complex, heterogeneous information spread across their systems. You can establish and manage high-quality data by using Information Server on Cloud Data Quality. You can cleanse data and monitor data quality on an ongoing basis to turn your data into trusted information.

Information Server on Cloud Data Quality helps you to do the following tasks:

- Deliver customizable data-cleansing functions in batch and in near real time
- Monitor and maintain data quality
- Provide a unified environment and support for advanced data validation rules

Information Server on Cloud Data Quality uses the characteristics of software-as-a-service (SaaS). You select the plan size based on your needs. IBM provisions the machines and deploys the Information Server on Cloud Data Quality software.

The following figure shows the topology of the server and client machines in a typical deployment.



As a hosted offering, you have the same control over your data in the cloud as in the on-premises system:

- Actively monitor and report any issues that you encounter with IBM Software as a Service (SaaS).
- Maintain the software platform of your cloud offering and the operating system to meet your security standards.
- Maintain software firewalls on servers that face the internet in a manner to provide required protection.
- Develop parallel jobs to transform and cleanse data, and develop server jobs to transform data. Establish connectivity between data sources and applications. Develop your own workload, business rules, monitoring, and scheduling for all jobs. You are responsible for the quality and performance of programs, applications, and jobs that you develop.
- Ensure the continuity, compatibility, and performance of the IBM SaaS platform by installing only permissible software, including any open source packages.
- Regularly upgrade the environment and operating system of your cloud offering.
- Create and maintain regular backups of data.
- Create and maintain high availability configurations.

The following managed add-on services are available to maintain and manage the infrastructure:

Jump start

This setup service provides up to 50 hours of remote consulting time for startup activities.

Accelerator

This setup service provides up to 50 hours of remote consulting time to perform various scoped activities.

Silver

This service provides monthly remote consulting time for operations and maintenance activities.

Gold

This service provides monthly remote consulting time for operations and maintenance activities. The service includes everything that is provided by the Silver service and delivers extra activities.

Note the following limitations and restrictions of Information Server on Cloud Data Quality:

- If your offering is designated as "Non-Production", Information Server on Cloud Data Quality can be deployed only as part of your development and test environments for internal non-production activities. These activities include, but are not limited to: testing, performance tuning, fault diagnosis, internal benchmarking, staging, quality assurance activity, developing internally used additions or extensions to the offering by using published application programming interfaces.
- Users must not modify the configuration file that is needed to run the job for parallel processing.

Available configurations

IBM® Information Server on Cloud Data Quality servers for the small and medium plans are virtual servers with dedicated CPUs. The servers in the large plan are in a bare metal environment.

Select the offering plan that fits your usage and environment needs.

Table 10: Offering sizes: small production and non-production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
Service metadata	16	4	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
Enterprise Search	32	8	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN
(Optional) Cognos	16	4	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	500 GB SAN

Table 11: Offering sizes: medium production and non-production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN
Service metadata	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN
Enterprise Search	64	16	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	1TB SAN

Table 11: Offering sizes: medium production and non-production (continued)

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
(Optional) Cognos	32	8	1 Gbps with 1000 GB bandwidth	100 GB SAN	1 TB SAN

Table 12: Offering sizes: large production

Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk
Engine	64	12	1 Gbps with 5000 GB bandwidth	1.7 TB SSD	1.7 TB SSD
Service metadata	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
Enterprise Search	64	12	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD
(Optional) Cognos	32	8	1 Gbps with 5000 GB bandwidth	960 GB SSD	960 GB SSD

Small and medium offerings include the following configuration:

- One virtual server machine with services and repository tiers.
- One virtual server machine with the engine tier.
- One virtual server machine with Enterprise Search Node
- Optional: One virtual machine with IBM Cognos® Business Intelligence.
- One client machine for small offerings and three client machines for medium offerings.

The number of client machines is based on the number of concurrent users. For small offerings, two concurrent users are allowed. For medium offerings, five concurrent users are allowed. By default, Microsoft Windows operating system allows two concurrent users to access the machine by using Remote Desktop Connection.

Large-size offering includes the following configuration:

- Two bare metal machines with Information Server on Cloud Data Quality. One machine has the services and repository tiers while the other machine has the engine tier.
- One bare metal machine with Enterprise Search Node
- Optional: One bare metal machine with IBM Cognos Business Intelligence.
- Five client machines based on 10 concurrent users.

Layout of IBM Information Server on Cloud Data Quality server and client disks

The layout of the Information Server on Cloud Data Quality server and client disks depends on the plan size of your system.

Virtual servers for small and medium plans

Information Server on Cloud Data Quality comes with three virtual servers. One server has the services and repository tiers, 2nd server has the engine tier, and third having enterprise search node. Optionally, you have another virtual server with IBM® Cognos® Business Intelligence.

The small and medium plans come with two Storage Area Network (SAN) disks. The Red Hat Enterprise Linux operating system is on the first SAN disk in both the small and medium plans. The second SAN disk is encrypted by using Linux Unified Key Setup (LUKS).

The encryption key details are provided in the Welcome letter from the IBM Sales Representative. It is recommended that you add your own key and remove the supplied key before you use the system.

The product is installed on the /opt directory. User data can be stored on /data directory. Both directories are on the partition /dev/xvdc1 that is encrypted.

Table 13: Small and medium disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	LUKS	/dev/xvdc1	500 GB for small plan. 1000 GB for medium plan	Secondary disk is /dev/xvdc	Product installation. Directories /data and /opt are created on this disk.
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging

Bare metal server for large plan

The large plan also comes with three servers. Two servers has a Solid State Drive (SSD) disk of about 960 GB for the services tier. The other server has an SSD disk of about 1.6 TB for the engine tier. Optionally, you can have additional server with IBM Cognos Business Intelligence with Solid State Drive (SSD) disk of about 960 GB. RAID level 1 implementation makes them appear as a single disk.

The disk is divided into four partitions. The Red Hat Enterprise Linux operating system is on a 10 GB partition. The boot data is on a 256 MB partition. The swap space is on a 2 GB partition. The remaining space is on another partition that is encrypted by using LUKS.

Table 14: Bare metal disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	256 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda2	2 GB	/dev/sda	Swap space for paging

Table 14: Bare metal disk layout (continued)

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/	None	/dev/sda3	10 GB	/dev/sda	Operating system data
/disk1	LUKS	/dev/sda5	About 900 GB for the services tier and enterprise search node, about 1.6 TB for the engine tier, and about 900 GB for the optional Cognos	/dev/sda	Product installation. Directories /data, /installables, and /opt are created on this disk.

Client for all plans

The Information Server on Cloud Data Quality client machine configuration is the same for all plans. The client machine has two SAN disks that are 100 GB each. One disk is drive C for the Microsoft Windows operating system. The other disk is drive F, and it is an empty disk.

Information roadmap

This roadmap lists information resources that are available for users who are new to the data quality products from IBM®. These products include IBM InfoSphere® DataStage® and QualityStage®, IBM InfoSphere Information Analyzer, and IBM InfoSphere Information Governance Catalog. These resources provide information about various subject areas, such as learning basic skills and troubleshooting.

IBM InfoSphere DataStage and QualityStage

- **Getting started with InfoSphere DataStage and InfoSphere QualityStage** Use tutorials to learn the skills that you need to develop parallel jobs that transform data and parallel jobs that cleanse data.
- **Designing DataStage and QualityStage jobs** You design InfoSphere DataStage and QualityStage jobs using the IBM InfoSphere DataStage and QualityStage Designer client.
- **Cleansing data with InfoSphere QualityStage jobs** The cleansing process can include, but is not limited to, eliminating redundant, obsolete, or inaccurate data. Clean data is a critical component for accurate information, reports, and analyses. Throughout your organization, people make business decisions based on data that is provided to them. By cleansing data, you provide high-quality data.
- **Standardizing address data with InfoSphere QualityStage** IBM InfoSphere QualityStage add-on modules standardize address records and match those records to postal validation reference files. These modules let you concentrate on the quality of your address data so that you do not have to become an expert in international postal standards.
- **Running InfoSphere DataStage and QualityStage jobs** You run your InfoSphere DataStage and QualityStage jobs from the InfoSphere DataStage and QualityStage Director client.
- **Monitoring jobs** You can use the IBM InfoSphere DataStage and QualityStage Operations Console to monitor the job runs, services, system resources, and workload management queues on several IBM InfoSphere Information Server engines.
- **Administering workload management** You can use the workload management queues to control the starting of parallel and server jobs.
- **Administering projects** IBM InfoSphere DataStage and QualityStage jobs are organized in projects, along with associated design items.

IBM InfoSphere Information Analyzer

- **Methodology and best practices** Use IBM InfoSphere Information Analyzer to understand the content, structure, and overall quality of your data at a given point in time.
- **Configuring InfoSphere Information Analyzer** After InfoSphere Information Analyzer is installed, you must complete additional configuration steps before you can use it. You should also validate any settings created during installation.
- **Managing metadata** After you import metadata into the metadata repository, you can add or modify information about the metadata such as the description of a table. You can also add information such as contacts, policies, and terms in IBM InfoSphere Information Governance Catalog to the imported schemas, directories, tables, files, and data fields.
- **Analyzing data with the on-premise client** Before you can begin analyzing data, you must understand the project environment, open a project, review the dashboard, and decide which analysis jobs or analysis reviews you want to conduct.
- **Analyzing data with InfoSphere Information Analyzer thin client** You can analyze data sets, and view and edit analysis results in a browser. You can also view data quality scores for data sets and columns, view data rules and quality rules associated with data sets, and create or delete quality rules.
- **Validating data by using data rules** You can define and run data rules, which evaluate or validate specific conditions associated with your data sources. Data rules can be used to extend your data profiling analysis, to test and evaluate data quality, or to improve your understanding of data integration requirements.
- **Publishing and transferring analysis results** You can view analysis results and publish it to the metadata repository. The analysis results can also be transferred into InfoSphere Optim™ and InfoSphere Guardium®.
- **Importing and exporting projects from the client** You can import and export projects and project assets and move them between instances of the IBM InfoSphere Information Server metadata repository by using the InfoSphere Information Analyzer import and export wizards.
- **Running and scheduling InfoSphere Information Analyzer jobs** You can run and schedule IBM InfoSphere Information Analyzer jobs outside of the InfoSphere Information Analyzer client.
- **Managing tables** You can manage reference tables and tables that are generated during analysis. You can open a table, delete a table, export a table to a file, add a table to the project, and view the metadata in a table.
- **Developing InfoSphere Information Analyzer applications with the HTTP API** A Hypertext Transfer Protocol (HTTP) application programming interface (API) is provided with IBM InfoSphere Information Analyzer. You can develop applications with this API to access and analyze InfoSphere Information Analyzer content.
- **Reports for information analysis** You can create reports that summarize analysis results and show details about your project. Reports are saved in the metadata repository and can be accessed by any user who is authorized to view them.

IBM InfoSphere Information Governance Catalog

- **Overview** Learn about the features and benefits of InfoSphere Information Governance Catalog.
- **Designing the catalog** A catalog is an authoritative dictionary of the assets that are used throughout the organization. One of the main benefits of a well-designed catalog is increased trust and confidence in organization information. Planning, designing, and publishing a catalog involves several tasks.
- **Administering the catalog** The catalog is composed of glossary assets (terms, categories, information governance policies, and information governance rules) and information assets. You can assign security roles and permissions to users to control access to the catalog. You can also assign workflow roles, assign users as stewards, define custom attributes, and define external asset types. In addition, you can configure the display of glossary assets.
- **Governing your data** You can create, edit, and delete catalog assets. Some assets are created when they are imported into the catalog. You can assign stewards and assets, define and assign custom

attributes, and extend data flows for lineage reports by importing assets. In addition, you can configure assets for lineage analysis reports. Assets of different types that have a common business purpose can be grouped in a collection.

- **Viewing catalog content** You can view, search, browse, and query the catalog to find catalog assets.
- **Expanding catalog capabilities**
- **Look up terms** while you work in other applications from the Microsoft Windows desktop by using Glossary Anywhere.
- **Develop and extend applications** by using InfoSphere Information Governance Catalog for Eclipse. You can access your glossary content from within your Eclipse-based development tool.
- **Write client applications to access and author catalog content** by using InfoSphere Information Governance Catalog REST API. You can integrate your catalog content into other software tools and portals.

New in Information Server 11.7

- Various new features introduced in Information Server 11.7 can be found [here](#)
- A key addition to Information Server 11.7 is **Information Server Enterprise Search**

Getting started and using IBM Information Server on Cloud Data Quality

You must set up your connection to Information Server on Cloud Data Quality. Information Server on Cloud Data Quality provides all of the functions of its on-premises counterparts, IBM® InfoSphere® Information Analyzer, IBM InfoSphere Information Governance Catalog, and IBM InfoSphere DataStage® and QualityStage®. It is in an IBM SoftLayer® hosted environment.

Prerequisite: You must know the IP address and the credentials of an account on the Information Server on Cloud Data Quality server and client computers. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM Sales Representative.

The Information Server on Cloud Data Quality client is on a Microsoft Windows virtual machine that is hosted on SoftLayer. When you connect to the client, you can access the client user interfaces. McAfee anti-virus software is installed on the client machine.

The Information Server on Cloud Data Quality servers are on Red Hat Enterprise Linux virtual or bare metal computers that are hosted on SoftLayer. The services and repository tiers are on one computer. The engine tier is on the other computer. Optionally, IBM Cognos® Business Intelligence is on a separate Red Hat Enterprise Linux virtual or bare metal computer. When you connect to the servers, you can access the IBM InfoSphere Information Server engine, services, and repository tiers. You can restart Information Server on Cloud Data Quality and do administrative tasks.

The default firewall configuration of server machines allows SSH connections only from client machines. You must first connect to a client machine by using a remote desktop connection, and then from the client machine you can connect to server machines by using SSH. After you log in to a server machine, you can change the firewall configurations to allow SSH connections from other machines. Communication between the server and client systems happens through a private IP. If you want to access the server from an on-premises client machine, you must modify the iptable rules.

Note: The Add Subscription window in Subscription Manager is disabled when you use a host name in the URL. To enable the window, use an IP address in the Subscription Manager URL.

When you connect for the first time, follow these steps:

1. SSH into the any server machine using unique user (order id) provided in the welcome letter using port 4362.
2. Change the password and su to root user.
3. On the Information Server on Cloud Data Quality servers, run the [ISALite tool](#). The `IS_install_path` for the server computers is `/opt/IBM/InformationServer`

4. Connect the Information Server on Cloud Data Quality client to the Information Server on Cloud Data Quality servers by following these steps:
 - **If your local computer is in a Microsoft Windows environment**
 - a. On your local computer, go to the **Start** menu. Click **Accessories > Remote Desktop Connection**.
 - b. Enter the IP address of the Microsoft Windows computer that hosts your Information Server on Cloud Data Quality client. Click **Connect**.
 - c. In the Windows Security window, enter the user name and password for the Information Server on Cloud Data Quality client. The user ID, password, and IP address of the client are in your Welcome letter. **Important:** Do not include the domain name with the user name.
 - d. In the Information Server on Cloud Data Quality client, open the file C:\Windows\System32\drivers\etc\hosts. Make sure that an entry with the private IP exists in the file for the Information Server on Cloud Data Quality servers that you are connecting to. **Important** The server IP must be a private IP. You cannot open other Information Server on Cloud Data Quality clients when you use a public IP.
 - **If your local computer is in an Apple Mac environment**
 - a. On your local computer, install Microsoft Remote Desktop from the Apple App Store.
 - b. Click the Microsoft Remote Desktop icon, and then click **Open**.
 - c. In the Microsoft Remote Desktop window, click **New**.
 - d. In the Edit Remote Desktops window, supply the following information:
 - In the PC name field, type in the IP address of the cloud client machine.
 - In the User name and Password fields, type in the Windows user name and password that are in the Welcome letter.
5. Verify the connection and installation on the Information Server on Cloud Data Quality client by following these steps:
 - [Run the ISALite tool](#). The *IS_install_path* for the client computer is C:\IBM\InformationServer.
 - [Test the installation](#) of the Information Server on Cloud Data Quality client.
6. Optional: Enable multiple users to open remote sessions to the Information Server on Cloud Data Quality client by following these steps on the client computer:
 - a. [Create user accounts](#).
 - b. [Give users permission](#) to do a remote desktop connection. The number of concurrent users in a remote session to the client is based on the [offering size](#) of your Information Server on Cloud Data Quality. The small offering size has a maximum of two concurrent sessions. The medium offering size has a maximum of five concurrent sessions. The large offering size has a maximum of 10 concurrent sessions.
7. [Reset the password](#) for users and administrators on the Information Server on Cloud Data Quality servers.
8. Optional: If you choose to include IBM Cognos Business Intelligence in your environment, you must configure it. For details, see:
 - [IBM InfoSphere Information Governance Dashboard](#)
 - [Configure the IBM InfoSphere Business Glossary URI](#)

After the initial connection, you can do any of the following tasks:

- [Open the InfoSphere DataStage® and QualityStage® Designer client](#)
- [Connect to an on-premises computer](#)
- [Connect to an IBM dashDB™ database](#)
- [Connect to an on-premises DB2® database instance](#)
- [Perform general administration and security tasks](#)

- Open Information Server on Cloud Enterprise Edition clients from the client machine by following either method.
 - For browser clients, open Microsoft Internet Explorer. In the Favorites Bar,
 1. click the **IIS_Launchpad** bookmark. From the Launchpad window, click the icon of the client of InfoSphere Information Server.
 2. click the **Kubernetes_Dashboard** bookmark to login to the Kubernetes Dashboard sitting on your Enterprise Search node.
 - For thick clients, on your desktop click **Start > Programs > IBM Infosphere Information Server > client-name**.

You can connect to any of the following InfoSphere Information Server clients that were provisioned with their server component:

- IBM InfoSphere Information Governance Catalog
- IBM InfoSphere Information Analyzer
- IBM InfoSphere Metadata Integration Bridges and the metadata interchange agent
- IBM InfoSphere Metadata Asset Manager
- IBM InfoSphere Information Server istool command-line utility
- IBM InfoSphere Information Server Manager client
- Multi-Client Manager
- IBM InfoSphere DataStage and QualityStage Administrator
- IBM InfoSphere DataStage and QualityStage Designer
- IBM InfoSphere DataStage and QualityStage Director.

Related information

- [Enhancing security of Information Server on Cloud computers](#)

Chapter 4. DataStage on Cloud

IBM® DataStage® on Cloud provides a hosted environment that you configure and control. You can use DataStage on Cloud to extend the reach of your business by leveraging cloud offerings, while you reduce the costs that are associated with providing these services. DataStage on Cloud provides several different plans so that any size business can access the powerful and scalable ETL platform by IBM.

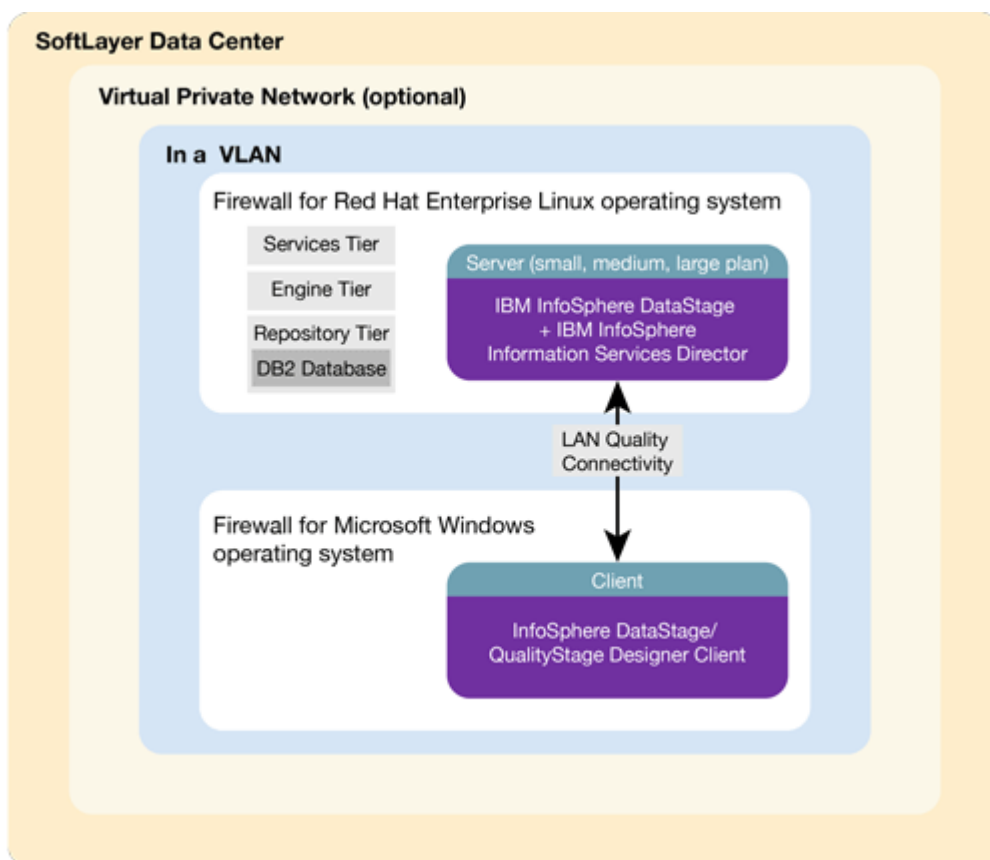
Overview

IBM® DataStage® on Cloud is a data integration tool for designing, developing, and running jobs that move and transform data. DataStage on Cloud provides all of the functions of its on-premises counterpart, IBM InfoSphere® DataStage.

DataStage on Cloud combines some of the characteristics of infrastructure-as-a-service (IaaS) and platform-as-a-service (PaaS). You select the plan size based on your needs. IBM provisions the machine and deploys the DataStage on Cloud software.

Overview of Information Server DataStage on Cloud High Availability

The following figure shows the topology of the server and client machines in a typical deployment without High Availability and backup Configurations.



As a hosted offering, you have the same control over your data in the cloud as in the on-premises system:

- Actively monitor and report any issues that you encounter with IBM Software as a Service (SaaS).
- Maintain the software platform of your cloud offering and the operating system to meet your security standards.
- Maintain software firewalls on servers that face the internet in a manner to provide required protection.

- Develop integration, transformation, and other jobs. Establish connectivity between data sources and applications. You can also develop your own workload, business rules, monitoring, and scheduling for all jobs. You are responsible for the quality and performance of programs, applications, and jobs that you develop.
- Ensure the continuity, compatibility, and performance of the IBM SaaS platform by installing only permissible software, including any open source packages.
- Regularly upgrade the environment and operating system of your cloud offering.
- Create and maintain regular backups of data.
- Create and maintain high availability configurations.

The following managed add-on services are available to maintain and manage the infrastructure:

Jump start

This setup service provides up to 50 hours of remote consulting time for startup activities.

Accelerator

This setup service provides up to 50 hours of remote consulting time to perform various scoped activities.

Silver

This service provides monthly remote consulting time for operations and maintenance activities.

Gold

This service provides monthly remote consulting time for operations and maintenance activities. The service includes everything that is provided by the Silver service and delivers extra activities.

Note the following limitations and restrictions of DataStage on Cloud:

- If your offering is designated as "Non-Production," DataStage on Cloud can be deployed only as part of your development and test environments for internal non-production activities. These activities include, but are not limited to: testing, performance tuning, fault diagnosis, internal benchmarking, staging, quality assurance activity, developing internally used additions or extensions to the offering by using published application programming interfaces.
- Users must not modify the configuration file that is needed to run the job for parallel processing.
- Usage of IBM InfoSphere Information Services Director with DataStage on Cloud has the following restrictions:
 - Batch jobs are restricted to Topology I batch jobs. For details, see [Designing IBM InfoSphere DataStage and QualityStage® jobs as services](#).
 - Restricted to SOAP over HTTP as service binding.

Available configurations

IBM® DataStage® on Cloud servers for the small and medium plans are virtual servers with dedicated CPUs. The servers in the large plan are in a bare metal environment.

Select the offering plan that fits your usage and environment needs.

Table 15: DataStage on Cloud available configurations

Offering	Environment	Memory (GB)	Number of Cores	Network speed (Gbps)	First disk	Second disk
Small	Production Non-production	16	4	1	100 GB Storage area network (SAN)	500 GB SAN
Medium	Production Non-production	32	8	1	100 GB SAN	1 TB SAN

Offering	Environment	Memory (GB)	Number of Cores	Network speed (Gbps)	First disk	Second disk
Large	Production	64	12	1	2-960 GB SSD	

Small and medium offerings include the following configuration:

- One virtual machine with DataStage on Cloud, IBM InfoSphere® Information Services Director
- One virtual machine with a Microsoft Windows operating system and all client utilities installed

Large-size offering includes the following configuration:

- One bare metal machine with DataStage on Cloud, IBM InfoSphere Information Services Director
- Two virtual machines with a Microsoft Windows operating system and all client utilities installed

Layout of IBM DataStage on Cloud server and client disks

The layout of the DataStage® on Cloud server and client disks depends on the plan size of your system.

Virtual server for small and medium plans

The small and medium plans come with two Storage Area Network (SAN) disks. The Red Hat Enterprise Linux operating system is on the first SAN disk in both the small and medium plans. The second SAN disk is encrypted by using Linux Unified Key Setup (LUKS) for production offerings of both these sizes.

The encryption key details are provided in the Welcome letter from the IBM® Sales Representative. It is recommended that you add your own key and remove the supplied key before you use the system.

The product is installed on the /opt directory. User data can be stored on /data directory. Both directories are on the partition /dev/xvdc1 that is encrypted for production grade parts.

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	LUKS	/dev/xvdc1	500 GB for small plan. 1000 GB for medium plan	Secondary disk is /dev/xvdc	Product installation. Directories /data and /opt are created on this disk.
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging

Bare metal server for large plan

The large plan comes with two Solid State Drive (SSD) disks that are 960 GB each. RAID level 1 implementation makes them appear as a single disk.

The disk is divided into four partitions. The Red Hat Enterprise Linux operating system is on a 10 GB partition. The boot data is on a 256 MB partition. The swap space is on a 2 GB partition. The remaining space is on another partition that is encrypted by using LUKS.

Table 17: Bare metal server disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	256 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda2	2 GB	/dev/sda	Swap space for paging
/	None	/dev/sda3	10 GB	/dev/sda	Operating system data
/disk1	LUKS	/dev/sda5	About 900 GB	/dev/sda	Product installation. Directories /data, /installables, and /opt are created on this disk.

Client for all plans

The DataStage on Cloud client machine configuration is the same for all plans sizes. The client machine has two SAN disks that are 100 GB each. One disk is C: for the Microsoft Windows operating system. The second disk is F: , and it is an empty disk.

DataStage on Cloud High Availability

IBM® DataStage® on Cloud is now available with High Availability in small, medium, and large sizes. As part of the High Availability offerings, additional machines are provided in the same data centre and VLAN to act as passive instances that can take over in case of failures to the active instances. The passive instances have pre-installed and configured parts of the application. Clustering is used between the active and passive instances for IBM WebSphere® Application Server. HADR is configured between the active and instances of IBM DB2® in active-passive mode. And, shared storage is used between them for the Engine, Service, and Metadata Repository Tiers of DataStage. [The High Availability offering also comes with backup functionality configured.](#)

Technologies and Concepts

HADR

Configuration of two instances of a software or hardware component for High Availability and Disaster Recovery. In the active-passive mode that is used for HADR configuration between the two instances of DB2, the passive instance keeps getting logs from the active one and apply the transactions on it's local database. So, the passive instance would be ready to take over whenever the active instance can't serve it's purpose.

Automatic Client Reroute

This term refers to the configuration of data sources to connect to the available database server in case the primary one goes down.

Portable IP

Portable IP is an IP that can be assigned for any of the active and passive instances and is initially assigned to the active instance or server.


Pacemaker® (PM)

A High availability software that helps detect a failure to an instance of software component and handles automatic assignment of a portable IP to another machine.

DataStage on Cloud High Availability comes with one active and one passive server machine, and a number of clients depending on the size of the offering chosen. Both the active and passive server machines have instances of DB2 and WebSphere Application Server installed on non-shared storage, and Information Server Engine, Metadata Repository and Services Tiers installed on the shared storage.

A portable IP, also called virtual IP is initially assigned to the active server machine. A virtual hostname associated to the virtual IP is used for installation of Information Server.

The WebSphere Application Server cluster is built using the instances of the same installed on both the machines, and DB2 HADR in active passive mode setup between the DB2 instances on them.



WebSphere Application Server on the passive server machine has a WebSphere node and HTTP Server installed while active server machine also has Deployment Manager installed along with these. HTTP Server takes care of distributing the load between the WebSphere Application Server nodes in a round robin fashion.

All the WebSphere Application Server data sources are configured for Automatic Client Rerouting so that when an active instance fails or is not reachable, the transactions are automatically redirected to passive instance of DB2. All the default databases that come with the product, DSODB, and XMETA are configured for HADR.

Passive Engine Tier on the passive server machine can be brought up with a few manual steps executed on it whenever the active Engine Tier doesn't work as all the Information Server Tiers are installed on the shared storage.

The dotted lines in the above diagram show the connections to the passive instances of different components while the solid lines show connections to the active instances.

The high availability software Pacemaker installed on both the servers helps automatically switch the virtual IP to the passive instances whenever the active machine goes down. The virtual IP can also be switched to the passive server machine by running a Pacemaker command to stop the specific resources on the active machine. Both public and private virtual IPs are used. These virtual IPs or the associated virtual host names should be used when connecting to the DataStage on Cloud machines, and when integrating with other applications for making sure that High Availability implementation works with the integration.

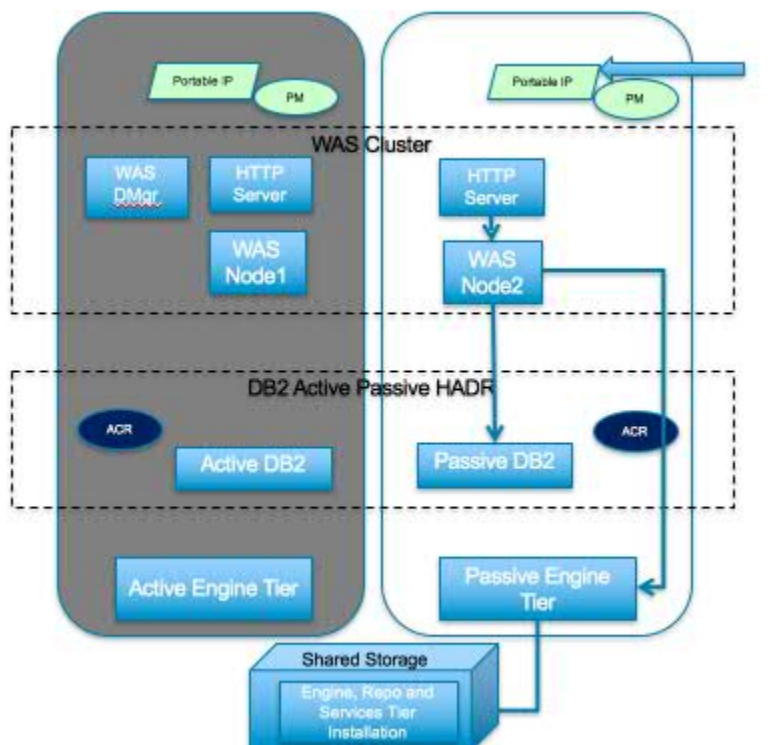
For switching over virtual IPs manually from the active server machine to passive server machine, run the script below on the active server machine:

```
/opt/IBM/ha_scripts/switchPortableIP.sh
```

While on the failures of the some of the components the passive instances take over automatically, some would need manual steps for the take over by the corresponding passive instance to happen. The following scenarios discuss some of the failure cases and the take over procedures in such cases.

Failure Scenarios

Active Server machine fails




Failure: The machine in dark grey colour has failed all the communication to that machine is stopped.

Take over procedure: Pacemaker detects the failure and assigns the portable IP to the passive machine. HTTP Server on the passive machine receives all the HTTP requests and sends them to the only available WebSphere Application Server node, Node2.

DB2 take over steps should be performed on the passive instance of DB2. Take over procedure is described in the failure case **Active DB2 fails** below.

DataStage Engine Tier should be started on the passive machine. [Run the steps for running DataStage Engine on the passive machine](#)


WebSphere Application Server Node1 fails



Failure: The WebSphere Application Server node, Node 1 shown in dark grey colour goes away, and all the connections between this and other processes are broken.

Take over procedure: HTTP Server detects that one of the WebSphere Application Server nodes is not available and directs all requests to the available node, Node2 automatically.

Active DB2 fails



Failure: Active DB2 shown in the dark grey colour goes away, and all database transactions fail.

Take over procedure: Administrator should run take over commands on the passive DB2. Automatic client re-route feature makes redirect all transactions to passive DB2 automatically.

Steps for DB2 take over on the passive machine:

1. Login as the user db2inst1
2. Change to the home directory of the user, db2inst1 by running the following command: `cd /home_/db2inst1`
3. Execute DB2 profile on the current shell: `. sqllib/db2profile`
4. Run the DB2 take over command for taking over on the passive/standby machine: `db2 takeover hadr on database <db_name>` . For example to take over XMETA on the passive machine, the command is "db2 takeover hadr on database xmeta".

Update the database URL in the file `/opt/IBM/InformationServer/Server/DSODB/DSODBConnect.cfg` to point to the host name of the new machine that is running DB2 Primary.

Available configurations

IBM® DataStage® on Cloud servers for the small and medium plans are virtual servers with dedicated CPUs. The servers in the large plan are in a bare metal environment.

Select the offering plan that fits your usage and environment needs.

Table 18: DataStage on Cloud available configurations

Offering	Environment	Memory (GB)	Number of Cores	Network speed (Gbps)	First disk	Second disk
Small	Production Non-production	16	4	1	100 GB Storage area network (SAN)	500 GB SAN
Medium	Production Non-production	32	8	1	100 GB SAN	1 TB SAN
Large	Production	64	12	1	2-960 GB SSD	

For the backup machine, we have the following configurations.

Offering	Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk	Performance disk
Small	Backup Server	64	8	1 Gbps with 250 GB bandwidth	100 GB storage area network (SAN)	3 TB SAN	1TB 4000 IOPS
Medium	Backup Server	64	8	1 Gbps with 1000 GB bandwidth	100 GB storage area network (SAN)	4 TB SAN	1TB 6000 IOPS

Offering	Tier	Memory (GB)	Number of cores	Network speed and bandwidth	First disk	Second disk	Performance disk
Large	Backup Server	128	12	1 Gbps with 5000 GB bandwidth	6.8TB SSD	8 TB SATA	

Small and medium offerings include the following configuration:

- One virtual machine with DataStage on Cloud, IBM InfoSphere® Information Services Director
- One virtual server machine used for storing backup artifacts
- One virtual machine with a Microsoft Windows operating system and all client utilities installed


Large-size offering includes the following configuration:

- One bare metal machine with DataStage on Cloud, IBM InfoSphere Information Services Director
- One bare metal machine used for storing backup artifacts
- Two virtual machines with a Microsoft Windows operating system and all client utilities installed

Available configurations of DataStage on Cloud HA

High Level Model of DataStage® on Cloud High Availability

The High Availability offering of IBM® DataStage on Cloud will have two server machines as shown in the below diagram, along with one or more client machines depending on the size chosen.



Virtual servers with dedicated CPUs are provided for the small and medium plans of IBM DataStage on Cloud High Availability. The servers in the large plan are bare metal. Both the server machines will have shared file storage along with non-shared storage.

Select the offering plan that fits your usage and environment needs.

Configurations

Offering	Environment	Memory (GB)	Number of Cores	Network speed (Gbps)	First disk	Second disk
Small	Production Non-production	16	4	1	100 GB Storage area network (SAN)	500 GB SAN
Medium	Production Non-production	32	8	1	100 GB SAN	1 TB SAN
Large	Production	64	12	1	2 x 960 GB SSD with RAID Level 1	

Small and medium offerings include the following configuration:

- Two virtual machines with DataStage on Cloud, IBM InfoSphere® Information Services Director
- One virtual machine with a Microsoft Windows operating system and all client utilities installed for small part and three for medium part.

Large-size offering includes the following configuration:

- Two bare metal machines with DataStage on Cloud, IBM InfoSphere Information Services Director
- Five virtual machines with a Microsoft Windows operating system and all client utilities installed

Concurrent users based on part size:

- The number of client machines is based on the number of concurrent users. For small offerings, two concurrent users are allowed, while five for medium and ten for large are allowed. By default, Microsoft Windows operating system allows two concurrent users to access the machine by using Remote Desktop Connection.

Layout of IBM DataStage on Cloud High Availability server and client disks

The layout of the DataStage® on Cloud High Availability server and client disks depends on the plan size of your system.

Virtual servers for small and medium plans

The small and medium plans come with two Storage Area Network (SAN) disks. The Red Hat Enterprise Linux operating system is on the first SAN disk in both the small and medium plans. The second SAN disk is encrypted by using Linux Unified Key Setup (LUKS) for production offerings of both these sizes.

The encryption key details are provided in the Welcome letter from the IBM® Sales Representative. It is recommended that you add your own key and remove the supplied key before you use the system.

Both of the active and passive server machines are mounted with a common shared storage for installing Information Server tiers.

The Information Server tiers are installed on the /opt directory which is bound to the shared storage, while DB2 and WebSphere Application Server are installed on the non-shared storage /opt2. User data can be stored on /data directory. Both directories are on the partition /dev/xvdc1 that is encrypted for production grade parts.

Small and medium server disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	Provided by Softlayer	Dynamic name	500 GB for small, 1000 GB for medium plans respectively.	Dynamic Name	This is the shared storage between the two server machines. Information Server tiers are installed on this. /opt is created on this storage.
/disk2	LUKS	/dev/xvdc1	500 GB for small plan. 1000 GB for medium plan	Secondary disk is /dev/xvdc	DB2 and WebSphere Application Server installation. Directory /opt2 is created on this disk.
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging

Bare metal server for large plan

The large plan comes with two Solid State Drive (SSD) disks that are 960 GB each. RAID level 1 implementation makes them appear as a single disk.

The disk is divided into four partitions. The Red Hat Enterprise Linux operating system is on a 10 GB partition. The boot data is on a 256 MB partition. The swap space is on a 2 GB partition. The remaining space is on another partition that is encrypted by using LUKS.

Bare metal server disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	256 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda2	2 GB	/dev/sda	Swap space for paging
/	None	/dev/sda3	10 GB	/dev/sda	Operating system data

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/disk1	Provided by Softlayer	Dynamic name	1000 GB	Dynamic Name	This is the shared storage between the two server machines. Information Server tiers are installed on this. /opt is created on this storage.
/disk2	LUKS	/dev/sda5	About 900 GB	/dev/sda	Product installation. Directories /installables, and /opt2 are created on this disk. DB2 and WebSphere Application Server are installed here.

Client for all plans

The DataStage on Cloud client machine configuration is the same for all plan sizes. The client machine has two SAN disks that are 100 GB each. One disk is C : for the Microsoft Windows operating system. The second disk is F : , and it is an empty disk.

Information roadmap

This roadmap lists information resources that are available for users who are new to the DataStage® products from IBM®. These resources provide information about various subject areas, such as learning basic skills and troubleshooting.

Product overview

- **InfoSphere® DataStage features** Learn about the features and benefits of InfoSphere DataStage.
- **InfoSphere DataStage product documentation** The InfoSphere Information Server Knowledge Center provides you with InfoSphere DataStage concepts and usage information to help new users prepare for using your new system.
- **InfoSphere DataStage developerWorks® forum** Use this forum to interact with other InfoSphere DataStage users to better understand how to design, build, debug and deploy jobs for information collection, integration and transformation.

Getting started

- **Overview of InfoSphere DataStage** This overview covers InfoSphere DataStage job life cycles, job designs, and how the product integrates with the rest of the InfoSphere Information Server suite.
- **Tutorial: Creating parallel jobs** This tutorial demonstrates how you can use InfoSphere DataStage to develop jobs that extract, transform, and load data. By transforming and cleansing the source data and applying consistent formatting, you enhance the quality of the data.

Using InfoSphere DataStage

- **[Designing InfoSphere DataStage jobs](#)** This topic explains how to design InfoSphere DataStage jobs using the IBM InfoSphere DataStage and QualityStage® Designer client. The Designer client gives you the tools that you need to create jobs that extract, transform, load, and check the quality of data.
- **[Developing InfoSphere DataStage parallel jobs](#)** This topic covers how to design parallel jobs to transform and to cleanse data. Parallel jobs are compiled and run on the InfoSphere Information Server engine.
- **[Developing server jobs](#)** Learn how server jobs are compiled and run on the InfoSphere Information Server engine. Such jobs connect to a data source, extract and transform data, and write data to a target database or file, such as a data warehouse.
- **[Deploying jobs](#)** Use the InfoSphere Information Server Manager to move InfoSphere DataStage assets between projects on the same engine, or on different engines. You can also use the InfoSphere Information Server Manager to move assets from one domain to another.
- **[Running InfoSphere DataStage jobs](#)** This topic covers how to run InfoSphere DataStage from the IBM InfoSphere DataStage and QualityStage Director client.
- **[Administering workload management](#)** Use the workload management queues to control the starting of parallel and server jobs.
- **[Monitoring jobs](#)** Use IBM InfoSphere DataStage and QualityStage Operations Console to monitor the job runs, services, system resources, and workload management queues on several InfoSphere Information Server engines.
- **[Parallel job reference](#)** Use this reference material to perform more advanced operations with parallel jobs.
- **[Data Flow Designer](#)** A new web-based tool, to create, edit, load, and run DataStage jobs is introduced in Information Server 11.7

Troubleshooting and support

- **[Troubleshooting InfoSphere DataStage](#)** Use the Troubleshooting page to find common troubleshooting topics.
- **[IBM Support Portal](#)** Use the Support Portal for InfoSphere DataStage to search for known problems and APARs.

Getting started and using IBM DataStage on Cloud

You must set up your connection to DataStage® on Cloud to design, develop, and run jobs that move and transform data. DataStage on Cloud provides all of the functions of its on-premises counterpart, IBM® InfoSphere® DataStage. DataStage on Cloud is in an IBM SoftLayer® cloud hosted environment that offers deployment without the cost, complexity, and risk of managing your own infrastructure.

Prerequisite: You must know the IP address and the credentials of an account on the DataStage on Cloud server and client computer. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM Sales Representative.

The DataStage on Cloud client is on a Microsoft Windows virtual machine that is hosted on SoftLayer. When you connect to the client, you can access the client user interfaces, create jobs, and run them. McAfee anti-virus software is installed on the client machine.

The DataStage on Cloud server is on a Red Hat Enterprise Linux virtual or bare metal computer that is hosted on SoftLayer. When you connect to the server, you can access the IBM InfoSphere Information Server engine, services, and repository tiers. You can restart InfoSphere DataStage, edit the dsenv file, and do administrative tasks.

The default firewall configuration of server machines allows SSH connections only from client machines. You must first connect to a client machine by using a remote desktop connection, and then from the client machine you can connect to server machines by using SSH. After you log in to a server machine, you can

change the firewall configurations to allow SSH connections from other machines. Communication between the server and client systems happens through a private IP. If you want to access the server from an on-premises client machine, you must modify the iptable rules.

When you connect for the first time, follow these steps:

1. SSH into the any server machine using unique user (order id) provided in the welcome letter using port 4362.
2. Change the password and su to root user.
3. On the DataStage on Cloud server, run the ISALite tool. The *IS_install_path* for the server computers is /opt/IBM/InformationServer.
4. {#connecte} Connect to the DataStage on Cloud client to the DataStage on Cloud server by following these steps
 - **If your local computer is in a Microsoft Windows environment**
 - a. On your local computer, go to the **Start** menu. Click **Accessories > Remote Desktop Connection**.
 - b. Enter the IP address of the Microsoft Windows computer that hosts your DataStage on Cloud client. Click **Connect**.
 - c. In the Windows Security window, enter the user name and password for the DataStage on Cloud client. The user ID, password, and IP address of the client are in your Welcome letter. **Important:** Do not include the domain name with the user name.
 - d. In the DataStage on Cloud client, open the file C:\Windows\System32\drivers\etc\hosts. Make sure that an entry with the private IP exists in the file for the DataStage on Cloud server that you are connecting to. **Important:** The server IP must be a private IP. You cannot open other client machines when you use a public IP.
 - **If your local computer is in an Apple Mac environment**
 - a. On your local computer, install Microsoft Remote Desktop from the Apple App Store.
 - b. Click the Microsoft Remote Desktop icon, and then click **Open**.
 - c. In the Microsoft Remote Desktop window, click **New**.
 - d. In the Edit Remote Desktops window, supply the following information:
 - 1) In the PC name field, type in the IP address of the cloud client machine.
 - 2) In the User name and Password fields, type in the Windows user name and password that are in the Welcome letter.
5. Verify the connection and installation on the DataStage on Cloud client by following these steps:
 - Run the ISALite tool. The *IS_install_path* for the client computer is C:\IBM\InformationServer.
 - Test the installation of the IBM InfoSphere DataStage and QualityStage® Administrator and InfoSphere DataStage and QualityStage Designer clients.
6. Optional: Enable multiple users to open remote sessions to the DataStage on Cloud client by following these steps on the client computer:
 - a. Create user accounts.
 - b. Give users permission to do a remote desktop connection. The number of concurrent users in a remote session to the client is based on the offering size of your DataStage on Cloud. The small offering size has a maximum of two concurrent sessions. The medium offering size has a maximum of five concurrent sessions. The large offering size has a maximum of 10 concurrent sessions.
7. Install Adobe Flash Player on the client computer to enable use of the Hierarchical Data stage in parallel job designs. The version of Adobe Flash Player to use is listed in the Web Browser Plug-Ins table of the System Requirements page.
8. Reset the password for users on the DataStage on Cloud client. Likewise, reset the password for administrators on the DataStage on Cloud server.

After the initial connection, you can do any of the following tasks:

- [Open the InfoSphere DataStage® and QualityStage® Designer client](#)
- [Connect to an on-premises computer](#)
- [Connect to an IBM dashDB™ database](#)
- [Connect to an on-premises DB2® database instance](#)
- [Perform general administration and security tasks](#)

Related information

- [Getting started and using IBM Information Governance Catalog on Cloud](#)


Backing up IIS DataStage on Cloud components

IBM Information Server DataStage on Cloud Premium and High Availability services provides software and hardware infrastructure for taking backups. Backup capability is based on IBM Spectrum Protect version 8.1.3 product (formerly known as Tivoli Storage Manager). One dedicated server machine with installation of IBM Spectrum Protect server version 8.1.3 is provided with each deployment of Information Server on Cloud Premium and High Availability service. Spectrum Protect Server version used is 8.1.3 and Spectrum Protect Client version used is 8.1.2.

Sample configuration templates are also provided. You can create new configurations or customize the sample templates to take regular backups.

For more information on IBM Spectrum Protect, you can check [IBM Spectrum Protect Knowledge Center](#).

Spectrum Protect setup



IBM Spectrum Protect Server installation

Two instances of IBM Spectrum Protect server are configured on a dedicated machine for each deployment of IIS on Cloud Premium and High Availability service. To store backup data on the storage attached with the server machine, the first instance is configured with Directory Container Storage Pool. The second instance is configured with Cloud Container Storage Pool to store data in Object Store. For description of different data pools types, check [Storage pool types](#).

By default, operating system IP table rules allows communication to the Spectrum Protect backup server only from those machines where IBM Spectrum Protect Client is installed. In order to use Spectrum Protect Operation Console (web application), you need to open 11090 port for specific IPs in the IP table firewall rules. The port 11090 will be open for the IIS Windows Client to communicate with the Spectrum Protect backup server.

For small & medium plans, Directory Container Storage Pool is mapped to SAN storage and DB2 (which is used by IBM Spectrum Protect Server) is installed on Performance Storage. SAN storage is encrypted using operating system's LUKS encryption. IBM SoftLayer provides default encryption for Performance Storage in most of the data centers, for remaining data centers encryption is done at operating system level using LUKS. For data center list you can check [IBM SoftLayer documentation](#).

Directory Container Storage Pool (Instance 1)

A Directory Container Storage Pool is configured in Instance 1, this pool is used to store backup data locally. To know more about Directory Container Storage Pool, check [Directory-container storage pools FAQs](#).

Domain configurations are created for all client machines; these domains are linked with Directory Container Storage pool. For details about policy domain configuration check [Creating a policy domain](#).

Cloud Container Storage Pool (Instance 2)

The Cloud Container Storage Pool configured on Instance 2 is used to store data in cloud storage. The cloud-container storage pools that are provided by IBM Spectrum Protect can store data to cloud storage that is object-based. By storing data in cloud-container storage pools, you can exploit the cost per unit advantages that clouds offer along with the scaling capabilities that cloud storage provides. IBM Spectrum Protect manages the credentials, security, read and write I/Os, and the lifecycle for data that is stored to the cloud. You can back up and restore data or archive and retrieve data directly from the cloud-container storage pool. To understand more, check [Cloud-container storage pools FAQs](#) and [Configuring a cloud-container storage pool](#) pages.

Before sending data to Object Store, Spectrum Protect server encrypts the data using encryption key. For encryption configurations details, check (https://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.0/srv.admin/t_cloud_encryption.html).

Like Directory Container Storage Pool, policy domain configurations are created for all Client machines; these domains are linked with Cloud Container Storage pool.

Node replication

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.

In the setup provided to you, replication is enabled for all clients and backed-up data is replicated from Spectrum Protect Server Instance 1 to Instance 2. If required you can change replication settings by following instructions available at [Replicating client data to another server](#).

Data is replicated from Spectrum Protect instance 1 to Spectrum Protect instance 2 using node replication. The administrative schedule is configured for this purpose. There are two schedules `replicate_nodes_weekend` and `replicate_nodes_weekday`.

`schedule replicate_nodes_weekday` replicates data from Spectrum Protect instance 1 to Spectrum Protect instance 2 for every 3 hours.

`schedule replicate_nodes_weekend` replicates data from Spectrum Protect instance 1 to Spectrum Protect instance 2 for every 12 hours.

By default this schedule is stopped, run the following.

```
update schedule replicate_nodes_weekday type=administrative expiration=never
update schedule replicate_nodes_weekend type=administrative expiration=never
```

Run above commands from the "Command Builder" of Spectrum protect operation center.

IBM Spectrum Protect client installation

IBM Spectrum Protect Client is installed on all the machines except the one on which IBM Spectrum Protect Server is installed and the IIS Windows Client machine in which the IIS Designer Client and the other thick clients are installed. IBM Spectrum Protect Client is configured to send backup data/ metadata/configuration files to Spectrum Protect Server over SSL. Client communicates with IBM Spectrum Protect server using server's private IP.

To learn about IBM Spectrum Protect Client, check [IBM Spectrum Protect Knowledge Center](#)

To enable access to IBM Spectrum Protect Client user interfaces, VNC server is installed on Client machines. By default, IP table firewall rules does not allow communication over 5901 port which is used by VNC server. To allow communication from the machine on which you want to access the user interfaces, you need to update IP Table firewall rules for port 5901. VNC server communication is not encrypted, if your organization mandates this communication to be secure, then you can use some other tool which support encryption.

Ports exposed

The following ports are opened to and from both the Spectrum Protect server and Spectrum Protect client machines:

- 1550
- 1552
- 1553
- 1650
- 1652
- 1653

Port 11090 is exposed from Spectrum Protect Server machine to the IIS Windows Client machine to access the Operations Center.

Port 4362 is also opened to access Spectrum Protect Server from the IIS Windows Client machines.

Apart from these ports, all the other ports are blocked for communication in the Spectrum Protect Server.

Getting started with IBM Spectrum Protect Operations Center console

IBM Spectrum Protect provides a web application called Operations Center for managing IBM Spectrum Protect environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

More details about Operations Center are available at [Managing the Operations Center](#).

In the setup provided to you, Operations Center is accessible using port 11090. Default IP table firewall rules on the IBM Spectrum Protect server machine does not allow communication with port 11090 from external machines. By default, the only machine which can access Spectrum Prtotect server using port 11090 is the IIS Windows Client machine. This is done for security purposes.

You can follow below steps to enable Operations Center access from an external machine.

1. Connect to Spectrum Protect server machine using putty or terminal. (Not required if accessing using IIS on Cloud Windows Client machine)
2. Go to scripts directory. (Not required if accessing using IIS on Cloud Windows Client machine)

```
cd /opt/IBM/scripts
```

3. Execute below command after replacing <IP_ADDRESS> with IP address of the machine from where you want to access Operations Console. (Not required if accessing using IIS on Cloud Windows Client machine)


```
./allow_port_from_ip.sh <IP_ADDRESS> 11090
```

When Operations Center is opened for the first time, it asks for some inputs. You must follow below steps to provide inputs when you open Operations Center for the first time

1. Open following URL in browser after replacing <Spectrum_Protct_Server_IP_Address> with your Spectrum Protect server machine IP.

```
https://<Spectrum_Protct_Server_IP_Address>:11090/oc
```

2. When you open Operations Console for the first time, it will ask for credentials.



3. Replace default details with correct values.

```

Localhost:1500      -- <Spectrum Protect server PUBLIC or PRIVATE IP>:1550
Administrator      -- tsminst
Password           -- Password for tsminst user is provided in welcome letter.


```

4. In the next page you will be asked to provide password (two times) for "Administrator ID". Provide password.
5. After providing password details, in the next page you need to specify how frequently you want to collect data. Depending on your requirement you can select 1 minute to 1 hour.
6. Follow instructions on the user interface to finish the wizard.

Note that "Instance 2" may be down and may take few minutes to start. You can check status of both instances under overview tab of Operations Center console.


Starting Command Builder

Open following URL in browser https://<Spectrum_Protct_Server_IP_Address>:11090/oc to access Operation Center. To open the command-line interface, hover over the globe icon in the Operations Center menu bar, and click Command Builder.



IBM Spectrum Protect processes administrator commands either in the foreground or in the background. Commands that process in the foreground must complete before you can issue another command. When commands are processing in the background, you can issue additional commands at any time. Most IBM Spectrum Protect commands process in the foreground.

Spectrum Protect server contains two instances of servers which are connected using node replication feature for the fail-over scenario. In Command builder left side down you can see both Spectrum Protect instance. By default first instance of Spectrum Protect server is selected. You need to select the second instance of Spectrum protect server if you need to execute any commands against the second instance of Spectrum protect server. In "Command Builder" you can select second instance which has name like <ORDER_ID>_/_s/m/l_bckp1. Select drop down menu from left downside corner, as shown in the image.



Retention Policy

Life cycle of backup data objects

A backup object exists in three states, active, inactive, and expired, before being purged from the Spectrum Protect server. The four steps involved in the life cycle of a backup data object are listed here.

1. A copy of the client data is sent to the Spectrum Protect server as a backup object. When a backup object is sent to the Spectrum Protect server, it becomes the active version.
2. It remains in an active state until the Spectrum Protect client program deletes the backup object manually, or a newer version of the backup object is sent. The backup object changes state from active to inactive.
3. The backup object remains inactive until it exceeds its retention settings. A backup object can exceed retention settings by either time or number of versions. The backup object changes state from inactive to expired.
4. The backup object remains in the expired state until expiration processing runs on the Spectrum Protect server. This process is invoked by a Spectrum Protect administrator with the expire inventory command. When expiration processing encounters a backup object in the expired state, it purges that object from the Spectrum Protect database and frees up the storage space where the backup object resided.

Spectrum Protect server sample domains for directory container storage pool (Spectrum Protect server local storage) are configured with backretention=30 archretention=30 Spectrum Protect server sample domains for cloud container storage pool (Object storage) are configured with backretention=365 archretention=365

BACKREtention :

Specifies the number of days (from the date the backup versions became inactive) to retain backup versions of files that are no longer on the client file system. This parameter is optional. You can specify an integer from 0 to 9999. The default value is 30. The server uses the backup retention value to manage

inactive versions of files when any of the following conditions occur: - A file is rebound to a new management class, but the new management class and the default management class do not contain a backup copy group. - The management class to which a file is bound no longer exists. The default management class does not contain a backup copy group. - The backup copy group is deleted from the management class to which a file is bound. The default management class does not contain a backup copy group.

ARCHREtention :

Specifies the number of days (from the date of archive) to retain archive copies. This parameter is optional. You can specify an integer from 0 to 30000. The default value is 365. The server uses the archive retention value to manage archive copies of files when either of the following conditions occur: - The management class to which a file is bound no longer exists. The default management class does not contain an archive copy group. - The archive copy group is deleted from the management class to which a file is bound. The default management class does not contain an archive copy group.

More details about domain configuration details are [here](#)

Below copygroup is defined for directory container storage pool (Spectrum Protect server local storage)
Spectrum Protect server sample copygroup defined with domain for backup is configured with
VEREXISTS=NOLimit VERDEL=NOLimit RETEXTRA=30 RETONLY=30

Below copygroup is defined for cloud container storage pool (Object storage) Spectrum Protect server
sample copygroup defined with domain for backup is configured with VEREXISTS=NOLimit
VERDEL=NOLimit RETEXTRA=365 RETONLY=365

Domain and copygroup created for each Spectrum Protect client machine have same settings.

VERExists :

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional. The default value is 2.

VEREXISTS=NOLimit Specifies that you want the server to retain all backup versions. The number of backup versions to retain is controlled by this parameter until versions exceed the retention time specified by the RETEXTRA parameter.

VERDeleted :

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using IBM Spectrum Protect. This parameter is optional. The default value is 1. If a user deletes a file from the client file system, the next incremental backup causes the server to expire the oldest versions of the file in excess of this number. The expiration date for the remaining versions is determined by the retention time specified by the RETEXTRA or RETONLY parameter.

VERDEL=NOLimit Specifies that you want the server to retain all backup versions for files that are deleted from the client file system after being backed up.

RETEtra :

Specifies the number of days to retain a backup version after that version becomes inactive. A version of a file becomes inactive when the client stores a more recent backup version, or when the client deletes the file from the workstation and then runs a full incremental backup. The server deletes inactive versions based on retention time even if the number of inactive versions does not exceed the number allowed by the VEREXISTS or VERDELETED parameters. This parameter is optional. The default value is 30 days.

RETOly :

Specifies the number of days to retain the last backup version of a file that has been deleted from the client file system. This parameter is optional. The default value is 60.

You can change sample retention policy values according to your requirement, keeping Spectrum Protect server storage space in mind.

More details about copygroup configuration details are [here](#)

Configuring Object storage

You can store deduplicated data and non-deduplicated data in a cloud-container storage pool and restore the data as required. You can configure IBM Spectrum Protect to temporarily store data in one or more local storage pool directories during data ingestion. The data is then moved from local storage to the cloud. In this way, you can improve data backup and archive performance.

After you define a storage pool directory, the IBM Spectrum Protect server uses that directory as a temporary landing spot for the data that you are transferring to cloud object storage. The server uses an automated background process to transfer data from local storage in the directory to cloud object storage. You do not need to take any additional steps to start or manage this transfer process. After the server successfully moves the data from local storage to cloud object storage, the server deletes the data from the directory and releases space for more incoming data.

If storage pool directories contain no more free space, backup operations stop prematurely. To avoid this situation, you can allocate more storage pool directories. You can also wait for the data to be automatically removed from the local directories after the data moves to the cloud.

Spectrum Protect server supports these cloud service providers.

- Amazon S3
- IBM Cloud Object Storage
- IBM SoftLayer
- OpenStack Swift

Amazon S3 API object storage has been used for the sample domains, policies and schedules. Object store is configured with dummy credentials and URL. Once the user has created an object storage with S3 API of their own, they can input the appropriate values for credentials and URLs and other necessities required to configure an object storage to Spectrum Protect server.

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. In the sample policy provided, each Spectrum Protect client has its own bucket, ie. its own subdirectory in the object storage. Bucket names and unique IDs (called "keys" in S3) are used to access data from object storage in Spectrum Protect.

Use Operation center in updating the details of Object store details or use "Command Center" in Operation center to update values.

In the sample policy provided, data retention policy for Object store is set to 365 days, which means data stored in Object store is available for 365 days. These are specified in the backretention and archretention parameters of a domain, as mentioned in the [Retention Policy](#) section.

Each machine contains 2 cloud pools, one used for backup and another for archive , where dummy credentials are provided.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

More details about SoftLayer Object store is [here](#)

More details about cloud object storage details are [here](#)

More details about configuring cloud-container storage pools for IBM SoftLayer is [here](#)

More details about encrypting data for cloud-container storage pools is [here](#)

Limitations and best practices

Spectrum Protect server setup and client installations are provided only with IIS on Cloud Premium and High Availability services. If you have also subscribed to Non-Production offerings then you need to develop your own backup artifacts for this offering.

Spectrum Protect server installation which is provided with IIS on Cloud Premium and High Availability offerings can only be used for taking backups of applications and files which are part of IIS on Cloud deployment.

IIS on Cloud setup consists of many applications and components. When you configure different policies & schedulers to take backup of different applications and components, backups are created at different timestamps. Hence after restoration of specific application backup, its data may not be in complete sync with related data elements in other applications or components.

Spectrum Protect server and Operation Center are installed in same system, so communication between Spectrum Protect server and Operation Center is through non-SSL. Spectrum Protect server has 2 instances which are connected through node replication. Both the instances are in same system, so communication between these 2 instances is through non-SSL.


Cloud storage is connected to 2nd instance of Spectrum Protect server, which acts as a fail-over server. There are some limitations in connecting to cloud storage from Spectrum Protect client. You can't connect to Cloud storage and retrieve data when 1st instance of Spectrum Protect server is up and running. If 1st instance of Spectrum Protect server is down, then you can connect to 2nd instance of Spectrum Protect server. You'll have only read only access, which means it's used only to retrieve data, you can't take backup and archive using this.

You should follow product or application specific documentation and best practices while developing artifacts to take backups, below are some examples: - IBM Information Server does not support hot backups. For details, check [Backing up IBM InfoSphere Information Server components](#). - WebSphere Application Server documentation suggests that servers are stopped while taking backup of node configurations. For details, check [backupConfig command](#)

IMPORTANT

- Spectrum Protect server's database backup files are not backed-up automatically. These files are mandatorily required to restore a SPectrum Protect Server in case of a failure scenario. Hence, it is recommended that the user backs the artifacts related to this to a secure location. Details about Spectrum Protect server database backup is available at "Spectrum Protect server database backup" section. [Spectrum Protect server database backup](#)
- Spectrum Protect server master encryption key is stored in the server password file, dsmserv.pwd. This file takes care of encryption and decryption of data being transferred for backup and restore. Hence, it is recommended that the user backs the artifacts related to this to a secure location. Details about master key is available at "Protecting the master encryption key" section. [Protecting the master encryption key](#)

IBM Spectrum Protect setup for IIS DataStage on Cloud Premium service



IIS DataStage on Cloud has the following machines:

- IIS Engine machine
- Spectrum Protect Backup server
- IIS Windows Client machine
- IIS Compute machines (Optional)
- BPM/Cognos machine (Optional)

Spectrum Protect server software is installed in the Spectrum Protect server machine and the Spectrum Protect Client software is installed in the rest except the IIS Windows Client machine.

By default all the sample schedulers are disabled by setting expiration value to -1. In order to use sample scheduler, you need to enable them by setting appropriate expiry date. This is discussed in detail for each of the machines in these offerings in the following sections.

Backup files will be available only for 30 days, later they are removed from Spectrum Protect server directory storage pool. If cloud container is configured backup files will be available for 365 days, later they are removed from Spectrum Protect server cloud storage pool. You can change these settings according to your requirement, more details on modifying these settings is available in "Retention Policy" section. Make sure storage in Spectrum Protect server is limited.

IIS Engine machine (main)

To take backups of the artifacts present in the IIS Engine machine, a sample policy domain configuration named as MAIN-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of WAS profile directory, WAS profile configuration, database full, database incremental online backups and other files which are located in IIS Engine machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine, execute commands to create backup archive files, copy files to specific location, executes commands to take db2 full and incremental backups and invoke Spectrum Protect commands to copy files to Spectrum Protect Server machine.

WebSphere Application Profiles Backup

A sample scheduler named as MAIN-WAS-WEEKLY is configured to take backups of WAS profile directories. This schedule invokes WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Engine Machine.

MAIN-WAS-WEEKLY scheduler is configured to execute weekly once, on Sunday at midnight.

WASProfileBackup_IIS.sh executes commands to create an archive file named IIS_AppServer_backup.tar which contains all the files from following directories.

- WAS InfoSphere profile backup : /opt/IBM/WebSphere/AppServer/profiles/*

IIS_AppServer_backup.tar contains content of WAS profile directory which is located at /opt/IBM/WebSphere/AppServer/profiles/

The WAS profile directory contains only one profile named as the "InfoSphere" profile

Archive files are stored in /home_/WAS_Backup/ folder. WASProfileBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServer_backup.tar files.

In order to enable MAIN-WAS-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-WAS-WEEKLY expiration=never
```

Above command enable MAIN-WAS-WEEKLY scheduler without expiry date. After enabling MAIN-WAS-WEEKLY schedule , start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

WebSphere Application Configuration Backup

A sample scheduler named as MAIN-WAS-DAILY is configured to take backups of WAS profile configuration. This schedule invokes WASProfileConfigBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Engine Machine. Before enabling MAIN-WAS-DAILY schedule, you must update WASProfileConfigBackup_IIS.sh and provide value for PASSWORD in the shell script.

MAIN-WAS-DAILY scheduler is configured to execute daily at midnight.

WASProfileConfigBackup_IIS.sh executes commands to create an archive file named IIS_AppServerConfig.zip, which contains all profile configurations.

- WAS InfoSphere profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/InfoSphere

IIS_AppServerConfig.zip contains InfoSphere profile configuration which is generated by using backupConfig command.

Archive files are stored in /home_/WAS_Conf/ folder. WASProfileConfigBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServerConfig.zip file.

In order to enable MAIN-WAS-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-WAS-DAILY expiration=never
```

Above command enable MAIN-WAS-WEEKLY scheduler without expiry date.

After enabling MAIN-WAS-WEEKLY schedule , start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine and run `service dsmcad restart` using root user.

IIS Database Backup

Database is configured with linear logging, which means all the transaction (archive) logs of database are stored in the IIS Engine machine. It is recommended that these logs should be stored in Spectrum Protect server itself. In case of a failover scenario, the logs are safe and secure in the Spectrum Protect server.

As transaction (archive) logs of database are stored in the IIS Engine machine, it is the user's responsibility to clean up unused logs, as these logs are not cleared by default. Once you move

transaction (archive) logs from local disk to Spectrum Protect server, it's the responsibility of Spectrum Protect server to clean unused logs. Use below commands in moving transaction (archive) logs from local disk to Spectrum Protect server:

Open Putty or terminal in the IIS Engine machine, switch to db2inst1 user.

```
db2 update database configuration for xmeta using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager
```

db2 update database configuration for dsosb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS db2 stop db manager force db2 start db manager

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. Hence before the sample schedules are started, you have to make sure that one full backup of the respective database is available. Verify available backups for a database by using 'db2adutl query db <DATABASE_NAME>'. Open Putty or terminal in the IIS Engine machine, switch to db2inst1 user to run this command. You have to execute the shell script db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in the IIS Engine Machine, incase there are no full backups available.

Before running db2FullBackup.sh you must consider, where to store database archive logs.

Online full database backups

A sample scheduler named as MAIN-DB-FULL-WEEKLY is configured to take backups of online full database. This schedule invokes db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Engine Machine.

MAIN-DB-FULL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

In order to enable MAIN-DB-FULL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DB-DOMAIN MAIN-DB-FULL-WEEKLY expiration=never
```

Above command enable MAIN-DB-FULL-WEEKLY scheduler without expiry date. After enabling MAIN-DB-FULL-WEEKLY schedule , start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

Online incremental database backups

A sample scheduler named as MAIN-DB-INCREMENT-DAILY is configured to take online incremental database backups. This schedule invokes db2IncrementalBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Engine Machine.

MAIN-DB-INCREMENT-DAILY scheduler is configured to execute from Monday to Saturday at midnights, except Sunday.

In order to enable MAIN-DB-INCREMENT-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DB-DOMAIN MAIN-DB-INCREMENT-DAILY expiration=never
```

Above command enable MAIN-DB-INCREMENT-DAILY scheduler without expiry date. After enabling MAIN-DB-INCREMENT-DAILY schedule , start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

IIS artifacts backup using ISTOOL

A sample scheduler named as MAIN-ISTOOL-WEEKLY is configured to take backups of ISTool export configuration. This schedule invokes `istool_assets.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder in IIS Engine Machine. Before enabling MAIN-ISTOOL-WEEKLY schedule, you must update `istool_assets.sh` and provide value for `PASSWORD` field.

MAIN-ISTOOL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

`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 configuration

Archive file generated by `istool.sh` is stored in `/home_/istool/`. `istool_assets.sh` executes Spectrum Protect server selective backup command to take backup of `istool` generated file.

In order to enable MAIN-ISTOOL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-ISTOOL-WEEKLY expiration=never
```

Above command enable MAIN-ISTOOL-WEEKLY scheduler without expiry date.

After enabling MAIN-ISTOOL-WEEKLY schedule, start '`dsmcad`' service from IIS machine. `dsmcad` provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

Files and Directories Backup

Two sample schedulers named as MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 is configured to take backups of files and folders. These schedule invokes `main_FilesDaily1.sh` and `main_FilesDaily2.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Engine Machine.

MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 scheduler are configured to execute everyday at midnight.

`main_FilesDaily1.sh` and `main_FilesDaily2.sh` executes commands to take backup of below files and folders.

- `/root/keyfile`
- `/home_/db2inst1/db2keystore.p12`
- `/opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/*`
- `/opt/IBM/InformationServer/ASBServer/apps/lib/iis/classes/*`
- `/opt/IBM/InformationServer/Server/Projects/*`
- `/etc/sysconfig/iptables`
- `/opt/IBM/InformationServer/Server/MsgHandlers/*`
- `/opt/IBM/InformationServer/Server/Configurations/*`
- `/opt/IBM/InformationServer/Updates/*`
- `/opt/IBM/InformationServer/Server/DSODB/*.cfg`
- `/opt/IBM/InformationServer/Server/DSEngine/dsenv`
- `/etc/services`
- `/etc/inittab`

`main_FilesDaily1.sh` and `main_FilesDaily2.sh` executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-FILES-DAILY1 expiration=never
update schedule MAIN-DOMAIN MAIN-FILES-DAILY2 expiration=never
```

Above command enable MAIN-FILES-DAILY1 and MAIN-FILES-DAILY1 scheduler without expiry date.

After enabling MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script `/opt/tivoli/tsm/client/ba/bin/main_FilesDaily1.sh` and `/opt/tivoli/tsm/client/ba/bin/main_FilesDaily2.sh` available in IIS Engine machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Engine machine is *main*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like `<ORDER_ID>-/s/m/l-bckp1`.

```
update stgpool main-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool main-cloud-pool identity=<USERNAME>
update stgpool main-cloud-pool password=<PASSWORD>

update stgpool main-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool main-arc-cloud-pool identity=<USERNAME>
update stgpool main-arc-cloud-pool password=<PASSWORD>
```

IIS Compute machine (cmpt)

To take backups of the artifacts present in the IIS Compute machine, a sample policy domain configuration named as COMPUTE-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Compute machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

Files and Directories Backup

A sample scheduler named as COMPUTE-FILES-DAILY is configured to take backups of files and folders. This schedule invokes `compute_FilesDaily.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Compute Machine.

COMPUTE-FILES-DAILY scheduler is configured to execute everyday at midnight.

`compute_FilesDaily.sh` executes commands to take backup of below files and folders.

- `/root/keyfile`
- `/etc/sysconfig/iptables`

- /etc/services
- /etc/inittab

compute_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable COMPUTE-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule COMPUTE-DOMAIN COMPUTE-FILES-DAILY expiration=never
```

Above command enable COMPUTE-FILES-DAILY scheduler without expiry date.

After enabling COMPUTE-FILES-DAILY schedule, start 'dsmcad' service from IIS Compute machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Compute machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script /opt/tivoli/tsm/client/ba/bin/main_FilesDaily.sh available in IIS Compute machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS compute machine is *compute*.

Use Operation center in updating the details of Object store details.

IIS Compute machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

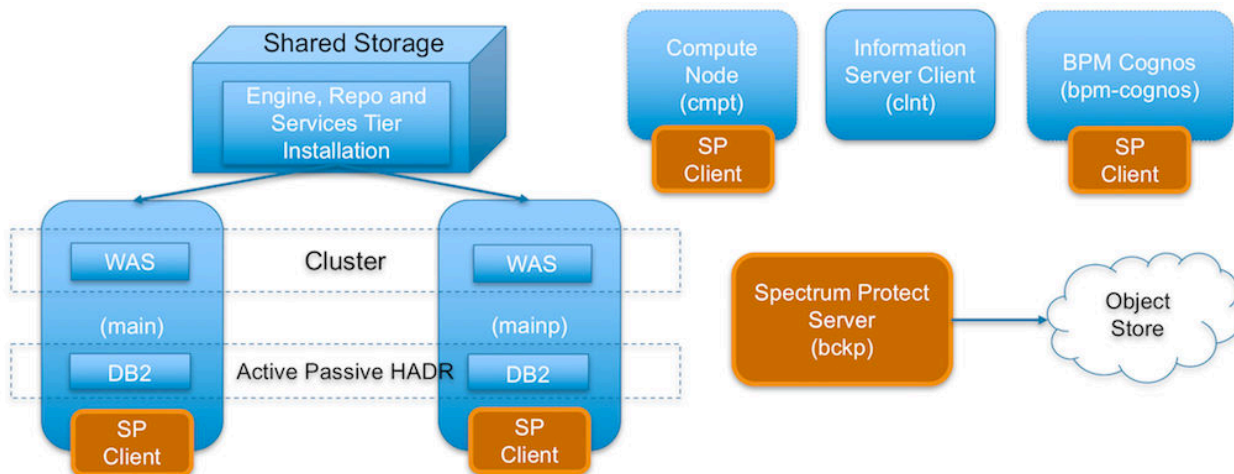
```
update stgpool compute-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool compute-cloud-pool identity=<USERNAME>
update stgpool compute-cloud-pool password=<PASSWORD>

update stgpool compute-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool compute-arc-cloud-pool identity=<USERNAME>
update stgpool compute-arc-cloud-pool password=<PASSWORD>
```

IBM Spectrum Protect setup for IIS on Cloud Enterprise Edition High Availability service

IIS Enterprise Edition on Cloud HA offering has the following machines:

- IIS Engine machine (Active)
- IIS Engine machine (Passive)
- Spectrum Protect Backup server
- IIS Windows Client machine
- IIS Compute machines (Optional)
- BPM/Cognos machine (Optional)



Spectrum Protect server software is installed in the Spectrum Protect server machine and the Spectrum Protect Client software is installed in the rest except the IIS Windows Client machine.

IIS Active Engine machine (main)

To take backups of the artifacts present in the IIS Engine machine, a sample policy domain configuration named as MAIN-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of WAS profile directory, WAS profile configuration, database full, database incremental online backups and other files which are located in IIS Engine machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine, execute commands to create backup archive files, copy files to specific location, executes commands to take db2 full and incremental backups and invoke Spectrum Protect commands to copy files to Spectrum Protect Server machine.

WebSphere Application Profiles Backup

A sample scheduler named as MAIN-WAS-WEEKLY is configured to take backups of WAS profile directories. This schedule invokes WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Engine Machine.

MAIN-WAS-WEEKLY scheduler is configured to execute weekly once, on Sunday at midnight.

WASProfileBackup_IIS.sh executes commands to create an archive file named IIS_AppServer_backup.tar which contains all the files from following directories.

- WAS InfoSphere profile backup : /opt/IBM/WebSphere/AppServer/profiles/*

IIS_AppServer_backup.tar contains content of WAS profile directory which is located at /opt/IBM/WebSphere/AppServer/profiles/

The WAS profile directory contains two profiles, namely "Custom01" and "DMGR01"

Archive files are stored in /home_/WAS_Backup/ folder. WASProfileBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServer_backup.tar files.

In order to enable MAIN-WAS-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-WAS-WEEKLY expiration=never
```

Above command enable MAIN-WAS-WEEKLY scheduler without expiry date. After enabling MAIN-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Engine machine. dsmcad provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

WebSphere Application Configuration Backup

A sample scheduler named as MAIN-WAS-DAILY is configured to take backups of WAS profile configuration. This schedule invokes WASProfileConfigBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Engine Machine. Before enabling MAIN-WAS-DAILY schedule, you must update WASProfileConfigBackup_IIS.sh and provide value for PASSWORD in the shell script.

MAIN-WAS-DAILY scheduler is configured to execute daily at midnight.

WASProfileConfigBackup_IIS.sh executes commands to create an archive file named IIS_AppServerConfig.zip, which contains all profile configurations.

- WAS Dmgr01 profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/Dmgr01
- WAS Custom01 profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/Custom01

IIS_AppServerConfig.zip contains InfoSphere profile configuration which is generated by using backupConfig command.

Archive files are stored in /home_/WAS_Conf/ folder. WASProfileConfigBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServerConfig.zip file.

In order to enable MAIN-WAS-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-WAS-DAILY expiration=never
```

Above command enable MAIN-WAS-WEEKLY scheduler without expiry date.

After enabling MAIN-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

IIS Database Backup

Database is configured with linear logging, which means all the transaction (archive) logs of database are stored in the IIS Engine machine. It is recommended that these logs should be stored in Spectrum Protect server itself. In case of a failover scenario, the logs are safe and secure in the Spectrum Protect server.

As transaction (archive) logs of database are stored in the IIS Engine machine, it is the user's responsibility to clean up unused logs, as these logs are not cleared by default. Once you move transaction (archive) logs from local disk to Spectrum Protect server, it's the responsibility of Spectrum Protect server to clean unused logs.

Database is configured with HADR scenario, verify HADR status before applying any database configuration changes.

Open Putty or terminal in the IIS Engine machine, switch to db2inst1 user.

```
db2pd -db xmeta -hadr db2pd -db dsodb -hadr
```

If it's proper then follow these steps.

Use below commands in moving transaction (archive) logs from local disk to Spectrum Protect server:

Open Putty or terminal in the IIS Engine machine, switch to db2inst1 user.

```
db2 update database configuration for xmeta using LOGARCHMETH1 TSM:SERVICEMGMTCLASS
db2 stop db manager force
db2 start db manager
```

```
db2 update database configuration for dsosb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS db2 stop
db manager force db2 start db manager
```

Wait for 2 or 3 minutes and check HADR status.

```
db2pd -db xmeta -hadr
db2pd -db dsodb -hadr
```


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. Hence before the sample schedules are started, you have to make sure that one full backup of the respective database is available. Verify available backups for a database by using 'db2adutl query db <DATABASE_NAME>'. Open Putty or terminal in the IIS Engine machine, switch to db2inst1 user to run this command. You have to execute the shell script db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in the IIS Engine Machine, incase there are no full backups available.

Before running db2FullBackup.sh you must consider, where to store database archive logs.

Online full database backups

A sample scheduler named as MAIN-DB-FULL-WEEKLY is configured to take backups of online full database. This schedule invokes db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Engine Machine.

MAIN-DB-FULL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

In order to enable MAIN-DB-FULL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DB-DOMAIN MAIN-DB-FULL-WEEKLY expiration=never
```

Above command enable MAIN-DB-FULL-WEEKLY scheduler without expiry date. After enabling MAIN-DB-FULL-WEEKLY schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

Online incremental database backups

A sample scheduler named as MAIN-DB-INCREMENT-DAILY is configured to take online incremental database backups. This schedule invokes db2IncrementalBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Engine Machine.

MAIN-DB-INCREMENT-DAILY scheduler is configured to execute from Monday to Saturday at midnights, except Sunday.

In order to enable MAIN-DB-INCREMENT-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DB-DOMAIN MAIN-DB-INCREMENT-DAILY expiration=never
```

Above command enable MAIN-DB-INCREMENT-DAILY scheduler without expiry date. After enabling MAIN-DB-INCREMENT-DAILY schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

IIS artifacts backup using ISTOOL

A sample scheduler named as MAIN-ISTOOL-WEEKLY is configured to take backups of ISTool export configuration. This schedule invokes istool_assets.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Engine Machine. Before enabling MAIN-ISTOOL-WEEKLY schedule, you must update istool_assets.sh and provide value for PASSWORD field.

MAIN-ISTOOL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

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 configuration

Archive file generated by istool.sh is stored in /home_/istool/. istool_assets.sh executes Spectrum Protect server selective backup command to take backup of istool generated file.

In order to enable MAIN-ISTOOL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-ISTOOL-WEEKLY expiration=never
```

Above command enable MAIN-ISTOOL-WEEKLY scheduler without expiry date.

After enabling MAIN-ISTOOL-WEEKLY schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

Files and Directories Backup

Two sample schedulers named as MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 is configured to take backups of files and folders. These schedule invokes `main_FilesDaily1.sh` and `main_FilesDaily2.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Engine Machine.

MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 scheduler are configured to execute everyday at midnight.

`main_FilesDaily1.sh` and `main_FilesDaily2.sh` executes commands to take backup of below files and folders.

- /root/keyfile
- /home_/db2inst1/db2keystore.p12
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/*
- /opt/IBM/InformationServer/ASBServer/apps/lib/iis/classes/*
- /opt/IBM/InformationServer/Server/Projects/*
- /etc/sysconfig/iptables
- /opt/IBM/InformationServer/Server/MsgHandlers/*
- /opt/IBM/InformationServer/Server/Configurations/*
- /opt/IBM/InformationServer/Updates/*
- /opt/IBM/InformationServer/Server/DSODB/*.cfg
- /opt/IBM/InformationServer/Server/DSEngine/dsenv
- /opt/IBM/HTTPServer/conf/*
- /etc/services
- /etc/inittab

`main_FilesDaily1.sh` and `main_FilesDaily2.sh` executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAIN-DOMAIN MAIN-FILES-DAILY1 expiration=never
update schedule MAIN-DOMAIN MAIN-FILES-DAILY2 expiration=never
```

Above command enable MAIN-FILES-DAILY1 and MAIN-FILES-DAILY1 scheduler without expiry date.

After enabling MAIN-FILES-DAILY1 and MAIN-FILES-DAILY2 schedule, start 'dsmcad' service from IIS Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Engine machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating scripts `/opt/tivoli/tsm/client/ba/bin/main_FilesDaily1.sh` and `/opt/tivoli/tsm/client/ba/bin/main_FilesDaily2.sh` available in IIS Engine machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Engine machine is *main*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

```
update stgpool main-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool main-cloud-pool identity=<USERNAME>
update stgpool main-cloud-pool password=<PASSWORD>

update stgpool main-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool main-arc-cloud-pool identity=<USERNAME>
update stgpool main-arc-cloud-pool password=<PASSWORD>
```

IIS Passive Engine machine (mainp)

To take backups of the artifacts present in the IIS Passive Engine machine, a sample policy domain configuration named as MAINP-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of WAS profile directory, WAS profile configuration, database full, database incremental online backups and other files which are located in IIS Passive Engine machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine, execute commands to create backup archive files, copy files to specific location, executes commands to take db2 full and incremental backups and invoke Spectrum Protect commands to copy files to Spectrum Protect Server machine.

As this machine acts as passive, the user must not run any schedulers or configuration changes provided for this machine until the machine behaves as active one.

Don't run below schedulers or configuration changes at the starting. Start these schedulers and make the configuration changes only if this machine is taken over as an Active Engine machine

WebSphere Application Profiles Backup

A sample scheduler named as MAINP-WAS-WEEKLY is configured to take backups of WAS profile directories. This schedule invokes WASProfileBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Passive Engine Machine.

MAINP-WAS-WEEKLY scheduler is configured to execute weekly once, on Sunday at midnight.

WASProfileBackup_IIS.sh executes commands to create an archive file named IIS_AppServer_backup.tar which contains all the files from following directories.

- WAS InfoSphere profile backup : /opt/IBM/WebSphere/AppServer/profiles/*

IIS_AppServer_backup.tar contains content of WAS profile directory which is located at /opt/IBM/WebSphere/AppServer/profiles/

The WAS profile directory contains one profile, namely "node01".

Archive files are stored in /home_/WAS_Backup/ folder. WASProfileBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServer_backup.tar files.

In order to enable MAINP-WAS-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAINP-DOMAIN MAINP-WAS-WEEKLY expiration=never
```

Above command enable MAINP-WAS-WEEKLY scheduler without expiry date. After enabling MAINP-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

WebSphere Application Configuration Backup

A sample scheduler named as MAINP-WAS-DAILY is configured to take backups of WAS profile configuration. This schedule invokes WASProfileConfigBackup_IIS.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in IIS Passive Engine Machine. Before enabling MAINP-WAS-DAILY schedule, you must update WASProfileConfigBackup_IIS.sh and provide value for PASSWORD in the shell script.

MAINP-WAS-DAILY scheduler is configured to execute daily at midnight.

WASProfileConfigBackup_IIS.sh executes commands to create an archive file named IIS_AppServerConfig.zip, which contains all profile configurations.

- WAS "node01" profile configuration backup : /opt/IBM/WebSphere/AppServer/profiles/node01

IIS_AppServerConfig.zip contains InfoSphere profile configuration which is generated by using backupConfig command.

Archive files are stored in /home_/WAS_Conf/ folder. WASProfileConfigBackup_IIS.sh executes Spectrum Protect server selective backup command to take backup of IIS_AppServerConfig.zip file.

In order to enable MAINP-WAS-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAINP-DOMAIN MAINP-WAS-DAILY expiration=never
```

Above command enable MAINP-WAS-WEEKLY scheduler without expiry date.

After enabling MAINP-WAS-WEEKLY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine and run `service dsmcad restart` using root user.

IIS Database Backup

As this machine acts as passive, the user must not run any database schedulers or configuration changes provided for this machine until the machine behaves as active one.

Don't run below schedulers or configuration changes at the starting. Start these database schedulers or configuration changes only if this machine is taken over as an Active Services machine

Database is configured with linear logging, which means all the transaction (archive) logs of database are stored in the IIS Passive Engine machine. It is recommended that these logs should be stored in Spectrum Protect server itself. In case of a failover scenario, the logs are safe and secure in the Spectrum Protect server.

As transaction (archive) logs of database are stored in the IIS Passive Engine machine, it is the user's responsibility to clean up unused logs, as these logs are not cleared by default. Once you move transaction (archive) logs from local disk to Spectrum Protect server, it's the responsibility of Spectrum Protect server to clean unused logs.

Open Putty or terminal in the IIS Passive Engine machine, switch to db2inst1 user.

```
db2pd -db xmeta -hadr  
db2pd -db dsodb -hadr
```

If it's proper then follow these steps.

Use below commands in moving transaction (archive) logs from local disk to Spectrum Protect server:

Open Putty or terminal in the IIS Passive Engine machine, switch to db2inst1 user.

```
db2 update database configuration for xmeta using LOGARCHMETH1 TSM:SERVICEMGMTCLASS  
db2 stop db manager force  
db2 start db manager
```

db2 update database configuration for dsodb using LOGARCHMETH1 TSM:SERVICEMGMTCLASS db2 stop db manager force db2 start db manager

Wait for 2 or 3 minutes and check HADR status.

```
db2pd -db xmeta -hadr db2pd -db dsodb -hadr
```

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. Hence before the sample schedules are started, you have to make sure that one full backup of the respective database is available. Verify available backups for a database by using 'db2adutl query db <DATABASE_NAME>'. Open Putty or terminal in the IIS Passive Engine machine, switch to db2inst1 user to run this command. You have to execute the shell script db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder in the IIS Passive Engine Machine, incase there are no full backups available.

Before running db2FullBackup.sh you must consider, where to store database archive logs.

Online full database backups

A sample scheduler named as MAINP-DB-FULL-WEEKLY is configured to take backups of online full database. This schedule invokes db2FullBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Passive Engine Machine.

MAINP-DB-FULL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

In order to enable MAINP-DB-FULL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAINP-DB-DOMAIN MAINP-DB-FULL-WEEKLY expiration=never
```

Above command enable MAINP-DB-FULL-WEEKLY scheduler without expiry date. After enabling MAINP-DB-FULL-WEEKLY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

Online incremental database backups

A sample scheduler named as MAINP-DB-INCREMENT-DAILY is configured to take online incremental database backups. This schedule invokes db2IncrementalBackup.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Passive Engine Machine.

MAINP-DB-INCREMENT-DAILY scheduler is configured to execute from Monday to Saturday at midnights, except Sunday.

In order to enable MAINP-DB-INCREMENT-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder us available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAINP-DB-DOMAIN MAINP-DB-INCREMENT-DAILY expiration=never
```

Above command enable MAINP-DB-INCREMENT-DAILY scheduler without expiry date. After enabling MAINP-DB-INCREMENT-DAILY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

IIS artifacts backup using ISTOOL

A sample scheduler named as MAINP-ISTOOL-WEEKLY is configured to take backups of ISTool export configuration. This schedule invokes `istool_assets.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder in IIS Passive Engine Machine. Before enabling MAINP-ISTOOL-WEEKLY schedule, you must update `istool_assets.sh` and provide value for PASSWORD field.

MAINP-ISTOOL-WEEKLY scheduler is configured to execute weekly once, on Sunday's at midnight.

`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 configuration

Archive file generated by `istool.sh` is stored in `/home_/istool/`. `istool_assets.sh` executes Spectrum Protect server selective backup command to take backup of `istool` generated file.

In order to enable MAINP-ISTOOL-WEEKLY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAINP-DOMAIN MAINP-ISTOOL-WEEKLY expiration=never
```

Above command enable MAINP-ISTOOL-WEEKLY scheduler without expiry date.

After enabling MAINP-ISTOOL-WEEKLY schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

Files and Directories Backup

Two sample schedulers named as MAINP-FILES-DAILY1 and MAINP-FILES-DAILY2 is configured to take backups of files and folders. These schedule invokes `main_FilesDaily1.sh` and `main_FilesDaily2.sh` which is available at `/opt/tivoli/tsm/client/ba/bin` folder inside IIS Passive Engine Machine.

MAINP-FILES-DAILY1 and MAINP-FILES-DAILY2 scheduler are configured to execute everyday at midnight.

`main_FilesDaily1.sh` and `main_FilesDaily2.sh` executes commands to take backup of below files and folders.

- `/root/keyfile`
- `/home_/db2inst1/db2keystore.p12`
- `/opt/IBM/InformationServer/ASBServer/apps/lib/iis/15properties/*`
- `/opt/IBM/InformationServer/ASBServer/apps/lib/iis/classes/*`
- `/opt/IBM/InformationServer/Server/Projects/*`
- `/etc/sysconfig/iptables`
- `/opt/IBM/InformationServer/Server/MsgHandlers/*`
- `/opt/IBM/InformationServer/Server/Configurations/*`
- `/opt/IBM/InformationServer/Updates/*`
- `/opt/IBM/InformationServer/Server/DSODB/*.cfg`

- /opt/IBM/InformationServer/Server/DSEngine/dsenv
- /opt/IBM/HTTPServer/conf/*
- /etc/services
- /etc/inittab

main_FilesDaily1.sh and main_FilesDaily2.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable MAINP-FILES-DAILY1 and MAINP-FILES-DAILY2 schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule MAINP-DOMAIN MAINP-FILES-DAILY1 expiration=never
update schedule MAINP-DOMAIN MAINP-FILES-DAILY2 expiration=never
```

Above command enable MAINP-FILES-DAILY1 and MAINP-FILES-DAILY1 scheduler without expiry date.

After enabling MAINP-FILES-DAILY1 and MAINP-FILES-DAILY2 schedule, start 'dsmcad' service from IIS Passive Engine machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Passive Engine machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating scripts /opt/tivoli/tsm/client/ba/bin/main_FilesDaily1.sh and /opt/tivoli/tsm/client/ba/bin/main_FilesDaily2.sh available in IIS Passive Engine machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS Engine machine is *mainp*.

Use Operation center in updating the details of Object store details.

IIS Services machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-s/m/l-bckp1.

```
update stgpool mainp-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool mainp-cloud-pool identity=<USERNAME>
update stgpool mainp-cloud-pool password=<PASSWORD>

update stgpool mainp-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool mainp-arc-cloud-pool identity=<USERNAME>
update stgpool mainp-arc-cloud-pool password=<PASSWORD>
```

IIS Compute machine (cmpt)

To take backups of the artifacts present in the IIS Compute machine, a sample policy domain configuration named as COMPUTE-DOMAIN is created on Spectrum Protect Server.

Sample schedulers are available to take backup of files which are located in IIS Compute machine. These schedulers when enabled notifies Spectrum Protect Client to execute specific scripts. Sample scripts which are available at Spectrum Protect Client machine which invokes Spectrum Protect commands to copy files to Spectrum Protect Server machine.

Files and Directories Backup

A sample scheduler named as COMPUTE-FILES-DAILY is configured to take backups of files and folders. This schedule invokes compute_FilesDaily.sh which is available at /opt/tivoli/tsm/client/ba/bin folder inside IIS Compute Machine.

COMPUTE-FILES-DAILY scheduler is configured to execute everyday at midnight.

compute_FilesDaily.sh executes commands to take backup of below files and folders.

- /root/keyfile
- /etc/sysconfig/iptables
- /etc/services
- /etc/inittab

compute_FilesDaily.sh executes Spectrum Protect server selective backup command to take backup of files and folders as mentioned above.

In order to enable COMPUTE-FILES-DAILY schedule you need to execute following command using "Command Builder". Details about Command Builder is available at "Starting command builder" section. [Starting Command Builder](#)

```
update schedule COMPUTE-DOMAIN COMPUTE-FILES-DAILY expiration=never
```

Above command enable COMPUTE-FILES-DAILY scheduler without expiry date.

After enabling COMPUTE-FILES-DAILY schedule, start 'dsmcad' service from IIS Compute machine. *dsmcad* provides a light-weight timer which automatically starts and stops the scheduler process as needed.

Open Putty or terminal in IIS Compute machine and run `service dsmcad restart` using root user.

You can modify existing files, folders or add new files, folders for backup by updating script /opt/tivoli/tsm/client/ba/bin/compute_FilesDaily.sh available in IIS Compute machine.

Object store configuration

Out of the box Object store in Spectrum protect server is configured with dummy credentials and URL. You need to update the details of Object store credentials, in cloud storage pools.

Bucket name used for IIS compute machine is *compute*.

Use Operation center in updating the details of Object store details.

IIS Compute machine contains 2 cloud pools, one used to store backup data and another used to store archive data.

Use private URL instead of public URL, if it's private URL data transfer between Spectrum Protect server and Object store is very fast compared to usage of public URL.

Run below commands from "Command Builder" of Spectrum protect operation center. As these are cloud container pools client needs to run from second instance of Spectrum Protect server. In "Command Builder" select second instance which has name like <ORDER_ID>-/s/m/l-bckp1.

```
update stgpool compute-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool compute-cloud-pool identity=<USERNAME>
update stgpool compute-cloud-pool password=<PASSWORD>

update stgpool compute-arc-cloud-pool clouduurl=https://<PRIVATE_URL>
update stgpool compute-arc-cloud-pool identity=<USERNAME>
update stgpool compute-arc-cloud-pool password=<PASSWORD>
```

Add new schedulers

Spectrum protect backup server comes with sample schedulers and the user can create new schedulers based on their requirement.

More information on creating new schedulers is available [here](#)

Update existing schedulers

Spectrum protect backup server comes with sample schedulers and the user can change existing schedulers according to their requirement or they can create a new scheduler.

The user can update the starting time of the scheduler, when to execute the scheduler or disable scheduler based on their requirements.

More information on updating schedulers is available [here](#)

Create new policies and domain

Spectrum protect backup server comes with default policies and domain for each machine. The user can new policies and domains based on their requirement.

More information about adding new policies and domain is available [here](#)

The user can specify their own rules on when to take backup, archive and data retention requirements. More information about these are available [here](#)

Start Spectrum Protect after server restart or reboot

If Spectrum Protect server is restarted or rebooted, follow these steps to start both Spectrum Protect instances and Operation center.

Open Spectrum Protect server using putty or terminal from the IIS Windows Client machine. Execute the following commands.

```
cd /bckp/opt/tivoli/tsm/server/bin/  
./bckp/tsminst1/sqllib/db2profile  
./dsmsevr -u tsminst1 -i /bckp/tsminst1 -q &
```

This will start Spectrum Protect server 1.

```
./bckp/tsminst2/sqllib/db2profile  
./dsmsevr -u tsminst2 -i /bckp/tsminst2 -q &
```

This will start Spectrum Protect server 2.

Starting Operation center

```
cd /bckp/opt/tivoli/tsm/ui/Liberty/bin  
service opscenter.rc status  
service opscenter.rc start
```

Initially Spectrum Protect server second instance will be down, it'll take 5 to 10 minutes to come up. You can check if both instances are up, under overview tab in operation center console.

https://<PUBLIC_IP>:11090/oc or https://<PRIVATE_IP>:11090/oc

If any of Spectrum protect client (like Services or Engine machine) is restarted run dsmcad service Open Putty or terminal in Spectrum protect client machine which is restarted and run `service dsmcad start` using root user.

Protecting the master encryption key

Data encryption and decryption is handled automatically by the Spectrum Protect server and does not require any user action apart from some initial configuration. To encrypt data for cloud-container storage pools, the server uses a master encryption key, which is created when the server password is set. The master encryption key is itself encrypted, and is stored as part of the server password file.

The master encryption key is stored in the server password file, `/bckp/tsminst1/dsmsevr.pwd` for the first server instance and `/bckp/tsminst2/dsmsevr.pwd` for the second server instance in the Spectrum Protect server machine. The master encryption key is encrypted by a different key, so the master encryption key is itself protected. The master encryption key is re-encrypted whenever the server password is set by the

SET SERVERPASSWORD command, so the user can issue this command periodically to further protect the key.

To decrypt data that was sent to encrypted cloud-container storage pools, the master encryption key is required. For this reason, it is important that the server password file is protected. If the server password file is lost or corrupted, the server cannot decrypt the data.

It is recommended that the user copy these files to Object Store or some secure location.

More details about master key is [here](#)

Spectrum Protect server database backup

In case of a scenario, where the Spectrum Protect server is no longer accessible, the Spectrum Protect server can be restored back to its latest state. To recover or restore any Spectrum Protect server, the following artifacts pertaining to that are required:

- Database used by Spectrum Protect for its functionality (.dbv)
- Metadata volume history file (volhist.out)
- Device configuration file (devconfig.out)

Since there are two Spectrum Protect server instances used in this offering, we need to backup two sets of the above mentioned artifacts. Scripts have been created to store these artifacts in a specific location locally. These artifacts are mandatorily required in the scenario where the Spectrum Protect server has to be restored. Hence, **it is recommended that the user should back transfer these files to a secure place.**

Scripts have been created within the Spectrum Protect server to take backup of these artifacts. Administrative schedulers are used to run these scripts. The details of these scripts and their relation to the two server instances can be found below.

First instance

An administrative schedule named *tminst1db*, created in the first instance, is used to call this script named *backuptsms*. The schedule calls this script everyday at 18:00:00. The script, *backuptsms*, creates the artifacts required to restore the Spectrum Protect Server in the */bckp/tsmdbbckps/* location in the Spectrum Protect server. The script also deletes the old backup artifacts and maintains only the latest two versions at any point of time. The script can be queried by using the following command in the command builder.

```
query script backuptsms format=lines
```

Second instance

An administrative schedule named *tminst1db*, created in the second instance, is used to call this script named *backuptsmt*. The schedule calls this script everyday at 19:00:00. The script, *backuptsmt*, creates the artifacts required to restore the Spectrum Protect Server in the */bckp/tsmdbbckpt/* location in the Spectrum Protect server. The script also deletes the old backup artifacts and maintains only the latest two versions at any point of time. The script can be queried by using the following command in the command builder.

```
query script backuptsmt format=lines
```

These scripts also clears the Spectrum Protect database archive logs.

Spectrum Protect server Inventory Expiration

As mentioned in the [Retention Policy](#) section, a file can be present in 3 states: active, inactive and expired. Once the file has reached the "expired" mode, it has to be manually deleted from the Spectrum Protect server to free up the space in order to take continuous backup. Inventory Expiration enables us to delete these expired artifacts.

Expire Inventory command might take several minutes to hours some times which results in slowness of the Spectrum Protect Server. This command should be scheduled to run only when the Spectrum Protect server is not busy.

Schedules have been created for both the Spectrum Protect server instances, which call scripts to execute the Inventory Expiration command. Information on these scripts and schedules have been given below.

First instance

An administrative schedule named *EXPIRES* is used to call the script, *EXPIRES*, every day at 07:00:00. It can be queried by using the following command in the command builder.

```
query script EXPIRES format=lines
```

Second instance

An administrative schedule named *EXPIRET* is used to call the script, *EXPIRET*, every day at 07:00:00. It can be queried by using the following command in the command builder.

```
query script EXPIRET format=lines
```

More details about expire inventory details are [here](#)

Restoring Backups

This section covers the steps involved to restore various IIS artifacts to IIS server. Before any steps to restore are executed, the services of IIS must be stopped. Follow the following steps to stop the services and then restore the artifacts. After restoring the artifacts, start the services again.

Steps to stop IIS services.

1. Login to IIS Engine machine as root and execute the following in Putty.

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSEngine
./dsenv; bin/uv -admin -stop;
```

2. Stop Node Agents

```
su - root
cd /opt/IBM/InformationServer/ASBNode/bin
./opt/IBM/InformationServer/Server/DSEngine/dsenv
./NodeAgents.sh stop
```

3. Stop Appwatcher process

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSODB/bin
./DSAppWatcher.sh -stop
```

4. Stop WAS

```
su - root
cd /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin
./stopServer.sh server1 -user wasadmin -password <PASSWORD>
```

Steps to start IIS services

1. Start WAS

Login to IIS Engine machine

```
su - root
cd /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin
./startServer.sh server1 -user wasadmin -password <PASSWORD>
```

1. Start DS Engine services

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSEngine
. ./dsenv; bin/uv -admin -start;
```

2. Start node Agents

```
su - root
cd /opt/IBM/InformationServer/ASBNode/bin
. /opt/IBM/InformationServer/Server/DSEngine/dsenv
./NodeAgents.sh start
```

3. Start Appwatcher process

```
su - dsadm
cd /opt/IBM/InformationServer/Server/DSODB/bin
./DSAppWatcher.sh -start
```

Restoring DB2 database

The following db2 databases are available in IIS DataStage on Cloud in IIS Engine machine:

- xmeta
- dsodb

The databases are referred to as <DATABASE_NAME> in the upcoming sections.

Follow the mentioned steps in order to restore a specific database back to a specific timestamps

1. Drop existing database.
2. Restore database.
3. Verify if the database is restored.

Drop existing database

1. Open terminal for IIS Engine machine, switch to db2inst1 user using command `su - db2inst1`
2. `db2 LIST APPLICATIONS` , to check if any applications are connected to this database.
3. Using `db2adutl` command , check available full or incremental backups. Note down timestamp which you are going to restore. Once database is dropped you can't see available full or incremental backups.
4. Drop the database using `db2 drop db <DATABASE_NAME>` , if you face any issues that means database is connected to applications , we need to close all connections to database before dropping it.

```
db2 connect to <DATABASE_NAME> user <DATABASE_USER> using <PASSWORD>
db2 quiesce db immediate force connections
db2 connect reset;
db2 LIST APPLICATIONS
db2 terminate
db2 force application all
db2 drop database <DATABASE_NAME>
```

5. In order to make sure , there is no database named <DATABASE_NAME> , execute list command. `db2 list db directory`

Restore Database

1. Using `db2adutl` command , check available full or incremental backups.
2. Select full or incremental backup timestamp which needs to be restored.
3. `db2 restore db <DATABASE_NAME> use tsm taken at 20170526063832 ENCRYPT`
4. In above statement '20170526063832' is timestamp when the DB backup was taken.
5. "use TSM" is used which means , we are restoring a database which is stored in Spectrum Protect server.

6. "ENCRYPT" is used as existing database is db2 native encrypted.
7. While taking backup "include logs" is used , so if needed you can use "include logs" option, when there is a need to extract transaction logs needed in restore scenario.
8. Select incremental backup timestamp if you want to restore till that date.


```
db2 restore db <DATABASE_NAME> incremental automatic use tsm taken at 20170605134603 1.
```

 On top of full backup , we are restoring incremental backup.
9. db2 rollforward db <DATABASE_NAME> to end of logs and complete.
 - a. Used to rollforward database till end of logs.

Verify database

1. Connect to <DATABASE_NAME> database using following command. db2 connect to <DATABASE_NAME> user <DATABASE_USER> using <PASSWORD>
2. Verify any table to check restore process is done and expected data is available.
3. Terminate database using command, db2 terminate
4. Verify if db2 native encryption is available.


```
db2 connect to <DATABASE_NAME> user db2inst1 using <PASSWORD>
db2 "SELECT * FROM TABLE (SYSPROC.ADMIN_GET_ENCRYPTION_INFO())"
```
5. Verify if AES is available in the output.

Restoring files and directories

As has been discussed in the Backup Section of IIS DataStage on Cloud, sample policies and schedules are created to backup specific files and directories to the Spectrum Protect Server from individual machines like the IIS Engine machine and IIS Compute machine.


There are two things to be particular about when a file/directory is being restored:

- Whichever file/directory is being restored, make sure that a copy of the same is created as a temporary file.
- The user privileges of the file/directory must be noted before restore. After restore, if there are any discrepancies, the user privileges must be set by the user for the restored files/directories so that they match with the ones of the current file/directory.


The method to restored has been explained for the IIS Engine machine. The procedure is the same for all other machines as well where Spectrum Protect Client is installed. The files/directories can be restored in the following way:

1. Log into the IIS Services machine as root in a GUI session.

2. Execute the following command: /opt/tivoli/tsm/client/ba/bin/dsmj



3. Under the "Restore" section, all the files and directories which have been backed up will be available to be selected for restore.



4. Select the files to be restored and click on the restore button.

Restoring WAS artifacts

The following WAS artifacts are backed up to the Spectrum Protect Server. WAS artifacts are to be restored in the IIS Engine machine.

- InfoSphere profile

- InfoSphere profile configurations

InfoSphere profile

The backup file of InfoSphere profile is stored in the following location: /home_/WAS_Backup/IIS_AppServer_backup.zip

The user can use the file current in this location, which is the backup taken on the previous Sunday. Or the user can choose different backup versions of this file. Different versions of backup file can be restored using the /opt/tivoli/tsm/client/ba/bin/dsmj. Once the required backup file has been made available, execute the following commands as "root" user to restore the InfoSphere profile.

```
/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin/stopServer.sh server1 -username wasadmin -password PASSWORD
```

Rename the "InfoSphere" profiles directory in /opt/IBM/WebSphere/AppServer/profiles to "InfoSphere1". Execute the following commands.

```
/opt/IBM/WebSphere/AppServer/bin/manageprofiles.sh -restoreProfile -backupFile <location_of_backup_file>  
/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/bin/startServer.sh server1
```

InfoSphere profile configurations

The backup file of InfoSphere profile is stored in the following location: /home_/WAS_Conf/IIS_AppServerConfig.zip

The user can use the file current in this location, which is the backup taken on the previous Sunday. Or the user can choose different backup versions of this file. Different versions of backup file can be restored using the /opt/tivoli/tsm/client/ba/bin/dsmj. Once the required backup file has been made available, follow the procedure to restore the InfoSphere profile configurations.

1. Open terminal for IIS Engine machine using root user.
2. Execute the following command: /opt/IBM/WebSphere/AppServer/bin/restoreConfig.sh <location_of_backup_file> -nostop -username wasadmin -password <PASSWORD> -profileName InfoSphere

Restoring ISTOOL asset

InfoSphere Information Server provides a backup utility called "ISTOOL" which enable us to backup assets related to Information Governance Catalog, Information Anaylzer etc. If these components are to be restored, they are backed up in the following location in the IIS Engine machine.

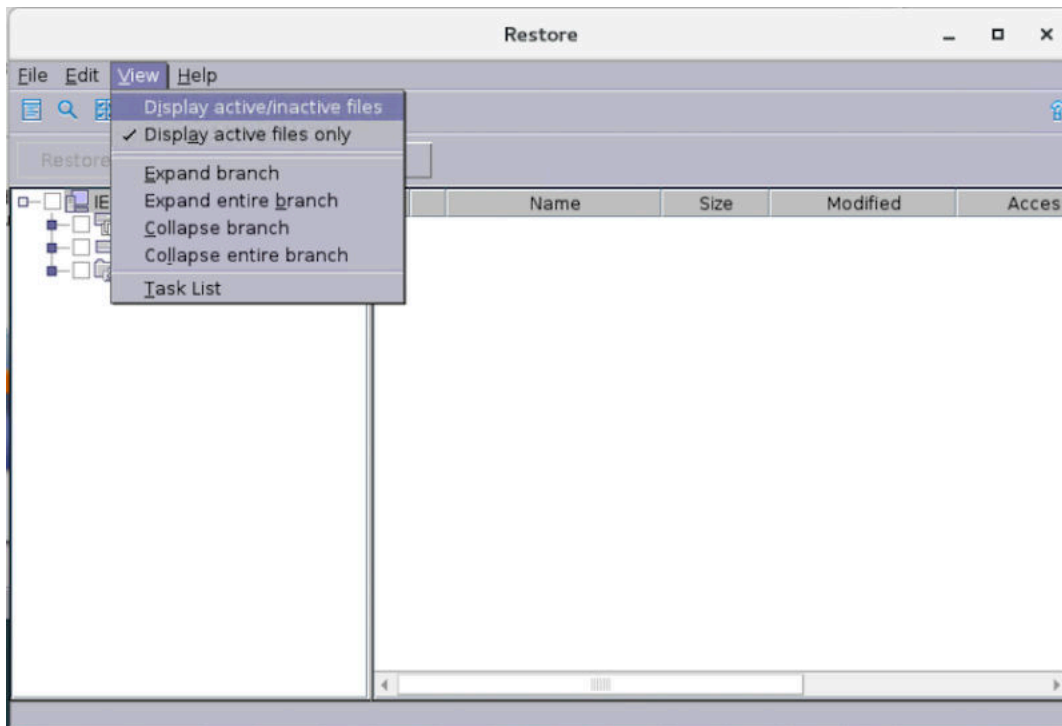
```
/home_/istool/istool.isx
```

Follow the procedure to restore these assets from the .isx archive file

```
cd /opt/IBM/InformationServer/Client/istools/cli  
import -dom <URL_IIS_ENGINE>:9443 -u isadmin -p <PASSWORD> -archive "/home_/istool/istool.isx" -all
```

Viewing and restoring multiple versions of a specific file

Spectrum Protect enables the user to restore not just the latest version of the file backup but older versions as well, if they are available in the Spectrum Protect server. In order to view the available versions, you can toggle the "Display active/inactive files" option in the "View" tab of Restore window.



The versions available in this option will not be more that 30 days old. If an older version is required, the user can use the cloud object storage to restore backup versions which are older that 30 days but not more than 365 days.

Restoring from Cloud Object Storage

As mentioned in the Spectrum Protect Setup, there are two Spectrum Protect Server instances. First instance is connected to local storage with the Spectrum Protect server and the second instance uses the Object Storage for storing backup.

Cloud Object storage stores backup of artifacts which are up to 365 days old. In order to restore the from the cloud object storage, follow the procedure.

Halt the first instance of the Spectrum Protect Server

Shut down the first instance of Spectrum Protect Server (tsminst1), so that the Spectrum Protect client can connect to the second instance of the Spectrum Protect server. The second instance of Spectrum Protect server is connected to the Cloud Object Storage.

1. Connect to the Operations Center. [Start the command builder.](#)
2. Execute the following commands.

```
DISABLE SESSIONS
QUERY SESSIONS
CANCEL SESSIONS
HALT
```

This stops the first server instance (tsminst1) and connects the Spectrum Protect Clients to the second server instance (tsminst2).

Change the SSL Certificate to connect to second instance of Spectrum Protect server

Certificate files of first instance of spectrum protect must be deleted so that the client machine can connect to the second instance. These files are located at the /opt/tivoli/tsm/client/ba/bin/ location of the Spectrum Protect client machine.

```
dsmcert.crl
dsmcert.kdb
```

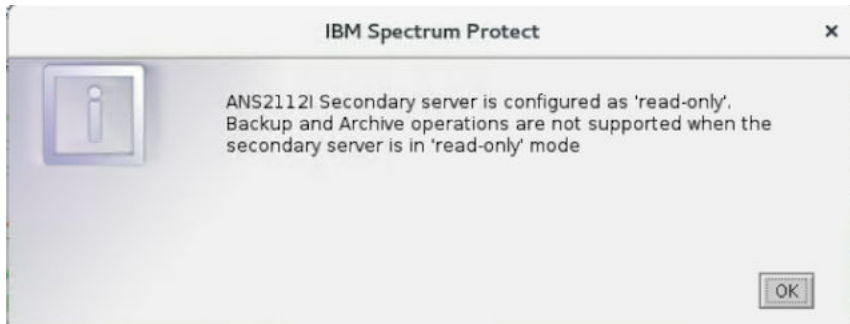


```
dsmcert.rdb  
dsmcert.sth
```

Execute the following commands to create the ssl certificate key for the second server instance.

```
cd /opt/tivoli/tsm/client/ba/bin/  
gsk8capicmd_64 -keydb -create -populate -db dsmcert.kdb -pw <Spectrum_Protect_Server_Password> -  
stash  
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed -label "<description>" -file /opt/tivoli/tsm/  
client/ba/bin/tsminst2/cert256.arm -format ascii  
chmod 777 dsmcert.*
```

Use the `/opt/tivoli/tsm/client/ba/bin/dsmj` in the client machine to restore the required files/directories. The GUI will be opened in read only format. This implies we can only restore the files that have already been backed up.



After the restoration is completed, to revert back to the original configuration (ie. connecting back to first instance of Spectrum Protect Server), follow the procedure in the next sub-section.

Connecting back to first instance of Spectrum Protect server

The first server instance (tsminst1) has to be started to revert back to the original configuration of backup and its schedules. Open the terminal and execute the following commands in Spectrum Protect Server machine.

```
./bckp/tsminst1/sql/lib/db2/profile  
/bckp//opt/tivoli/tsm/server/bin/dsmserve -u tsminst1 -i /bckp/tsminst1
```

The above mentioned commands starts the first server instance. Next, the SSL certificates in the Spectrum Protect client machine has to be made compliant to the first server instance.

In the Spectrum Protect client machines, execute the following steps.

1. Delete the following files in `/opt/tivoli/tsm/client/ba/bin/`

```
dsmcert.crl dsmcert.kdb dsmcert.rdb dsmcert.sth
```

2. Execute the following commands to change the ssl certificate key.

```
cd /opt/tivoli/tsm/client/ba/bin/ gsk8capicmd_64 -keydb -create -populate -db dsmcert.kdb -pw  
<Spectrum_Protect_Server_Password> -stash gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed -  
label "<description>" -file /opt/tivoli/tsm/client/ba/bin/tsminst1/cert256.arm -format ascii chmod  
777 dsmcert.*
```

Chapter 5. DataStage on Cloud Designer Client

IBM® DataStage® on Cloud Designer Client provides a hosted environment that you configure and control. You can use DataStage on Cloud Designer Client is included in IBM Information Server on Cloud Enterprise Edition, but it can be used as an independent service. Use DataStage on Cloud Designer Client to create, design, and develop DataStage jobs.

Overview


IBM® DataStage® on Cloud Designer Client is a data integration software platform that helps organizations derive more value from the complex, heterogeneous information spread across their systems. DataStage on Cloud Designer Client provides a single unified platform to understand, cleanse, transform, and deliver trustworthy and context-rich information.

DataStage on Cloud Designer Client includes the following IBM Information Server on Cloud clients:

- IBM InfoSphere® Information Governance Catalog
- IBM InfoSphere Information Analyzer
- IBM InfoSphere Metadata Integration Bridges and the metadata interchange agent
- IBM InfoSphere Metadata Asset Manager
- IBM InfoSphere Information Server istool command-line utility
- IBM InfoSphere Information Server Manager client, Multi-Client Manager
- IBM InfoSphere DataStage and QualityStage® Administrator
- IBM InfoSphere DataStage and QualityStage Designer
- IBM InfoSphere DataStage and QualityStage Director

DataStage on Cloud Designer Client uses the characteristics of software-as-a-service (SaaS). You select the plan size based on your needs. IBM provisions the machine and deploys the DataStage on Cloud Designer Client software.

The following figure shows the topology of the client machine in a typical deployment.



As a hosted offering, you have the same control over your data in the cloud as in the on-premises system:

- Actively monitor and report any issues that you encounter with IBM Software as a Service (SaaS).
- Maintain the software platform of your cloud offering and the operating system to meet your security standards.
- Maintain software firewalls on servers that face the internet in a manner to provide required protection.
- Develop integration, transformation, and other jobs. Establish connectivity between data sources and applications. You can also develop your own workload, business rules, monitoring, and scheduling for all jobs. You are responsible for the quality and performance of programs, applications, and jobs that you develop.
- Ensure the continuity, compatibility, and performance of the IBM SaaS platform by installing only permissible software, including any open source packages.
- Regularly upgrade the environment and operating system of your cloud offering.
- Create and maintain regular backups of data.
- Create and maintain high availability configurations.

The following managed add-on services are available to maintain and manage the infrastructure:

Jump start

This setup service provides up to 50 hours of remote consulting time for startup activities.

Accelerator

This setup service provides up to 50 hours of remote consulting time to perform various scoped activities.

Silver

This service provides monthly remote consulting time for operations and maintenance activities.

Gold

This service provides monthly remote consulting time for operations and maintenance activities. The service includes everything that is provided by the Silver service and delivers extra activities.

Restriction: All Information Server on Cloud clients are installed, but you can use only the clients of components that are installed at the server side.

Available configuration

IBM® DataStage® on Cloud Designer Client is on virtual machine with dedicated CPUs.

The following table lists the configuration of the client machine.

<i>Table 19: Designer Client available configurations</i>	
Configuration	Size
Memory (GB)	4
Number of cores	4
Network speed (Gbps)	1 Gbps with 1000 GB bandwidth
First disk	100 GB Storage Area Network (SAN)
Second disk	100 GB SAN

The virtual machine has a Microsoft Windows operating system environment. All IBM Information Server on Cloud client utilities are installed on DataStage on Cloud Designer Client.

Layout of IBM® DataStage® on Cloud Designer Client disks

The DataStage® on Cloud Designer Client disk comes in one size for all plans.

DataStage on Cloud Designer Client is a client virtual machine with two Storage Area Network (SAN) disks that are 100 GB each. One disk is drive C for the Microsoft Windows operating system. The other disk is drive F, and it is an empty disk.

DataStage on Cloud Designer Client can be used in combination with any of the servers that are offered with IBM® Information Server on Cloud and IBM DataStage on Cloud offerings.

Information roadmap

This roadmap lists information resources that are available for users who are new to the DataStage® products from IBM®. These resources provide information about various subject areas, such as learning basic skills and troubleshooting.

Product overview

- **[InfoSphere® DataStage features](#)** Learn about the features and benefits of InfoSphere DataStage.
- **[InfoSphere DataStage product documentation](#)** The InfoSphere Information Server Knowledge Center provides you with InfoSphere DataStage concepts and usage information to help new users prepare for using your new system.
- **[InfoSphere DataStage developerWorks® forum](#)** Use this forum to interact with other InfoSphere DataStage users to better understand how to design, build, debug and deploy jobs for information collection, integration and transformation.

Getting started

- **[Overview of InfoSphere DataStage](#)** This overview covers InfoSphere DataStage job life cycles, job designs, and how the product integrates with the rest of the InfoSphere Information Server suite.
- **[Tutorial: Creating parallel jobs](#)** This tutorial demonstrates how you can use InfoSphere DataStage to develop jobs that extract, transform, and load data. By transforming and cleansing the source data and applying consistent formatting, you enhance the quality of the data.

Using InfoSphere DataStage

- **[Designing InfoSphere DataStage jobs](#)** This topic explains how to design InfoSphere DataStage jobs using the IBM InfoSphere DataStage and QualityStage® Designer client. The Designer client gives you the tools that you need to create jobs that extract, transform, load, and check the quality of data.
- **[Developing InfoSphere DataStage parallel jobs](#)** This topic covers how to design parallel jobs to transform and to cleanse data. Parallel jobs are compiled and run on the InfoSphere Information Server engine.
- **[Developing server jobs](#)** Learn how server jobs are compiled and run on the InfoSphere Information Server engine. Such jobs connect to a data source, extract and transform data, and write data to a target database or file, such as a data warehouse.
- **[Deploying jobs](#)** Use the InfoSphere Information Server Manager to move InfoSphere DataStage assets between projects on the same engine, or on different engines. You can also use the InfoSphere Information Server Manager to move assets from one domain to another.
- **[Running InfoSphere DataStage jobs](#)** This topic covers how to run InfoSphere DataStage from the IBM InfoSphere DataStage and QualityStage Director client.
- **[Administering workload management](#)** Use the workload management queues to control the starting of parallel and server jobs.
- **[Monitoring jobs](#)** Use IBM InfoSphere DataStage and QualityStage Operations Console to monitor the job runs, services, system resources, and workload management queues on several InfoSphere Information Server engines.
- **[Parallel job reference](#)** Use this reference material to perform more advanced operations with parallel jobs.
- **[Data Flow Designer](#)** A new web-based tool, to create, edit, load, and run DataStage jobs is introduced in Information Server 11.7

Troubleshooting and support

- **Troubleshooting InfoSphere DataStage** Use the Troubleshooting page to find common troubleshooting topics.
- **IBM Support Portal** Use the Support Portal for InfoSphere DataStage to search for known problems and APARs.

Getting started and using IBM DataStage on Cloud Designer Client

You must set up your connection to DataStage® on Cloud Designer Client to design, develop, and run jobs that move and transform data. DataStage on Cloud Designer Client includes all clients of IBM® InfoSphere® Information Server. It is in an IBM SoftLayer® cloud hosted environment.

Prerequisite: You must know the IP address and the credentials of an account on DataStage on Cloud Designer Client computer. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM Sales Representative.

DataStage on Cloud Designer Client is on a Microsoft Windows virtual machine. When you connect to DataStage on Cloud Designer Client, you can access client user interfaces, create jobs, and run them. McAfee anti-virus software is installed on the client machine.

Before using DataStage on Cloud Designer Client with any of the IBM Information Server on Cloud, IBM Information Server on Cloud Data Quality, and IBM DataStage on Cloud offerings, required connections must be set up between the client and servers.

When you connect for the first time, follow these steps:

1. Connect DataStage on Cloud Designer Client to the Information Server on Cloud server by following these steps:
 - **If your local computer is in a Microsoft Windows environment**
 - a. On your local computer, go to the **Start** menu. Click **Accessories > Remote Desktop Connection**.
 - b. Enter the IP address of the Microsoft Windows computer that hosts your DataStage on Cloud Designer client. Click **Connect**.
 - c. In the Windows Security window, enter the user name and password for the DataStage on Cloud Designer client. The user ID, password, and IP address of the client are in your Welcome letter. **Important:** Do not include the domain name with the user name.
 - d. In the DataStage on Cloud Designer client, open the file C:\Windows\System32\drivers\etc\hosts. Make sure that an entry with the private IP exists in the file for the Information Server on Cloud server that you are connecting to. **Important:** The server IP must be a private IP. The connection from DataStage on Cloud Designer Client fails when you use a public IP.
 - **If your local computer is in an Apple Mac environment**
 - a. On your local computer, install Microsoft Remote Desktop from the Apple App Store.
 - b. Click the Microsoft Remote Desktop icon, and then click **Open**.
 - c. In the Microsoft Remote Desktop window, click **New**.
 - d. In the Edit Remote Desktops window, supply the following information:
 - In the PC name field, type in the IP address of the cloud client machine.
 - In the User name and Password fields, type in the Windows user name and password that are in the Welcome letter.
2. Verify the connection and installation on DataStage on Cloud Designer Client. On the client computer, follow these steps:
 - Run the ISALite tool. The *IS_install_path* for the client computer is C:\IBM\InformationServer.

- [Test the installation](#) of the IBM InfoSphere DataStage and QualityStage® Administrator and InfoSphere DataStage and QualityStage Designer clients.
3. Optional: On the client computer, enable multiple users to open remote sessions to DataStage on Cloud Designer Client by following these steps.
 - a. [Create user accounts](#).
 - b. [Give users permission](#) to do a remote desktop connection.
 4. [Reset the password](#) for users on DataStage on Cloud Designer Client.

After the initial connection, you can do any of the following tasks:

- You can [add additional DataStage on Cloud Designer Client machines](#) to your environment.
- Depending on the clients that you purchased with the offering, you can create jobs, create data rules, and other tasks. DataStage on Cloud Designer Client includes the following InfoSphere Information Server clients:
 - IBM InfoSphere Information Governance Catalog
 - IBM InfoSphere Information Analyzer
 - IBM InfoSphere Information Governance Dashboard
 - IBM InfoSphere Information Services Director
 - IBM InfoSphere Data Click
 - IBM InfoSphere FastTrack
 - IBM InfoSphere Metadata Integration Bridges and the metadata interchange agent
 - IBM InfoSphere Information Server istool command-line utility
 - IBM InfoSphere Information Server Manager client
 - Multi-Client Manager
 - IBM InfoSphere DataStage and QualityStage Administrator client
 - IBM InfoSphere DataStage and QualityStage Designer client
 - IBM InfoSphere DataStage and QualityStage Director client
- [Open the InfoSphere DataStage® and QualityStage® Designer client](#)
- [Connect to an on-premises computer](#)
- [Connect to an IBM dashDB™ database](#)
- [Connect to an on-premises DB2® database instance](#)
- [Perform general administration and security tasks](#)

As your computing needs grow, you can add and connect more DataStage on Cloud Designer Client machines in your Information Server on Cloud environment.

Related tasks

- [Enhancing security of Information Server on Cloud computers](#)

Adding and connecting extra DataStage on Cloud Designer Client machines

You can add extra IBM® DataStage® on Cloud Designer Client machines in your IBM Information Server on Cloud environment as your computing needs increase.

Prerequisite: You must be logged in to an administrative account on the Information Server on Cloud server.

You use scripts to add more DataStage on Cloud Designer Client machines to your Information Server on Cloud environment.

The scripts are in directory `/opt/IBM/scripts` on the Information Server on Cloud server.

Information Server on Cloud servers come with iptables firewall setup. You change the iptables rules to allow traffic between the new DataStage on Cloud Designer Client machine and the servers.

1. To use the new DataStage on Cloud Designer Client machine with IBM Information Server on Cloud Data Quality and IBM Information Server on Cloud Enterprise Edition, follow these steps:
 - a. In the services tier machine, do the following steps in a command-line window:
 - 1) Open the file `/etc/hosts`, and then add a host entry for the new DataStage on Cloud Designer Client machine.
 - 2) Run the script `allow_services_tier_from_ip.sh` with the private IP address of the new DataStage on Cloud Designer Client machine. For example, the command might be:
`allow_services_tier_from_ip.sh 10.xxx.xxx.xxx.`
 - b. In the engine tier machine, run the script `allow_engine_tier_from_ip.sh` with the private IP address of the new DataStage on Cloud Designer Client machine. This script permits traffic from the DataStage on Cloud Designer Client machine to the engine tier machine.
 - c. In the new DataStage on Cloud Designer Client machine, open the Microsoft Windows PowerShell command-line tool. Run the following command: `netsh advfirewall firewall add rule name="Open HTTPS port for IS" protocol=tcp localport=19443 dir=in remoteip=<Private_ip_of_the_services_tier> profile=any action=allow enable=Yes`. The `<Private_ip_of_the_services_tier>` value is the IP address of the services tier of the Information Server on Cloud machine. This command adds a rule to the iptables firewall to permits traffic from the DataStage on Cloud Designer Client machine to the services tier machine.
2. To use the new DataStage on Cloud Designer Client machine with IBM DataStage on Cloud server, follow these steps:
 - a. In the new DataStage on Cloud Designer Client machine, open the `C:\Windows\System32\drivers\etc\hosts` file. Add a host entry for the DataStage on Cloud server.
 - b. In the DataStage on Cloud server, open the file `/etc/hosts`. Add a host entry for the DataStage on Cloud Designer Client machine.
 - c. Update the iptables firewall rules by following these steps:
 - 1) On the engine tier machine of the DataStage on Cloud server, run the script `allow_datastage_from_ip.sh` with the private IP address of the new DataStage on Cloud Designer Client machine. For example, the command might be:
`allow_datastage_from_ip.sh 10.xxx.xxx.xxx. Tip: Check the list of open ports that are opened by using this script. Depending on the list of services that you are using from the DataStage on Cloud Designer Client machine, you might want to close some ports.`
 - 2) In the new DataStage on Cloud Designer Client machine, open the Windows PowerShell command-line tool. Run the command `netsh advfirewall firewall add rule name="Open HTTPS port for IS" protocol=tcp localport=19443 dir=in remoteip=Private_ip_of_datastage_server profile=any action=allow enable=Yes`. This command opens HTTPS port 19443 on the DataStage on Cloud Designer Client machine to the DataStage on Cloud server.

Chapter 6. Information Governance Catalog on Cloud

IBM® Information Governance Catalog on Cloud provides a hosted environment that you configure and control. You can use Information Governance Catalog on Cloud to extend the reach of your business by leveraging cloud offerings, while you reduce the costs that are associated with providing these services. Information Governance Catalog on Cloud provides several different plans so that any size business can access the powerful and scalable ETL platform by IBM.

Overview

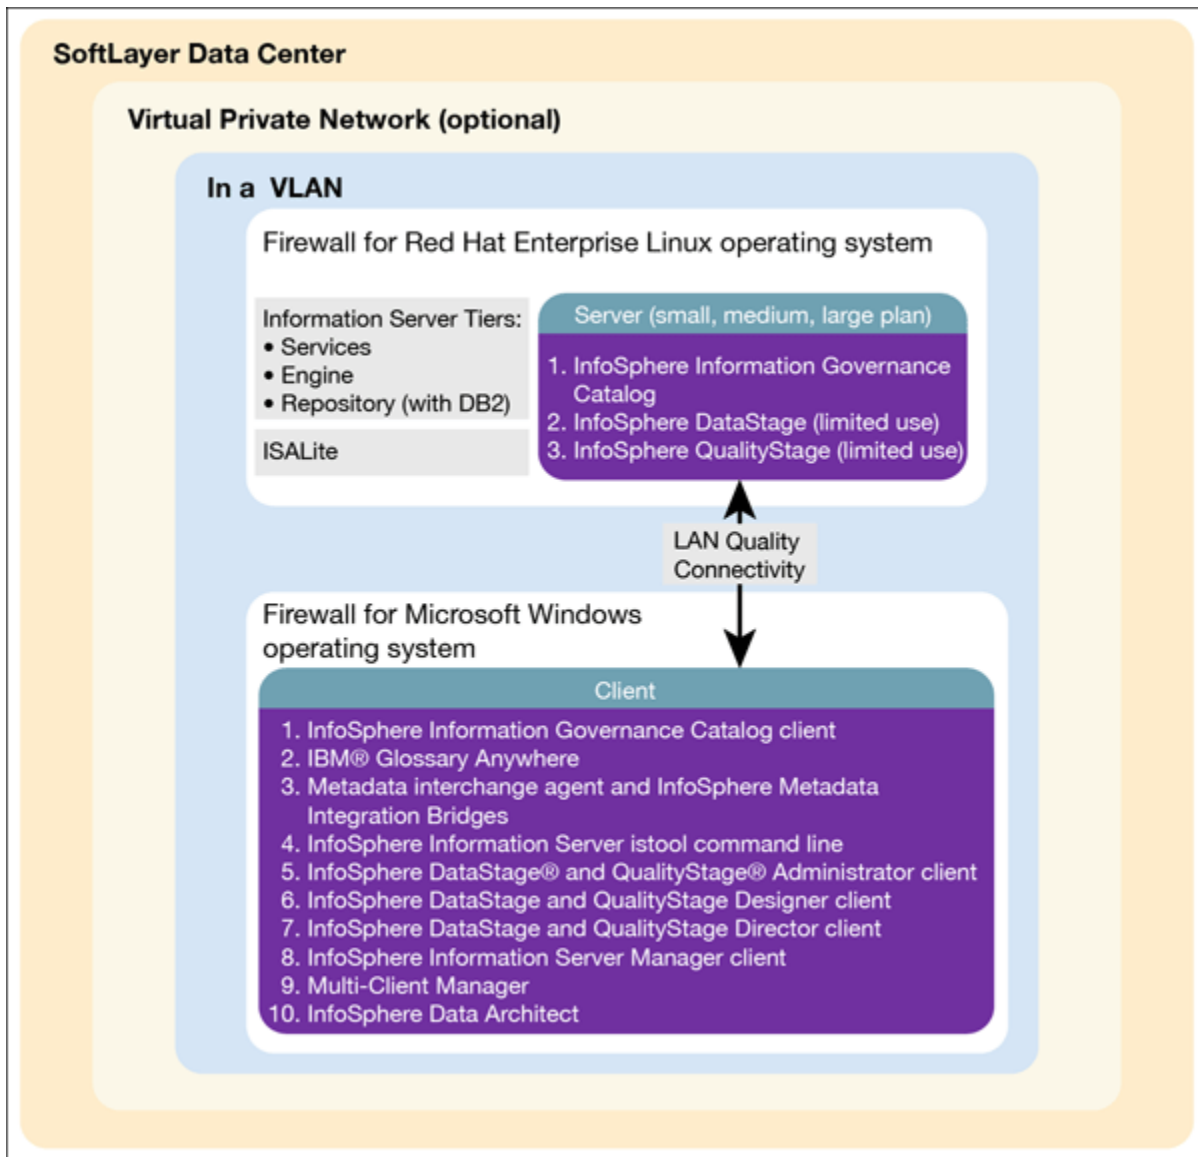
IBM® Information Governance Catalog on Cloud is an interactive, web-based tool in the cloud that enables users to create, manage, and share an enterprise vocabulary and classification system in a central catalog. Information Governance Catalog on Cloud provides all of the functions of its on-premises counterpart, IBM InfoSphere® Information Governance Catalog.

You can do the following tasks by using Information Governance Catalog on Cloud:

- Create a well-documented, end-to-end information blueprint to align your business requirements with your enterprise and reference architectures before you start a strategic project
- Establish a common business language and manage business perspectives about information to align those views with the IT perspective
- Manage and explore data lineage to create trusted information that supports data governance and compliance efforts
- Provide a solid foundation for different types of information integration and governance projects, including information integration, lifecycle management, and security initiatives

Information Governance Catalog on Cloud uses the characteristics of software-as-a-service (SaaS). You select the plan size based on your needs. IBM provisions the machine and deploys the Information Governance Catalog on Cloud software.

The following figure shows the topology of the server and client machines in a typical deployment.



As a hosted offering, you have the same control over your data in the cloud as in the on-premises system:

- Actively monitor and report any issues that you encounter with IBM Software as a Service (SaaS).
- Maintain the software platform of your cloud offering and the operating system to meet your security standards.
- Maintain software firewalls on servers that face the internet in a manner to provide required protection.
- Develop integration, transformation, and other jobs. Establish connectivity between data sources and applications. You can also develop your own workload, business rules, monitoring, and scheduling for all jobs. You are responsible for the quality and performance of programs, applications, and jobs that you develop.
- Ensure the continuity, compatibility, and performance of the IBM SaaS platform by installing only permissible software, including any open source packages.
- Regularly upgrade the environment and operating system of your cloud offering.
- Create and maintain regular backups of data.
- Create and maintain high availability configurations.

The following managed add-on services are available to maintain and manage the infrastructure:

Jump start

This setup service provides up to 50 hours of remote consulting time for startup activities.

Accelerator

This setup service provides up to 50 hours of remote consulting time to perform various scoped activities.

Silver

This service provides monthly remote consulting time for operations and maintenance activities.

Gold

This service provides monthly remote consulting time for operations and maintenance activities. The service includes everything that is provided by the Silver service and delivers extra activities.

Restriction: With Information Governance Catalog on Cloud, users are not authorized to use any of the following components or functions:

- Run a job function of IBM InfoSphere DataStage®
- Design a new asset of IBM InfoSphere DataStage and QualityStage® Designer

Available configurations

IBM® Information Governance Catalog on Cloud servers for the small and medium plans are virtual servers with public CPUs. The servers in the large plan are in a bare metal environment.

Select the offering plan that fits your usage and environment needs.

Table 20: Configurations of Information Server machine

Offering	Memory (GB)	Number of Cores	Network speed and bandwidth	First disk	Second disk
Small	12	4	1 Gbps with 250 GB bandwidth	100 GB Storage area network (SAN)	100 GB SAN
Medium	16	8	1 Gbps with 250 GB bandwidth	100 GB SAN	500 GB SAN
Large	64	12	1 Gbps with 1000 GB bandwidth	960 GB SSD	960 GB SSD

Table 21: Configurations of Information Server Enterprise Search node

Offering	Memory (GB)	Number of Cores	Network speed and bandwidth	First disk	Second disk
Small	32	8	1 Gbps with 250 GB bandwidth	100 GB Storage area network (SAN)	100 GB SAN
Medium	64	16	1 Gbps with 250 GB bandwidth	100 GB SAN	500 GB SAN
Large	64	12	1 Gbps with 1000 GB bandwidth	960 GB SSD	960 GB SSD

A client machine is provided with all offering sizes. The client is in a Microsoft Windows environment with 6 GB of memory, two cores, and two 100 GB SAN storage. The network speed is 1 Gbps, with 250 GB bandwidth.

The following IBM Information Server on Cloud client tier components are installed with Information Governance Catalog on Cloud on client machine:

- IBM InfoSphere® Information Governance Catalog client
- IBM Glossary Anywhere
- IBM InfoSphere Metadata Integration Bridges and the metadata interchange agent
- IBM InfoSphere Metadata Asset Manager
- IBM InfoSphere Information Server istool command-line utility
- IBM InfoSphere Information Server Manager client
- Multi-Client Manager
- IBM InfoSphere DataStage® and QualityStage® Administrator client
- IBM InfoSphere DataStage and QualityStage Designer client
- IBM InfoSphere DataStage and QualityStage Director client
- IBM InfoSphere Data Architect client

New consoles available from Information Server version 11.7

- IBM InfoSphere® Information Governance Catalog clientNew
- IBM InfoSphere® InformationGovernance Monitor
- IBM InfoSphere® InformationServer Enterprise Search

These client tier components import or create assets and their metadata in the metadata repository. The assets and metadata can then be accessed by Information Governance Catalog on Cloud.

Layout of IBM Information Governance Catalog on Cloud server and client disks

The layout of the Information Governance Catalog on Cloud servers and client disks depends on the plan size of your system.

Virtual servers for small and medium plans

The small and medium plans come with two Storage Area Network (SAN) disks. The Red Hat Enterprise Linux operating system is on the first SAN disk in both the small and medium plans. The second SAN disk is encrypted by using Linux Unified Key Setup (LUKS) for both these sizes.

The encryption key details are provided in the Welcome letter from the IBM® Sales Representative. It is recommended that you add your own key and remove the supplied key before you use the system.

The product is installed on the /disk1 directory. User data can be stored on /data directory. Both directories are on the partition /dev/xvdc1 that is encrypted.

Table 22: Small and medium disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/xvda1	256 MB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/	None	/dev/xvda2	99.8 GB	Primary disk is /dev/xvda	Kernel boot area and operating system data
/disk1	LUKS	/dev/xvdc1	100 GB for small plan. 500 GB for medium plan	Secondary disk is /dev/xvdc	Installation files and creation of /data and /opt directories.

Table 22: Small and medium disk layout (continued)

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
SWAP	None	/dev/xvdb1	2 GB	/dev/xvdb	Swap space for paging

Bare metal servers for large plan

The servers in large plan come with two Solid State Drive (SSD) disks that are 960 GB each. RAID level 1 implementation makes them appear as a single disk.

The disk is divided into four partitions. The Red Hat Enterprise Linux operating system is on a 10 GB partition. The boot data is on a 248 MB partition. The swap space is on a 2 GB partition. The remaining space is on another partition that is encrypted by using LUKS.

Table 23: Bare metal disk layout

Mount point	Encryption status	Partition name	Partition size	Disk	Used for
/boot	None	/dev/sda1	248 MB	/dev/sda	Kernel boot data
SWAP	None	/dev/sda3	2 GB	/dev/sda	Swap space for paging
/	None	/dev/sda2	10 GB	/dev/sda	Operating system data
/opt	LUKS	/dev/mapper/cryptData	About 900 GB	/dev/sda	Installation files and creation of /opt directory.

Client for all plans

The Information Governance Catalog on Cloud client machine configuration is the same for all plans sizes. The client machine has two SAN disks that are 100 GB each. One disk is C: for the Microsoft Windows operating system. The second disk is F:, and it is an empty disk.

Information roadmap

This roadmap lists information resources that are available for users who are new to the IBM® InfoSphere® Information Governance Catalog products. These resources provide information about various subject areas.

- **Overview** Learn about the [features and benefits](#) of InfoSphere Information Governance Catalog.
- **Designing the catalog** A catalog is an authoritative dictionary of the assets that are used throughout the organization. One of the main benefits of a well-designed catalog is increased trust and confidence in organization information. Planning, designing, and publishing a catalog involves several tasks.
- **Administering the catalog** The catalog is composed of glossary assets (terms, categories, information governance policies, and information governance rules) and information assets. You can assign security roles and permissions to users to control access to the catalog. You can also assign workflow roles, assign users as stewards, define custom attributes, and define external asset types. In addition, you can configure the display of glossary assets.
- **Governing your data** You can create, edit, and delete catalog assets. Some assets are created when they are imported into the catalog. You can assign stewards and assets, define and assign custom

attributes, and extend data flows for lineage reports by importing assets. In addition, you can configure assets for lineage analysis reports. Assets of different types that have a common business purpose can be grouped in a collection.

- **Viewing catalog content** You can view, search, browse, and query the catalog to find catalog assets.
- **Expanding catalog capabilities**
- Look up terms while you work in other applications from the Microsoft Windows desktop by using Glossary Anywhere.
- Develop and extend applications by using InfoSphere Information Governance Catalog for Eclipse. You can access your glossary content from within your Eclipse-based development tool.
- Write client applications to access and author catalog content by using InfoSphere Information Governance Catalog REST API. You can integrate your catalog content into other software tools and portals.
- **Redbooks publication and Support**
- Information Governance Principles and Practices for a Big Data Landscape Published in 2014, this publication describes how the IBM Big Data Platform provides the integrated capabilities that are required for the adoption of Information Governance in the big data landscape. The IBM Big Data Platform, which is coupled with a framework for Information Governance, provides an approach to build, manage, and gain significant value from the big data landscape.
- IBM Support Portal Use the Support Portal for IBM InfoSphere Information Server to search for known problems and APARs.

New in Information Server 11.7

- Various new features introduced in Information Server 11.7 can be found [here](#)
- A key addition to Information Server 11.7 is [Information Server Enterprise Search](#)

Getting started and using IBM Information Governance Catalog on Cloud

You must set up your connection to Information Governance Catalog on Cloud. Information Governance Catalog on Cloud provides all of the functions of its on-premises counterpart, IBM® InfoSphere® Information Governance Catalog. It is deployed in an IBM SoftLayer® hosted environment.

Prerequisite: You must know the IP address and the credentials of an account on the Information Governance Catalog on Cloud server and client computers. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM Sales Representative.

Prerequisite: If your DataStage® jobs use the Hierarchical Data stage in parallel job designs, you must install Adobe Flash Player on the client computer. The version of Adobe Flash Player to use is listed in the Web Browser Plug-Ins table of the [System Requirements](#) page.

The Information Governance Catalog on Cloud client is on a Microsoft Windows virtual machine that is hosted on SoftLayer. When you connect to the client, you can access the client user interfaces. McAfee anti-virus software is installed on the client machine.

The Information Governance Catalog on Cloud server is on a Red Hat Enterprise Linux virtual or bare metal computer that is hosted on SoftLayer. When you connect to the server, you can access the IBM InfoSphere Information Server engine, services, and repository tiers. You can restart Information Governance Catalog on Cloud and do administrative tasks.

The Information Governance Catalog on Cloud license does not permit you to run DataStage jobs, but you can import those jobs by using DataStage on Cloud Designer Client. Information Governance Catalog on Cloud uses the job definitions to draw the lineage graph.

The default firewall configuration of server machines allows SSH connections only from client machines. You must first connect to a client machine by using a remote desktop connection, and then from the client machine you can connect to server machines by using SSH. After you log in to a server machine, you can change the firewall configurations to allow SSH connections from other machines. Communication

between the server and client systems happens through a private IP. If you want to access the server from an on-premises client machine, you must modify the iptable rules.

When you connect for the first time, follow these steps:

1. SSH into the any server machine using unique user (order id) provided in the welcome letter using port 4362.
2. Change the password and su to root user.
3. On the Information Governance Catalog on Cloud server, run the ISALite tool. The *IS_install_path* for the server computers is `/opt/IBM/InformationServer`
4. Connect to the Information Governance Catalog on Cloud client to the Information Governance Catalog on Cloud server by following these steps:
 - **If your local computer is in a Microsoft Windows environment**
 - a. On your local computer, go to the **Start** menu. Click **Accessories > Remote Desktop Connection**.
 - b. Enter the IP address of the Microsoft Windows computer that hosts your Information Governance Catalog on Cloud client. Click **Connect**.
 - c. In the Windows Security window, enter the user name and password for the Information Governance Catalog on Cloud client. The user ID, password, and IP address of the client are in your Welcome letter. **Important:** Do not include the domain name with the user name.
 - d. In the Information Governance Catalog on Cloud client, open the file `C:\Windows\System32\drivers\etc\hosts`. Make sure that an entry with the private IP exists in the file for the Information Governance Catalog on Cloud server that you are connecting to. **Important:** The server IP must be a private IP. You cannot open other client machines when you use a public IP.
 - **If your local computer is in an Apple Mac environment**
 - a. On your local computer, install Microsoft Remote Desktop from the Apple App Store.
 - b. Click the Microsoft Remote Desktop icon, and then click **Open**.
 - c. In the Microsoft Remote Desktop window, click **New**.
 - d. In the Edit Remote Desktops window, supply the following information:
 - In the PC name field, type in the IP address of the cloud client machine.
 - In the User name and Password fields, type in the Windows user name and password that are in the Welcome letter.
5. Verify the connection and installation on the Information Governance Catalog on Cloud client by following these steps:
 - Run the ISALite tool. The *IS_install_path* for the client computer is `C:\IBM\InformationServer`.
 - Test the installation of the Information Governance Catalog on Cloud client.
6. Optional: Enable multiple users to open remote sessions to the Information Governance Catalog on Cloud client by following these steps on the client computer:
 - a. Create user accounts.
 - b. Give users permission to do a remote desktop connection.
7. Reset the password for users and administrators on the Information Governance Catalog on Cloud server.

After the initial connection, you can do any of the following tasks:

- Open the InfoSphere DataStage® and QualityStage® Designer client
- Connect to an on-premises computer
- Perform general administration and security tasks

Related information

- [Getting started and using IBM DataStage on Cloud](#)
- [Enhancing security of Information Server on Cloud computers](#)

Chapter 7. Connecting to other systems

You can connect your IBM Information Server cloud offerings to on-premises systems and to cloud systems. Not all cloud offerings can connect to every system.

The following table lists a sample of which system can connect to Information Server cloud offerings. For information about how to connect, click the links in the **Connect to these systems** column.

Note that with network connection available, cloud offerings can connect to the same data sources as their on-premise equivalents. For details, see https://www.ibm.com/support/knowledgecenter/SSZJPZ_11.5.0/com.ibm.swg.im.iis.connect.nav.doc/containers/cont_iisinfoconnect.html

Table 24: Sample of systems that can connect to cloud offerings

Connect to these systems	IBM® Information Server on Cloud Enterprise Edition	IBM DataStage® on Cloud	IBM Information Governance Catalog on Cloud	IBM Information Server Data Quality on Cloud
IBM InfoSphere DataStage and QualityStage® Designer client	Yes	Yes	Yes	Yes
On-premises computer	Yes	Yes	Yes	Yes
IBM dashDB™ database	Yes	Yes	No	Yes
On-premises DB2® database instance	Yes	Yes	No	Yes

Connecting to the IBM InfoSphere DataStage and QualityStage Designer client

You connect to the InfoSphere® DataStage® and QualityStage® Designer client of IBM® DataStage on Cloud from IBM Information Server on Cloud offerings.

Prerequisite: You must know the IP address, host name, and credentials of the DataStage on Cloud server. This information is in the Welcome letter from your IBM Sales Representative.

You can connect from these cloud offerings:

- IBM Information Server on Cloud Enterprise Edition
- IBM DataStage on Cloud
- IBM Information Governance Catalog on Cloud
- IBM Information Server on Cloud Data Quality

1. On the Microsoft Windows machine where the DataStage on Cloud client is installed, click **Start > Programs > IBM InfoSphere DataStage and QualityStage Designer client**.

2. Enter the following information:

- **Host name of the services tier:** The host name or IP address of the services tier on the DataStage on Cloud server in the following format: <IP_hostname_services_tier>:<port_number>. For example, 169.XX.XX.XX:9446.
- **User name:** The user name of the IBM InfoSphere Information Server administrator user on the DataStage on Cloud server. For example, isadmin.

- **Password:** The password of the InfoSphere Information Server administrator user.
 - **Project:** The name of the default project is dstage1. For example, <host_name>/dstage1.
3. Click **Login**, and then accept and install the certificate that is generated. You can now use the InfoSphere DataStage and QualityStage Designer client on the cloud.

Connecting to an on-premises computer

You can connect to an on-premises computer from IBM® Information Server on Cloud cloud offerings.

Prerequisite: You must know the IP address and credentials of an administrator user of the on-premises computer.

Prerequisite: You must know the VPN client that you want to use in your DataStage® on Cloud server or client machines.

You can connect to an on-premises computer from these cloud offerings:

- IBM Information Server on Cloud Enterprise Edition
- IBM DataStage on Cloud
- IBM Information Governance Catalog on Cloud
- IBM Information Server on Cloud Data Quality

A VPN client must be installed on your DataStage on Cloud server or client machine only if your on-premise machine is behind a firewall. Then, after the VPN client is installed, you can access the on-premises system.

You can set up the VPN client to always run on the DataStage on Cloud server. Alternatively, you can run the VPN client only when you run a DataStage job that accesses computers behind a firewall.

Your on-premise system might be a server that is running in a UNIX or Microsoft Windows operating system environment. The VPN client enables a job that is running in the DataStage on Cloud to read from and write to an on-premise system.

1. If the on-premise machine that you need to connect to is behind a firewall, do these steps before you connect for the first time:
 - a. Log in to the DataStage on Cloud server machine that you want to connect from.
 - b. Install a VPN client that is appropriate for the operating system of the cloud machine. The DataStage on Cloud server is on a Linux operating system.
 - c. On your server machine, ping the on-premises computer. This step verifies that the connection is valid. After the connection is available, you can connect to the on-premise machines by using DataStage jobs to read and write data.
2. To run the VPN client only during job execution, design the job to open up the firewall at the beginning of the job execution and to close the firewall when the job execution is done. Use a command-line or API of your VPN client in an OSH script or in a stage to open and close the firewall.

Connecting to IBM dashDB

You can connect some IBM® Information Server cloud offerings to your IBM dashDB™ instance on the cloud.

Prerequisite: You must know the IP address or host name, port number, and credentials for a user to connect to IBM InfoSphere® DataStage® and QualityStage® Designer on the client machine.

Prerequisite: You must know the JDBC URL and credentials to connect to the dashDB instance.

You can connect to your IBM dashDB instance from these cloud offerings:

- IBM Information Server on Cloud Enterprise Edition

- IBM DataStage on Cloud
- IBM Information Server on Cloud Data Quality

Then, you can read data from dashDB or write to dashDB from other sources, such as a flat file or RDBMS stages.


You can use the JDBC Connector stage to connect to dashDB. For more information about how to use this stage, see [JDBC data sources](#).

For example, suppose that you have a dashDB service that runs on IBM Bluemix®. You can connect to that dashDB instance by using the JDBC Connector stage when you run a job on your DataStage on Cloud server. You use the credentials that are required to connect to the dashDB service.

1. On your DataStage on Cloud client machine, log in with the credentials of a user in the InfoSphere DataStage and QualityStage Designer client.
2. In the InfoSphere DataStage and QualityStage Designer client, design a job that includes a JDBC Connector stage by following these steps:
 - a. Click **File > New** from the menu.
 - b. In the New window, select the Parallel Job icon, and then click **OK**.
 - c. Add the JDBC Connector stage to the job by following these steps:
 - 1) In the palette, select the Database category.
 - 2) Locate JDBC Connector in the list of available stage types.
 - 3) Drag the JDBC Connector stage icon to the canvas.
 - 4) Rename the JDBC Connector stage. Choose a name that indicates the role of the stage in the job.
 - d. Create the necessary links and more stages for the job:
 - For a job that reads data from a JDBC data source, create the next stage in the job, and then create an output link from the JDBC Connector stage to the next stage.
 - For a job that writes data to a JDBC data source, create one or more links from other stages in the job to the JDBC Connector stage.
 - e. Save the job.
3. Get the JDBC URL, user name, and password of your dashDB instance by following these steps:
 - a. Go to your dashboard in Bluemix.
 - b. From the list of services, select dashDB. The connection parameters are available in JSON format under the topic "Service Credentials".
4. [Enter the connection parameters](#) of your dashDB instance in the JDBC connector stage.
5. Compile the job, and then run it.

Example

You design a job that generates data by using the Row Generator stage. This stage writes to your dashDB instance by using a JDBC connector. You provide the connection information to the JDBC Connector stage, compile the job, and then run the job.



Connecting on an on-premises DB2 database instance

You can connect some IBM® Information Server cloud offerings to your DB2® database instance on an on-premises computer.

Prerequisite: You must know the IP address or host name, port number, and credentials for a DB2 user to connect to the DB2 database on the on-premises computer.

Prerequisite: You must know the IP address or host name, port number, and credentials of the DataStage® on Cloud server to connect from the IBM InfoSphere® DataStage and QualityStage® Designer client.

Prerequisite: If the on-premises DB2 database is installed on a system that is behind a firewall, follow the steps to [connect to an on-premises computer](#).

You can connect to your on-premises DB2 database instance from these cloud offerings:

- IBM Information Server on Cloud Enterprise Edition
- IBM DataStage on Cloud
- IBM Information Server Data Quality on Cloud

Then, you can move data from your on-premises database to the cloud, and from the cloud to your on-premises database.

You can use the DB2 Connector stage from the DataStage on Cloud server to read data from the on-premises DB2 database. For more information about how to use this stage, see [DB2 connector](#).

You catalog the on-premises database to the DataStage on Cloud to connect them with each other.

1. Verify the connection between the on-premises computer and the DataStage on Cloud server by using the **ping** command.
2. On the DataStage on Cloud server, [configure the DB2 database connector](#). The DB2 connector requires that you catalog each on-premises database to DB2 on the DataStage on Cloud server.
3. On the DataStage on Cloud server, log in as a DB2 instance user.
4. Catalog the on-premises database on the DataStage on Cloud server by doing these commands:

```
db2 catalog tcpip node node_name remote remote_ip_address db2_port_number
db2 catalog database database_name at node node_name
db2 terminate
```

If you do the commands inside the DB2 prompt, do not include "db2" in the command.


5. At the DB2 prompt, confirm that the on-premises DB2 database is cataloged:

```
list database directory
```

Example

You create a job to transfer data from the on-premises DB2 database to a sequential file on your DataStage on Cloud server. You design a job by using the DB2 connector and Sequential File stages.

In DB2 Connector stage properties, you provide the database name, user name, and password details of your on-premises DB2 database.



Complete the job design with other details, compile the job, and then run the job.

Chapter 8. Administering cloud offerings on the server

You administer your IBM® Information Server on Cloud Enterprise Edition offerings on the server with the same utilities and tools that you use for your on-premises systems.

Information Server on Cloud Enterprise Edition offerings are based on IBM InfoSphere® Information Server. As a result, administrative tasks for your Information Server on Cloud server are managed by the same utilities and tools as for your on-premises systems. See [Administering](#) for detailed information about administering your system.

You can do the following administrative tasks:

- **Start services and application server services** For information, see [Starting services \(Linux, UNIX\)](#).
- **Shut down services and application server services** For more information, see [Shutting down services \(Linux, UNIX\)](#).
- **Change passwords** For more information about changing passwords, see [Changing the password](#) for the administrator and database accounts.
- **Encrypt DB2® database with native encryption** This encryption is done on all machines. For more information, see [DB2 native encryption](#).
- **Back up and restore the services tier, engine tier, and repository tier** For more information, see [Backing up IBM InfoSphere Information Server components](#) and [Performing a manual migration](#).

Open ports on server and client machines

At the operating system level, the firewall is implemented on Linux-based server machines by using iptables. On the Windows client machines, the firewall is implemented by using netsh. Although most ports are blocked, some ports are open by default on the server and client machines.

Open ports for incoming traffic on the server machine

Ports that are required for Information Server on Cloud to work are opened on private IP addresses. For details about these ports, see [Configuring your network](#).

For [Information Server DataStage® on Cloud High Availability](#) and [Information Server on Cloud Enterprise Edition High Availability](#) offerings, and an additional set of ports that are required for IBM® WebSphere® Application Server components on different server machines are opened on private IP addresses. For details about these ports, see [WebSphere Application Server ports](#).

Note the following changes in port configuration for Information Server on Cloud:

- The default ORB Listener port number is changed from 9100 to 9108
- ASB Agent's object port is fixed to 33311
- The default Information Server console port is changed from 9443 to 9446
- The default WebSphere® Application Server Bootstrap port is changed from 2809 to 2825

To connect to Information Server on Cloud from any client machine other than the one that is bundled with your part number, you must also open ports from that machine.

Table 25: Open ports on other Linux server machines in the offering

Category	Protocol	Port/Type	Reason that the port is open	From
HTTP+SSL	TCP	443	To ensure that HTTP communication works	All
SSH	TCP	Given in the Welcome letter	To allow SSH connections	Client Machines
Bootstrap	UDP	68	To receive messages from BOOTP configuration server at boot time	All
PING	ICMP	Echo-request	To allow ping	All

Open ports for incoming traffic on client machines

Table 26: Open ports on Information Server on Cloud clients

Port	Protocol	Profile	Reason that the port is open	Open from which machine
5986	TCP	Domain	Install the product by using Microsoft Windows Remote Management	All machines
3389	TCP	Domain	Connect to the client machine by using Microsoft Remote Desktop Connection	All machines
19443	TCP	Any	HTTPS port for service tier	All
Echo-request	ICMP	Any	Allow ping	All

(Optional) Open ports for incoming traffic on IBM Business Process Manager on Cloud

Table 27: Open ports if you decide to use IBM® Business Process Manager Standard

Component	Port	From which machine
BPM Process Centre	WAS_Secure_port: 9043. WC_defaulthost_secure: 9443. DCS_UNICAST_ADDRESS: 9354. SIB_ENDPOINT_SECURE_ADDRESS: 7286. SIP_DEFAULTHOST_SECURE: 5061. IPC_CONNECTOR_ADDRESS: 9633.	Client machines. Information Server Services tier machine

Table 27: Open ports if you decide to use IBM® Business Process Manager Standard (continued)

Component	Port	From which machine
BPM Process Server	WAS_Secure_port: 9045. WC_defaulthost_secure: 9444. DCS_UNICAST_ADDRESS: 9357. SIB_ENDPOINT_SECURE_ADDRESSES: 7287. SIP_DEFAULTHOST_SECURE: 5063. IPC_CONNECTOR_ADDRESS: 9635.	Client machines. Information Server Services tier machine
SSH	Given in the Welcome letter	Client machines
Bootstrap	68	All
PING	Echo-request	All

(Optional) Open ports for incoming traffic on IBM Cognos Business Intelligence on Cloud

Table 28: Open ports if you decide to use IBM Cognos® Business Intelligence

Component	Port	From which machine
IBM Cognos Business Intelligence	Dispatcher port: 9300. Log Server port: 9362	Client machines; Information Server Services tier machine
SSH	Given in the Welcome letter	Client machines
Bootstrap	68	All
PING	Echo-request	All

(Optional) Open ports on Information Server Services tier machine for incoming traffic from BPM and Cognos machine

Table 29: Open ports if you decide to use IBM Business Process Manager Standard and IBM Cognos Business Intelligence

Component	Port	From which machine
Information Server Services tier	IIS Console port: 9446. DB2 port: 50000	BPM. Cognos

Starting services after the Information Server on Cloud server machine reboots

When the Information Server on Cloud server machine restarts, services are started automatically. In some cases, you might need to start the services manually.

It takes several minutes for all services on the server machine to complete their startup. If you try to log in to any Information Server on Cloud client tools as soon as the server machine restarts, the services on the server machine might not yet be started. As a result, you cannot log in.

1. Wait a few minutes for all services to start, and then log in to either the IBM® InfoSphere® Information Server Web console.
2. If you still cannot log in, follow the steps to start the services manually.

Setting up firewall security

You can change the security level of your IBM® Information Server on Cloud server by configuring the iptables firewall. On the client, you can use the netsh command from Microsoft Windows to show the firewall profile and to list open ports.

Prerequisite: You must know the credentials of a user and the IP address of the Information Server on Cloud client and servers. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM® Sales Representative.

About these tasks

When the Information Server on Cloud server and client are delivered, the network access to the Linux server and the Microsoft Windows client is secured.

If you need to access ports of the Information Server on Cloud server from your on-premises application, you must modify the iptables rules.

On the Information Server on Cloud client, you can show the firewall profile, list the open ports, and close an open port.

To change the security level or to manage rules for the Information Server on Cloud server firewall

1. Log in to the root account of the Information Server on Cloud server.
2. Open a command-line window, and then type the command sudo.
3. Do any of following commands from the prompt:

List the rules in iptables

You can print the rules with the line numbers. This command is useful when you need to insert a rule at a specific location in iptables.

```
iptables -L --line-numbers
```

Restrict access to the Information Server on Cloud server from a range of IP addresses

You can restrict access to a subnet so that only those IP addresses can use ssh to access your Information Server on Cloud server.

```
iptables -A INPUT -i eth0 -p tcp -s <ip_subnet_address> --dport 22 -m state --state NEW,ESTABLISHED -j ACCEPT
```

where <ip_subnet_address> is the range of IP addresses in the subnet. For example, the following rule allows incoming ssh connections from only the 192.xxx.xxx.xxx subnet:

```
iptables -A INPUT -i eth0 -p tcp -s 192.xxx.xxx.xxx/24 --dport 22 -m state --state NEW,ESTABLISHED -j ACCEPT
```

Restrict access to the Information Server on Cloud server from a single IP address

This restriction is the most stringent.

```
iptables -A INPUT -i eth0 -p tcp -s <ip_address> --dport 22 -m state --state NEW,ESTABLISHED -j ACCEPT
```

where <ip_address> is the single IP address that is permitted to access the server. For example, the following rule allows incoming ssh connections from a single port on 192.xxx.xxx.xxx:

```
iptables -A INPUT -i eth0 -p tcp -s 192.xxx.xxx.xxx --dport 22 -m state --state NEW,ESTABLISHED -j ACCEPT
```

Open a port from another machine

To connect to the Information Server on Cloud server from any other machine, you must open the IBM InfoSphere® Information Server ports from that machine. For a list of ports that are open by default for incoming traffic on the server machine, see the topic [Open ports on server and client machines](#).

```
iptables -A INPUT -p tcp -s <ip_address> --dport 9446 -j ACCEPT
```

where <ip_address> is the IP address of the machine for which you want to open port 9446 on the Information Server on Cloud server machine. For example, the following rule opens port 9446 on the server for incoming traffic from port 9446 on 192 . xxx . xxx . xxx:

```
iptables -A INPUT -p tcp -s 192.xxx.xxx.xxx --dport 9446 -j ACCEPT
```

Open the port for a fix pack or patch installation

Fix packs and patches can be installed to the existing Information Server on Cloud server environment. By default, port 8446 is used for the installation. To do the installation in GUI mode from a remote machine, 8446 port must be opened.

```
iptables -A INPUT -p tcp -s <ip_address> --dport 8446 -j ACCEPT
```

where <ip_address> is the IP address for which you want to open port 8446 on the Information Server on Cloud server machine. For example, the following rule opens port 8446 on the server for incoming traffic from port 8446 on 192 . xxx . xxx . xxx:

```
iptables -A INPUT -p tcp -s 192.xxx.xxx.xxx --dport 8446 -j ACCEPT
```

If you use a port other than 8446 for installation, use the other port in the iptables command.

Add a rule at a specific line in the iptables

The order of the rules is important. Rules are appended to the end of existing rules. For example, if the rule that you added follows a DROP rule that matches incoming traffic, the DROP rule is considered before the newly added rule. You can add a rule at a specific line in the iptables to push subsequent rules down in the list.

```
iptables -I INPUT line_number -p tcp -s <ip_address> --dport 22 -j ACCEPT
```

For example, the following rule adds a rule at line 6 to accept incoming traffic on port 22 from 192 . xxx . xxx . xxx:

```
iptables -I INPUT 6 -p tcp -s 192.xxx.xxx.xxx --dport 22 -j ACCEPT
```

Delete all rules in iptables

Important: Before you remove all rules, be sure that the default INPUT policy is ACCEPT. Otherwise, you cannot connect to the Information Server on Cloud server from anywhere.</div>

```
iptables -F
```

Save rules in iptables permanently

To persist the modifications to the iptables rules after the system restarts, you must update the file /etc/ sysconfig/iptables.

```
iptables-save > /etc/sysconfig/iptables
```

To show the Microsoft Windows firewall profiles on the Information Server on Cloud client

1. Access the Information Server on Cloud client by using Microsoft Remote Desktop Connection.

2. Open a command-line window on the client.
3. Run the command `netsh advfirewall show allprofiles`.

To show a list of all open ports on the Information Server on Cloud client

1. Access the Information Server on Cloud client by using Microsoft Remote Desktop Connection.
2. On the Information Server on Cloud client, click **Start > Administrative Tools > Server Manager**.
3. In the upper-right corner of the Server Manager window, click **Tools**. From the list, select **Windows Firewall with Advanced Security**.
4. In the left pane, click **Inbound Rules** to see which ports are open for incoming network traffic.

To block an open port on the Information Server on Cloud client

Do the previous task. Right-click the port number, and then select the appropriate action.

Scripts to change the iptables firewall settings

You can change the settings of the iptables firewall to suit the needs of your organization.

Use scripts to help manage the iptables firewall settings on the cloud server. The following scripts are located in directory `/opt/IBM/scripts/` on the server. To use these scripts, you must be logged in to an administrative account on the server.

Refer to [Adding and connecting extra DataStage on Cloud Designer Client machines](#) to modify the iptable rules to access the Information Server on Cloud Enterprise Edition server and Information server on Cloud Data Quality server from an IBM® InfoSphere® Information Server client other than the Information Server on Cloud client that is provided as part of the offering.

allow_datastage_from_ip.sh

Use this script to access the DataStage® on Cloud server from an IBM® InfoSphere® DataStage client other than the Information Server on Cloud client that is provided as part of the offering.

The command line parameter `IP_address` is the IP address of the computer that you want to allow access to the DataStage on Cloud server.

For example, the following command enables `192.xxx.xxx.xxx` to access the DataStage on Cloud server.

```
allow_datastage_from_ip.sh 192.xxx.xxx.xxx
```

allow_ping_from_ip.sh

Use this script to be able to ping the Information Server on Cloud server from a specific computer.

The command line parameter `IP_address` is the IP address of the computer that you want to allow to ping the Information Server on Cloud server.

For example, the following command enables `192.xxx.xxx.xxx` to ping the server.

```
allow_ping_from_ip.sh 192.xxx.xxx.xxx
```

allow_port_from_ip.sh

Use this script to allow access to a specific port on the Information Server on Cloud server from a specific computer.

The first command line parameter `IP_address` is the IP address of the specific computer. The second parameter `Port_number` is the port on the server to allow the connection from `IP_address`.

For example, the following command 8446 enables 192 . xxx . xxx . xxx to connect to the server at port 8446.

```
allow_port_from_ip.sh 192.xxx.xxx.xxx
```

delete_rule_with_linenumber.sh

Use this script to delete a rule at a specific line of the iptables firewall on the Information Server on Cloud server.

The command line parameter *line_number_of_rule* is the line number of the rule that you want to delete.

Tip: To print the line numbers of iptables rules, use the command: `iptables -L --line-numbers`.

For example, the following command deletes the rule at line 20.

```
delete_rule_with_linenumber.sh 20
```

save_iptables_changes_redhat.sh

Use this script to persist the changes to iptables rules that you made on the Information Server on Cloud server. If the changes are not persisted, the changes might be lost when you restart the server.

allow_services_tier_from_bpm_cognos_ip.sh

Use this script to access the Information Server on Cloud server services tier for Information Server on Cloud Data Quality and Information Server on Cloud Enterprise Edition from the machine where IBM Business Process Manager and IBM Cognos® Business Intelligence are installed.

The command line parameter *IP_address* is the IP address of the computer that you want to allow access to the Information Server on Cloud services tier machine.

For example, the following command enables 192 . xxx . xxx . xxx to access the Information Server on Cloud services tier server.

```
allow_services_tier_from_bpm_cognos_ip.sh 192.xxx.xxx.xxx
```

Related information:

- [Getting started and using DataStage on Cloud](#)
- [Getting started and using IBM Information Governance Catalog on Cloud](#)

Enhancing security of Information Server on Cloud computers

You can improve the security of your Information Server on Cloud computers by using SSH key pairs, changing the default SSH port, and restricting SSH access to specific ports.

Related tasks

- [Getting started and using IBM Information Server on Cloud Enterprise Edition](#)
- [Getting started and using IBM Information Server on Cloud Data Quality](#)
- [Getting started and using IBM DataStage on Cloud Designer Client](#)
- [Getting started and using IBM Information Governance Catalog on Cloud](#)

Setting up SSH keys

You can use SSH keys to restrict access to your Information Server on Cloud server from specific computers.

Prerequisite: You must be logged in to the Information Server on Cloud server in the root account.

You can restrict access to the Information Server on Cloud server by using SSH keys. You create the key pairs on the local computer or virtual machine from which to connect to the Information Server on Cloud server, and then copy the keys to the server.

Next, you disable SSH login password authentication. Because this step disables connections to the cloud server that use login and password authentication, you cannot connect to the Information Server on Cloud server if you cannot connect to your local computer that generated the SSH keys.

1. On your local computer, type the command **ssh-keygen** to create SSH key pairs. Store the key pair in the default location. By default, private keys are stored in your home directory of `/home/username/.ssh/id_rsa`. Public keys are stored in `id_rsa.pub`.
2. Copy the public SSH key to the Information Server on Cloud server that you want to connect to from your local computer. In the Information Server on Cloud server, append the public key to the file `$HOME/.ssh/authorized_keys`. The directory `$HOME` might be `/home/<user_name>`, where `<user_name>` is your login name. For the root user, `$HOME` might be `/root`.
3. Disable SSH logon by following these steps:
 - a. At the command-line prompt on your local computer, open the file `/etc/ssh/sshd_config`.
 - b. Add the line `PasswordAuthentication no`. If this line exists but is commented out, uncomment the line.
 - c. Run the command `service sshd restart` to restart the SSHD service. **Important:** After you restart the SSHD service, you cannot connect to the Information Server on Cloud server by using login and password. To avoid being locked out of the Information Server on Cloud server, test your connection to the server by using SSH keys before you restart the SSHD service.

Allowing SSH access to specific IP addresses

You can allow SSH access to the IBM® Information Server on Cloud server from specific IP addresses to prevent intrusion and hacking.

Prerequisite: You must be logged in to the Information Server on Cloud server in the root account.

Prerequisite: You must know the port that is used for SSH. The port number is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM Sales Representative. It is also in the file `/etc/ssh/sshd_config` on the Information Server on Cloud server.

The SSH port of the Information Server on Cloud server is open only to client machines that are provisioned with the Information Server on Cloud Enterprise Edition offering. This procedure opens the Information Server on Cloud server to other machines.

Run the **allow_port_from_ip.sh** script. The first argument is the IP of the computer you want to give access to. The second argument is the SSH port number.

For example, the following command allows port 1234 from the IP address 192.xx.xx.xx to access the Information Server on Cloud server.

```
allow_port_from_ip.sh 192.xx.xx.xx 1234
```

Managing LUKS keys on the IBM Information Server on Cloud server

Any partition that is encrypted by using Linux Unified Key Setup (LUKS) can have eight different keys. You can use any of the keys to open the encrypted partition. In addition, you can add new keys or remove existing ones according to your needs.

Prerequisite: You must know the credentials of the root user and the IP address of the Information Server on Cloud server. This information is in the "Access Credentials and Initial Setup" section of the Welcome letter that you received from your IBM® Sales Representative.

One LUKS key is specified in the "Encryption Details" section of the Welcome letter that you received from your IBM Sales Representative. This key file is used from `/etc/crypttab` file to unlock the partition

when the cloud server reboots. For details about encrypted partitions on the Information Server on Cloud server, see the following topics:

- [Layout of Information Server Enterprise Edition on Cloud server and client disks](#)
- [Layout of Information Server Data Quality on Cloud server and client disks](#)
- [Layout of DataStage on Cloud server and client disks](#)
- [Layout of DataStage Designer Client on Cloud server and client disks](#)
- [Layout of Information Governance Catalog on Cloud server and client disks](#)

1. On the Information Server on Cloud server, open a command-line window.
2. Do any of the following tasks to manage your keys. The commands can be run from any directory.

Add new LUKS key

```
cryptsetup luksAddKey <partition_name>
```

The parameter <partition_name> is the partition that is encrypted.

For example, the partition for virtual servers is /dev/xvdc1. For bare metal servers, the partition is /dev/sda5.

Type in the existing key that is given in the Welcome letter. When prompted, enter the name of the new key.

Remove existing LUKS keys

```
cryptsetup luksRemoveKey <partition_name>
```

When prompted, enter the name of the specific key that you want to remove.

Add a key from a file

```
cryptsetup luksRemoveKey <partition_name> <keyfile_name>
```

The parameter <keyfile_name> must include the full file path of the file.

The command asks for an existing key and then reads the new key from the file.

Chapter 9. Troubleshooting cloud offerings

These topics contain troubleshooting information for Information Server on Cloud.

Cannot install a patch on services tier machine

When you try to install a patch for Information Server on Cloud Enterprise Edition and Information Server on Cloud Data Quality on services tier machine, you are unable to open the installation wizard.

Symptoms

During the installation, you receive a URL, which you can open on a Windows machine and run the installation by using the wizard (graphical mode). When you open the URL, you are unable to proceed.

The issue occurs on the Internet Explorer browser.

Resolving the problem

To resolve the issue, complete the following steps:

1. In the URL, replace the host name with the private IP address of service tier.
2. Configure the Compatibility View settings of Internet Explorer browser:
 - a. Click **Tools > Compatibility View settings**.
 - b. Specify the URL with private IP address in the Add to this website field and then click Add.
 - c. Click **Save**.

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 30: 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 30: 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, LACS^{Link}, ZIP, ZIP + 4, ZIP Code, Post Office, Postal Service, USPS and United States Postal Service. IBM Corporation is a non-exclusive DPV and LACS^{Link} licensee of the United States Postal Service.

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

