

Ready-to-use Virtual Appliance for Hands-on IBM Spectrum Archive Evaluation

Hiroyuki Miyoshi

Hiroshi Araki

Takeshi Ishimoto





Introduction

IBM® Spectrum Archive Enterprise Edition for the IBM TS4500, IBM TS3500, IBM TS4300, and IBM TS3310 tape libraries provides seamless integration of IBM Linear Tape File System (LTFS) with IBM Spectrum® Scale by creating an LTFS tape tier. You can run any application that is designed for disk files on tape by using IBM Spectrum Archive. IBM Spectrum Archive can play an important role in reducing the cost of storage for data that does not need the access performance of primary disk.

For more information about IBM Spectrum Archive, see the product information at:

<https://www.ibm.com/us-en/marketplace/data-archive>

The IBM Spectrum Archive Virtual Appliance can be deployed in minutes and key features can be tried along with this user guide. The virtual machine (VM) has a pre-configured IBM Spectrum Scale and a virtual tape library that allows to quickly test the IBM Spectrum Archive features without connecting to a physical tape library. The virtual appliance is provided as a VirtualBox .ova file.

In this IBM Redpaper publication, we show you how to set up a virtual appliance and include typical use cases with instructions. This publication includes the following topics:

- ▶ How to migrate files from disk to tape manually and automatically.
- ▶ How to recall files from tape to disk.
- ▶ Briefly introduces tape management commands, such as reconcile, reclaim, export, and import.
- ▶ How REST API is supported in IBM Spectrum Archive.

This version of the trial virtual appliance includes the following software:

- ▶ IBM Spectrum Archive Enterprise Edition 1.3
- ▶ IBM Spectrum Scale 5.0.0

Note: The web page location to download the trial ova file was updated August 2020 in Table 1 on page 3.

Note: Consider the following points:

- ▶ Because this virtual appliance uses the emulated tape library, some of the command output can be slightly different from the physical tape library.
- ▶ The capacity of each virtual tape cartridge is 10 GB, which is significantly smaller than the physical tape cartridges.
- ▶ The virtual appliance is not an appropriate environment for performance testing.
- ▶ IBM Spectrum Scale protocol related service (NFS/CIFS/Object) is currently disabled by default.

Requirements

The requirements to run the IBM Spectrum Archive Virtual Appliance are listed in Table 1.

Table 1 IBM Spectrum Archive Virtual Appliance requirements

VM host platform	<ul style="list-style-type: none">▶ Windows 10 Version 1903▶ Mac OS 10.14.5▶ RHEL 7.4 Note: Supports 64-bit only
Software	Tested with VirtualBox 6.0.12
CPU	1 Core
Virtual Appliance OVA file temporary download space	3 GB
Virtual Appliance OVA file	<ul style="list-style-type: none">▶ SpectrumArchive_1.3.0.5_TrialVM.ova▶ Note: Download the Virtual Appliance OVA file from: https://www.ibm.com/products/data-archive▶ Select Free 90-day trial and the web page will guide you through the registration of an IBM user ID to download the OVA file.
Minimum system memory allocated to the VM	4 GB
System disk capacity used by VM	7 GB - 100 GB Note: 100 GB is used when all disk space (including IBM GPFS and ten 10 GB tapes) are full. For test cases that are described in this document, 10 GB should be enough.
Machine virtualization acceleration setup	Must be set to ENABLED.
IP addresses	127.0.0.1 Note: You can connect the VM from the host via ssh -p 2022 root@localhost.

Installation

Installation of IBM Spectrum Archive Virtual Appliance consists of the following steps:

- ▶ Configure VirtualBox
- ▶ Import the OVA file
- ▶ Correct the network adapter

Configure VirtualBox

Complete the following steps to install the IBM Spectrum Archive Virtual Appliance:

1. Start VirtualBox.
2. Click **File** → **Host Network Manager**.
3. Ensure that there is a Host-only Adapter with IPv4 Address of 192.168.56.1 and IPv4 Network Mask of 255.255.255.0. DHCP Server must be enabled. If such a network does not exist, create a new Host-only Adapter, as shown in Figure 1.

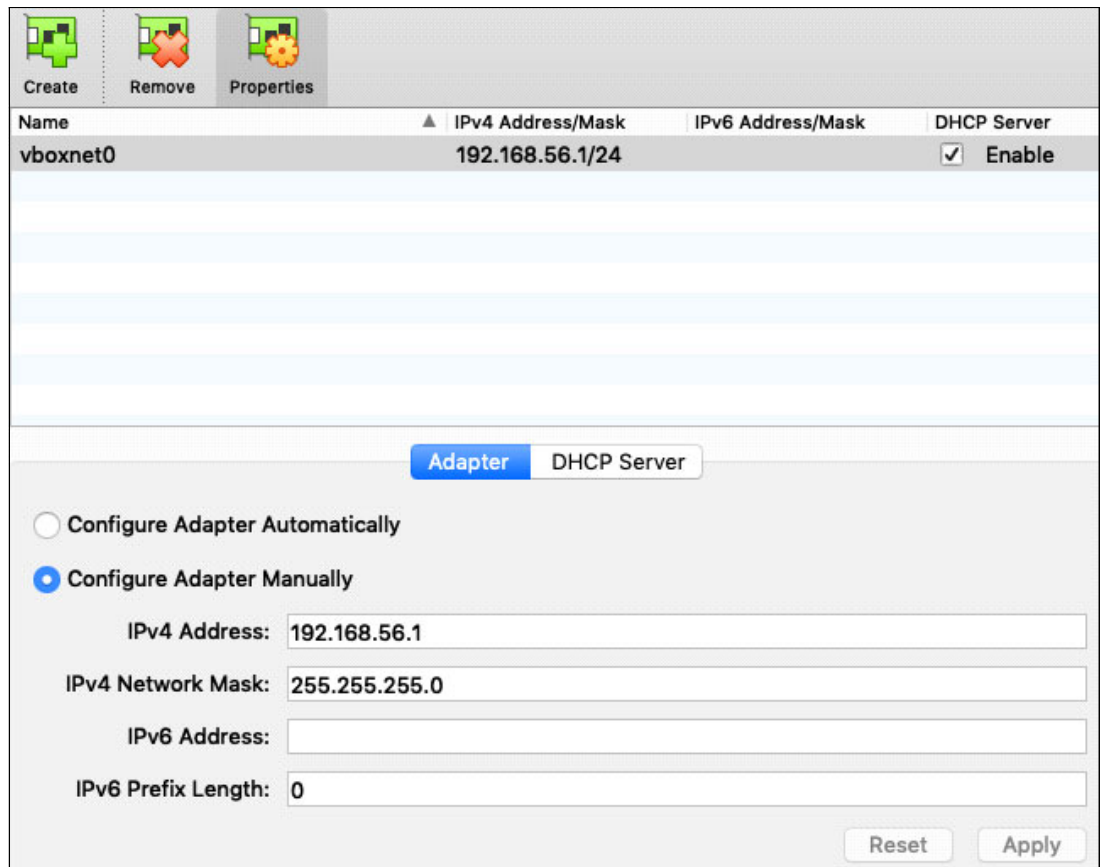


Figure 1 Host-only adapter settings

Remember the adapter name for the 192.168.56.1 network.

Import .OVA file

1. Create your own instance of the virtual appliance by clicking **File** → **Import Appliance**.
2. Select the SpectrumArchive_1.3.0.5_TrialVM.ova file.

Note: During the import process, select the **Include only NAT network adapter MAC address** for **MAC Address Policy** option.

3. During import, click **Agree** for the license statement, as shown in Figure 2.

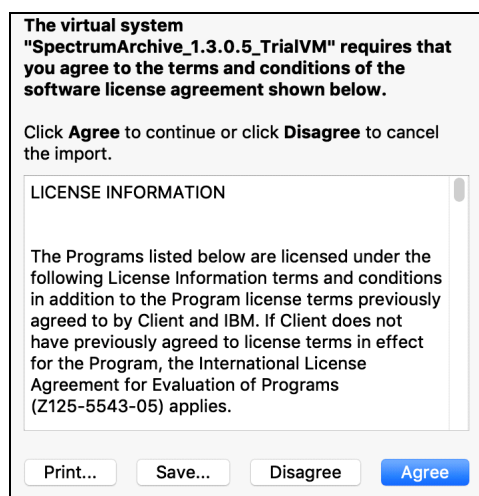


Figure 2 License statement

Correct network adapter

1. After the import process is complete, Select the Network settings of the virtual appliance. Ensure that adapter 2 is attached to the Host-only Adapter as shown in Figure 3 (adapter 2 was configured in step 3 of “Configure VirtualBox” on page 4).

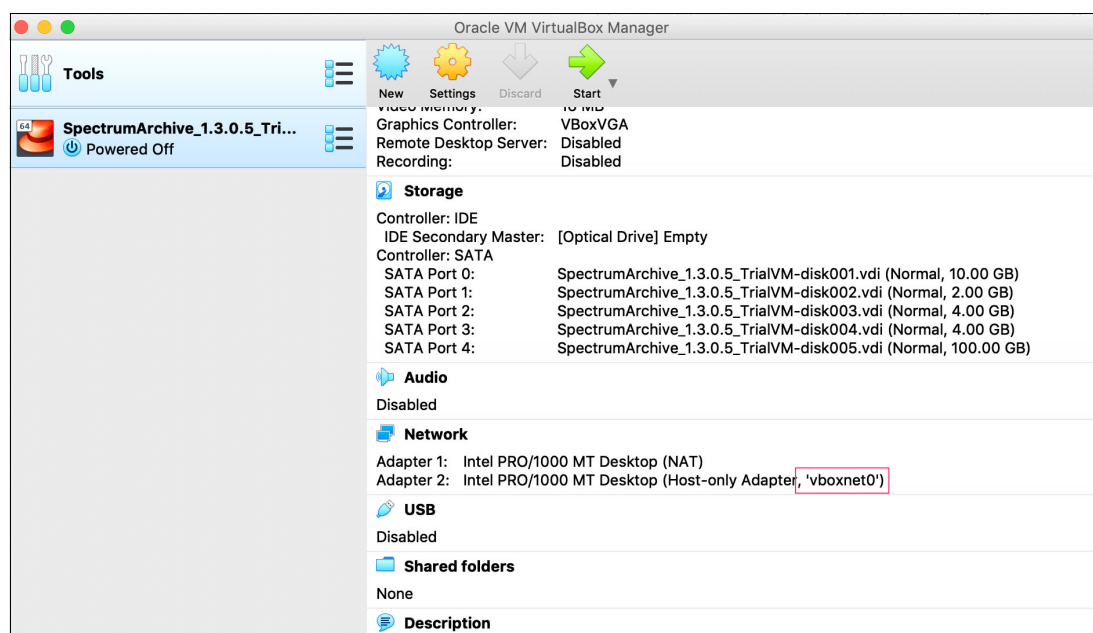


Figure 3 Network name attached to the Adapter 2

2. If the Host-only Adapter does not match, click **Settings** icon from the VirtualBox window. Click **Network** in the **Settings** dialog, select **Adapter 2**, and change the Name field of the adapter to the one configured in step 2 in “Configure VirtualBox” on page 4. Press **OK** to apply the change.
3. Start the virtual machine by clicking the **Start** icon.

Setting up IBM Spectrum Archive Virtual Appliance

In this section, we describe how to set up IBM Spectrum Archive Virtual Appliance after the installation is complete.

When the virtual machine boots, login by using the following credentials:

- ▶ User: root
- ▶ Password: ibm

The VM can also be connected from the host via `ssh -p 2022 root@localhost`.

IBM Spectrum Archive then can be configured by completing the initial setup process that is described next. This set up must be completed whenever the virtual machine starts.

Complete the following steps:

1. Ensure that the VM is a member of an IBM Spectrum Scale cluster, as shown in Example 1.

Example 1 Ensuring VM is member of an IBM Spectrum Scale Cluster

```
[root@spectrumscale ~]# mmfsccluster
```

GPFS cluster information

=====

```
GPFS cluster name:      spectrumscale.gpfscluster
GPFS cluster id:        10588301405112176324
GPFS UID domain:        spectrumscale.gpfscluster
Remote shell command:   /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:        CCR
```

Node	Daemon node name	IP address	Admin node name	Designation
1	spectrumscale	10.0.2.15	spectrumscale	quorum-manager-perfmon

2. Ensure that the IBM Spectrum Scale service is active and `/ibm/fs1` is mounted, as shown in Example 2. It may take approximately five minutes for `fs1` to be mounted.

Example 2 IBM Spectrum Scale service started

```
[root@spectrumscale ~]# mmgetstate
```

Node number	Node name	GPFS state
1	spectrumscale	active

```
[root@spectrumscale ~]# df /ibm/fs1
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
fs1	10485760	1647872	8837888	16%	/ibm/fs1

3. Start IBM Spectrum Archive by issuing the command shown in Example 3.

Example 3 Configuring and starting IBM Spectrum Archive

```
[root@spectrumscale ~]# eeadm cluster start
```

Library name: lib1, library serial: 0000VLIBRARY_LL0, control node (ltfsee_md) IP address: 10.0.2.15.

Starting - sending a startup request to lib1.

Starting - waiting for startup completion : lib1.
 Starting - opening a communication channel : lib1.
 .
 Starting - waiting for getting ready to operate : lib1.

 Started the IBM Spectrum Archive EE services for library lib1 with good status

4. Verify the status of IBM Spectrum Archive by issuing the commands that are shown in Example 4.

Example 4 Verifying status of IBM Spectrum Archive

```
[root@spectrumscale ~]# df /lufs
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
lufs:/filebackend/0000VLIBRARY_LL0	2199023255040	0	2199023255040	0%	/lufs

```
[root@spectrumscale ~]# eeadm node list
```

Node ID	State	Node IP	Drives	Ctrl Node	Library	Node Group	Host Name
1	available	10.0.2.15	4	yes(active)	lib1	G0	spectrumscale

```
[root@spectrumscale ~]# eeadm library list
```

Library Name	Model	Serial Number	Ctrl Node
lib1	03584L22	0000VLIBRARY_LL0	spectrumscale

```
[root@spectrumscale ~]# eeadm drive list
```

Drive S/N	State	Type	Role	Library	Node ID	Tape	Node Group	Task ID
VDRIVE0000	not_mounted	LT05	mrg	lib1	1	-	G0	-
VDRIVE0001	not_mounted	LT05	mrg	lib1	1	-	G0	-
VDRIVE0002	not_mounted	LT05	mrg	lib1	1	-	G0	-
VDRIVE0003	not_mounted	LT05	mrg	lib1	1	-	G0	-

```
[root@spectrumscale ~]# eeadm tape list
```

Tape ID	Status	State	Usable(GiB)	Used(GiB)	Available(GiB)	Reclaimable%	Pool	Library	Location	Task
VTAP00L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP01L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP02L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP03L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP04L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP05L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP06L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP07L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP08L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-
VTAP09L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot	-

A virtual tape library (lib1) with four virtual tape drives (VDRIVE0000 - VDRIVE0003) and 10 virtual tape cartridges (10 GB capacity for each) VTAP00L5 - VTAP09L5 are attached. The drives are pre-configured.

5. Create a tape pool for data migration and assign tapes to the pool, as shown in Example 5. This process is required only once when the virtual appliance is first created and set up.

Example 5 Creating a tape pool for data migration

```
[root@spectrumscale ~]# eeadm pool create pool1
```

2019-06-27 03:51:20 GLESL782I: Successfully created pool pool1.

```
[root@spectrumscale ~]# eeadm pool list
```

Pool Name	Usable(TiB)	Used(TiB)	Available(TiB)	Reclaimable%	Tapes	Type	Library	Node Group
pool1	0.0	0.0	0.0	0%	0	-	lib1	G0

```
[root@spectrumscale ~]# eeadm tape assign VTAP00L5 VTAP01L5 VTAP02L5 -p pool1 -f
2019-06-27 03:51:30 GLESL700I: Task tape_assign was created successfully, task id is 1000.
2019-06-27 03:51:31 GLESL087I: Tape VTAP01L5 successfully formatted.
2019-06-27 03:51:31 GLESL360I: Assigned tape VTAP01L5 to pool pool1 successfully.
2019-06-27 03:51:31 GLESL087I: Tape VTAP00L5 successfully formatted.
2019-06-27 03:51:31 GLESL360I: Assigned tape VTAP00L5 to pool pool1 successfully.
2019-06-27 03:51:32 GLESL087I: Tape VTAP02L5 successfully formatted.
2019-06-27 03:51:32 GLESL360I: Assigned tape VTAP02L5 to pool pool1 successfully.
```

```
[root@spectrumscale ~]# eeadm pool list
```

Pool Name	Usable(TiB)	Used(TiB)	Available(TiB)	Reclaimable%	Tapes	Type	Library	Node Group
pool1	0.0	0.0	0.0	0%	3	LTO	lib1	G0

```
[root@spectrumscale ~]# eeadm tape list
```

Tape ID	Status	State	Usable(GiB)	Used(GiB)	Available(GiB)	Reclaimable%	Pool	Library	Location
Task ID									
VTAP00L5	ok	appendable	9	0	9	0%	pool1	lib1	homeslot -
VTAP01L5	ok	appendable	9	0	9	0%	pool1	lib1	homeslot -
VTAP02L5	ok	appendable	9	0	9	0%	pool1	lib1	homeslot -
VTAP03L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -
VTAP04L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -
VTAP05L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -
VTAP06L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -
VTAP07L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -
VTAP08L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -
VTAP09L5	ok	unassigned	0	0	0	0%	-	lib1	homeslot -

Note: If a tape status is not OK and the state indicates errors such as `check_tape_library`, it is most likely a temporary error caused by the tape library emulation. In such a case, the following workaround might fix the tape status:

1. `eeadm tape unassign <tape IDs> -p <pool ID>`
2. `eeadm library rescan`
3. `eeadm tape assign -<tape IDs> -p <pool ID> -f`
4. `eeadm cluster stop`
5. `eeadm cluster start`

IBM Spectrum Archive features

The basic operations for IBM Spectrum Archive are described in this section.

Manual migration

Manual migration includes the following objectives:

- ▶ Check the file status to determine whether the file data is on disk or tape
- ▶ Understand the migration policy rule file
- ▶ Manually run the migration policy and verify that the file data was migrated

Complete the following steps:

1. Create test files under the `/ibm/fs1/archive` test directory, as shown in Example 6.

Example 6 Creating test files in /ibm/fs1/archive directory

```
[root@spectrumscale ~]# mkdir /ibm/fs1/archive
[root@spectrumscale ~]# cd /ibm/fs1/archive/
[root@spectrumscale archive]# dd if=/dev/urandom of=10Mfile_1 count=1 bs=10M
```

```
[root@spectrumscale archive]# for i in $(seq 2 3); do cp 10Mfile_1 10Mfile_$i;
done
[root@spectrumscale archive]# yes 'This is test' | dd of=text bs=2M count=1
[root@spectrumscale archive]# ls -l
total 28672
-rw-r--r--. 1 root root 10485760 Jun 27 03:52 10Mfile_1
-rw-r--r--. 1 root root 10485760 Jun 27 03:52 10Mfile_2
-rw-r--r--. 1 root root 10485760 Jun 27 03:52 10Mfile_3
-rw-r--r--. 1 root root 2097152 Jun 27 03:52 text
```

2. To determine whether file data is on disk or tape, run the **dsm1s** command or **eeadm file state** command.

Example 7 shows the **dsm1s** command output. The “r” indicates that the files are resident (on disk).

Example 7 Confirming whether file data is on disk or tape

```
[root@spectrumscale archive]# dsm1s
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 7.0
  Client date/time: 06/27/2019 03:53:04
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.
```

ActS	ResS	ResB	FSt	Prvd	FName
/ibm/fs1/archive:					
10485760	10485760	10240	r	-	10Mfile_1
10485760	10485760	10240	r	-	10Mfile_2
10485760	10485760	10240	r	-	10Mfile_3
2097152	2097152	2048	r	-	text

Example 8 shows the **eeadm file state** command output. The command also has a shorter output version. All files are resident. The file data is on disk only.

Example 8 File data is on disk only

```
[root@spectrumscale archive]# eeadm file state *
Name: /ibm/fs1/archive/10Mfile_1
State: resident

Name: /ibm/fs1/archive/10Mfile_2
State: resident

Name: /ibm/fs1/archive/10Mfile_3
State: resident

Name: /ibm/fs1/archive/text
State: resident

[root@spectrumscale archive]# eeadm file state * -s
R 0 - - - -- /ibm/fs1/archive/10Mfile_1
R 0 - - - -- /ibm/fs1/archive/10Mfile_2
R 0 - - - -- /ibm/fs1/archive/10Mfile_3
R 0 - - - -- /ibm/fs1/archive/text
```

The following file statuses are also available:

- Premigrated (p): Data is on disk and tape
- Migrated (m): Data is on tape only

3. Define a policy rule file to migrate these test files to tapes. The policy file is found at:

/root/policy/policy_migArchiveDir

The contents of the policy file are shown in Example 9.

Example 9 Migrate policy contents

```
define(  
    exclude_list,  
    PATH_NAME LIKE '/ibm/fs1/.ltfsee/%'  
    OR PATH_NAME LIKE '/ibm/fs1/.SpaceMan/%'  
    OR PATH_NAME LIKE '/ibm/fs1/.snapshots/%'  
    OR ((CURRENT_TIMESTAMP - MODIFICATION_TIME) < INTERVAL '121' SECONDS)  
    OR (FILE_SIZE < 1024)  
)  
  
define(  
    is_migrated,  
    MISC_ATTRIBUTES LIKE '%V%'  
)  
  
RULE EXTERNAL POOL 'Archive_files'  
    EXEC '/opt/ibm/ltfsee/bin/eedm'  
    OPTS '-p pool1@lib1'  
  
RULE 'Mig-test' MIGRATE FROM POOL 'system'  
    TO POOL 'Archive_files'  
    WHERE (  
        PATH_NAME LIKE '/ibm/fs1/archive/%'  
        AND NOT (exclude_list)  
        AND NOT (is_migrated)  
    )
```

When this policy is run, it migrates all resident and pre-migrated files that are in the /ibm/fs1/archive directory from disk to tape.

Note: For best practice, the following files are excluded from migration:

- ▶ Files that are less than 1 MB in size
- ▶ Files that have been modified within the last two minutes

4. Wait two minutes after the test file creation. Run the `mmapplypolicy` command to run the policy, as shown in Example 10.

Example 10 Running mmapplypolicy command

```
[root@spectrumscale archive]# mmapplypolicy /ibm/fs1/archive -P  
/root/policy/policy_migArchiveDir  
[I] GPFS Current Data Pool Utilization in KB and %  
Pool_Name          KB_Occupied      KB_Total  Percent_Occupied  
system             3952640        4194304   94.238281250%  
[I] 59017 of 107520 inodes used: 54.889323%.
```

[I] Loaded policy rules from /root/policy/policy_migArchiveDir.
Evaluating policy rules with CURRENT_TIMESTAMP = 2019-06-27@08:04:52 UTC
Parsed 2 policy rules.

```
RULE EXTERNAL POOL 'Archive_files'
  EXEC '/opt/ibm/ltfsee/bin/eedm'
  OPTS '-p pool1@lib1'
```

```
RULE 'Mig-test' MIGRATE FROM POOL 'system'
  TO POOL 'Archive_files'
  WHERE (
    PATH_NAME LIKE '/ibm/fs1/archive/%'
    AND NOT (PATH_NAME LIKE '/ibm/fs1/.ltfsee/%'
    OR PATH_NAME LIKE '/ibm/fs1/SpaceMan/%'
    OR PATH_NAME LIKE '/ibm/fs1/snapshots/%'
    OR ((CURRENT_TIMESTAMP - MODIFICATION_TIME) < INTERVAL '121' SECONDS)
    OR (FILE_SIZE < 1024)
  )
  AND NOT (MISC_ATTRIBUTES LIKE '%V%')
)
```

[I] 2019-06-27@08:04:52.493 Directory entries scanned: 5.

[I] Directories scan: 4 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.

[I] 2019-06-27@08:04:52.557 Parallel-piped sort and policy evaluation. 5 files scanned.

[I] 2019-06-27@08:04:52.588 Piped sorting and candidate file choosing. 4 records scanned.

[I] Summary of Rule Applicability and File Choices:

Rule#	Hit_Cnt	KB_Hit	Chosen	KB_Chosen
KB_I11	Rule			
0	4	32768	4	32768

0 RULE 'Mig-test' MIGRATE FROM POOL 'system' TO POOL 'Archive_files'
WHERE(.)

[I] Filesystem objects with no applicable rules: 1.

[I] GPFS Policy Decisions and File Choice Totals:

Chose to migrate 32768KB: 4 of 4 candidates;

Predicted Data Pool Utilization in KB and %:

Pool_Name	KB_Occupied	KB_Total	Percent_Occupied
system	3919872	4194304	93.457031250%

[I] Messages tagged with <1> are from node localhost. dispatched. \.....

<1> 2019-06-27 04:04:52 GLESL700I: Task migrate was created successfully, task id is 1004.

<1> 2019-06-27 04:04:53 GLESM896I: Starting the stage 1 of 3 for migration task 1004 (qualifying the state of migration candidate files).

<1> 2019-06-27 04:04:53 GLESM897I: Starting the stage 2 of 3 for migration task 1004 (copying the files to 1 pools).

<1> 2019-06-27 04:04:53 GLESM898I: Starting the stage 3 of 3 for migration task 1004 (changing the state of files on disk).

<1> 2019-06-27 04:04:53 GLESL038I: Migration result: 4 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too early for migration.

[I] 2019-06-27@08:04:55.851 Policy execution. 4 files dispatched.

[I] A total of 4 files have been migrated, deleted or processed by an EXTERNAL EXEC/script;
0 'skipped' files and/or errors.

When the **mmapplypolicy** command completes, the files should be migrated to tapes, as shown in Example 11.

Example 11 Verifying migrated files

```
[root@spectrumscale archive]# dsmls | tail -n +10
10485760      0      0 m      ltfs      10Mfile_1
10485760      0      0 m      ltfs      10Mfile_2
10485760      0      0 m      ltfs      10Mfile_3
2097152       0      0 m      ltfs      text

[root@spectrumscale archive]# eeadm file state * -s
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_3
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/text
```

If the files are read, they are recalled from the tapes, as shown in Example 12. The recall is triggered by file access.

Example 12 Files recalled from tapes

```
[root@spectrumscale archive]# tail -5 /ibm/fs1/archive/text
This is test
This is test
This is test
This is test
This

[root@spectrumscale archive]# dsmls * | grep text
2097152      2097152      2048 p      ltfs      text

[root@spectrumscale archive]# eeadm file state text
Name: /ibm/fs1/archive/text
State: premigrated
ID: 17428432867258979361-5537376248122048639-1043068732-93711-0
Replicas: 1
Tape 1: VTAP02L5@pool1@lib1 (tape state=appendable)

[root@spectrumscale archive]# eeadm file state text -s
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/text
```

The file status is p (premigrated). After a read of the migrated file, the data is recalled on to disks. The data on the tape is still valid.

If you modify the migrated (or premigrated) file, the status is resident (r), as shown in Example 13.

Example 13 Resident status

```
[root@spectrumscale archive]# echo aaa >> 10Mfile_2

[root@spectrumscale archive]# dsmls * | grep 10Mfile_2
10485764      10485764      10248 r      -      10Mfile_2
```

```
[root@spectrumscale archive]# eeadm file state 10Mfile_2
Name: /ibm/fs1/archive/10Mfile_2
State: resident
```

5. The migration can also be invoked by directly passing file names to the **eeadm migrate** command, as shown in Example 14.

Example 14 Running eeadm migration command directly

```
[root@spectrumscale archive]# sleep 120
[root@spectrumscale archive]# cat * >/dev/null
[root@spectrumscale archive]# eeadm file state * -s
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_1
R 0 - - - -- /ibm/fs1/archive/10Mfile_2
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_3
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/text

[root@spectrumscale archive]# find . -type f | eeadm migrate -p pool1
2019-06-27 04:11:19 GLESL277I: The "eeadm migrate command" is called without
specifying an input file. Waiting for standard input.
If necessary press ^D to exit.
2019-06-27 04:11:19 GLESL700I: Task migrate was created successfully, task id
is 1010.
2019-06-27 04:11:20 GLESM896I: Starting the stage 1 of 3 for migration task
1010 (qualifying the state of migration candidate files).
2019-06-27 04:11:20 GLESM897I: Starting the stage 2 of 3 for migration task
1010 (copying the files to 1 pools).
2019-06-27 04:11:20 GLESM898I: Starting the stage 3 of 3 for migration task
1010 (changing the state of files on disk).
2019-06-27 04:11:20 GLESL038I: Migration result: 4 succeeded, 0 failed, 0
duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for
migration, 0 too early for migration.

[root@spectrumscale archive]# eeadm file state * -s
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_3
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/text
```

6. The **eeadm task** commands show the results of the completed commands and the status of the running commands. For instance, the completed commands are shown in Example 15.

Example 15 Seeing results of completed commands

```
[root@spectrumscale archive]# eeadm task list -c
TaskID Type Result CreatedTime(-0400) StartedTime(-0400)
CompletedTime(-0400)
1000 tape_assign succeeded 2019-06-27_03:51:30 2019-06-27_03:51:30
2019-06-27_03:51:32
1004 migrate succeeded 2019-06-27_04:04:52 2019-06-27_04:04:52
2019-06-27_04:04:53
1006 transparent_recall succeeded 2019-06-27_04:06:34 2019-06-27_04:06:34
2019-06-27_04:06:34
1007 transparent_recall succeeded 2019-06-27_04:07:38 2019-06-27_04:07:38
2019-06-27_04:07:38
```

```

1008 transparent_recall succeeded 2019-06-27_04:10:33 2019-06-27_04:10:33
2019-06-27_04:10:33
1009 transparent_recall succeeded 2019-06-27_04:10:33 2019-06-27_04:10:33
2019-06-27_04:10:33
1010 migrate succeeded 2019-06-27_04:11:19 2019-06-27_04:11:19
2019-06-27_04:11:20

```

```

[root@spectrumscale archive]# eeadm task show 1010
=== Task Information ===
Task ID:          1010
Task Type:        migrate
Command Parameters: eeadm migrate -p pool1
Status:          completed
Result:          succeeded
Accepted Time:    Thu Jun 27 04:11:19 2019 (-0400)
Started Time:     Thu Jun 27 04:11:19 2019 (-0400)
Completed Time:   Thu Jun 27 04:11:20 2019 (-0400)
Workload:         4 files. 1 replicas.
                  10485764 bytes to copy. 1 copy tasklets on pool1@lib1.
Progress:         -
                  1/1 copy tasklets completed on pool1@lib1.
Result Summary:   4 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong
pool, 0 not found, 0 too small, 0 too early.
                  (GLESM899I) All files have been successfully copied on
pool1/lib1.

```

```

[root@spectrumscale archive]# eeadm task show 1010 -v
=== Task Information ===
Task ID:          1010
Task Type:        migrate
Command Parameters: eeadm migrate -p pool1
Status:          completed
Result:          succeeded
Accepted Time:    Thu Jun 27 04:11:19 2019 (-0400)
Started Time:     Thu Jun 27 04:11:19 2019 (-0400)
Completed Time:   Thu Jun 27 04:11:20 2019 (-0400)
Workload:         4 files. 1 replicas.
                  10485764 bytes to copy. 1 copy tasklets on pool1@lib1.
Progress:         -
                  1/1 copy tasklets completed on pool1@lib1.
Result Summary:   4 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong
pool, 0 not found, 0 too small, 0 too early.
                  (GLESM899I) All files have been successfully copied on
pool1/lib1.
Messages:
2019-06-27 04:11:20.453474 GLESM896I: Starting the stage 1 of 3 for migration
task 1010 (qualifying the state of migration candidate files).
2019-06-27 04:11:20.460617 GLESM897I: Starting the stage 2 of 3 for migration
task 1010 (copying the files to 1 pools).
2019-06-27 04:11:20.637450 GLESM898I: Starting the stage 3 of 3 for migration
task 1010 (changing the state of files on disk).
2019-06-27 04:11:20.898493 GLESL038I: Migration result: 4 succeeded, 0 failed,
0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for
migration, 0 too early for migration.

```



```

--- Subtask(level 1) Info ---
Task ID:          1011
Task Type:        copy_replica
Status:           completed
Result:           succeeded
Accepted Time:    Thu Jun 27 04:11:20 2019 (-0400)
Started Time:     Thu Jun 27 04:11:20 2019 (-0400)
Completed Time:   Thu Jun 27 04:11:20 2019 (-0400)
Workload:         10485764 bytes to copy. 1 copy tasklets on
pool1@lib1.
Progress:         1/1 copy tasklets completed on pool1@lib1.
Result Summary:   (GLESM899I) All files have been successfully copied
on pool1/lib1.
Messages:
  2019-06-27 04:11:20.482769 GLESM825I: The Copy tasklet (0x20delc0) will be
  dispatched on drive VDRIVE0000 from the write queue (tape=VTAP02L5).
  2019-06-27 04:11:20.509717 GLESM031I: A list of 1 file(s) has been added
  to the migration and recall queue.
  2019-06-27 04:11:20.631374 GLESM134I: Copy result: 1 succeeded, 0 failed,
  0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for
  migration, 0 too early for migration.
  2019-06-27 04:11:20.631466 GLESM899I: All files have been successfully
  copied on pool1/lib1.

```

```

[root@spectrumscale archive]# eeadm task show 1010 -r
Result   Failure Code  Failed time           Node File name
Success  -                -                      - /ibm/fs1/archive/text
Success  -                -                      -
/ibm/fs1/archive/10Mfile_1
Success  -                -                      -
/ibm/fs1/archive/10Mfile_2
Success  -                -                      -
/ibm/fs1/archive/10Mfile_3

```

7. Clean up the test environment, as shown in Example 16.

Example 16 Cleaning up the test environment

```

[root@spectrumscale archive]# cd
[root@spectrumscale ~]# rm -rf /ibm/fs1/archive

```

Automatic migration

Automatic migration includes the following objectives:

- Understand the policy rule for automatic migration
- Configure the automatic migration

Complete the following steps:

1. Create test files under a test directory called /ibm/fs1/archive and run manual migration, as shown in Example 17. Ensure all test files are migrated by running the **dsm1s** command or **eeadm file state** command.

Example 17 Test directory

```

[root@spectrumscale ~]# mkdir /ibm/fs1/archive

```

```
[root@spectrumscale ~]# cd /ibm/fs1/archive
[root@spectrumscale archive]# dd if=/dev/urandom of=10Mfile_1 count=1 bs=10M
[root@spectrumscale archive]# for i in $(seq 2 3); do cp 10Mfile_1 10Mfile_$i;
done
[root@spectrumscale archive]# sleep 120; mmapplypolicy /ibm/fs1/archive -P
/root/policy/policy_migArchiveDir

[root@spectrumscale archive]# eeadm file state * -s
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_3
```

2. Add a callback to GPFS, as shown in Example 18.

Example 18 Adding a callback to GPFS

```
[root@spectrumscale archive]# mmaddcallback MIGRATION --command
/usr/lpp/mmfs/bin/mmstartpolicy --event lowDiskSpace --parms "%eventName
%fsName --single-instance"
mmaddcallback: mmsdrfs propagation completed.

[root@spectrumscale archive]# mmlscallback MIGRATION
MIGRATION
    command      = /usr/lpp/mmfs/bin/mmstartpolicy
    event         = lowDiskSpace
    parms         = %eventName %fsName --single-instance
```

3. Define a policy rule for the auto migration. The policy file to be applied is found at /root/policy/policy_migAuto. Unlike the policy_migArchiveDir file, the policy_migAuto does not pick up files in premigrate state, as shown in Example 19. The **THRESHOLD** statement is described later in this section.

Example 19 Policy file for auto migration

```
define(
    exclude_list,
    PATH_NAME LIKE '/ibm/fs1/.ltfsee/%'
    OR PATH_NAME LIKE '/ibm/fs1/.SpaceMan/%'
    OR PATH_NAME LIKE '/ibm/fs1/.snapshots/%'
    OR ((CURRENT_TIMESTAMP - MODIFICATION_TIME) < INTERVAL '121' SECONDS)
    OR (FILE_SIZE <1024)
)

define(
    is_premigrated,
    MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE '%V%'
)

define(
    is_migrated,
    MISC_ATTRIBUTES LIKE '%V%'
)

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'Archive_files'
EXEC '/opt/ibm/ltfsee/bin/eeadm'
```

```

OPTS '-p pool1@lib1'

RULE 'ARCHIVE_FILES' MIGRATE FROM POOL 'system'
  TO POOL 'Archive_files'
  THRESHOLD (0, 0)
  WHERE (
    PATH_NAME LIKE '/ibm/fs1/archive/%'
    AND NOT (exclude_list)
    AND NOT (is_migrated)
    AND NOT (is_premigrated)
  )

```

4. Apply the policy by running the **mmchpolicy** command, as shown in Example 20.

Example 20 Applying the policy

```

[root@spectrumscale archive]# mmchpolicy fs1 /root/policy/policy_migAuto -t
"System policy for LTFS"
Validated policy `System policy for LTFS': Parsed 3 policy rules.
Policy `System policy for LTFS' installed and broadcast to all nodes.

[root@spectrumscale archive]# cp 10Mfile_1 10Mfile_AutoMigTest
[root@spectrumscale archive]# eeadm file state * -s
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_3
R 0 - - - -- /ibm/fs1/archive/10Mfile_AutoMigTest

```

The migration is started for the newly created files after approximately 2 minutes. The file status changes from resident (R) to migrated (M), as shown in Example 21.

Example 21 Verifying file status change

```

[root@spectrumscale archive]# eeadm file state * -s
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_3
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/10Mfile_AutoMigTest

```

Notice that premigrated file is not migrated as defined in the policy.

The policy rule that is used for this test had THRESHOLD (0, 0) in the RULE MIGRATE statement so that the lowDiskSpace event always occurs, regardless of the disk space. These values can be adjusted, such as THRESHOLD (80, 60), when deploying at a customer's system so that the migration occurs only when 80% or more disk space is being used. The migration continues to run until the used disk space is lower than 60%.

5. To clean up, delete the callback, set a default policy, and delete the test files, as shown in Example 22.

Example 22 Cleaning up

```

[root@spectrumscale archive]# mmdelcallback MIGRATION
mmdelcallback: mmsdrfs propagation completed.

[root@spectrumscale archive]# mmchpolicy fs1 DEFAULT
Validated policy 'DEFAULT':
Policy `DEFAULT' installed and broadcast to all nodes.

```

```
[root@spectrumscale archive]# cd
[root@spectrumscale ~]# rm -rf /ibm/fs1/archive
```

Selective Recall

Selective recall includes running the **eeadm recall** command to start selective recalls.

Opposed to the transparent recall, a selective recall is an explicit recall of a specified group of files. The specified files will be recalled in a performance optimized manner.

Complete the following steps:

1. Create test files and migrate them, as shown in Example 23.

Example 23 Creating test files

```
[root@spectrumscale ~]# mkdir /ibm/fs1/archive
[root@spectrumscale ~]# cd /ibm/fs1/archive
[root@spectrumscale archive]# dd if=/dev/urandom of=2Mfile_1 count=1 bs=2M
[root@spectrumscale archive]# for i in $(seq 2 5); do cp 2Mfile_1 2Mfile_${i};
done
[root@spectrumscale archive]# sleep 120; mmappolicy /ibm/fs1/archive -P
/root/policy/policy_migArchiveDir

[root@spectrumscale archive]# eeadm file state * -s
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_3
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_4
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_5
```

2. Run the **eeadm recall** command by passing the file names, as shown in Example 24. Verify the file status as shown in the same example.

Example 24 Running the eeadm recall command

```
[root@spectrumscale archive]# find /ibm/fs1/archive -type f | eeadm recall
2019-06-27 04:34:13 GLESL277I: The "eeadm recall command" is called without
specifying an input file. Waiting for standard input.
If necessary press ^D to exit.
2019-06-27 04:34:13 GLESL268I: 5 file name(s) have been provided to recall.
2019-06-27 04:34:14 GLESL700I: Task selective_recall was created successfully,
task id is 1019.
2019-06-27 04:34:14 GLESL263I: Recall result: 5 succeeded, 0 failed, 0
duplicate, 0 not migrated, 0 not found, 0 unknown.

[root@spectrumscale archive]# eeadm file state /ibm/fs1/archive/* -s
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_1
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_2
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_3
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_4
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_5
```

3. The **eeadm recall** command also accepts a file list. To test the command using a file list, run a manual migration to migrate all test files, as shown in Example 25 on page 19.

Example 25 Preparing test files

```
[root@spectrumscale archive]# mmapplypolicy /ibm/fs1/archive -P
/root/policy/policy_migArchiveDir
[root@spectrumscale archive]# eeadm file state /ibm/fs1/archive/* -s
```

4. Create a file list of the migrated test files and pass it to the **eeadm recall** command, as shown in Example 26.

Example 26 Running eeadm recall by passing a file list

```
[root@spectrumscale archive]# find /ibm/fs1/archive -type f -name "2Mfile*"
>/tmp/file_list
[root@spectrumscale archive]# cat /tmp/file_list
/ibm/fs1/archive/2Mfile_1
/ibm/fs1/archive/2Mfile_3
/ibm/fs1/archive/2Mfile_2
/ibm/fs1/archive/2Mfile_4
/ibm/fs1/archive/2Mfile_5

[root@spectrumscale archive]# eeadm recall /tmp/file_list
2019-06-27 04:36:22 GLES268I: 5 file name(s) have been provided to recall.
2019-06-27 04:36:23 GLES2700I: Task selective_recall was created successfully,
task id is 1022.
2019-06-27 04:36:23 GLES263I: Recall result: 5 succeeded, 0 failed, 0
duplicate, 0 not migrated, 0 not found, 0 unknown.

[root@spectrumscale archive]# eeadm file state /ibm/fs1/archive/* -s
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_1
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_2
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_3
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_4
P 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_5
```

If a customer has millions of files to migrate, the **find** command can be slow. The **mmapplypolicy** command can be used to generate the file list faster. As shown in Example 27, the following policy rule can be used to find the files and run the **eeadm recall** command: `/root/policy/policy_recallArchiveDir`.

Example 27 Recall policy rule

```
RULE EXTERNAL POOL 'pool1'
  EXEC '/opt/ibm/ltfsee/bin/eeadm'
RULE 'ee_recall' MIGRATE FROM POOL 'pool1' TO POOL 'system'
  WHERE
  (
    PATH_NAME LIKE '/ibm/fs1/archive/%'
  )
```

Note: The rule **ee_recall** indicates data movement from pool1 (LTFS/tape) to system (GPFS/disk).

5. To test the **eeadm recall** command by using **mmappolicy**, run a manual migration to migrate all test files, as shown in Example 28.

Example 28 Running a manual migration

```
[root@spectrumscale archive]# mmappolicy /ibm/fs1/archive -P
/root/policy/policy_migArchiveDir
[root@spectrumscale archive]# eeadm file state /ibm/fs1/archive/* -s
```

6. Run the policy for selective recall and verify the status, as shown in Example 29.

Example 29 Running mmappolicy for selective recall

```
[root@spectrumscale archive]# mmappolicy /ibm/fs1/archive -P
/root/policy/policy_recallArchiveDir
....
[I] Summary of Rule Applicability and File Choices:
Rule#      Hit_Cnt      KB_Hit      Chosen      KB_Chosen
KB_I11     Rule
   0         5         0         5         10240
10240      RULE 'ee_recall' MIGRATE FROM POOL 'pool1' TO POOL 'system' WHERE(.)
....
<1> 2019-06-27 04:37:55 GLESL268I: 5 file name(s) have been provided to recall.
<1> 2019-06-27 04:37:55 GLESL700I: Task selective_recall was created
sucessfully, task id is 1025.
<1> 2019-06-27 04:37:56 GLESL263I: Recall result: 5 succeeded, 0 failed, 0
duplicate, 0 not migrated, 0 not found, 0 unknown.
[I] 2019-06-27@08:37:57.987 Policy execution. 5 files dispatched.
[I] A total of 5 files have been migrated, deleted or processed by an EXTERNAL
EXEC/script;
    0 'skipped' files and/or errors.

[root@spectrumscale archive]# eeadm file state /ibm/fs1/archive/* -s
P 1  VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_1
P 1  VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_2
P 1  VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_3
P 1  VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_4
P 1  VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_5
```

7. By executing the **selective recall** command, the status of the specified files have changed from migrated to premigrated. There is also an option to change the status from migrated or premigrated to resident status. To test the resident status option, run **eeadm recall** with the **--resident** option, as shown in Example 30.

Example 30 Running eeadm recall with --resident

```
[root@spectrumscale archive]# find /ibm/fs1/archive -type f | eeadm recall
--resident
2019-06-27 04:40:08 GLESL277I: The "eeadm recall command" is called without
specifying an input file. Waiting for standard input.
If necessary press ^D to exit.
2019-06-27 04:40:08 GLESL268I: 5 file name(s) have been provided to recall.
2019-06-27 04:40:08 GLESL700I: Task selective_recall was created
successfully, task id is 1027.
2019-06-27 04:40:09 GLESM750I: All 5 files were successfully changed to a
resident state.
```

```
[root@spectrumscale archive]# eeadm file state /ibm/fs1/archive/* -s
R 0 - - - -- /ibm/fs1/archive/2Mfile_1
R 0 - - - -- /ibm/fs1/archive/2Mfile_2
R 0 - - - -- /ibm/fs1/archive/2Mfile_3
R 0 - - - -- /ibm/fs1/archive/2Mfile_4
R 0 - - - -- /ibm/fs1/archive/2Mfile_5
```

8. Run **eeadm task** commands to see the results, as shown in Example 31.

Example 31 Running eeadm task commands

```
[root@spectrumscale archive]# eeadm task list -c
```

TaskID	Type	Result	CreatedTime(-0400)	StartedTime(-0400)	CompletedTime(-0400)
1000	tape_assign	succeeded	2019-06-27_03:51:30	2019-06-27_03:51:30	2019-06-27_03:51:32
1004	migrate	succeeded	2019-06-27_04:04:52	2019-06-27_04:04:52	2019-06-27_04:04:53
1006	transparent_recall	succeeded	2019-06-27_04:06:34	2019-06-27_04:06:34	2019-06-27_04:06:34
1007	transparent_recall	succeeded	2019-06-27_04:07:38	2019-06-27_04:07:38	2019-06-27_04:07:38
1008	transparent_recall	succeeded	2019-06-27_04:10:33	2019-06-27_04:10:33	2019-06-27_04:10:33
1009	transparent_recall	succeeded	2019-06-27_04:10:33	2019-06-27_04:10:33	2019-06-27_04:10:33
1010	migrate	succeeded	2019-06-27_04:11:19	2019-06-27_04:11:19	2019-06-27_04:11:20
1012	migrate	succeeded	2019-06-27_04:17:17	2019-06-27_04:17:17	2019-06-27_04:17:18
1014	transparent_recall	succeeded	2019-06-27_04:22:06	2019-06-27_04:22:06	2019-06-27_04:22:06
1015	migrate	succeeded	2019-06-27_04:25:43	2019-06-27_04:25:43	2019-06-27_04:25:44
1017	migrate	succeeded	2019-06-27_04:32:10	2019-06-27_04:32:10	2019-06-27_04:32:11
1019	selective_recall	succeeded	2019-06-27_04:34:14	2019-06-27_04:34:14	2019-06-27_04:34:14
1020	migrate	succeeded	2019-06-27_04:35:23	2019-06-27_04:35:23	2019-06-27_04:35:24
1022	selective_recall	succeeded	2019-06-27_04:36:23	2019-06-27_04:36:23	2019-06-27_04:36:23
1023	migrate	succeeded	2019-06-27_04:37:29	2019-06-27_04:37:29	2019-06-27_04:37:30
1025	selective_recall	succeeded	2019-06-27_04:37:55	2019-06-27_04:37:55	2019-06-27_04:37:56
1026	selective_recall	succeeded	2019-06-27_04:38:49	2019-06-27_04:38:49	2019-06-27_04:38:49
1027	selective_recall	succeeded	2019-06-27_04:40:08	2019-06-27_04:40:08	2019-06-27_04:40:09

```
[root@spectrumscale archive]# eeadm task show 1006
=== Task Information ===
Task ID: 1006
Task Type: transparent_recall
Command Parameters: dsmrecall
Status: completed
Result: succeeded
Accepted Time: Thu Jun 27 04:06:34 2019 (-0400)
Started Time: Thu Jun 27 04:06:34 2019 (-0400)
Completed Time: Thu Jun 27 04:06:34 2019 (-0400)
Workload: 2097152 bytes of file (name: /ibm/fs1/archive/text, inode: 93711)
Progress: -
Result Summary: Transparent Recall result: Succeeded
```

```
[root@spectrumscale archive]# eeadm task show 1027
=== Task Information ===
Task ID: 1027
Task Type: selective_recall
Command Parameters: eeadm recall --resident
Status: completed
Result: succeeded
Accepted Time: Thu Jun 27 04:40:08 2019 (-0400)
Started Time: Thu Jun 27 04:40:08 2019 (-0400)
Completed Time: Thu Jun 27 04:40:09 2019 (-0400)
Workload: 5 files, 0 bytes in total to recall in this task.
5 files, 0 bytes in total to recall in this task. [subtask:1028]
```

Progress: 5 completed (or failed) files / 5 total files.
5 completed (or failed) files / 5 total files. [subtask:1028]
Result Summary: Selective Recall result: 5 succeeded, 0 failed, 0 duplicate, 0 not migrated, 0 not found.
- [subtask:1028]

```
[root@spectrumscale archive]# eeadm task show 1027 -r
```

Result	Failure Code	Failed time	Node	File name
Success	-	-	-	/ibm/fs1/archive/2Mfile_1
Success	-	-	-	/ibm/fs1/archive/2Mfile_3
Success	-	-	-	/ibm/fs1/archive/2Mfile_2
Success	-	-	-	/ibm/fs1/archive/2Mfile_4
Success	-	-	-	/ibm/fs1/archive/2Mfile_5

9. Clean up the test environment, as shown in Example 32.

Example 32 Cleaning up the test environment

```
[root@spectrumscale archive]# cd  
[root@spectrumscale ~]# rm -rf /ibm/fs1/archive
```

Preserving file system objects on tape

Symbolic links, empty regular files, and empty directories are some of the file system objects that do not contain data or content. These objects cannot be migrated to tapes by using the **eeadm migrate** command. IBM Spectrum Archive provides a dedicated command to save these objects in tapes. The command is called the **eeadm save** command.

For additional information refer to the following documents:

- ▶ [IBM Knowledge Center for IBM Spectrum Archive](#)
- ▶ IBM Redbooks® publication *IBM Spectrum Archive Enterprise Edition V1.3.0: Installation and Configuration Guide*, SG24-8333, section “Preserving file system objects on tape”.

Tape Management Commands

IBM Spectrum Archive EE provides tape management commands to de-fragment data on tapes, unassign/assign tapes from a pool, export tapes from a library and import tapes into a library.

Table 2 lists the commands.

Table 2 Tape management commands

Feature	eeadm command
Reconciliation	eeadm reconcile
Reclamation	eeadm reclaim
Making tape cartridges offline	eeadm tape offline
Making offline tape cartridges to online	eeadm tape online
Exporting tape cartridges from the system	eeadm tape export
Importing tape cartridges into the system	eeadm tape import

For more information, see the following documents:

- ▶ [IBM Knowledge Center for IBM Spectrum Archive](#)
- ▶ IBM Redbooks publication *IBM Spectrum Archive Enterprise Edition V1.3.0: Installation and Configuration Guide*, SG24-8333.

REST API

IBM Spectrum Archive EE provides REST API support to obtain system configuration in JSON format via HTTP/HTTPS. By default, port 7100 is used for the REST API. In the current version, only GET method is supported and the configuration cannot be changed via REST API. On this virtual appliance, the REST API is ready to use without additional configuration.

REST API URL is in the format:

`http://<host>:7100/ibmsa/v1/<resource>/<id>`

Valid values for <resource> are “libraries”, “nodegroups”, “nodes”, “drives”, “pools”, “tapes” and “tasks”.

<id> is optional. If it is given, a JSON object representing <resource> specified by <id> is shown. If it is omitted, an array of all the JSON object representing <resource> is shown.

Complete the following steps:

1. Ensure that the httpd server is running, as shown in Example 33.

Example 33 Verifying status of httpd

```
[root@spectrumscale ~]# systemctl status httpd
httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; enabled; vendor
  preset: disabled)
   Active: active (running) since Thu 2019-06-27 03:46:46 EDT; 1h 12min ago
  ...
```

2. Ensure that the REST API is working, as shown in Example 34.

Example 34 Verifying status of Spectrum Archive EE REST API

```
[root@spectrumscale ~]# curl -i -XGET 'http://localhost:7100/ibmsa/v1'
HTTP/1.1 200 OK
Date: Thu, 27 Jun 2019 09:00:08 GMT
Server: Apache/2.4.6 (CentOS) OpenSSL/1.0.2k-fips mod_wsgi/3.4 Python/2.7.5
Content-Length: 83
Content-Type: application/json

{"message":"IBM Spectrum Archive REST API server is
working.", "status_code":"200"}
```

3. Get JSON array of all pools via REST API, as shown in Example 35. A parameter **pretty** can be used to make the output in a more human-readable format.

Example 35 Getting a JSON array of pools via REST API

```
[root@spectrumscale ~]# curl -XGET
'http://localhost:7100/ibmsa/v1/pools?pretty'
[
```

```
{
  "active_space": 9,
  "capacity": 29796335616,
  "device_type": "LT0",
  "fill_policy": "Default",
  "format_class": "default",
  "free_space": 29786898432,
  "id": "0f68d0cb-3a99-4289-8ed3-f36cb6efea61",
  "library_id": "1a5e8ded-ff95-4375-89a9-0c086f0cdcc8",
  "library_name": "lib1",
  "low_space_warning_enable": false,
  "low_space_warning_threshold": 0,
  "media_restriction": "^.{8}$",
  "mode": "normal",
  "mount_limit": 0,
  "name": "pool1",
  "no_space_warning_enable": false,
  "nodegroup_name": "G0",
  "non_appendable_space": 0,
  "num_of_tapes": 3,
  "owner": "System",
  "reclaimable_space": 9437175,
  "used_space": 9437184,
  "worm": "no"
}
]
```

4. Similar JSON based output can be obtained by **eeadm pool list** command with the **-json** option, as shown in Example 36.

Example 36 Getting a JSON array of pools via eeadm CLI command

```
[root@spectrumscale ~]# eeadm pool list --json | python -m json.tool
{
  "payload": [
    {
      "active_space": 9,
      "capacity": 29796335616,
      "device_type": "LT0",
      "fill_policy": "Default",
      "format_class": "default",
      "free_space": 29786898432,
      "id": "0f68d0cb-3a99-4289-8ed3-f36cb6efea61",
      "library_id": "1a5e8ded-ff95-4375-89a9-0c086f0cdcc8",
      "library_name": "lib1",
      "low_space_warning_enable": false,
      "low_space_warning_threshold": 0,
      "media_restriction": "^.{8}$",
      "mode": "normal",
      "mount_limit": 0,
      "name": "pool1",
      "no_space_warning_enable": false,
      "nodegroup_name": "G0",
      "non_appendable_space": 0,
      "num_of_tapes": 3,
      "owner": "System",
```

```

        "reclaimable_space": 9437175,
        "used_space": 9437184,
        "worm": "no"
    }
],
"version": 1
}

```

For more information about the REST API feature, see IBM Knowledge Center [here](#).

Policy rule tips

The **mmapplypolicy** policy rule is a powerful interface provided by IBM Spectrum Scale. Various parameters or build-in features can be used to select files or to make a list of files.

This section describes tips for using the **mmapplypolicy** command for IBM Spectrum Archive.

Migration policy tips

The recommended policy rule tips for migration are as follows:

- ▶ Exclude IBM Spectrum Archive internal directories and .snapshot directory
- ▶ Exclude files that have been updated within two minutes
- ▶ Exclude migrated files
- ▶ Exclude files less than 1 KB in size

These are all defined in `/root/policy/policy_migArchiveDir` which are used in previous migration/recall sections:

- ▶ Do not run multiple **mmapplypolicy** commands concurrently that would select the same files.

If the same files are passed to two different **eeadm migrate** commands at the same time, one will fail as a duplicate.

In case of automatic **mmapplypolicy** execution, it is recommended to specify the **--single-instance** option to **mmstartpolicy** when calling **mmaddcallback**, as shown in “Automatic migration” on page 15.

Other tips

This section describes other migration tips:

- ▶ How to write a rule to migrate to multiple pools:

Up to 3 pools can be specified by adding to OPTS definition in the policy rule file as shown in Example 37.

Example 37 Adding OPTS definition in the policy rule file

```

RULE EXTERNAL POOL 'Archive_files'
EXEC '/opt/ibm/ltfsee/bin/eeadm'
OPTS '-p pool1@lib1 pool2@lib1 pool3@lib1'

```

- ▶ How to filter/select files by directory/file names, timestamp or size:

As shown in `/root/policy/policy_migArchiveDir`, `PATH_NAME LIKE` can be used to filter by names:

```
PATH_NAME LIKE '/ibm/fs1/.ltfsee/%'
```

The file size can be obtained by FILE_SIZE:

```
FILE_SIZE < 1024
```

If you want to select files that have not been accessed in the last 30 days, the following condition can be used:

```
DAYS(CURRENT_TIMESTAMP) - DAYS(ACCESS_TIME) > 30
```

- How to control the order of files to migrate, especially for auto migration:

WEIGHT can be used to control the order of files to migrate. Example 38 makes the migration run from the files with the oldest POSIX access time.

Example 38 Using WEIGHT to control the order of files to migrate

```
RULE 'Mig-test' MIGRATE FROM POOL 'system'
  WEIGHT (CURRENT_TIMESTAMP - ACCESS_TIME)
  TO POOL 'Archive_files'
  WHERE (
    PATH_NAME LIKE '/ibm/fs1/archive/%'
    AND NOT (exclude_list)
    ....
```

- How to list migrated or premigrated files

Migrated files can be selected by the following **define** macro:

```
define(is_migrated, (MISC_ATTRIBUTES LIKE '%V%'))
```

Pre-migrated files can be selected by the following **define** macro:

```
define(is_premigrated, (MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE '%V%'))
```

- How to list saved objects:

Saved empty files, symbolic links, and empty directories can be selected by the following **define** macro example:

```
define(is_savedfile, (MISC_ATTRIBUTES LIKE '%F%' AND (NOT
XATTR('dmapi.IBMSTIME') IS NULL AND NOT XATTR('dmapi.IBMSPATH') IS NULL) ))
define(is_savedlink, (MISC_ATTRIBUTES LIKE '%L%' AND (NOT
XATTR('dmapi.IBMSTIME') IS NULL AND NOT XATTR('dmapi.IBMSPATH') IS NULL) ))
define(is_saveddir, (MISC_ATTRIBUTES LIKE '%D%' AND (NOT
XATTR('dmapi.IBMSTIME') IS NULL AND NOT XATTR('dmapi.IBMSPATH') IS NULL) ))
```

- How to list files that are migrated or premigrated to certain tapes:

The following expression can be used to find the files that point to a specified tape (VTAP00L5 in the following sample):

```
XATTR('dmapi.IBMTPS') LIKE '%VTAP00L5@%'
```

Note: In physical multi-library configuration, if tapes with the same barcodes exist in different libraries, the library ID also needs to be specified in the previous expression.

The policy rule can accept input arguments from **mmappolicy -M <key>=<value>**.

The /root/policy/policy_listFilesByTape is the full policy rule that accept the tape name from the **mmappolicy** command. The executed example is shown in Example 39.

Example 39 Policy scan to list up files migrated to a tape

```
[root@spectrumscale archive]# find /ibm/fs1/archive/ -type f | sort | xargs
eadm file state -s
```

```

M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_1
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_2
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_3
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_4
M 1 VTAP02L5@pool1@lib1 - - -- /ibm/fs1/archive/2Mfile_5

[root@spectrumscale archive]# mmapplypolicy /ibm/fs1/archive -P
/root/policy/policy_listFilesByTape -I defer -f /tmp/out -M
target_tape=VTAP02L5
...
Rule#      Hit_Cnt      KB_Hit      Chosen      KB_Chosen
KB_I11     Rule
0          0          0          0          0
0  RULE 'rule_premig' LIST 'files_premig' SHOW(.) WHERE(.)
1          5          0          5          0
0  RULE 'rule_mig' LIST 'files_mig' SHOW(.) WHERE(.)
2          0          0          0          0
0  RULE 'rule_savedfile' LIST 'files_saved' SHOW(.) WHERE(.)
3          0          0          0          0
0  RULE 'rule_savedlink' LIST 'links_saved' DIRECTORIES_PLUS SHOW(.)
WHERE(.)
4          0          0          0          0
0  RULE 'rule_saveddir' LIST 'dirs_saved' DIRECTORIES_PLUS SHOW(.) WHERE(.)
...

[root@spectrumscale archive]# cat /tmp/out.list.files_mig | awk '{print $1"
"$8}'
55814 /ibm/fs1/archive/2Mfile_5
55847 /ibm/fs1/archive/2Mfile_2
55848 /ibm/fs1/archive/2Mfile_3
55849 /ibm/fs1/archive/2Mfile_1
55855 /ibm/fs1/archive/2Mfile_4

```

For more information about the specific policy rule syntax and the built-in features, see the IBM Spectrum Scale Knowledge Center [here](#).

Stopping and restarting IBM Spectrum Archive

In this section, we describe how to stop and restart IBM Spectrum Archive. This process can help if any issue can be seen in your instance of the Virtual Appliance.

Stopping IBM Spectrum Archive

Run the following commands to stop IBM Spectrum Archive:

```
eeadm cluster stop
```

Restarting IBM Spectrum Archive

Run the following commands to restart IBM Spectrum Archive:

```
eeadm cluster start
```

Re-formatting tapes

You may quickly run out of tape space because the capacity of each virtual tape is much less than a physical tape. When tapes become full, take the following steps to re-format the tapes:

1. Delete some GPFS files that are migrated to the pool.
2. Run the **eeadm reconcile** command.
3. Run the **eeadm reclaim** command.

Furthermore, if all GPFS files can be deleted, the following steps would be faster:

1. Delete all GPFS files that are migrated to the pool.
2. Run the **eeadm tape unassign <tapes>** with **-E** option.
3. Run the **eeadm tape assign <tapes>** with **-f** option.

Detailed logs

The following files include detailed messages and log information:

- ▶ `/var/log/messages`
- ▶ `/var/log/ltfsee_trc.log`

This log file is encoded. Run following to see the decoded log:

```
cat /var/log/ltfsee_trc.log | /opt/ibm/ltfsee/bin/ltfsee_catcsvlog
```

In addition, adding **-L 6** to the **mmapp1ypolicy** command gives more information about the policy execution.

More information

For more information about IBM Spectrum Archive, see the following resources:

- ▶ IBM Knowledge Center:
https://www.ibm.com/support/knowledgecenter/ja/ST9MBR_1.3.0/ltfs_ee_ichome.html
- ▶ IBM Redbooks publication *IBM Spectrum Archive Enterprise Edition V1.3.0: Installation and Configuration Guide*, SG24-8333.
<http://www.redbooks.ibm.com/abstracts/sg248333.html>

For more information about IBM Spectrum Scale, see the following resource:

- ▶ IBM Knowledge Center:
https://www.ibm.com/support/knowledgecenter/ja/STXKQY_5.0.3/ibmspectrumscale503_welcome.html

Limitations

The virtual appliance has IBM Spectrum Archive 1.3 running on CentOS. However, CentOS is not officially supported by the product. This virtual appliance is provided as is for demonstration and training purposes. No official support by IBM is provided.

Authors

This Redpaper was produced by a team of specialists from around the world working with the International Technical Support Organization, Tucson Center.

Hiroyuki Miyoshi is an IBM software developer in Tokyo, who is currently engaged in IBM Spectrum Archive development. He joined IBM in 2001 with a master's degree in Electrical Science at Waseda University and has been developing various storage products, such as RAID subsystems, BladeCenter storage and switch, TS7700, SONAS, and IBM Spectrum Scale. His areas of expertise include HSM (Hierarchical Storage Management) with tapes, distributed file system, DR solution for scalable storage systems, and RAS features, such as resource monitoring and callhome.

Hiroshi Araki is an IBM software developer in Tokyo, who is currently engaged in IBM Spectrum Archive development. He joined IBM in 2008 with a master's degree in Computer Science at Keio University and worked in development of storage products, such as TS7700, SONAS, V7000 Unified and IBM Spectrum Scale. In 2015, he joined IBM Spectrum Archive development and has been working on various components, especially CLI command, REST API, and Monitoring server. His areas of expertise include RAS and DR solutions for enterprise storage systems.

Takeshi Ishimoto is the lead developer of Linear Tape File System (LTFS) and IBM Spectrum Archive products since 2009, and he oversees the software architecture, usability, and performance of IBM Spectrum Archive software. He is a contributor to the SNIA LTFS Format Specification, and is currently co-chair of the SNIA LTFS Technical Working Group. He is located in Tokyo and helping tape clients in Asia Pacific and other regions. He was the speaker at IBM Systems Technical University (TechU) in 2017 and 2018, and speaker at SNIA Storage Developer Conference 2019. Before storage development, he was developing firmware and device drivers for mobile devices and their operating systems, such as ThinkPad, WorkPad, Apple's PowerBook, Windows, and Linux for IBM PowerPC®.

Thanks to the following people for their contributions to this project:

Larry Coyne

IBM Systems Worldwide Client Experience Centers

Yuka Sasaki

IBM Systems

Now you can become a published author, too

Here's an opportunity to spotlight your skills, grow your career, and become a published author, all at the same time. Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the US. This material might 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.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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

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 actual people or business enterprises 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
IBM®

IBM Spectrum®
PowerPC®

Redbooks®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

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



REDP-5384-01

ISBN 0738457922

Printed in U.S.A.

Get connected

