

IBM Spectrum Protect™ Snapshot for Unix® 4.1.4 Installation Planning Sheet

The purpose of this document is to provide a list of considerations that should be reviewed before installing and configuring Spectrum Protect™ Snapshot.

Before going through this document, the Pre-Installation Checklist should be completed.

Tip: Beginning with Version 4.1.3, IBM Tivoli Storage FlashCopy Manager is now IBM Spectrum Protect™ Snapshot. Some applications such as the software fulfillment systems and IBM License Metric Tool use the new product name. However, the software and its product documentation continue to use the Tivoli Storage Manager product name. To learn more about the rebranding transition, see <http://www.ibm.com/support/docview.wss?uid=swg21963634>.

General considerations

During the installation process, the Spectrum Protect Snapshot setup script requests the following information. The decision on how to answer them should be made beforehand during the planning phase with the Pre-Installation Checklist:

•Select one of these configurations:

- (1) backup only
- (2) cloning only
- (3) backup and cloning

This selection determines the mode of operation for Spectrum Protect Snapshot. “Backup only” is the standard mode known from previous releases; it is the only supported choice for DB2 pureScale environments. In this mode databases and custom applications are protected by snapshot backups with optional combination with offloaded tape backups.

“Cloning only” is the cloning mode introduced by Spectrum Protect Snapshot 2.2. This mode is used to clone databases (DB2 and Oracle) without snapshot backup functionality. It is not available for DB2 pureScale environments. Finally, in environments not involving DB2 pureScale, both modes mentioned can be used in combination. This combined mode allows cloning and protecting a database at the same time.

•Are you going to perform offloaded backups to Tivoli Storage Manager? [Y/N]

The answer to this question basically switches support for off-loaded tape backups on or off. If answered with “Yes” this definitively results in the need for an auxiliary server (backup server). However, even if you do not use offloaded backups an auxiliary server might be still required. For more information, see

Chapter 2 → *Preparation of the backup and cloning server in the Installation and User's Guide.*

• **Do you want offloaded tape backups being triggered right after snapshot? [Y/N]**

If answered “Yes” the offload agent (tsm4acs) is added to the /etc/inittab and snapshot copies are offloaded to tape immediately. If answered “No” the agent is not added to the /etc/inittab so that offloaded backups need to be triggered manually or by custom scripts. For high-availability (HA) environments the offload agent is usually added to the individual HA scripts so answer “No” in this case.

- **a) Select the Backup System to update or delete:**
- **b) Select the Clone instance to update or delete:**

These options allow to manage Spectrum Protect Snapshot installations on backup and/or clone servers centrally from the production system via SSH. Please check if the use of SSH is a valid option in your IT environment. If yes, this is the recommended way of installing and configuring Spectrum Protect Snapshot on backup and clone servers. If SSH is not an option in your environment the installation and configuration procedures need to be executed separately on backup and clone servers. Alternatively, a NFS share can be established between production and backup/clone servers (not recommended).

What to read?

The intention of this document is to provide the relevant information for an initial Spectrum Protect Snapshot installation and configuration, in particular listing the most important parameters that must be provided for an initial configuration. The set of parameters depends on the selected applications and storage types. Hence, the document is split into sections each providing the right set of parameters applicable to your specific environment.

For Oracle SAP databases, refer to section **Oracle SAP**.

For plain Oracle (non-SAP) and Oracle ASM databases, refer to **Oracle plain / ASM**.

For DB2 and DB2 SAP databases, refer to section **DB2 / DB2 SAP**.

For Custom Application environments, refer to section **Custom Applications**.

If offloaded backups to tape are used, refer to the application-specific **Offload** section.

If cloning is used for DB2 or Oracle environments, refer to section **Cloning**.

For DS, SVC, and Storwize V7000 storage systems, refer to section **DS8000 / SVC / Storwize V7000**.

For IBM XIV storage systems, refer to section **XIV**.

For IBM N Series and NetApp storage systems, refer to section **IBM N Series and NetApp**

For DB2 pureScale environments based on GPFS, refer to section **GPFS**.

Oracle SAP

In case of SAP Oracle, the following parameters should be considered in the init<SID>.sap file (located in the \$ORACLE_HOME/dbs directory):

For online backups:

| | |
|--------------------------------|--|
| backup_type | online |
| backup_dev_type | util_vol_online |
| TARGET_DATABASE_SUSPEND | specified in the init<SID>.utl file, set to YES |

For offline backups:

| | |
|--------------------------------|-----------------------------|
| backup_type | offline |
| backup_dev_type | util_vol |
| TARGET_DATABASE_SUSPEND | ignored for offline backups |

util_par_file

For snapshot only backups (without TSM for ERP), needs to be set to the fully qualified path of the FCM profile.

For offloaded backups with TSM for ERP, needs to be set to the fully qualified path of the ERP profile (.utl file).

util_path

For snapshot only backups (without TSM for ERP), set to the path of backint_volume.

If TSM for ERP is in use, this parameter is not required to be defined.

util_vol_unit

Specifies the smallest unit that can be backed up with a snapshot or clone, and also determines restore granularity. The most suitable value for this parameter depends on your concrete environment. In principle, it should be set to the smallest possible value. For more details please refer to the User's Manual, chapter 'SAP® BR*Tools configuration profile (.sap)'.

util_vol_access

Specifies the accessibility of snapshot backup volumes, must be set to **none** on the production system

The following parameters queried by the setup wizard should be considered. For parameters not listed here, the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|--------------|---|----------|
| DEVICE_CLASS | For simple scenarios with a single storage device configuration you can keep the default value. Advanced arguments allow to define multiple device classes being applied for distinct time frames or distinct DB partitions only. Refer to the users guide for details. | STANDARD |

Oracle SAP offload

| Name | Description | Value |
|---------|---|-------|
| PROFILE | Fully qualified path to the external SAP Backint profile. | |

Oracle SAP offload with RMAN

If Oracle RMAN is used to perform the offloaded backups on the backup server the following parameters for the TSM for ERP profile (on both the production and backup server) should be considered.

| | | |
|------------------------------------|---|--|
| INCREMENTAL | INCREMENTAL or CUMULATIVE | |
| INCREMENTAL_CHANNELS | Number of RMAN channels to use for the backup | |
| INCREMENTAL_CATALOG_USER | User to connect to the catalog database | |
| INCREMENTAL_CATALOG_CONNECT_STRING | Name of the catalog database | |

Oracle plain / ASM

The following parameters should be considered for plain or ASM Oracle databases. For parameters not listed here, the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|---------------------------------|--|----------|
| DEVICE_CLASS | For simple scenarios with a single storage device configuration you can keep the default value. Advanced arguments allow to define multiple device classes being applied for distinct time frames or distinct DB partitions only. Refer to the users guide for details. | STANDARD |
| VOLUME_MGR | ASM LVM If ASM is selected, the existing option LVM_FREEZE_THAW must not be specified since there is no file system. If LVM is specified, the ASM-related options in the ORACLE section are not allowed in the profile. | |
| CATALOG_DATABASE_CONNECT_STRING | Specifies the connect string of the Recovery Catalog database to be used to catalog backup information. This value must correspond to the value defined in the \$ORACLE_HOME/network/admin/tnsnames.ora file. When in doubt, ask the Oracle administrator. | |
| CATALOG_DATABASE_USERNAME | Specifies a user name that has Oracle system database administrator privileges on the Recovery Catalog database. When in doubt, ask the Oracle administrator. | |
| TARGET_DATABASE_PARAMETER_FILE | Specifies the fully resolved path and file name of the Oracle parameter file (init<SID>.ora by default) for the target database. Note that this file must be a text-based Oracle parameter file (PFILE) and not an Oracle server file. When in doubt, ask the Oracle administrator. | |
| DATABASE_BACKUP_SCRIPT_FILE | Name of the RMAN backup script that contains the Data Protection for Oracle environment variables. When in doubt, ask the Oracle administrator or the TSM administrator. | |
| Oracle ASM only | | |
| ASM_INSTANCE_ID | SID of the ASM instance It is not really recommended by Oracle but possible to have a SID for the ASM instance other than '+ASM'. In such environments, this profile parameter can be used to specify the ASM instance SID. | +ASM |

| Name | Description | Value |
|-------------------|--|-------|
| ASM_ROLE | <p>Specifies the role that should be used when connecting to the ASM instance.</p> <p>The 'sysdba' role must be specified when using Oracle 10g. Specify 'sysasm' role when using Oracle 11g,</p> <p>However, when the database administrator role is separated from the ASM administrator role when using Oracle(R) 11g, the instance owner cannot use the 'sysasm' role.</p> | |
| ASM_INSTANCE_USER | <p>Specifies the user name of the ASM instance owner. Use this parameter when the target database and the ASM instance are running under different user IDs. When this parameter is set to 'AUTO', the ID of the user running the process is used.</p> <p>Note: The ASM instance user must have sysadm permission.</p> | |

Oracle plain / ASM offload

For plain Oracle databases, the defaults for all offload-related parameters can be accepted. Only for ASM the following parameters should be considered if they differ from the production system. If these values are left empty their corresponding values out of the Oracle section are used.

| Oracle ASM only | | |
|------------------------|---|------|
| ASM_INSTANCE_ID | <p>SID of the ASM instance</p> <p>It is not really recommended by Oracle but possible to have a SID for the ASM instance other than '+ASM'. In such environments, this profile parameter can be used to specify the ASM instance SID.</p> | +ASM |

DB2 / DB2 SAP

For DB2 and DB2 SAP databases, the following parameters queried by the setup wizard should be considered. For parameters not listed here, the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|--------------|---|----------|
| DEVICE_CLASS | For simple scenarios with a single storage device configuration you can keep the default value. Advanced arguments allow to define multiple device classes being applied for distinct time frames or distinct DB partitions only. Refer to the users guide for details. | STANDARD |

For DB2 pureScale protection the following additional parameter should be considered.

| Name | Description | Value |
|--------------|---|-------|
| MAX_VERSIONS | Number of GPFS file system snapshot backup versions to be kept before the oldest snapshot backup is deleted. Please note that GPFS file system snapshots can grow over time as the database is modified. Administrators must assess how many snapshots can be kept without risking a file-system-full condition and set the parameter MAX_VERSIONS accordingly. In addition, the administrators are required to keep monitoring free space in the file systems during operation. Setting this parameter to ADAPTIVE is not allowed for DB2 pureScale. | |

For DB2 Standby server protection the following additional parameters should be considered. To configure FCM for DB2 Standby server protection, the Setup script needs to be started in the advanced mode.

| Name | Description | Value |
|-----------------------|---|----------------|
| DB2_PRODUCTION_SERVER | Specify the hostname and port number of the DB2 server where the HADR primary server is running (production system) | |
| DB2_ALIAS | Specify the alias name of the DB2 database running on the HADR primary server. | |
| DB2_USERNAME | Specify the DB2 user that is used to connect from the HADR standby server to the DB2 database running on the HADR primary server. | |
| DB2_AUTH_TYPE | Specify the value of the DB2 instance AUTHENTICATION parameter on the DB2 HADR primary server. | SERVER_ENCRYPT |

| Name | Description | Value |
|----------------|---|----------|
| PRE_FLASH_CMD | This parameter identifies the command script or executable file that is used to quiesce the DB2 standby or DB2 HADR secondary immediately before the snapshot operation begins. | |
| POST_FLASH_CMD | This parameter identifies the command script or executable file that is used to resume the DB2 standby or DB2 HADR secondary immediately after snapshot creation. | |
| DEVICE_CLASS | This parameter is evaluated instead of the parameter specified in the CLIENT section when the DB2 system acts as DB2 standby server or as HADR secondary. | STANDARD |

DB2 / DB2 SAP offload

| Name | Description | Value |
|---------------------|--|-------|
| NUM_SESSIONS | Number of I/O sessions to be created between DB2 and Spectrum Protect server. When in doubt, check with the Spectrum Protect server administrator. | |
| DB2 SAP only | | |
| OPTIONS | If IBM Spectrum Protect for Enterprise Resource Planning is being used, the IBM Spectrum Protect for Enterprise Resource Planning DB2 vendor options file (vendor.env) must be specified. Must not be used when the DB2 native agent handles offloaded backups | |

If **DB2 Standby server protection is in use**, the profile contains a second offload section which has the same parameters and semantics as the standard offload section for custom applications. This second section is used whenever an offloaded backup from a DB2 HADR secondary occurs (instead as from the DB2 HADR primary). For details please refer to the section *Custom Applications Offload* below.

In DB2 pureScale environments, the configuration of a standby server not supported.

Custom applications

For custom application environments, the following parameters queried by the setup wizard should be considered. For parameters not listed here, the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|--------------|--|----------|
| DEVICE_CLASS | For simple scenarios with a single storage device configuration you can keep the default value. Advanced arguments allow to define multiple device classes being applied for distinct time frames or distinct DB partitions only. Refer to the <i>Installation and User's Guide</i> for details. | STANDARD |

Custom applications offload

| Name | Description | Value |
|---------------|---|-------|
| MODE | This parameter determines which of the Spectrum Protect Backup Archive client backup functions to use when creating a Spectrum Protect offloaded backup. In first instance, you can keep the default. Please refer to the <i>Installation and User's Guide</i> for advanced configurations. | FULL |
| ASNODENAME | This parameter identifies the name of the node where data is stored during a Spectrum Protect offloaded backup. | |
| VIRTUALFSNAME | This parameter identifies the virtual file space name of a backup group. It is available when the MODE parameter specifies a value of FULL or DIFF. | |

Cloning

Cloning can be applied for DB2 and Oracle environments (both with or without SAP), but not in environments involving DB2 pureScale. If cloning is activated an additional section CLONING is added to the profile.

| Name | Description | Value |
|--------------|---|-------|
| DEVICE_CLASS | Associates a DEVICE_CLASS section with the cloning operation. A DEVICE_CLASS that is already in use for backup cannot be used for cloning at the same time. The value "STANDARD" is preserved for backup, it cannot be used for cloning purposes. In contrast to the DEVICE_CLASS parameter of the CLIENT section the cloning variant prescribes the specification of at least one database name. | |

DS8000 / SVC / Storwize V7000

The storage device related parameters listed here should be checked with the storage administrator. For parameters not listed here the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|--------------------------------------|---|-------|
| COPYSERVICES_ HARDWARE_TYPE | DS8000 (For IBM DS8000 series) SVC (For IBM SAN Volume Controller and IBM Storwize V7000) | |
| COPYSERVICES_ PRIMARY_ SERVERNAME | Defines the TCP/IP address or hostname of the host running the CIM Agent for DS Open API (which can manage the primary and secondary Copy Services servers of the DS8000 cluster), the SVC master console, or embedded CIM Agent. | |
| COPYSERVICES_ USERNAME | cim user - to connect to CIM Agent for DS Open API svc user - to connect to the primary copyservices server | |
| FLASHCOPY_TYPE | COPY INCR NOCOPY Once flashcopy relations of a certain kind (COPY, INCR, NOCOPY) are established, this value cannot be changed without performing a withdrawal of the previous relation. Therefore, this parameter should be carefully considered using the detailed information in the manual. | |

| Name | Description | Value |
|------------------|--|-------|
| TARGET_SETS | Specify the target volumes to be used in the FlashCopy operation using one of these values: <ul style="list-style-type: none"> •VOLUMES_DIR •VOLUMES_FILE •<list of target set names> (see also TARGET_NAMING) | |
| VOLUMES_FILE | Only applicable if TARGET_SETS is set to VOLUMES_FILE. Specify the name of the target volumes file (.fct). | |
| TARGET_NAMING | Only applicable if TARGET_SETS lists the target set names. Using this parameter, target volume names can be derived from source volume names according to a certain naming pattern specified by this parameter: <string with wildcards %SOURCE and %TARGETSET> For more details please refer to the User's Manual. | |
| BACKUP_HOST_NAME | If your setup requires a backup server specify the name of the backup host as configured in the storage subsystem that is used during forced mount and offloaded tape backup operations. Specify 'PREASSIGNED_VOLUMES' if you are using static volume mapping on SVC / Storwize V7000. For more details please refer to the <i>Installation and User's Guide</i> . | NONE |

XIV

The storage device related parameters listed here should be checked with the storage administrator. For parameters not listed here the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|--------------------------------|--|--------------|
| COPYSERVICES_ HARDWARE_TYPE | IBM XIV® Storage System | XIV |
| COPYSERVICES_ SERVERNAME | Specify the hostname of the IBM XIV® Storage System. | |
| COPYSERVICES_ USERNAME | Username to log in to the XIV® system. | |
| BACKUP_HOST_NAME | Required only in environments involving a backup server. Specify the name of the host as defined in the XIV for the backup system. | |
| PATH_TO_XCLI | specifies the filepath where the IBM XIV® command line interface (XCLI) is installed. | |

IBM N Series and NetApp

The storage device related parameters listed here should be checked with the storage administrator. For parameters not listed here the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|------------------------------|---|----------------------------------|
| COPYSERVICES_ HARDWARE_ TYPE | IBM N Series and NetApp. Set the value to NSERIES_SAN if the N Series or NetApp is connected via SAN or iSCSI. Set the value to NSERIES_NAS if the N Series or NetApp is connected via NFS | NSERIES_SAN or NSERIES_NAS |
| COPYSERVICES_ SERVERNAME | Specify the TCP/IP hostname of the storage system. | |
| COPYSERVICES_ USERNAME | Username to log in to the storage system. | |
| BACKUP_ HOST_ NAME | Required only in environments involving a backup server. For SAN attached storage: Specify the name of the initiator group as defined in the storage system for the backup system. | |

GPFS

For parameters not listed here the default value can be accepted. If there is an entry in the "Value" column, this means that the entry is the recommended value for the parameter.

| Name | Description | Value |
|----------------------------------|---|--------------|
| COPYSERVICES_ HARDWARE_TYPE | GPFS file system snapshots | GPFS |
| NUMBER_GPFS_CONCURRENT_ TASKS | Number of threads that are used to parallelize tasks. The default is 3. | |

Passwords

The setup wizard prompts you to enter the following passwords:

- Device sections: One password for each individual storage box / CIM agent.
For GPFS snapshots, a password is not requested.
- (Oracle only)** Oracle catalog database password used to authenticate with the Oracle catalog DB.
- (DB2 standby setup only)** Password used to authenticate with the DB2 production system (HADR primary node).