

**Hands-on Lab:
Backup and Restore Manager for z/VM
Lab Exercises
SHARE, August 2014
Session 15731**



Reference information for all labs:

You each have a copy (in PDF format) of the Administration Guide and the User Guide for Backup and Restore Manager for z/VM. The following chapter and page numbers will help you find documentation on each of the commands or options used in these labs:

status	<i>Admin Guide</i>	<i>Chapter 6, page 97</i>
bkrlist	<i>User Guide</i>	<i>Appendix A, page 36</i>
bkruser	<i>User Guide</i>	<i>Appendix A, page 39</i>
bkrjob	<i>User Guide</i>	<i>Appendix A, page 33</i>
bkrvol	<i>User Guide</i>	<i>Appendix A, page 42</i>
review	<i>Admin Guide</i>	<i>Chapter 6, page 106</i>
submit	<i>Admin Guide</i>	<i>Chapter 6, page 107</i>
config	<i>Admin Guide</i>	<i>Chapter 5, page 71</i>
console	<i>Admin Guide</i>	<i>Chapter 5, page 83</i>

All Backup and Restore Manager commands that you issue from your user ID to the any of the Backup and Restore Manager service machines must be preceded by:

SMSG <server userid>

For example, to issue the STATUS command from your user ID, type:

SMSG BKRBKUP STATUS

In these labs, we'll be using the following user IDs:

- USER01 through USER15 – you'll each be assigned one of these user IDs to use during class.
- CLASSADM – this will be used by the instructor

- BKRBKUP – the master server for Backup and Restore Manager
- BKRCATLG – the catalog server for Backup and Restore Manager
- BKRWRK01 through BKRWRK04 – the worker service machines for Backup and Restore Manager

We'll also be using Operations Manager for z/VM to view the backup server consoles.

Lab 1: Become familiar with the panel interfaces for finding data that has been backed up and is available for restore

The backup catalog is housed in Shared File System (SFS). You can use the panel interfaces provided by Backup and Restore Manager to review the catalog contents from various starting points:

- **BKRLIST** is what end users will typically use to find their own data. For performance reasons, this command should not be used from a user ID that has administrator privileges on large systems.
- **BKRUSER** is what administrators will typically use to find and restore individual files for a specific user ID.
- **BKRJOB** is what administrators will typically use to find and restore data associated with a specific job and instance.
- **BKRVOL** is what administrators will typically use to find and restore data associated with a specific DASD volume.

In addition, if you are already familiar with SFS, you can use SFS commands such as **DIRLIST** to review the catalog content directly. This would typically only be needed if you are trying to debug a problem. You should not modify catalog content unless directed to do so by the IBM Support team.

You do not need to use or understand **DIRLIST** in order to use Backup and Restore Manager.

Continue with the step by step instructions for the lab on the following page.

Steps for Lab 1:

1. Using the **BKRUSER** command, find any data owned by your user ID in the backup catalog. Use F11 to select items from the lists presented and drill down to the next level.

You will drill down several levels. You can also use the filter at the top of most panels to narrow down the list of items displayed on that panel. Wildcards are allowed in the filter fields.

Use F3 to exit.

We are not doing a restore in this step. We're just browsing the catalog.

2. Using the **BKRLIST** command, find any data owned by your user ID in the backup catalog. Use one of the following methods:
 - Issue **BKRLIST** then use the Owner field in the top right of the panel to filter on your userid.
 - Issue **BKRLIST * * * USERnn** replacing the nn with your user number.

Again use F11 to select items from the lists presented and drill down to the next level. You will drill down several levels.

Use F3 to exit.

We are not doing a restore in this step. We're just browsing the catalog.

3. Using the **BKRJOB** command, review the backup jobs that have already been run on your system. Typically you would use this interface to restore an entire job. For the lab exercise, however, we'll just restore one file from a job:
 - a. Put the cursor on a specific job/instance/object and use F11 to see the details.
 - b. Select one of the files listed in the job and hit F10 to restore. You can pick any file since you will just restore to your reader – you won't actually replace the file on a disk.
 - c. You have multiple options on this panel:
 - i. Restore directly to disk (we will not do this)
 - ii. Restore to RDR (**we'll choose this one**). Fill in your user ID in the "Or to RDR of userid" field. You don't need to specify a node.
 - iii. Restore to an SFS directory (we will not do this)
 - d. Use F10 to submit the restore request.
 - e. Use F3 to exit
 - f. You should now have two files in your reader:
 - i. The file you requested in the restore
 - ii. A copy of the worker console (in case there were errors)

iii.

4. Using the **BKRVOL** command, review the volumes that contain data that has been backed up on your system. Typically you would use this interface to find and restore data on a volume.

Use F11 multiple times to navigate through the volume backup details and find the files on a device on that volume that has been backed up.

Use F3 to exit.

We are not doing a restore in this step. We're just browsing the catalog.

***** **End of Lab 1** *****

Lab 2: Perform a backup of your A-disk

We'll use a sample job template to create your own backup job that will only back up your A-disk. The backup will be written to disk (rather than tape.)

Many options will already be selected for you. (We'll review those options in future labs.) In this lab you'll only specify which disk is backed up.

For the class, all active job templates will be stored in an SFS directory. This allows Backup Manager to always have access to the latest files and allows each student to update his/her own file in the directory.

You have this SFS directory accessed as filemode E.

Steps for Lab 2:

1. Using FILELIST, verify your A-disk does not contain any large files. You can create one or more small files if you like.
2. Copy the sample job template we've created for the class to your personal job template. Both will be in the SFS directory we have set up for job templates. Use your user ID or user number as the filename to ensure uniqueness:

```
COPYFILE LABSAMP TEMPLATE E <userid> TEMPLATE E
```

3. XEDIT your job template:

```
XEDIT <userid> TEMPLATE E
```

4. Scroll down (around line 44) or find the specification for the target media:

```
CONFIG BKR_OUTPUT_SPEC.
```

You'll notice that the value is **CMSFILE BACKUP DISKPOOL ***.

5. From the XEDIT command line use **FILELIST** to find the file **BACKUP DISKPOOL** on any accessed disk.
6. Browse its contents. You'll notice a list of **<userid> <diskaddress>** pairs. This tells Backup Manager where to write the backup data, since we are not writing to tape.
7. Type **QQUIT** in XEDIT and hit F3 from FILELIST until you are back in the XEDIT session for your job template.
8. Scroll down or find the specification for the job name: **CONFIG BKR_JOB_NAME**. Change the value to match the filename of your job template: **USERnn**
9. Similarly, scroll down or find the specification for the filename of the catalog granules: **CONFIG BKR_CATALOG_GRANULE_FN**. Change the value to match the filename of your job template: **USERnn**
10. You can optionally review other specified options. We'll work with them later.
11. Scroll down or find the **Job_Header** section that contains **INCLUDE** and **EXCLUDE** statements.
12. Modify the **INCLUDE** statement to specify your user ID in the **INCLUDE** column and **0191** in the **VDEV** column. This will include only your 191 disk (or A-disk) in the backup job.

File your changes.

13. Submit your job for review

```
MSG BKRBKUP REVIEW <job template filename>
```

Notice the messages you receive. Correct any errors.

14. You should also receive at least one file from BKRBKUP: the resolved job template. Use the **RDRLIST** command and the **PEEK** function key to review the resolved job.

15. Find the **Job_Header** section and verify that you have only the following statement:

```
DUMPDYN <userid> 0191 $$DRIVER$$
```

If you do not see this statement, return to step 3 and correct any errors.

16. Submit your job for backup:

```
MSG BKRBKUP SUBMIT <job template filename>
```

Notice you receive the same messages as you received in Step 6, plus additional information about which worker the job has been sent to.

17. Using Operations Manager, view the console of that worker to see your job being processed:

```
VIEWCON BKRWRKnn
```

Keep in mind that the job may complete quickly and others in the class are also submitting jobs. You can scroll up (F7) or use the reverse find (\) command when viewing the console to find your user ID.

Exit the console using F3.

18. Using Operations Manager, view the console of the catalog server to verify your backup data has been cataloged.

```
VIEWCON BKRCATLG
```

You can search for messages like the following:

```
BKRBU19047I --- Owner - VDEV: <userid> - 0191
```

Exit the console using F3.

19. Using the **BKRUSER** command, find your data that has just been backed up. Remember to select an item and use F11 to drill down. You may have to drill down several levels to find the list of files you can restore.

20. Choose a file and request a restore of the file to your reader. Verify the file is in your reader.

***** End of Lab 2 *****

Lab 3: Practice INCLUDE/EXCLUDE options

In this exercise, you won't perform any actual backups. Instead, you'll practice specifying multiple **INCLUDE/EXCLUDE** statements to see how they affect each other. You will also specify various values for each of the available fields on the **INCLUDE/EXCLUDE** statements.

You'll also specify multiple workers for the job to see how Backup Manager divides the job amongst the workers. You'll then use the REVIEW function to verify your results. Again, we won't actually submit any backup jobs for execution.

To understand the various options on the **INCLUDE/EXCLUDE** statements, it's important to remember how wildcard characters work. Backup Manager supports the following z/VM wildcard characters:

* means 0 to n characters

% means one and only one character. You can use multiple of these to represent, for example, exactly 2 characters

means one and only one numeric character (0-9)

@ means one and only one hex character (0-9, A-F, a-f)

& means one and only one alphabetic character (A-Z, a-z)

Steps for Lab 3:

1. XEDIT your job template:
XEDIT <userid> TEMPLATE E
2. Find the specification for the number of workers: **CONFIG BKR_JOB_WORKERS**. Change the value to **2**.
3. Find the **Job_Header** section that contains **INCLUDE** and **EXCLUDE** statements.

4. Modify the **INCLUDE** and **EXCLUDE** statements to do any or all of the following.

Note: To specify SFS and DASD volume information, uncomment and use the **INCLUDE/EXCLUDE SFS** and **INCLUDE/EXCLUDE RDEVVOL** statements instead of **INCLUDE/EXCLUDE MINIDISK**.

- Exclude all minidisks on the system
- Include all of MAINT's minidisks
- Include all users' minidisks with a virtual address ending in 91
- Exclude all minidisks that are 50 cylinders or more in size
- Exclude everything in the SFS filepool BKRSFS (note, use the **EXCLUDE SFS** statement)
- Include everything in the SFS filepool BKRSFS owned by BKRADMIN (Backup Manager job templates)
- Include all DASD volumes with a volser that starts with TVn
- Exclude all DASD volumes with a volser that starts with TVx, where x is letter (not a number)

5. File your changes.

6. Submit your job for review:

```
MSG BKRBKUP REVIEW <job template filename>
```

Notice the messages you receive. Correct any errors.

7. Use the **RDRLIST** command and the **PEEK** function key to review the resolved job templates. You'll have one file for each worker you specified.
8. Find the **Job_Header** section in the first file (the filename ends in **00**) and notice the **DUMPDYN** statements, based on your **INCLUDE** and **EXCLUDE** specifications. Remember, this will only list every other disk since you specified 2 workers.

If you do not see these statements, return to step 1 and correct any errors.

9. Find the **Job_Header** section in the second file (the filename ends in **01**) and notice the **DUMPDYN** statements for the remaining disks and SFS filespace.

If you do not see these statements, return to step 1 and correct any errors.

10. As time permits, repeat the above steps specifying different options on your **INCLUDE** and **EXCLUDE** statements. You can also change the number of workers.

***** **End of Lab 3** *****