IBM Tivoli Storage FlashCopy® Manager 4.1.x for UNIX®

Objective of this document

This document provides comprehensive information about the environment required to support FlashCopy® Manager solutions, including:

- How to prepare for implementation: Complete this checklist before installing the software to avoid issues.
- Identify stakeholders required for successful implementation.
- Identify supported release levels of ancillary software.
- Get the most current tips and hints for installation.

Note: This document refers to the following storage systems: IBM System Storage DS8000, IBM System Storage SAN Volume Controller (SVC), IBM Storwize V7000, IBM Flex System™ v7000, IBM FlashSystem™, IBM XIV® Storage System, IBM Spectrum Accelerate™, IBM N Series, and NetApp®, and also for filesystem snapshots with the IBM General Parallel File System™ (GPFS™) in the context of DB2® pureScale® environments.

Note: In AIX and Linux environments, some checks are covered by the IBM FlashCopy® Manager Prerequisite Checker v4.1.3. For such checks, you find a grey check mark (✔) in the corresponding check column, and a reference to the corresponding property name as it appears in the output of the prerequisite checker is listed, denoted by a double arrow (⇒).
On AIX and Linux environments, all checks that are not marked with a grey check mark (✔) must be performed manually.
Run the IBM FlashCopy® Manager Prerequisite Checker on AIX and Linux in order to perform checks automatically. If the result of the prerequisite checker is "pass" the checks with a grey check mark were performed automatically and successfully. Otherwise, they must be performed manually.
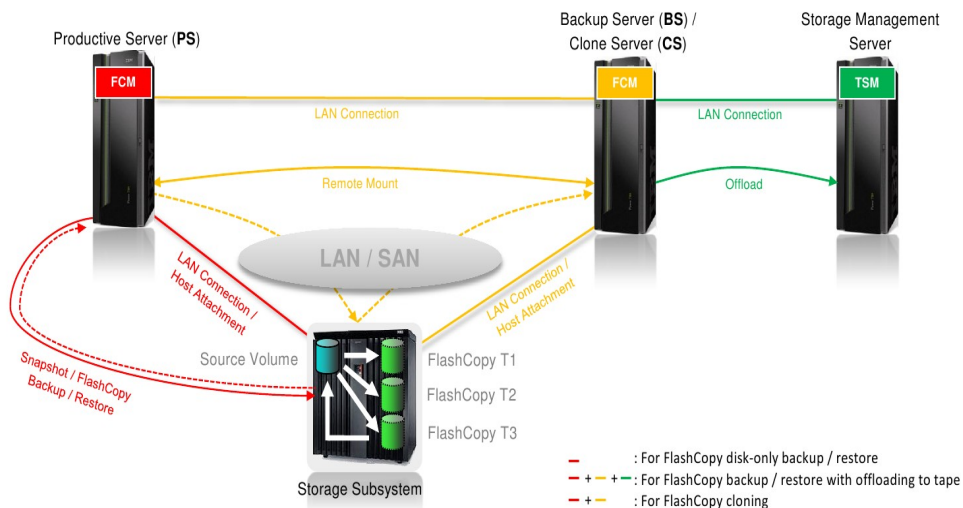On Solaris and HP-UX environments, all checks must be performed manually.

Table of Contents

1. Hardware Requirements

1.1 General Hardware Requirements

	Prerequisite	**checked**
1.1.1	For the use of FlashCopy® Manager (hereafter FCM) backup, offloaded tape backup and cloning operations, different sets of systems are required: **For FlashCopy disk-only backup / restore**: • one or more UNIX or Linux production system(s) (hereafter PS) • each PS must have a data path to the storage system (DS/SVC/XIV/ IBM N Series/NetApp) via FC or iSCSI for SAN or LAN for NAS type storage **For FlashCopy backup / restore with offloading to tape**: • one or more UNIX or Linux production systems (PS) • one or more UNIX or Linux auxiliary system(s) for forced mount and offload (hereafter BS) • the OS release level on PS and BS must be the same • all of these PS/BS must have a data path to the **same** storage system (via FC or iSCSI for SAN, or LAN for NAS type storage) **For FlashCopy cloning**: • one or more production system(s) (PS) • one or more auxiliary system(s) for cloning (hereafter CS) • BS and CS can be running on a single machine or on separate, dedicated machines. • the OS release level on PS and CS must be the same • all of these PS/CS must have a data path to the **same** storage system (via FC or iSCSI for SAN, or LAN for NAS type storage)	

	Prerequisite	checked
	Note 1: Offloaded backup and cloning operations can run on the same auxiliary system or on separate, dedicated auxiliary systems. **Note 2:** In an HA environment that uses AIX PowerHA or HACMP, the FCM auxiliary system cannot act as the standby/takeover server for the production server. For HP-UX, Linux, and Solaris, HA cluster environments are not supported. **For filesystem snapshots of a DB2 pureScale database:** • one or more UNIX servers of the same operating system (AIX, xLinux) **For filesystem snapshots of a DB2 pureScale database with offloading to tape:** • one or more UNIX servers of the same operating system (AIX, xLinux) • one or more servers are for PS, one or more servers act as auxiliary server (BS) **Note:** For DB2 pureScale environments, the term "server" refers to a GPFS cluster as part of a DB2 pureScale instance which consists of one or more servers. In DB2 pureScale environments GPFS file system snapshots are taken. These are independent of the underlying storage subsystem.	
1.1.2	For the following tasks, a LAN connection is required between PS and the storage management server (like TSM): • tape restore operations • backups and restores of the log files For tape backup operations on the BS/CS a LAN connection to the storage management server is required as well unless PS and BS/CS are installed on the same system. **For FlashCopy disk-only backup / restore**: • LAN connection between PS and the storage system **Note:** Under some conditions an auxiliary system (BS) is mandatory even for disk-only backups. For more information, see the *Installation and User's Guide*. In chapter 2, the section "Preparing backup and cloning servers" provides details. If a BS is mandatory it requires: o LAN connection between PS and BS o LAN connection between BS and the storage system *see Note 2 **For FlashCopy backup / restore with offloading to tape**: • LAN connection between PS and the storage system • LAN connection between BS and the storage system *see Note 2 below • LAN connection between PS and BS • LAN connection between the BS and the storage management server for offloaded (tape) backups	

	Prerequisite	checked
	For FlashCopy cloning: • LAN connection between PS and the storage system • LAN connection between CS and the storage system *see Note 2 below • LAN connection between PS and CS **Note 1:** The PS has a LAN connection to the CIM Agent for DS, SVC cluster, IBM Storwize V7000, XIV or IBM N Series or NetApp. **Note 2:** (XIV) The BS/CS also has a LAN connection to the XIV management port to allow FlashCopy Manager to issue XCLI commands. (SVC or IBM Storwize V7000 with dynamic host assignment) The BS/CS also have a LAN connection to the SVC cluster or IBM Storwize V7000. **Note 3:** For simplified installation and configuration of FlashCopy Manager a LAN connection between production and backup/cloning servers is recommended. **Note 4:** In DB2 pureScale environments, maximum performance for tape backup operations can be achieved by either using a high-speed interconnect between production and backup cluster or by using a SAN connection between backup cluster and the GPFS disks itself. For details on GPFS remote mount functionality see GPFS Advanced Administration Guide, Chapter 1	

1.1.3	LUNs of the production database must not be distributed over multiple disk storage subsystems / SVC storage clusters.
	
	In some environments multiple storage systems must be configured:
	• AIX LVM mirroring environments:
	  Each mirror set of the production database must be located completely on a single storage cluster of the storage system (DS / SVC / XIV / IBM N series / NetApp)
	
	• Oracle ASM support:
	  • All failure groups representing a consistent database image must be located on a single disk storage subsystem DS / SVC storage cluster / XIV / IBM N series / NetApp
	  • For DS just one LUN per disk storage subsystem is supported
	  • (Linux) For DS only one target set is allowed to be used
	  • (Linux) For SVC the setting PREASSIGNED_VOLUMES is not supported for the profile parameter BACKUP_HOST_NAME in the device class section of the profile
	⇒ fcm.same_storage_device, fcm.warn_same_storage
		✔

1.1.4	The LUNs of the storage system for the production application must be attached/accessible to the PS and must be visible to FCM. They are used as the source volumes for the FlashCopy / snapshot process. **Note:** For performance reasons, keep the number of LUNs allocated to one application low.	

		Prerequisite	checked
1.1.5		For performing FCM FlashCopy operations, so called target volumes / target LUNs must be created / available. For the different scenarios, consider the following: **For FlashCopy disk-only backup / restore**: • The target volumes must exist but must not be mapped to any host on the storage system. **For FlashCopy backup / restore with offloading to tape**: • The target volumes must be accessible to the BS. **For FlashCopy cloning**: • The target volumes must be accessible to the CS. **Note:** Accessible in this context means the correct zoning/LUN masking (SAN) / network configuration (NAS) is configured to ensure the target volumes can be mapped to BS/CS. (DS/SVC/V7000) **Note:** This applies to DS storage system and also to SVC and IBM Storwize V7000 if preassigned volumes are used. It is recommended to use dynamic target volume mapping instead. (XIV/N Series/NetApp) **Note:** For XIV, IBM N Series or NetApp the target LUNs are created automatically during snapshot process and are assigned to the BS/CS . See documentation of profile parameter *BACKUP_HOST_NAME* in the *Installation and User's Guide* for more details.	
1.1.6		The sizes of the target volumes must be the same as the source volumes. (DS) DS requires source and target volume pairs be located on the same storage system. (SVC/V7000) For SVC and IBM Storwize V7000 those can be located in different storage devices of one and the same SVC cluster.	

1.1.7		**FCM operations in VMware® Virtual Machines**
		Tivoli Storage FlashCopy Manager can back up, restore, and perform clone operations on VMware virtual machines that run Linux guest operating systems.
		
		All source LUNs that participate in backup or clone operations must be attached to the virtual machine by using one of the following attachment methods:
		  • VMware physical mode raw device mapping (pRDM)
		  • iSCSI
		  • Network File System (NFS)
		Tivoli Storage FlashCopy Manager can perform a FlashCopy restore operation from a snapshot to an existing pRDM disk but does not create a virtual machine or pRDM definition as part of the restore process.
		
		If the backup or clone system is a virtual machine then all target LUNs that	

	Prerequisite	checked

	participate in backup or clone operations must be attached to the virtual machine by using one of the following attachment methods:
	   • iSCSI
	   • Network file system (NFS)	

1.1.8	**FCM operations in Linux KVM Virtual Machines**
	Tivoli Storage FlashCopy Manager can back up, restore, and perform clone operations on Linux KVM virtual machines that run Linux guest operating systems.
	
	All source LUNs that participate in backup or clone operations must be attached to the virtual machine by using one of the following attachment methods:
	   • Block device mapping (BDM)
	   • iSCSI
	   • Network File System (NFS)
	   • PCI Passthrough
	
	For the Tivoli® Storage FlashCopy® Manager KVM setup using Block Device Mapping (BDM), the LUNs must be mapped to the KVM guest as multipath devices. The LUNs must be visible inside the KVM guest as multipath devices.
	
	From within the KVM guest, run the multipath command. The output from running this command looks similar to the following example:
	
	`kvm-guest:~ # multipath -ll`
	`mpathat (360050768018205de4000000000001949) dm-7 IBM    ,2145`
	`size=2.0G features='1 queue_if_no_path' hwhandler='0' wp=rw`
	`` `-+- policy='service-time 0' prio=50 status=active ``
	`` `- 3:0:0:3 sdf 8:80  active ready running ``
	
	The LUN identifier is a unique number like this 360050768018205de4000000000001949, and must not be overwritten by the KVM stack. The product storage identifier like this IBM ,2145 must be visible inside the KVM guest.
	
	If the backup or clone system is a virtual machine then all target LUNs that participate in backup or clone operations must be attached to the virtual machine by using one of the following attachment methods:
	   • iSCSI
	   • Network file system (NFS)	

1.2 Hardware Requirements for DS8000

	Prerequisite	checked
1.2.1	IBM System Storage DS8000 (DS8100, DS8300, DS8700, DS8800 or DS8870) with the Point in Time Copy feature enabled. With DS8000 Release 4.1 (bundle 64.1.16.0) and later the embedded CIM agent in the HMC is enabled and configured by default and can be used with FlashCopy Manager. **Note:** IBM CIM Agent for DS Open API version 5.4.2 has a known problem with refresh of an incremental FlashCopy using FCM. Therefore make sure that CIM Agent 5.4.2.xx is not used. Please upgrade the DS microcode to a level that provides CIM Agent >= 5.4.3 (for example: DS8000 bundle version 64.30.78.0 comes with CIM Agent 5.4.3.52)	
1.2.2	DS8000 LIC level for FlashCopy to work with FlashCopy Manager: FlashCopy Manager supports all DS8000 Releases >= R3.1. For older DS8000 releases, the following minimum microcode levels are required: mcode 6.1.600.52/DSCLI 5.0.5.17, mcode 6.2.400.7x/DSCLI 5.2.400 or higher mcode	
1.2.3	Support is provided for Fibre Channel attached volumes	
1.2.4	The user defined for FCM operations in DS8000 needs to be assigned to the "Administrator" user group.	

1.3 Hardware Requirements for SVC / Storwize

	Prerequisite	checked
1.3.1	(SVC) IBM System Storage SAN Volume Controller versions 6.1, 6.2, 6.3, 6.4, 7.1, 7.2, 7.3 and 7.4 (Storwize) IBM Storwize V7000 v1 Disk System IBM Storwize V7000 Version 6.4, 7.1, 7.2, 7.3 and 7.4 IBM Storwize v3700 Version 6.4, 7.1, 7.2, 7.3 and 7.4 IBM Storwize v5000 7.1, 7.2, 7.3 and 7.4 IBM Storwize V7000 Unified v1.3 Disk System (block level support only) IBM Storwize V7000 Unified v1.4 Disk System (block level support only) IBM Storwize V7000 Unified v1.5 Disk System (block level support only) IBM Flex System™ v7000 Version 6.4, 7.1 and 7.2 IBM FlashSystem™ V840 IBM FlashSystem™ V9000 **Note:** FlashCopy Manager communicates with the SVC/Storwize only. There is no communication between FlashCopy Manager and storage systems attached to the SVC.	
1.3.2	Support is provided for Fibre Channel and iSCSI attached volumes	
1.3.3	Support for Metro Mirror and Global Mirror: The target volumes of the Metro/Global Mirror relationships are usually located on a secondary storage system in a remote site. It can be the Disaster Recovery site. For Remote FlashCopy backups, FCM requires that	

	the remote relationships are created, started and added to a unique consistency group. For Metro Mirror, the relationships must be synchronized. For Global Mirror the relationships must already be copied for the first time, otherwise FCM could wait too long and fail running out of time. If you are going to offload the FlashCopy disk backup to TSM, an additional backup host is needed on the secondary site. Prior to the offload backup, FCM will attach the FlashCopy target volumes on the secondary site to this backup host. If you are using the FlashCopy backup on the secondary site for restores, you need a takeover host for the restore containing the same instance structure as in the production host on the production site.	
1.3.4	SVC and Storwize must be configured for https or http communication. By default, SVC/Storwize is configured for HTTPS communication.	
1.3.5	For SVC and Storwize, the user defined for FCM operations in SVC/Storwize needs to be assigned to the "Administrator" user group.	

Note: Starting with SVC 6.1, the CIM agent can be restarted using the Service Assistant for 6.1.0. At http://clusterip/service, go to the Restart Service navigation item and select the CIMOM.

1.4 Hardware Requirements for XIV

	Prerequisite	checked
1.4.1	**IBM XIV Storage System Hardware:** IBM XIV 2nd Generation Storage System v10.0.0.b, and later 10.x levels IBM XIV 3rd Generation Storage System 11.x levels **IBM XIV Storage System Software:** IBM XIV Management Tools 3.x for all IBM XIV® Generations IBM XIV Management Tools 4.2.x for all IBM XIV® Generations or XIV CLI 2.3.1 and later 2.x and 3.x levels for IBM XIV 2nd Generation XIV CLI 3.0.1 and later 3.x levels for IBM XIV 3rd Generation XIV CLI 4.2 and later 4.2.x levels for IBM XIV 3rd Generation **IBM Spectrum Accelerate™:** IBM Spectrum Accelerate™ Version 11.5.x FCM requires the usage of mirror consistency groups for synchronous and asynchronous remote mirroring on XIV. This feature is first implemented with XIV 10.2	
1.4.2	Support is provided for Fibre Channel and iSCSI attached volumes	
1.4.3	Support for Synchronous and Asynchronous Mirroring: The target volumes of the Remote Mirroring relationships are usually located on a secondary storage system in a remote site. It can be the Disaster Recovery site. For Remote Snapshot backups, FCM requires that the remote relationships are created, started and added to a unique consistency group. For Synchronous Mirroring, the relationships must be synchronized. For Asynchronous Mirroring the relationships must already be copied for the first time, otherwise FCM could wait too long and fail running out of time. If you are going to offload the FlashCopy disk backup to TSM, an additional backup host is needed on the secondary site. Prior to the offload backup, FCM will attach the remote snapshots on the secondary site to this backup host. If you are using the FlashCopy backup on the secondary site for restores, you need a takeover host for the restore containing the same instance structure as in the production host on the production site.	
1.4.4	The user account that will be used by FlashCopy Manager to log onto the storage system must have "Storage Administrator" privileges	

1.5 Hardware Requirements for N Series/NetApp

	Prerequisite	checked
1.5.1	Supported ONTAP versions are 7.3 or later 7.3.x levels and 8.1 or later 8.1.x levels	
1.5.2	Volumes must be attached to the production server via SAN, iSCSI or NAS.	
1.5.3	The user account that will be used by FlashCopy Manager to log onto the storage system must have the permissions / capabilities:	

api-system-*, api-volume*, login-http-admin, api-lun-list*, api-lun-get-serial-number, api-cg-start, api-cg-commit, and snapshot-list-info	

2. Software Requirements

2.1 Software Requirements – Operating System

The following software - if not otherwise specified - is required on the PS and if applicable on the BS/CS:

AIX

	Prerequisite	checked PS	checked BS/CS
2.1.1	AIX 6.1 (64 bit) AIX 7.1 (64 bit) ⇒ *os.version and os.architecture*	✔	
	The following minimum C++ runtime levels are required: (AIX 6.1): xlC.rte 9.0.0.1 and xlC.aix61.rte 9.0.0.1 (or later) filesets (AIX 7.1): xlC.rte 11.1.0.0 and xlC.aix61.rte 11.1.0.0 (or later) filesets ⇒ *os.lib.xlC.aix61.rte and os.lib.xlC.rte*	✔	
	Note 1: The operating system levels must be the same on the PS and BS/CS servers.		
	Note 2: Virtual I/O support for SAN, see 2.1.8		
2.1.2	File systems: • JFS or JFS2, see 2.1.3 ⇒ *fcm.fileSystemType*	✔	✖
	• IBM General Parallel Filesystem (GPFS) 3.5.0.14 as a minimum level in combination with DB2 pureScale		
	• (FCM 4.1.1.2 only) IBM General Parallel Filesystem (GPFS) 4.1.0.5 as a minimum level in combination with Custom Applications		
	• NFS in combination with IBM N Series or NetApp storage. ⇒ *fcm.fileSystemType (not checked: combination with IBM N Series or NetApp)*	(✔)	✖
	• Oracle ASM		

2.1.3	JFS2 inline logs and concurrent I/O are supported.
	
	For concurrent I/O make sure that the option 'cio' is in the /etc/filesystems (check with command 'lsfs') and in the LVCB (logical volume control block) of the filesystem. This can be checked with the command getlvcb:
	
	`x1:root:/#getlvcb -AT fslv01`
	`AIX LVCB`
	`intrapolicy = m`
	`copies = 1`
	`interpolicy = m`
	`lvid = 0059d79a00004c00000000fb53c1c4d1.9`
	`lvname = fslv01`		

	Prerequisite	checked PS	checked BS/CS
	``label = /oracle/C01/sapdata7
machine id = 9D7AA4C00
number lps = 13
relocatable = y
strict = y
stripe width = 0
stripe size in exponent = 0
type = jfs2
upperbound = 32
fs =
vfs=jfs2:log=/dev/loglv10:mount=true:*options=cio,*rw:account=false
time created = Wed Mar 10 13:50:33 2004
time modified = Tue Aug 16 17:11:15 2005```		
2.1.5	AIX Base System Locale ISO Code Set - U.S. English (the command 'locale -a' must show that the locale en_US.ISO8859-1 is available) ⇒ *os.localeInstalled.en_US.ISO8859-1*	✔	
2.1.6	For N series and NetApp storage systems only: The GCC runtime library libgcc version 4 must be installed. The package can be downloaded from the IBM AIX toolbox at: • http://www.ibm.com/systems/power/software/aix/linux/toolbox/download.html ⇒ *storage.netapp*	✔	
2.1.7	If SDD is installed, the correct vpath device configuration must be checked with lsvpcfg. On the production and backup systems, all pvids containing storage-system volumes must be converted to vpath (hd2vp command, only in case of SDD). SDDPCM with AIX MPIO can be used instead of SDD. **Note:** When AIX LVM mirroring is used, then SDD must not be used. Instead AIX MPIO and SDDPCM must be used.		

2.1.8	Virtual I/O for HBA is supported through
	 - native VIOS on p5 / p6 / p7 with
	   SVC, Storwize V7000 and DS8000 with target disks mapped statically to the backup host (parameter *BACKUP_HOST_NAM*E set to *PREASSIGNED_VOLUMES*)
	 - N_Port ID Virtualization (NPIV) on p6 and p7 with
	   all storage systems and with target disks mapped statically to the backup host (parameter *BACKUP_HOST_NAM*E set to *PREASSIGNED_VOLUMES*) or mapped dynamically to the backup host (parameter *BACKUP_HOST_NAM*E set to *<hostname>*)
	6
	
	NPIV is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. This allows multiple Fibre Channel initiators to share a single physical port.		

2.1.9	Disks used with Oracle ASM need the `reserve_policy` attribute set to `no_reserve`. Change it with: `chdev -l hdisk`*n* `-a reserve_policy=no_reserve`		

Linux

	Prerequisite	checked PS	checked BS/CS
2.1.20	Red Hat Enterprise Linux 5 x64 Version 5.1 up to 5.8 Red Hat Enterprise Linux 6.1 up to 6.6 SUSE Enterprise Linux 11 x64 ⇒ *os.version and os.architecture*	✔	
	Note: The operating system levels must be the same on the PS and BS/CS.		

2.1.21	A volume manager must be installed, configured, and running.
	Volume Managers supported:
	  • LVM2
	    ⇒ *os.package.lvm2 and os.fileInfo.lvm.conf*
		✔	

	• Oracle ASM		
	Note: Refer to the following URL for *Known Issues and Limitations* • *http://www-01.ibm.com/support/docview.wss?uid=swg21650880*		

2.1.22	Filesystems:
	  • ext2,ext3 or ext4
	    ⇒ *fcm.fileSystemType*
		✔	✖

	• IBM General Parallel Filesystem (GPFS) 3.5.0.14 as a minimum level in combination with DB2 pureScale		
	• (FCM 4.1.1.2 only) IBM General Parallel Filesystem (GPFS) 4.1.0.5 as a minimum level in combination with Custom Applications		
	• NFS in combination with IBM N Series or NetApp storage. ⇒ *fcm.fileSystemType (not checked: combination with IBM N Series or NetApp)*	(✔)	✖
	• Oracle ASM		
2.1.23	Only entire disks are part of the volume groups / disk groups managed by FCM. **Note:** Disk partitions within volume groups / disk groups **are not** supported.		✖

2.1.24	Multipathing needs to be installed and configured. The according package names are:
	  • multipath-tools on SLES systems
	    ⇒ `os.package.multipath-tools`
		✔	

	• device-mapper-multipath on RHEL systems ⇒ `os.package.device-mapper-multipath`	✔	
2.1.25	Native DM-MPIO is installed. ⇒ `os.lsmod.dm_multipath`	✔	
2.1.26	The sg3_utils package needs to be installed.		

Solaris

	Prerequisite	checked PS	checked BS/CS
2.1.30	Solaris 10 SPARC Update 4, or later **Note 1:** The operating system levels must be the same on the PS and BS/PS. **Note 2:** For a proper operation of FlashCopy Manager in conjunction with DB2 it is crucial to install the following Solaris patch or newer patches accumulating it: **Patch Id:** 142900-09		
2.1.31	Volume Manager: Veritas Volume Manager (VxVM) version 5.0 or 5.1 or Oracle ASM **Note:** Refer to following URL for *Known Issues and Limitations* • *http://www.ibm.com/support/docview.wss?uid=swg21650880*		
2.1.32	File systems: UFS, Veritas File System (VxFS), Oracle ASM		✖
2.1.33	FlashCopy Manager needs to run in the global zone		
2.1.34	Native MpxIO is installed.		

HP-UX

	Prerequisite	checked PS	checked BS/CS
2.1.40	HP-UX 11iv3 (11.31) Update 3 release, or later **Note 1:** The operating system levels must be the same on the PS and BS/PS. **Note 2:** For a proper operation of FlashCopy Manager it is crucial to install the following HP-UX patch or newer patches accumulating it: **Patch Id:** PHCO_41072		
2.1.41	Volume Manager: HP LVM (BaseLVM) version B.11.31.1009 to B.11.31.1403 or Oracle ASM		
2.1.42	File systems:Veritas File System (VxFS) 4.1, or later, Oracle ASM		✖
2.1.43	Volume Group Versions 1.0, 2.0, 2.1 and 2.2 are supported		
2.1.44	Native MPIO or VxDMP is installed.		
2.1.45	Supported in combination with HP Integrity servers (Intel® Itanium®)		
2.1.46	The LVM volume group configurations for LUNs must be using the agile (persistent) naming model.		

15

	Prerequisite	checked PS	checked BS/CS
2.1.47	Prior installing FlashCopy Manager update the /etc/dld.sl.conf to include the following paths: <INSTANCE_DIR>/jre/jre/lib/IA64N <INSTANCE_DIR>/jre/jre/lib/IA64N/server <INSTANCE_DIR>/jre/jre/lib/IA64N/jli <INSTANCE_DIR>/jre/jre/lib/IA64W <INSTANCE_DIR>/jre/jre/lib/IA64W/server <INSTANCE_DIR>/jre/jre/lib/IA64W/jli Examples of the INSTANCE_DIR: DB2: <Instance owner $HOME directory>/sqllib/acs/ Oracle: <Instance owner $HOME directory>/acs/ See IBM Tivoli Storage FlashCopy Manager Installation and User's Guide for UNIX and Linux to find correct INSTANCE_DIR for specific application.		

2.2 Software Requirements - Storage Subsystems

XIV

	Perequisite	checked
2.2.1	XIV CLI 2.3.x or 3.x. or 4.2.x XIV CLI 4.1 is not supported **XCLI download page:** **http://www-01.ibm.com/support/docview.wss?rs=1319&context=STJTAG&dc=D400&q1=ssg1*&uid=ssg1S4000813&loc=en_US&cs=utf-8&lang=en**	

2.3 Other Software Requirements

Note: For DB2 pureScale environments the term server refers to a GPFS cluster as part of a DB2 pureScale instance which consists of one or more servers.

	Prerequisite	checked PS	checked BS/CS
2.3.1	For the installation of FCM, the FCM package according to the application to be protected must be available.		
2.3.2	**For FlashCopy only backup/restore:** **-** **For FlashCopy backup / restore with offloading into a storage management server:** • An IBM Tivoli Storage Manager Server (min. V6.3.x) • Alternative to TSM - a third party storage management server. FCM makes the data available on the BS. The transfer into the storage management server must be performed by third party components • Optional: TSM Storage Agent for LANfree data transfer on PS (for logs) and BS **For FlashCopy cloning:** **-**		

3. Requirements for DB2

Note: For DB2 pureScale environments, the term "server" refers to a GPFS cluster as part of a DB2 pureScale instance which consists of one or more servers.

3.1 Software requirements

The following software - if not otherwise specified - is required on the PS and BS/CS.

	Prerequisite	checked PS	checked BS/CS
3.1.1	64-bit DB2 Enterprise Server Edition V9.7; for Solaris, Fix Pack 2 or a later 9.7 Fix Pack is required. 64-bit DB2 Enterprise Server Edition V10.1 Fix Pack 2 and later 10.1.x Fix Packs 64-bit DB2 Enterprise Server Edition V10.5 and later 10.5.x Fix Packs ⇛ *fcm.db_Version*	✔	
	(Solaris) • For a proper operation of FlashCopy Manager in conjunction with DB2 it is crucial to install the following Solaris patch or newer patches accumulating it: **Patch Id:** 142900-09		
	On the BS a DB2 instance with the same name as on PS needs to be created with the same configuration as on PS.		
	(DB2 pureScale) • DB2 Enterprise Server Edition V10.5 FP4 or a later 10.5 Fix Pack is required. • Use the GPFS package included within this DB2 package. Other installed GPFS versions should be uninstalled before the DB2 package is installed. • In the backup cluster a DB2 pureScale instance with the same name and configuration like in the production cluster needs to be created.		

	Prerequisite	checked PS	checked BS/CS
3.1.2	**For FlashCopy only backup/restore:** **-** **For FlashCopy backup / restore with offloading into a storage management server:** • (SAP) an IBM Tivoli Storage Manager for Enterprise Resource Planning V6.3 can be used to backup the database redo logs on the PS and to backup the database on the BS/CS or the DB2 native TSM agent can be used to backup the database redo logs on the PS and to backup the database on the BS/CS • (non-SAP or DB2 pureScale) the DB2 native TSM agent must be used to backup the database redo logs on the PS and to backup the database on the BS/CS **For FlashCopy cloning:** **-**		

3.2 Environmental requirements

	Prerequisite	checked PS	checked BS/CS
3.2.1	All DB2 database files must reside on source volumes with supported file systems, see 2.1.2, 2.1.22, 2.1.32, 2.1.42 (raw devices are not supported) •all tablespace containers •the local database directory •the online log directory ⇒ *fcm.fileSystemType*	✔	✖
3.2.2	All table spaces must be DMS table spaces except system temporary table spaces or user temporary table spaces which can be SMS table spaces. DB2 automatic storage table spaces are also supported. ⇒ *fcm.tablespace_format*	✔	✖
3.2.3	The DB2 database must be in log retain mode.		✖
3.2.4	Make sure that the same unique UID can be used for the DB2 instance owner user on PS and BS		
3.2.5	Make sure that the same unique GID can be used for the primary group of the DB2 instance owner on PS and BS		
3.2.6	For hardware snapshots, soft links to DB2 tablespace files are supported as long as they are pointing to files residing in volume groups that can be flashed. The normal SAP setup of DB2 tablespaces files is not using links. In combination with DB2 pureScale, only soft links pointing to files in the same GPFS file set are supported.		✖
3.2.7	In DB2 pureScale environments, the backup cluster needs to be authenticated against the production cluster. On the production cluster use `mmauth add` and `mmauth grant` to allow for remote mount and write access to all DB2 GPFS database file systems from the backup cluster. On the backup cluster use `mmremotecluster` add to add the production cluster. Do not use `mmremotefs add` and `mmmount`. This will be done by FCM. For details on GPFS remote mount functionality see GPFS Advanced Administration Guide, Chapter 1.		
3.2.8	In DB2 pureScale environments, the backup cluster needs the same cluster topology as the production cluster. The same number of members and also the same member numbers are required. Cluster caching facilities (CF) do not count in this manner. Multiple DB2 members can be created on a single host as logical members, for example through the DB2 setup wizard db2isetup.		
3.2.9	In case the backup cluster of a DB2 pureScale environment consists of two or more machines the backup server identified to FlashCopy Manager must be the system where DB2 pureScale member 0 resides. Member numbers and the machines members reside on can be found in db2nodes.cfg.		

3.3 Volume group layout requirements

Note: For Solaris, the term volume group (VG) corresponds to disk groups. The following prerequisites related to volume groups are applicable to Solaris disk groups in the same way.

	Prerequisite	**checked PS**
3.3.1	Make sure that the VGs listed in 3.3.2/3/4 do **not** contain any other file system (FS) than listed below. Only the FS listed below need to be and are allowed to be included in the snapshot backup. • the FS of VG1 and VG2 can be put together into one VG • the FS of VG3 must be separated from all other listed FS •⇒ *fcm.volume_group_layout*	✔
	Important: For database instances managing multiple databases each of the databases need to be located on its own volume groups according to this layout specification.	
3.3.2	Volume group layout: **Note:** In the following rows <SID> means the database name in capital letters and <sid> means the database name in small letters. Also the FS names listed are an example for a SAP database where the DB2 instance owner is db2<sid> and its $HOME directory is /db2/<SID>. The VG names VG1, VG2 and VG3 are samples. The VGs can have any name. • VG1 (can be more than 1 VG): /db2/<SID>/sapdata1..N all FS containing data tablespaces /db2/<SID>/sapdatat or /db2/<SID>/saptemp1..N all FS containing temporary tablespaces Location: Supported storage subsystems ⇒ *fcm.volume_group_layout*	✔
3.3.3	• VG2: /db2/<SID>/db2<sid> the FS containing the local database directory (which includes /db2/<SID>/db2<sid>/NODE0000/SQL00001) Location: Supported storage subsystems ⇒ *fcm.volume_group_layout*	✔
3.3.4	• VG3: /db2/<SID>/log_dir the FS containing the online log files. Location: Supported storage subsystems ⇒ *fcm.volume_group_layout*	✔

		Prerequisite	checked PS

3.3.5		For multipartitioned databases (DPF) a dedicated set of VGs must be defined for each partition for a partition-based source- target relation, e.g.
		VGdata0000: /db2/<SID>/sapdata1..N/NODE0000
		                        /db2/<SID>/sapdatat/NODE0000
		                        /db2/<SID>/saptemp1/NODE0000
		                        /db2/<SID>/db2<sid>/NODE0000
		VGlog0000: /db2/<SID>/log_dir/NODE0000
		.
		.
		.
		VGdata<NNNN>: /db2/<SID>/sapdata1..N/NODE<NNNN>
		                        /db2/<SID>/sapdatat/NODE<NNNN>
		                        /db2/<SID>/saptemp1/NODE<NNNN>
		                        /db2/<SID>/db2<sid>/NODE<NNNN>
		VGlog<NNNN>: /db2/<SID>/log_dir/NODE<NNNN>
		
		For DB2 automatic storage make sure that the storage path resides in dedicated file systems on dedicated VGs for each partition.
		For example:
		Database AS1 (instance db2as1) is created with storage path
		/db2/AS1/sapdata1. In this case the following filesystems/VGs must exist:
		VGdata0000: /db2/AS1/sapdata1/db2as1/NODE0000
		.
		.
		.
		VGdata<NNNN>: /db2/AS1/sapdata1/db2as1/NODE<NNNN>
		⇒ *fcm.no_shared_volumes*
			✔

3.3.6		Other FS like /db2/<SID>/sqllib /db2/<SID> /db2 /opt/IBM/db2/V10.x and other application specific FS are not allowed to be on the same VGs VG1, VG2 and VG3. They are allowed to reside on local storage or on any storage system. ⇒ *fcm.volume_group_layout*	✔
3.3.7		For performance reasons the number of LUNs should be kept as low as possible.	

3.4 DB2 pureScale: GPFS File system layout requirements

Note: For DB2 pureScale, the snapshot granularity is an independent fileset. The following prerequisites are related to independent filesets. Note that each GPFS file system has a default independent fileset named `root`. That means either multiple GPFS file systems and/or independent filesests within one or more file systems can be used for separating data according to the rules listed in the table below.

	Prerequisite	**checked PS**
3.4.1	Make sure that the filesets listed in 3.4.2/3/4 do **not** contain any other directories than listed below. Only the directories listed below need to be and are allowed to be included in the snapshot backup. Fileset1 and Fileset2 can be combined into one fileset. Fileset3 must be separated from all other listed filesets. **Note:** In the following rows <SID> means the database name in capital letters and <sid> means the database name in small letters. Also the directory names listed are an example for a SAP database where the DB2 instance owner is db2<sid> and its $HOME directory is /db2/<SID>. The fileset names Fileset1, Fileset2 and Fileset3 are samples. The filesets can have any name.	
3.4.2	Fileset1 (can be more than one fileset): /db2/<SID>/sapdata1..N all directories containing data tablespaces /db2/<SID>/sapdatat or /db2/<SID>/saptemp1..N all directories containing temporary tablespaces Location: Any storage supported by GPFS	
3.4.3	Fileset2: /db2/<SID>/db2<sid> the directory containing the local database directory (which includes /db2/<SID>/db2<sid>/NODE0000/SQL00001) Location: Any storage supported by GPFS	
3.4.4	Fileset3: /db2/<SID>/log_dir the directory containing the online log files. Location: Any storage supported by GPFS	
3.4.5	For DB2 automatic storage make sure that the storage path resides in dedicated filesets.	
3.4.6	Other directories like /db2/<SID>/sqllib /db2/<SID> /db2 /opt/IBM/db2/V10.x and other application specific directories are not allowed to be on the same filesets Fileset1, Fileset2 and Fileset3.	
3.4.7	GPFS snapshots are located in the GPFS file systems. Due to the copy-on-write technology used, their sizes grow over time as the database is modified. Database administrators must keep monitoring free space in the GPFS file systems in order to avoid a file-system-full condition for the productive database. The number of GPFS snapshots kept in the file systems can be controlled with FlashCopy Manager profile parameter MAX_VERSIONS.	

3.5 NFS File system layout requirements

Note: For NFS support fo IBM N Series or NetApp storage, the snapshot granularity are independent NFS filesystems on dedicated NetApp volumes. FCM does not support the export of sub-directories from the same NetApp volume. The following prerequisites are related to independent NFS filesystems.

	Prerequisite	checked PS
3.5.1	Make sure that the NFS filesystems listed in 3.5.2/3/4 do **not** contain any other directories than listed below. Only the directories listed below need to be and are allowed to be included in the snapshot backup. ⇛ *fcm.volume_group_layout*	✔

3.5.2	Volume group layout:
	    Note: In the following rows <SID> means the database name in capital letters and <sid> means the database name in small letters. Also the directory names listed are an example for a SAP database where the DB2 instance owner is db2<sid> and its $HOME directory is /db2/<SID>. The NFS filesystems names NFS-filesystem1 , NFS-filesystem2 and NFS-filesystem3 are samples. The NFS filesystem can have any name.
	    • NFS-filesystem1 (can be more than one NFS filesystem):
	    /db2/<SID>/sapdata1..N all directories containing data tablespaces
	    /db2/<SID>/sapdatat or /db2/<SID>/saptemp1..N all directories containing temporary tablespaces
	    Location: Supported storage subsystems
	    ⇛ *fcm.volume_group_layout*
		✔

3.5.3	• NFS-filesystem2:
	    /db2/<SID>/db2<sid> the directory containing the local database directory (which includes /db2/<SID>/db2<sid>/NODE0000/SQL00001)
	    Location: Supported storage subsystems
	    ⇛ *fcm.volume_group_layout*
		✔

3.5.4	• NFS-filesystem3:
	    /db2/<SID>/log_dir the directory containing the online log files.
	    Location: Supported storage subsystems
	    ⇛ *fcm.volume_group_layout*
		✔

	Prerequisite	**checked PS**

3.5.5	For multipartitioned databases (DPF) a dedicated set of NFS filesystems must be defined for each partition for a partition-based source- target relation, e.g.
	NFS FS for partition 0: /db2/<SID>/sapdata1..N/NODE0000
	                          /db2/<SID>/sapdatat/NODE0000
	                          /db2/<SID>/saptemp1/NODE0000
	                          /db2/<SID>/db2<sid>/NODE0000
	                          /db2/<SID>/log_dir/NODE0000
	.
	.
	.
	NFS FS for other partitions: /db2/<SID>/sapdata1..N/NODE<NNNN>
	                          /db2/<SID>/sapdatat/NODE<NNNN>
	                          /db2/<SID>/saptemp1/NODE<NNNN>
	                          /db2/<SID>/db2<sid>/NODE<NNNN>
	                          /db2/<SID>/log_dir/NODE<NNNN>
	
	For DB2 automatic storage make sure that the storage path resides in dedicated file systems on dedicated VGs for each partition.
	For example:
	Database AS1 (instance db2as1) is created with storage path
	/db2/AS1/sapdata1. In this case the following NFS filesystems must exist:
	NFS FS for partition 0: /db2/AS1/sapdata1/db2as1/NODE0000
	.
	.
	.
	NFS FS for other partitions: /db2/AS1/sapdata1/db2as1/NODE<NNNN>
	⇒ *fcm.no_shared_volumes*
		✔

3.5.6	Other directories like /db2/<SID>/sqllib /db2/<SID> /db2 /opt/IBM/db2/V10.x and other application specific directories are not allowed to be on the same NFS filesystems NFS-filesystem-1, NFS-filesystem-2 and NFS-filesystem-3. They are allowed to reside on local storage or on any storage system. ⇒ *fcm.volume_group_layout*	✔

4. Requirements for SAP on Oracle

4.1 Software requirements

Note: The following software - if not otherwise specified - is required on PS and BS/CS by FlashCopy Manager.

	Prerequisite	checked PS	checked BS/CS
4.1.1	Oracle Enterprise Server Edition 11gR1 or 11gR2 - 64-bit at a level certified by the respective SAP release The following Oracle release at a level certified by the respective SAP release is currently supported: • Oracle Enterprise Server Edition 11gR2 – 64-bit ⇛ *fcm.db_Version*	✔	
	Note: If off-loading or cloning shall be used, an Oracle DBMS must be installed on BS server.	✘	
4.1.2	The following BR*Tools releases are currently supported: • BR*Tools release 7.10 (patch level 14 or higher) • BR*Tools release 7.20 (patch level 6 or higher) • BR*Tools release 7.40 **Hint:** For Oracle 11gR2 release 7.20 or 7.40 of BR*Tools is required. **Note:** FlashCopy-Manager-only interfaces with the SAP BR*Tools through the backint interface (BC-BRI volume based).		
4.1.3	**For FlashCopy only backup/restore:** - **For FlashCopy backup / restore with offloading into a storage management server:** • an IBM Tivoli Storage Manger for Enterprise Resource Planning v6.3 on PS and BS **For FlashCopy cloning:** -		

4.2 Environmental requirements

	Prerequisite	checked PS	checked BS
4.2.1	All Oracle database files must reside on source volumes in supported file systems, see 2.1.2, 2.1.22, 2.1.32, 2.1.42 (raw devices are not supported) • all tablespace data files • all control files • all redo log files (origlogA/B, mirrlogA/B) ⇒ *fcm.fileSystemType*	✔	✖
4.2.2	The SAP environment files (.sapenv_<hostname>.csh, .sapenv_<hostname>.sh) of the user ora<sid> **must not** contain statements like 'tset -I -Q'. If this statement exists in the environment files, the FlashBack Restore does not work. The software reports that a source volume is no longer attached to the production system.		
4.2.3	Make sure that the same unique UID can be used for user ora<sid> on PS and BS		
4.2.4	Make sure that the same unique GID can be used for group dba on PS and BS		

4.3 Volume group layout requirements

As a general rule: Layouts which were supported with previous releases also work with this one.

	Prerequisite	**checked PS**
4.3.1	Make sure that the VGs listed do **not** contain any other FS than those listed below. Only the FS listed below need to be and are allowed to be included in the snapshot backup. ⇒ *fcm.volume_group_layout*	✔
4.3.2	Volume group layout: **Note:** In the following rows <SID> means the database name in capital letters. The VG names VG1 and VG2 are samples. The VGs can have any name. • VG1 (can be more than 1 VG): /oracle/<SID>/sapdata1..N all FS containing tablespaces and the control files Location: Supported storage subsystems ⇒ *fcm.volume_group_layout*	✔
4.3.3	• VG2 (can be more than 1 VG): /oracle/<SID>/origlogA the FS containing the online redo logs and optionally a copy of the control file /oracle/<SID>/origlogB the FS containing the online redo logs and optionally a copy of the control file /oracle/<SID>/mirrlogA the FS containing the online redo logs and optionally a copy of the control file /oracle/<SID>/mirrlogB the FS containing the online redo logs and optionally a copy of the control file Location: Allowed to reside on local storage or on any storage system as long as only FCM backup but not FCM cloning is desired. The DB control files can be stored in the origlogA/origlogB and mirrlogA/mirrlogB as well as in sapdata1..N. ⇒ *fcm.volume_group_layout*	✔
4.3.4	**Note:** For performance reasons the number of LUNs should be kept as low as possible.	
4.3.5	Other FS like /oracle/<SID>/oraarch /oracle/<SID>/saparch /oracle/<SID>/sapbackup /oracle/<SID>/sapreorg /oracle/<SID> /oracle and other application specific FS are not allowed to be on the same VGs VG1 and VG2. They are allowed to reside on local storage or on any storage system. ⇒ *fcm.volume_group_layout*	✔

4.4 NFS File system requirements

Note: For NFS support fo IBM N Series or NetApp storage, the snapshot granularity are independent NFS filesystems on dedicated NetApp volumes. FCM does not support the export of sub-directories from the same NetApp volume. The following prerequisites are related to independent NFS filesystems.

	Prerequisite	**checked PS**
4.4.1	Make sure that the NFS filesystems listed do **not** contain any other FS than those listed below. Only the FS listed below need to be and are allowed to be included in the snapshot backup. ⇛ *fcm.volume_group_layout*	✔

4.4.2	Volume group layout:
	    Note: In the following rows <SID> means the database name in capital letters. The NFS filesystems names NFS-filesystem1 and NFS-filesystem2 are samples. The NFS filesystems can have any name.
	  • NFS-filesystem1 (can be more than 1 NFS filesystem):
	    /oracle/<SID>/sapdata1..N all FS containing tablespaces and the control files
	    Location: Supported storage subsystems
	    ⇛ *fcm.volume_group_layout*
		✔

4.4.3	• NFS-filesystem2 (can be more than 1 NFS filesystem):
	    /oracle/<SID>/origlogA the FS containing the online redo logs and optionally a copy of the control file /oracle/<SID>/origlogB the FS containing the online redo logs and optionally a copy of the control file /oracle/<SID>/mirrlogA the FS containing the online redo logs and optionally a copy of the control file /oracle/<SID>/mirrlogB the FS containing the online redo logs and optionally a copy of the control file
	
	    Location: Allowed to reside on local storage or on any storage system as long as only FCM backup but not FCM cloning is desired.
	
	    The DB control files can be stored in the origlogA/origlogB and mirrlogA/mirrlogB as well as in sapdata1..N.
	
	    ⇛ *fcm.volume_group_layout*
		✔

4.4.4	The FS listed here are not allowed to be located on the same NFS filesystems NFS-filesystem1 and NFS-filesystem2:
	  • /oracle/<SID>/oraarch
	  • /oracle/<SID>/saparch
	  • /oracle/<SID>/sapbackup
	  • /oracle/<SID>/sapreorg
	  • /oracle/<SID>
	  • /oracleVGs
	and other application specific FS. They are allowed to reside on local storage or on any storage system.
	⇛ *fcm.volume_group_layout*
		✔

5. Requirements for Oracle (non-SAP)

5.1 Software requirements

The following software - if not otherwise specified - is required on both the production and backup systems by DataProtection for

FlashCopy:

	Prerequisite	checked PS	checked BS/CS
5.1.1	The following Oracle releases are currently supported: • Oracle Database Standard Edition 11gR1 – 64-bit • Oracle Database Standard Edition 11gR2 – 64-bit • Oracle Database Enterprise Server Edition 11gR1 – 64-bit • Oracle Database Enterprise Server Edition 11gR2 – 64-bit • Oracle Real Application Clusters (RAC) option on supported Oracle Database releases ⇒ *fcm.db_Version*	✔	
	Note 1: Oracle Real Application Clusters (RAC) is only supported in combination with Oracle ASM only. Please also consider *Known Issues and Limitations* on following URL: • http://www.ibm.com/support/docview.wss?uid=swg21650880		
	Note 2: An Oracle DBMS must be installed on the backup server for offloaded tape backups. An Oracle DBMS does not need to be installed on the BS for disk-only backups.	✖	
5.1.2	**For FlashCopy only backup / restore**: - **For FlashCopy backup / restore with offloading to tape**: • Tivoli Storage Manager for Databases - Oracle version v6.3 or later on PS and BS **For FlashCopy cloning**: -		
5.1.3	For ASM on AIX, the following patches are required: Oracle Enterprise Server 11gR1: • Patch 7225720 (included in 11.1.0.7.1) • Patch 8471165 Oracle Enterprise Server 11gR2: • Patch 11785938 (included in 11.2.0.3) • Patch 11666137 (included in 11.2.0.3) • Patch 13366202	✔	

⇒ *fcm.oracleASMPatches*			

5.2 Environmental requirements

	Prerequisite	checked PS	checked BS/CS

5.2.1	All Oracle database files must reside
	  • on source volumes in supported file systems, see 2.1.2, 2.1.22, 2.1.32, 2.1.42
	    ⇒ *fcm.fileSystemType*
		✔	✖

	• or Oracle ASM storage (AIX only)		
5.2.2	Make sure that the same unique UID can be used for the Oracle instance owner user on PS and BS. One dedicated instance owner user is required per database instance.		
5.2.3	Make sure that the same unique GID can be used for the primary group of the Oracle instance owner on PS and BS		
5.2.4	RMAN catalog database needs to be configured (even for disk only backups). It should not reside on the PS (because it would be lost in the case of a disaster) but must be accessible from both PS and BS. ⇒ *fcm.catalog_db_connection*	✔	
5.2.5	(Linux) On Linux FCM supports only ASM environments setup with Oracle ASMLib since Oracle recommends ASMLib over raw devices for ease-of-use and performance reasons. Raw devices are not supported. ASMLib 2.0 is required by FCM.		
5.2.6	(Linux) For multipathing to work properly the parameters ORACLEASM_SCANORDER and ORACLEASM_SCANEXCLUDE need to be set in the ASMLib configuration. For example exclude single paths from the scan by setting ORACLEASM_SCANEXCLUDE="sd" and force multipath devices by setting ORACLEASM_SCANORDER="dm-*". For details see section "Configuring the Scan Order" on this page: http://www.oracle.com/technetwork/topics/linux/multipath-097959.html		
5.2.7	(Linux) In Oracle ASM environments with ASMLib on Linux, the ASM_DISKSTRING parameter is required to be set to 'ORCL:*' on production and backup system. For further information see: http://docs.oracle.com/cd/E11882_01/install.112/e41961/storage.htm#CWLIN293		
5.2.8	(AIX, Solaris, HPUX) For Oracle ASM the setting of the ASM_DISKSTRING parameter has to be left to the default for the backup server.		
5.2.9	For Oracle ASM diskgroups are created on entire disks. (Setups on disk partitions are not supported.)		✖
5.2.10	For Oracle Real Application Clusters (RAC) databases, the database instance owner user is required to be member of the group: asmadmin		
5.2.11	For Oracle Real Application Clusters (RAC) databases, a shared file system across all RAC nodes of the the production database is required. It is used to share the FCM binaries, configuration and repository.		✖
5.2.12	Database file names must not contain blanks, tab or new line characters.		✖

5.3 Volume group layout requirements

	Prerequisite	checked PS
5.3.1	One or more Oracle control file copies **can** be placed in the same VG as the data files. If they are, follow the recovery procedure described in the user's guide chapter "Restoring Oracle snapshots" after a restore operation. Make sure that the VGs listed do **not** contain any other FS than those listed below. Only the FS listed below need to be and are allowed to be included in the snapshot backup. ⇒ *fcm.volume_group_layout*	✔
5.3.2	Volume group layout: **Note:** In the following rows <SID> means the database name in capital letters. The VG names VG1 and VG2 are samples. The VGs can have any name. • VG1 (can be more than 1 VG): /oracle/<SID>/data1..N all FS containing tablespaces Location: Supported storage subsystems ⇒ *fcm.volume_group_layout*	✔
5.3.3	• VG2 (can be more than 1 VG): containing the online redo logs and optionally a copy of the control file Location: Allowed to reside on local storage or on any storage system as long as only FCM backup but not FCM cloning is desired. ⇒ *fcm.volume_group_layout*	✔
5.3.4	For performance reasons the number of LUNs should be kept as low as possible.	
5.3.5	The FS listed here are not allowed to be located on the VGs (VG1,VG2) where data, log and control files are located and must be explicitly separated: • /oracle/<SID>/oraarch • /oracle/<SID> • /oracle They are allowed to reside on local storage or on any storage system. ⇒ *fcm.volume_group_layout*	✔

5.4 Disk Group layout requirements (Oracle ASM)

	Prerequisite	checked
5.4.1	One or more Oracle control file copies **can** be placed in the same disk group as the data files. If they are, follow the recovery procedure described in the user's guide chapter "Restoring Oracle snapshots" after a restore operation. ⇒ *fcm.volume_group_layout*	✔
5.4.2	Disk Group DG1 (can be more than 1 DG): One or multiple disk groups containing the tablespace files. Location: Supported storage subsystems ⇒ *fcm.volume_group_layout*	✔
5.4.3	DG2 (can be more than 1 DG): One or more additional disk groups containing the online redo logs and control files. Location: Allowed to reside on local storage or on any storage system and also in regular file systems rather than ASM disk groups as long as only FCM backup but not FCM cloning is desired. ⇒ *fcm.volume_group_layout*	✔
5.4.4	Disk groups must not be shared between multiple database instances. ⇒ *fcm.extrafiles*	✔
5.4.5	On the backup system (if used) all disk groups that are specified with parameter `control_files` as control file locations in the database profile need to be created manually for the control file restore that is part of the TSM backup operation.	
5.4.5	Note. For performance reasons the number of LUNs should be kept as low as possible.	

5.5 NFS File system layout requirements

Note: For NFS support fo IBM N Series or NetApp storage, the snapshot granularity are independent NFS filesystems on dedicated NetApp volumes. FCM does not support the export of sub-directories from the same NetApp volume. The following prerequisites are related to independent NFS filesystems.

	Prerequisite	checked
5.5.1	One or more Oracle control file copies **can** be placed in the same NFS filesystems as the data files. If they are, follow the recovery procedure described in the user's guide chapter "Restoring Oracle snapshots" after a restore operation. Make sure that the NFS filesystems listed do **not** contain any other FS than those listed below. Only the FS listed below need to be and are allowed to be included in the FlashCopy to the backup system. ⇒ *fcm.volume_group_layout*	✔

5.5.2	Volume group layout:
	Note: In the following rows <SID> means the database name in capital letters.
	The NFS filesystems names NFS-filesystem1 and NFS-filesystem2 are samples. The NFS filesystems can have any name.
	   • NFS-filesystem1 (can be more than 1 NFS filesystem):
	     /oracle/<SID>/data1..N all FS containing tablespaces
	     Location: Supported storage subsystems
	     ⇒ *fcm.volume_group_layout*
		✔

5.5.3	• NFS-filesystem2 (can be more than 1 NFS filesystem):
	     containing the online redo logs and optionally a copy of the control file
	     Location: Supported storage subsystems
	     ⇒ *fcm.volume_group_layout*
		✔

5.5.4	The FS listed here are not allowed to be located on the NFS filesystems NFS-filesystem1 and NFS-filesystem2: • /oracle/<SID>/oraarch • /oracle/<SID> • /oracle They are allowed to reside on local storage or on any storage system ⇒ *fcm.volume_group_layout*	✔

6. Requirements for Custom Applications

6.1 Software requirements

In principle, FlashCopy Manager for Custom Applications just creates snapshot copies of plain filesystems. However, the data residing on these filesystems is usually owned by some kind of application (e.g. a DBMS, a FTP server, or a Web-Server). To get a transaction-consistent snapshot copy these custom applications may need to be suspended before and resumed after the snapshot. This can be realized by customized scripts specific to your application which are invoked by FlashCopy Manager at the right time.

	Prerequisite	checked PS	checked BS/CS
6.1.1	Your custom application is properly installed.		✖
6.1.2	If applicable: Pre- and post processing scripts to suspend / resume your custom application for the snapshot copy process are prepared, tested and accessible by FlashCopy Manager.		✖
6.1.3	**For FlashCopy only backup / restore**: - **For FlashCopy backup / restore with offloading to tape**: • Tivoli Storage Manager Backup Archive Client version v6.3 or later on PS and BS • Alternative to TSM - a third party storage management server. FCM makes the data available on the BS. The transfer into the storage management server must be performed by third party components **For FlashCopy cloning**: • not supported		

6.2 Environmental requirements

	Prerequisite	checked PS	checked BS/CS
6.2.1	All data files must reside on source volumes in supported file systems, see 2.1.2, 2.1.22, 2.1.32, 2.1.42		✖
6.2.2	Make sure that the same unique UID can be used on PS and BS to access the application data.		
6.2.3	Make sure that the same unique GID can be used on PS and BS to access the application data.		

6.3 Volume group layout requirements

	Prerequisite	checked PS
6.3.1	Ensure that no alien files are located on the filesystems being subject to the snapshot backup.	
6.3.2	Ensure that the corresponding volumes of the protected filesystems are sole members of a dedicated volume group. Alien volumes **must not** be included in the same VG.	

6.4 NFS File system layout requirements

Note: For NFS support fo IBM N Series or NetApp storage, the snapshot granularity are independent NFS filesystems on dedicated NetApp volumes. FCM does not support the export of sub-directories from the same NetApp volume. The following prerequisites are related to independent NFS filesystems.

	Prerequisite	checked PS
6.4.1	Ensure that no alien files are located on the NFS filesystems being subject to the snapshot backup.	

6.5 GPFS File system layout requirements

FlashCopy Manager for Unix and Linux V4.1.1.2 adds support for GPFS file system snapshots for Custom Applications. This support is limited to local snapshot backup and restore only. Offloaded backups to TSM are not supported with FCM V4.1.1.2.

Note: For Custom Applications, the snapshot granularity is an independent fileset. The following prerequisites are related to independent filesets. Note that each GPFS file system has a default independent fileset named root. That means either multiple GPFS file systems and/or independent filesests within one or more file systems can be used for separating data according to the rules listed in the table below.

	Prerequisite	checked PS
6.5.1	Ensure that no alien files are located on the GPFS filesets being subject to the snapshot backup.	
6.5.2	GPFS snapshots are located in the GPFS file systems. Due to the copy-on-write technology used, their sizes grow over time as the database is modified. Database administrators must keep monitoring free space in the GPFS file systems in order to avoid a file-system-full condition for the productive database. The number of GPFS snapshots kept in the file systems can be controlled with FlashCopy Manager profile parameter MAX_VERSIONS.	

7. Requirements for FlashCopy Cloning

To create clones of databases the same requirements apply as mentioned in the previous chapters. The following chapter provides additional requirements especially for FlashCopy Cloning operations.

Note 1: FlashCopy cloning is not available in combination with DB2 pureScale.

Note 2: FlashCopy cloning is not supported for NFS filesystems.

Note 3: FlashCopy cloning for Oracle ASM is not supported on HP-UX and Solaris platforms

Note 4: FlashCopy cloning for Oracle ASM is not supported on Oracle 11gR1

7.1 Environmental requirements (DB2)

	Prerequisite	**checked PS**	**checked BS/CS**
7.1.1	On the cloning system a DB2 instance with the same name like on PS needs to be created with the same configuration like on PS (backup instance is named in the following section). Even if a cloning system is used for cloning only the DB2 database instance of the backup instance must exist. Make sure the DB2 database manager configuration parameter DFTDBPATH points to the correct location.	✖	
7.1.2	In addition to the DB2 instance with the same name like on PS a second DB2 instance with the clone instance name needs to be created on the clone system. Make sure the DB2 database manager configuration parameter DFTDBPATH points to the correct location.	✖	
7.1.3	(SAP only) On the cloning server the same SAP release level needs to be installed and a SAP clone instance needs to be created prior to the first FlashCopy Cloning.	✖	
7.1.4	The clone database instance owner must be a member of the primary group of the backup instance owner (see 7.1.1).	✖	
7.1.5	For DPF environments the clone database instance must be configured with the same amount of nodes like the production and backup database instances.	✖	
7.1.6	(SAP only) For SAP Kernel releases >= 6.10: On the cloning system the environment variable • dbs_db6_schema must be set correctly for user db2<sid> and <sid>adm to the production system database schema. Normally this should be 'SAP<SID>'.		

7.2 Environmental requirements (Oracle SAP and non-SAP)

	Prerequisite	checked PS	checked BS/CS
7.2.1	Oracle DBMS needs to be installed on the cloning server. Also an Oracle clone target database instance owner user must be created. It must be in the same groups as the instance owner user of the production instance.	✖	
7.2.2	If a system is used for backups and clones the instance owner of the backup and clone instance must be in the same groups or the clone database instance owner must be a member of the primary group of the backup instance owner.	✖	
7.2.3	(SAP only): For SAP Kernel releases >= 6.10: On the cloning system the environment variable • dbs_ora_schema must be set correctly for user ora<sid> and <sid>adm to the production system database schema. Normally this should be 'SAP<SID>'.		
7.2.4	Make sure that the instance owner of the cloning instance on the cloning system uses ksh or bash as login shell. Otherwise the cloning process will fail because of problems in environment variable substitution.	✖	
7.2.5	If an Oracle spfile does exist on the clone server it must be removed. Such file may be created when a full SAP instance is installed on the clone server to prepare the environment for the clone. The default name of the spfile is ${ORACLE_HOME}/dbs/spfile$ {SID}.ora.	✖	
7.2.6	If the database is configured to use a pfile, ensure this pfile does not contain any IFILE statements. Such statements are not supported.		✖
7.2.7	Oracle ASM: The clone target database instance owner must be member of the group: asmdba.	✖	
7.2.8	Oracle ASM: Cloning is not supported on Oracle 11gR1		

7.3 Volume group layout requirements for SAP on Oracle

	Prerequisite	checked PS	checked BS/CS
7.3.1	The volume group(s) listed in 4.3.3 containing online redo logs are required to be located on supported disk subsystems.		✖

7.4 Volume group layout requirements for Oracle (non-SAP)

	Prerequisite	checked PS
7.4.1	The volume group(s) listed in 5.3.3 containing online redo logs are required to be located on supported disk subsystems.	
7.4.2	At least one control file must be located on the same volume groups as the online redo log files. ⇒ *fcm.volumes_for_cloning*	✔

7.5 Disk Group layout requirements for Oracle ASM

	Prerequisite	checked PS
7.5.1	FlashCopy Cloning of an Oracle database that resides on Oracle ASM is supported for AIX and Linux starting with FCM version 4.1.1. There is no support for HP-UX and Solaris.	
7.5.2	The disk group(s) listed in 5.4.3 containing online redo logs are required to be located on supported disk subsystems.	
7.5.3	All control file copies must be located on ASM disk groups on supported storage subsystems.	

7.6 DB2 pureScale

	Prerequisite	checked PS
7.6.1	FlashCopy Cloning of an DB2 pureScale database is not supported	

7.7 Pre- and Post-Processing Scripts Requirements

	Prerequisite	checked PS	checked BS/CS

| 7.7.1 | All pre- and post-processing configuration files and the scripts that are defined in the configuration files must meet the following requirements:

• They must exist on the clone system.
• They must have a minimum permission of read and execute access for the clone database instance owner. The preprocessing and postprocessing scripts must have a minimum permission of read and execute access for the user who owns the scripts and who is the operator of the scripts. Furthermore, if the scripts are intended to be run by any user registered on the system, the scripts must be owned by the root user and the must have a minimum permission of read and execute for the User, Group, and World user group.
• They must have a minimum permission of read access for the production database instance owner.
• They must exists in the same directory.
• They must be valid files. Symbolic links are not allowed.
• They must be specified with a fully qualified path and name.
• SQL scripts (DB2 SQL or Oracle SQL) must have the extension „.sql".
• Shell scripts must have the extension „.sh". | ✖ | |

8. Environmental Requirements

	Prerequisite	checked PS	checked BS/CS

8.0.1	The ulimits of the DB2 instance owner (DB2 UDB), Oracle instance owner (Oracle) user or custom application owner user and root on the production and backup system should at a minimum be set to (check with ulimit -a):
	
	`data seg size (kbytes)     unlimited`
	`max memory size (kbytes)   131000`
	`stack size (kbytes)        131000`
	
	Depending on the user's shell and OS level, the output of the command ulimit -a can vary.		

8.0.2	The environment files (e.g. .profile, .login, etc.) of the database instance user **must not** contain statements like 'tset -I -Q', any tty-statements like 'stty', 'tty' or 'ttytype' and statements like 'tabs'.		
8.0.3	The time on all hosts must be synchronized.		

8.1 AIX LVM Mirroring Requirements

	Prerequisite	checked PS
8.1.1	IBM FlashCopy Manager for AIX supports AIX LVM mirrored environments with the following specifications: • Only two copy (mirror) sets are allowed • One complete copy set must reside in one single storage subsystem (or SVC cluster). It is not allowed to spread a copy set across multiple subsystems with LVM mirroring. • The other complete copy set must reside in the second storage subsystem (or SVC cluster).	
8.1.2	For all logical volumes in one of the volume groups involved in the FlashCopy, check the following: • mirror write consistency must be set on and must be active • mirror write scheduling policy must be set to parallel write (with parallel, sequential or round robin read) or striped. In addition, for all jfslog or jfs2log volumes in one of the involved volume groups the mirror write scheduling policy must be parallel.	
8.1.3	For each of the involved volume groups check the following: • quorum must be set off	
8.1.4	For performance reasons the number of LUNs should be kept small.	

Trademarks

The following terms are trademarks or registered trademarks of IBM corporation in the United States, other countries, or both:

AIX, DB2, GPFS, DS6000, DS8000, XIV, SVC, IBM Storwize V7000, eServer, FlashCopy, HACMP, IBM, Passport Advantage, POWER, pSeries, System p, Tivoli, TotalStorage, WebSphere.

Oracle and Solaris are registered trademarks of Oracle Corporation and/or its affiliates.

SAP, SAP NetWeaver, and mySAP are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Symantec and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries.

UNIX is a registered trademark of The Open Group.

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

HP-UX is a product name of Hewlett-Packard Company.

Intel and Itanium are trademarks of Intel Corporation in the U.S. and/or other countries.

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