

ARKEIA 5 USER MANUAL



Document version:

Title: Arkeia 5 User Manual	Version: 2002-11-29 – 1.25 C
-----------------------------	------------------------------

TABLE OF CONTENTS

Introduction	2
Typographical conventions	3
Who should read this manual?	4
Copyright notice.....	5
1: Installation pre-requisites	6
Overview	7
Available clients and servers	7
Hardware requirements	9
Memory	9
Network bandwidth, configuration and adapters.....	9
SCSI adapters, tape drives and libraries	10
Platform-specific requirements	14
Overview	14
Compaq Tru64 UNIX/Digital Unix DEC OSF	14
Hewlett-Packard HP-UX.....	15
IBM AIX	16
LINUX.....	17
SGI IRIX	18
Sun Solaris	19
2: Installation procedure	21
Installation under the Linux operating system.	22
Installation using the “.rpm” file	22
Installation using the “.tar.gz” file.....	24
Installation check	27
Post-installation root password	29
Installation of Arkeia for a client-only machine	30
Overview	30
Configuration of an Arkeia client	30
Arkeia license management	32
Overview	32
Entering the license number through the GUI	32
Additional functionalities that require a separate license	33
3: Hardware installation	35
Overview of hardware management.....	36
Interface and host adapter supported	36
Supported backup media	36
Supported tape drives.....	38
Handling and replacement of backup media	38
Advanced backup hardware	39
Basic SCSI configuration	41
Overview of the SCSI interfaces	41
SCSI device management under Linux	41
Quick SCSI setup and troubleshooting	42
Basic installation procedure	46

4: Basic Concepts	50
Overview	51
Architectural concepts.....	52
Client/Server architecture	52
Three-tier architecture	53
Heterogeneous networks.	54
Arkeia Internal Architecture	57
Modular architecture.....	57
Multiple data flows	57
Savepacks.....	58
Tapes and tape pools	58
Drivepacks	60
5: GUI Concepts.	62
Overview	63
The Window functions	64
Window controls.....	64
Information fields	65
The list functions.....	67
A short example: the main menu	69
6: Objects and users creation	71
Creating tapes	72
Log on the Arkeia server	72
Create the tape(s)	74
Tape types.....	77
Creating tape pools	79
Log on the Arkeia server	79
How to create a tape pool	79
Checking a tape pool statistics.....	82
Creating savepacks	83
Overview	83
Creating an empty Savepack	83
Adding trees to a Savepack	85
Defining options for the Savepack	91
Setting up the options of a tree.....	95
Creating users	99
Overview	99
Arkeia users categories	99
How to create users under Arkeia	100
Creating Drivepacks.....	103
Overview	103
How to create Drivepacks.....	103
7: Interactive backup	106
Overview	107
How to access the Interactive backup	108
Main functions of the interactive backup.....	110
Main functions: Archive backup only	113
Main functions: Incremental backup only	114
Monitoring a running interactive backup	116
Monitoring a backup through the main screen	116

Monitoring a backup through the list of operations.....	118
Modifying a running backup	120
Adding a Savepack to a running backup	120
Setting the operating mode of a running backup	121
8: Periodic backup	123
How to access the Periodic backup menu	124
Overview: configuration of a periodic backup.....	126
Creating a periodic backup	127
Step one – creating the backup itself.....	127
Step 2 – Defining the periodicity of the backup	128
Step 3 – Configuration of the backup	132
Step 4 – Defining exceptions and advanced options.....	135
9: Data restoration.....	138
Accessing the Restoration function	139
The Restoration menu	140
File selection	142
File selection with the Database Navigator	142
File selection through the file search	145
Setting the options of a Restoration	149
Accessing the options	149
Configuring the options	149
Applying the options to the Restoration	150
Redirecting a Restoration.....	151
Accessing the redirection function.....	151
Options of the redirection.....	151
Arkeia path syntax	152
Applying the redirection	154
Tape analysis and launching a Restoration	155
Tape analysis overview	155
Using the Tape analysis	155
Launching a Restoration	156
10: Advanced Hardware.....	158
Overview	159
Differences between libraries and other backup devices.	159
Libraries managed by Arkeia	159
Advantages of the library over other backup devices	160
Creation of a library under Arkeia version 5	161
Creation of the library.....	161
Adding and cleaning drives into a library	163
Adding tapes to a library (slot assignment)	165
Managing a library under Arkeia version 5	167
Starting and stopping a library	167
Using a library for a backup	167
Tape operations under Arkeia version 5	168
A special case: the file libraries.....	168
11: Optimization.....	169
Overview	170
Average network performance.....	170
Use of multiple flows.....	171

Overview	171
Parallelism of a single machine and of multiple machines ...	171
How to limit backup speed.....	172
Why limit the backup speed?	172
How to limit backup speed.....	172
Priority	174
Overview	174
Usage	174
Chaining	175
Overview	175
Usage	175
Configuring compression in Arkeia.....	176
Overview	176
Compression settings	176
12: Arkeia Utilities	178
Extracting data from tape with readarkeia.....	179
Overview	179
Obtaining readarkeia.....	179
Using readarkeia	180
List the content of a given tape.....	180
Extract all the files on a tape.....	180
Extract a specific file.....	181
Recovering the Arkeia database with arkrstdb	182
Overview	182
General usage and syntax	182
Options	182
Automatic hardware detection with arkdetect	184
Overview	184
Running arkdetect	184
Analyzing the results of arkdetect.....	186
Additional options of "arkdetect"	190
Tape duplication with arkc.....	191
Overview	191
Overview of the arkc syntax	192
Testing the Arkeia server connectivity with chknlp	194
Overview	194
Usage	194
chknlp test description.....	194
Renaming servers/clients with arkmvhost	198
Overview	198
Syntax of arkmvhost	198
Listing all the (modified) clients in the Arkeia database ..	198
Complete update of the Arkeia database with arkmvhost ...	199
Renaming a single client in the Arkeia database	199
Common options of arkmvhost.....	200
Removing a host with arkrmhost	201
Overview	201
Syntax of "arkrmhost"	201
Checking the database integrity with arkdbchk	202

Overview	202
Usage and syntax.....	202
Examples.....	203
Stopping and starting Arkeia with arkboot.....	205
Overview	205
Usage, configuration and syntax.....	205
13: Disaster Recovery	207
Introduction to Disaster Recovery	208
Overview	208
Disaster Recovery issues	209
Preparing for a Disaster Recovery.....	211
Creating a specific Disaster Recovery configuration	211
How to backup a client/server parameters	213
Setting up a comprehensive backup procedure	216
Disaster Recovery procedure for a client	217
Hardware limitations.....	217
Disaster Recovery – step one	218
Disaster Recovery – step two	221
Disaster Recovery – step three	225
Disaster Recovery – step four.....	226
Disaster Recovery procedure for a server	228
Overview	228
Disaster Recovery – step one	228
Disaster Recovery – step two	229
Disaster Recovery – step three	229
Disaster Recovery – step four.....	230
14: Advanced network configuration.....	231
Name resolution and servers with multiple NICs	232
Overview	232
Configuration for an Arkeia client.....	232
Configuration for an Arkeia server.....	232
Hosts file used by Arkeia	233
Setting up “NLP_HOSTFILE” and “NLP_ONLYHOSTFILE”	233
Syntax of the Arkeia-specific “hosts” file.....	234
Usage of “NLP_HOSTNAME”	234
Configuring Arkeia not to use FQDN	234
How to use different TCP ports with Arkeia	236
Configuring Arkeia to work from behind a firewall	237
Introduction	237
Standard configuration.....	237
Using the port ranges.....	238
SSH configuration.....	239
Using Arkeia and DHCP	241
Overview	241
Using the variable “VIRTUAL_HOSTNAME”	241
Appendix 1: MySQL Backup.....	242
Overview	243
Configuration of the Arkeia plug-in	244
The “arkmysql.pr#” preference file.	244

Configuration of the log verbosity.	244
Configuration of the UNIX socket used by the plug-in ...	245
Configuration of the host name or TCP/IP address.	245
Configuration of the MySQL user ID.	245
Configuration of the MySQL password.	245
Configuration of the TCP/IP port.	245
Using the plug-in to backup and restore MySQL	246
Backup with the MySQL plug-in.....	246
Restoration with the MySQL plug-in.....	247
Appendix 2: DB2 Backup	248
Introduction	249
How can Arkeia be used to backup an IBM DB2 Database?	249
Arkeia for IBM DB2 – requirements.	249
Offline Backup Configuration	250
“Offline” (Cold) Backup of the DB2 Database	250
Prerequisite for an “Offline” backup.....	250
Arkeia Configuration for “Offline” backup.....	251
“Hot” Backup Configuration	253
Online Backup of DB2 Database	253
Setting the DB2 instance logging to “ON”	253
Creating the Backup Scripts.....	253
The “create_lock.sh”	254
The “delete_lock.sh” script	255
The “db2_backup.sh” script.....	256
The “db2_backup_cat.sh” script.....	257
The “db2_restore.sh” script.....	258
The “db2_restore_cat.sh” script.....	259
The “db2_check.sh” script	260
Arkeia Configuration for “Online” Backup.....	262
Advanced Tree Configuration.....	263
Appendix 3: Migrating to version 5	265
Overview of the migration tool.....	266
Obtaining the migration tool	267
Running the migration tool on a server	268
Keep a copy of the Arkeia version 4.2 configuration	268
Required operations before the migration.....	268
Install Arkeia version 5 (“.rpm” only)	269
Running the migration tool.	269
Tasks performed by the migration tool	270
Tasks to be performed after the migration	271
Testing the migrated information.	271
Migrating a client machine.	272
Known problems and issues.	274
Obtaining help on the migration tool.	276
Appendix 4: Microsoft Windows backup	277
Overview	278
Installation	279
Backup	280
Backup of the Registry through the “sysinfo” plug-in	280

Backup of the Windows Active directory	281
Backup of the Certificates through the “sysinfo” plug-in.	282
Backup of the Event Logs through the “sysinfo” plug-in.	283
Restoration of the Windows system information	284
Restoring the Windows Registry with the “sysinfo” plug-in. .	284
Restoring the Active Directory with the “sysinfo” plug-in.	284
Restoring the Certificates with the “sysinfo/” plug-in.	285
Restoring the Event Logs with the “sysinfo/” plug-in.	285
Index	286

INTRODUCTION

Welcome to the Arkeia 5 User Manual!

Thank you for purchasing Arkeia, the leading software for heterogeneous network backup! We hope that Arkeia version 5 will answer your needs for fast and reliable backup under many Operating Systems.

This manual is designed to help you install, configure and use Arkeia on the different machines of your network. Because Arkeia is a full-featured software, reading this manual before any installation and operation is highly recommended for all users.

For your convenience, an Index and a Table of Contents are available at the end of this manual. If you downloaded this manual from the Internet, in the form of a PDF (*“Portable Document Format”*) file, the Table of Contents should also be displayed on the left of this page. Clicking on a Table of Contents item will then take you to the relevant page of the manual.

You will find, on the very first page of this manual, a version number for this document. If the manual you are using appears to be outdated or incomplete, download the latest version by connecting to the Arkeia Corporation web site, at the following URL:

`http://www.arkeia.com/download.html`

You may notice that this document contains many links like the one above, as well as many reference to other chapters of the manual. If you click on of these links, your PDF viewer should either take you to the correct section of the manual, or the default world-wide web browser of your system should open the correct web page (provided you are connected to the Internet).

Finally, we would like to hear from you! In order to improve the overall quality of the documentation, an email address is available. If you have any remarks or corrections, feel free to send an email to: `doc@arkeia.com`. Please note that this address is for the documentation-related questions only. All technical inquiries should be sent to the Arkeia Corporation Technical Support staff.

1 • Typographical conventions

You may want to take note of the following typographical conventions, which will help you find the information you require in a more convenient way.

In this manual, the following paragraph contains important information, that applies to all versions of Arkeia. You should always read this paragraph and determine if the information applies to your Arkeia installation and/or to your network:



Please note: this information is important, and should be read carefully. It may be necessary to apply the information contained in this paragraph to your network, or Arkeia server, in order to successfully perform a backup or a restoration.

The following paragraph contains information that may, or may not, apply to your situation. Nevertheless, the information it contains should be read carefully.



Please note: the information contained in this paragraph may or may not apply to your installation. Nevertheless, it should be reviewed carefully and understood before proceeding with the rest of the user manual.

The following font is used to indicate the result of a command that should be entered on a command line, as well as the results of said command:

```
$ ls -Al
drwxr-xr-x  2 root    root          2048 Dec  6 23:03 bin
drwxr-xr-x  3 root    root          1024 Sep  1  1999 boot
drwxr-xr-x  5 root    root        34816 Feb 25 04:02 dev
drwxr-xr-x 31 root    root          3072 Feb 27 17:28 etc
drwxr-xr-x 19 root    root          1024 Feb  7 18:49 home
drwxr-xr-x  4 root    root          3072 Sep  1  1999 lib
drwxr-xr-x  2 root    root        12288 Sep  1  1999 lost+found
drwxr-xr-x  9 root    root          1024 Sep 29 18:26 mnt
dr-xr-xr-x 61 root    root           0 Feb  6 11:51 proc
drwxr-xr-x  7 root    root          1024 Nov 29 12:29 root
drwxr-xr-x  3 root    root          2048 Aug  9  2000 sbin
drwxrwxrwt 11 root    root          1024 Feb 28 17:01 tmp
drwxr-xr-x 23 root    root          1024 Nov 23  1999 usr
drwxr-xr-x 18 root    root          1024 Sep  1  1999 var
```

In the above example, the command “`ls -Al`” has been entered in the root (“/”) directory of a UNIX machine. The lines under the command represent its output, which is displayed on the screen of the machine itself.

Finally, if you’d like to print this documentation to keep a copy for reference, please use the “Fit to page” option of your PDF software if you are not using the standard European “A4” page format.

2 • Who should read this manual?

This manual is designed, first and foremost, for the network and system administrators who have to install, manage and maintain the machines and the network(s) backed-up by Arkeia.

Most administrators will therefore benefit greatly by reading it carefully before any installation or configuration of Arkeia. Most errors and problems that may occur during the use of Arkeia can be avoided (or solved) by reading the relevant section or chapter of this manual.

Many (non-administrator) users of Arkeia should read the chapter “4: Basic Concepts” on page 50 to understand the main concepts that are used by the software, especially the distinction between a server, a client and an interface. Also of interest to these users are the explanations found in “5: GUI Concepts.” on page 62, on the functions offered by the graphical interface, as well as chapters “6: Objects and users creation” on page 71, “7: Interactive backup” on page 106, and “9: Data restoration” on page 138.

3 • Copyright notice

Microsoft, Exchange, “*MS Exchange*”, Microsoft Office, Windows, Windows 98, Windows 95, Windows NT, Windows ME, Windows 2000 and Windows XP are trademarks and Copyright © by Microsoft Corporation Inc., All Rights Reserved.

Sun, Sun Solaris, SPARC, Sun Microsystems and the “*Sun*” logo are trademarks and Copyright © by Sun Microsystems, All Rights Reserved.

Linux is Copyright © by Linus Torvalds, All Rights Reserved.

MySQL is a registered trademark of MySQL AB in Sweden, and a trademark of MySQL AB in the United States of America and other countries, All Rights Reserved.

All other logos and trademarks are Copyright by their respective owners.

This manual, including all the figures and screenshots it contains, is Copyright © 2001, 2002 by Arkeia Corporation, All Rights Reserved.

CHAPTER 1: INSTALLATION PRE-REQUISITES

- *Summary:*

This chapter details the configuration required to install Arkeia version 5 on different machines and operating systems. Make sure you read this chapter and check that the configuration of your servers matches the requirements of Arkeia.

1 • Overview

Arkeia is designed to perform backup and restore tasks with numerous operating systems and hardware platforms. New platforms are added regularly, and you may want to check the most up-to-date list, which is available on the World Wide Web, at the following address:

<http://www.arkeia.com/platforms.html>

1.1 Available clients and servers

You will find below a list of clients and servers available for Arkeia version 5. This list contains the following information:

- **Server:** an Arkeia version 5 server exists for this operating system. This server is able to manage backup and restore operations on this particular platform.
- **Client:** an Arkeia version 5 client exists for this operating system. By communicating with an Arkeia server, the client is able to save (and restore) the data of its host machine to the server.
- **GUI:** an Arkeia version 5 Graphical User Interface (GUI) exists for this operating system. The GUI allows you to use a machine of this type as a management workstation, which can connect to different servers and configure their operations.
- **ARKC:** an Arkeia version 5 “*arkc*” command-line interface exists for this operating system. ARKC is a simplified interface to control a server through a standard UNIX command line.

For more information on the Arkeia version 5 architecture (client/server/interface), refer to chapter “5: GUI Concepts.” on page 62.

Available clients and servers.

Operating System	Server	Client	GUI	ARKC
AIX 3.2				
AIX 4.x	•	•	•	
BSDi 3.0, 4.0		•		
DGUX AviiON		•		
DRSNX 7 (ICL, Intel x86)		•		
DRSNX 7 (ICL, Sparc)		•		
FreeBSD 2.2.6 (Intel x86)		•		
FreeBSD 3.2 (Intel x86)		•		
HP-UX 9		•		

Available clients and servers.

Operating System	Server	Client	GUI	ARKC
HP-UX 10	•	•	•	•
HP-UX 11	•	•	•	•
IRIX 4		•		
IRIX 5.3		•		
IRIX 6.x	•	•	•	•
Linux 2.x (Intel x86)	•	•	•	•
Linux 2.x (MIPS, Cobalt)		•		
Linux 2.x (Alpha)		•		
Linux 2.x (ARM, Netwinder)		•		
MaxOS 4.2 (MIPS)		•		
MaxOS 4.2 (PowerPC)		•		
Novell Netware 5/6 (Intel x86)		•		•
DEC Alpha UNIX 3.2		•		
SCO v5 (Intel x86)		•		
Solaris 2.5, 2.6, 7, 8, 9 (Sparc)	•	•	•	•
SunOS 4.1		•		
Compaq Tru64	•	•	•	•
Unixware 2.x		•		
Unixware 7.1	•	•		
OpenUnix 8	•	•		
Windows 95/98/ME		•		
Windows NT Server 4.0 (Alpha)		•		•
Windows NT Server 4.0 (Intel x86)		•		•
Windows NT Workstation 4.0 (Intel x86)		•		•
Windows 2000		•		•
Windows XP		•		•

2 • Hardware requirements

Like all software, Arkeia has a number of requirements. Unlike classical office and productivity software (for instance, word processing), Arkeia is designed to work on a network, and therefore requires a certain network and server capacity. For more information on the Arkeia client/server architecture, refer to chapter “5: GUI Concepts.” on page 62.

The following sections of this chapter are designed to inform you of the various requirements of Arkeia. Make sure you read and understand these before installing the software on the machines you are responsible for.

2.1 Memory

Arkeia uses large amounts of memory, both on the client that needs to be backed-up and on the server that controls the backup operations.

For a client, the minimum amount of memory is 64 MB of RAM. It is recommended to install up to 128 MB of RAM in the client.

For a server, the minimum amount of memory is 128 MB of RAM. It is recommended to install 256 MB of RAM or more, depending on the number of clients.



Please note: Arkeia needs a large amount of memory on a server. This amount may vary, according to the number of clients that are going to be backed-up by the server. If your Arkeia server is managing a large network, 256 MB of RAM may be the minimum required for adequate performance.

2.2 Network bandwidth, configuration and adapters

Since Arkeia is designed to work exclusively on a network, proper attention should be given to the quality of the network adapters. Equally important is a proper configuration of the network.

Before installing Arkeia, it is therefore recommended to check the configuration of your network. To do this, you may want to verify the information detailed in the steps below:

- Check the IP addresses of your network, as well as the machine (client) names.
- Check the entries in the `/etc/hosts` file and in the DNS server.
- Ping several machines from the Arkeia server, using both their IP addresses and their names. Check the Ping results to make sure the IP addresses and names are correct.

- Check the auto-negotiation, the speed and the duplex configuration of the network adapter(s) as these parameters have a lot of influence on the overall performance of the network.



Please note: an Arkeia backup server communicates constantly with the clients it is supposed to back-up. Therefore, it is essential that the network configuration itself does not prove to be a bottleneck to the server operations.

2.3 SCSI adapters, tape drives and libraries

- *Overview*

The typical amount of data processed by an Arkeia backup server requires high-speed connections to both the network and the peripheral(s) used for the backup itself. For this reason, a lot of care should be paid to the SCSI configuration, and to the SCSI adapters used on the Arkeia server.

- *SCSI Configuration*

On an Intel i386 (or compatible) PC-class machine, make sure you have checked the following items in your configuration, and that the following options are set in the BIOS of your SCSI adapter:

- Your tape drive/library should be plugged in a separate SCSI controller, and not in the same SCSI adapter as hard disk drives or CD-ROM drives.
- SCSI peripherals connected to a SCSI adapter should be of the same generation.
- If you are using, under Linux, a library that uses multiple LUN, the “*Multiple LUN support*” option of the SCSI adapter BIOS should be set to: “**e n a b l e d**”.



Please note: Arkeia only supports SCSI tape drives. IDE tape drives are not supported. You should also note that Travan tape drives have very poor performance during restore operations. This is due to specific features supported only by the manufacturer’s drivers and not by Arkeia.

- *Differences between tape drives and libraries/autoloaders*

In general, Arkeia uses one device for a standalone tape drive and two (or more) devices for libraries or autoloaders. The main differences between these two types of peripherals are the following:

- A standalone tape drive is usually controlled through the standard SCSI tape drive device of the Operating System of your server. For instance, this device is `/dev/st0` under Linux.
- A library, or autoloader, uses one standard tape device for each tape drive and one generic SCSI device for the library media changer. For instance, this is `/dev/sg0` under Linux.
- Within a library, the tape drive and the media changer have different SCSI IDs, but use the same LUN.
- Within most autoloaders, the tape drive(s) and the media changer have the same SCSI IDs, but use different LUNs.

Despite the differences indicated above, the tape drives of the libraries and of the autoloader are defined and detected as standalone tape drives.



Please note: standalone tape drives require that your kernel is configured to support them. On some operating systems this may require recompiling the kernel.

Furthermore, media changers within libraries or autoloaders are considered generic SCSI devices, and your kernel should also be configured to support these generic devices.

Finally, if you are using a library or an autoloader that support this function, you should make sure that your kernel is configured to probe all LUNs on a SCSI bus – since your backup peripheral may be using more than one LUN. For more information, refer to the documentation that was supplied with your backup device.

- *Kernel configuration*

Make sure the following options are activated in the kernel you are running:

- SCSI tape drives support.
- Generic SCSI support.
- Probe all LUNs support.

For more information about specific configurations and operating systems, refer to the following sections: “*Mounting and probing tape devices*” and “*Library managements*”.

- *Mounting and probing tape devices*

For more information on the detection of SCSI tape devices, refer to chapter “3: Hardware installation” on page 35 of this manual, and to the section “Automatic hardware detection with arkdetect” on page 184, in chapter “12: Arkeia Utilities” on page 178.

A tape device installation is usually simple, provided the kernel used on the system supports such devices.

The correct device can be checked (“*probed*”) by using the following command:

```
mt -f [tape device] status
```

In the above command, the [tape device] section depends on the operating system you are using. You will find below a short list of the possible names for a tape device under several operating systems:

List of tape devices per Operating Systems

Operating systems	Tape device
AIX	/dev/rmtX
Digital UNIX/Compaq Tru64 UNIX	/dev/rmtXh
HP-UX	/dev/rmt/Xmb (where “b” = Berkeley mode)
Linux v2.x	/dev/stX
SGI Irix	/dev/rmt/tpsXd3
Sun Solaris	/dev/rmt/Xhb (where “b” = Berkeley mode)



Please note: in the table above, the “x” corresponds to the number of the device. Replace this letter with the number of the device you would like to mount or probe.

You can find below an example of a correct “mt” command output under Linux:

```
# mt -f /dev/st0 status
SCSI 2 tape drive
File number=0, block number=0, partition=0.
Density code 0x24 (DDS-2).
Soft error count since last status=0
General status bits on (41010000):
BOT ONLINE IM_REP_EN
```

- *Library management*

As stated above, the main difference between a stand-alone tape drive and a library, from an “*hardware*” point of view, is that a library (or autoloader) contains a robotic

arm. This arm is managed as a generic SCSI device. In order for this arm to be correctly managed by your system, you should therefore make sure that the generic SCSI devices are correctly supported by your kernel.

You will find below a list of generic SCSI devices under different operating systems:

List of generic SCSI devices per operating systems

Operating system	Control device for library media changer
AIX	/dev/pthruX (provided by Arkeia Corporation)
Digital UNIX/Compaq Tru64	/dev/bXtXlX (where "b" = bus, "t" = target, "l" = LUN)
HP-UX	/dev/scsiX /dev/scsi/cXtXdX (where "c" = instance, "t" = target ID, d = LUN)
Linux v2.x	/dev/sgX
SGI Irix	/dev/rmt/tpsXd3
Sun Solaris	/dev/rsstX (provided by Arkeia Corporation)

Note that in the examples above, the "x" should be replaced by the relevant numerical value for the library arm you are trying to probe/control.

For more information on the detection of SCSI tape devices, refer to chapter "3: Hardware installation" on page 35 of this manual, and to the section "Automatic hardware detection with arkdetect" on page 184, in chapter "12: Arkeia Utilities" on page 178.

3 • Platform-specific requirements

3.1 Overview

Since Arkeia runs several tasks at the same time on the machine it is installed on, it may be necessary to optimize the shared memory available, as well as the message queues.

This section details the IPC (*“Inter-process communication”*) optimization procedures for several operating systems. Using these parameters should enable Arkeia to function properly in a conventional environment.



Please note: the values given below are average values. As such, they may be insufficient if the server where Arkeia is installed is used for other tasks than backup. For instance, if the machine being configured is also used as an Oracle database server, it is recommended that you increase the values given below.

3.2 Compaq Tru64 UNIX/Digital Unix DEC OSF

- *How to edit the IPC settings*

To modify the IPC settings on a Compaq Tru64 UNIX or a Digital UNIX DEC/OSF machine, follow the steps indicated below:

- Edit the “`/etc/sysconfigtab`” file, by entering the following commands:

```
cd /etc
vi sysconfigtab
```

- Go to the “`ipc`” section (create this section if it does not exist in the original file) and enter the following lines:

```
ipc:
msg-max = 8192
msg-mnb = 65535
msg-mni = 64
msg-tql = 1500
shm-max = 4194304
shm-min = 1
shm-mni = 512
shm-seg = 512
sem-mni = 128
sem-msl = 25
sem-opm = 10
sem-ume = 10
sem-vmx = 32767
sem-aem = 16384
num-of-sems = 60
max-kernel-ports = 22487
port-hash-max-num = 1124350
port-reserved-max-num = 22487
set-max-num = 1029
```

- Go to the “`proc`” section (create it if it does not exist in the original file) and enter the following lines:

```
proc:
max-proc-per-user = 64
max-threads-per-user = 256
per-proc-stack-size = 2097152
max-per-proc-stack-size = 33554432
per-proc-data-size = 134217728
max-per-proc-data-size = 1073741824
max-per-proc-address-space = 1073741824
per-proc-address-space = 1073741824
```

- *How to change the IPC settings*

Certain setting changes do not need the reboot of the machine to take effect. Enter this command instead:

```
/sbin/sysconfigdb -s
```

Otherwise, restart the machine with this command:

```
/sbin/reboot
```

- *How to view the IPC and process settings*

To view the IPC and process setting on a Tru64 UNIX machine, or Digital UNIX machine, enter the following commands:

- To view the “`ipc`” configuration, enter the command:

```
/sbin/sysconfig -q ipc
```

- To view the “`process`” configuration, enter the command:

```
/sbin/sysconfig -q proc
```

3.3 Hewlett-Packard HP-UX

- *How to edit the IPC and process settings*

Under Hewlett-Packard HP-UX, the SAM management tool configures IPCs. You can run this utility either in a text or in a graphical mode by entering the following command: “`sam`”.

- *How to access the IPC and process settings via the SAM menu*

- Message Queue configuration:

```
[Kernel configuration]      [Configurable Parameter]
[msgmap] 102
[msgmax] 8192
```

```
[msgmnb] 65535
[msgmni] 200
[msgseg] 2048
[msgssz] 32
[msgtql] 100
```

- Semaphore configuration:

```
[Kernel configuration] [Configurable Parameter]
[sema] 1
[semaem] 16384
[semmap] 66
[semnmi] 64
[semnms] 200
[semnmu] 30
[semume] 10
[semvmx] 32767
```

- Shared memory segments configuration:

```
[Kernel configuration] [Configurable Parameter]
[shmem] 1
[shmmax] 67108864
[shmmni] 200
[shmseg] 120
```

- Number of process configuration:

```
[Kernel configuration] [Configurable Parameter]
[nproc] 316
```

To modify each of these parameters, use the menu: “[Kernel configuration] / [Configurable Parameter] / [Action] / [Modify configurable parameter...]”.

- *How to change the settings*

To have these settings take effect, use the menu: “[Action] / [New Kernel]”.

The machine will then reboot automatically to execute the new kernel.

- *How to view the IPC and process settings*

The IPC and process configuration can be viewed under HP/UX by using the SAM management tool.

3.4 IBM AIX

There are no specific IPC or process settings that need to be applied to the AIX operating system (version 3.2 and version 4.1).

3.5 LINUX

There is no specific IPC configuration that should be applied to the versions of Linux using a kernel with a version number smaller than 2.4.x.

However, the following configuration should be applied to systems running a kernel with a version number equal, or greater than, 2.4.



Please note: The following information should be entered in the file:

```
/etc/sysctl.conf
```

If this file does not exist on your system, please create it using the text editors supplied with your distribution of Linux (see example below).

If the file exists on your system, please add the following information to it, or replace the existing information if need be.

To set the 2.4.x kernel parameters at boot time, follow the instructions below:

Login as the “root” user, with the command (assuming, of course, that you have the rights to login as “root” and know the correct password):

```
su root
```

Enter the “/etc” directory:

```
cd /etc
```

Edit the “sysctl.conf” file with the command:

```
vi sysctl.conf
```

Enter the following lines in the file and save it:

```
# Sets maximum number of message queues
# to 200 (by default it is 16)

kernel.msgmni = 200

# Sets maximum size of message queue to 65536
# (by default the size of the message queue is 16348 since 2.2.14)

kernel.msgmnb = 65536
```

You should then enter the following command to apply the changes to the kernel configuration:

```
sysctl -p
```


On some distributions, you may need to add the command “`sysctl -p`” to the system initialization files so that the kernel changes are applied after each reboot. These files are usually named “`boot.local`” or “`rc.local`”.

When these changes have been applied, you can check the changes to the IPC configuration by entering the following command:

```
ipcs -l
```

3.6 SGI IRIX

- *How to modify the IPC settings (IRIX Kernel version 6.4 and earlier)*

The default values are all located in the following read-only files:

- IPC configuration: “`/var/sysgen/mtune/shm`”.
- Semaphores: “`/var/sysgen/mtune/sem`”.
- Message queues: “`/var/sysgen/mtune/msg`”.
- Number of processes: “`/var/sysgen/mtune/kernel`”.

If you need to change these values, create the following file: “`/var/sysgen/stune`” (if it does not already exist), by entering the following commands:

```
cd /var/sysgen
vi stune
```

Insert the following lines (if they do not already exist) to modify the shared memory segments:

```
shmmax=0x20000000
shmin=1
shmmni=100
shmseg=512
shmall=512
```

Insert the following lines (if they do not already exist) to modify semaphores:

```
semnmi=10
semnms=60
semnmu=30
semmsl=25
semopm=10
semume=10
semvmx=32767
semaem=16384
```

Insert the following lines (if they do not already exist) to modify the message queues:

```
msgmax=8192
msgmnb=65535
msgmni=64
msgssz=8
```

```
msgtql=1000
msgseg=1536
```

The configuration processing is automatic.

- *How to apply the new settings (IRIX Kernel version 6.4 and earlier)*

Enter the command: “`/usr/sbin/autoconfig`”. Then, reboot the machine with the command: “`reboot`”.

- *How to view the IPC settings (IRIX Kernel version 6.4 and earlier)*

Enter one of the following commands:

- To view the IPC configuration: “`/usr/sbin/systune | grep shm`”.
 - To view the process configuration: “`/usr/sbin/systune | grep proc`”.
 - To view the semaphore configuration: “`/usr/sbin/systune | grep sem`”.
 - To view the message configuration: “`/usr/sbin/systune | grep msg`”.
- *How to modify the IPC settings (Irix Kernel version 6.5.X)*

For local backups to run with good performance, follow the procedure outlined below:

- Change the following parameters in the kernel:

```
tcp_recvspace = 184320
tcp_sendspace = 61440
```

These parameters should be entered either in the file “`/var/sysgen/master.d/bsd`” or in the file “`/var/sysgen/mtune/bsd`”. It is possible to dynamically list and modify these parameters with the command: “`systune`”.

3.7 Sun Solaris

- *How to modify the IPC and process settings*

If you need to change the IPC or process configuration, create the file “`/etc/system`” (if it does not already exist), with the commands:

```
cd /etc
vi system
```

In this file, insert the following lines (if they do not already exist) to modify the shared memory segments:

```
set shmsys:shminfo_shmmax=2097152
```

```
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=30
set shmsys:shminfo_shmseg=100
```

Insert the following lines (if they do not already exist) to modify the message queues:

```
set msgsys:msginfo_msgmap=500 New
set msgsys:msginfo_msgmax=8192
set msgsys:msginfo_msgmnb=65536
set msgsys:msginfo_msgssz=8 New
set msgsys:msginfo_msgseg=8192 New
set msgsys:msginfo_msgmni=100
set msgsys:msginfo_msgtql=500
```

- *How to apply the new configuration*

Reboot the machine by entering the command: “`reboot`”.

- *How to display the IPC and process settings*

Enter the following command: “`sysdef -i`”.

CHAPTER 2: INSTALLATION PROCEDURE

- *Summary*

This chapter will detail the installation procedure that should be followed in order to install Arkeia version 5 on a backup server. This chapter will offer a separate section for all the different operating systems available.



Please note: Reading this chapter, and understanding the procedures described within is essential for the people in charge of installing Arkeia on a backup server.

Please make sure you understand the different installation options and requirements before proceeding with the installation itself!



If Arkeia version 4.x is already installed on a backup server, it is possible to “migrate” this 4.x installation to Arkeia version 5. To do this, please refer to “Appendix 3: Migrating to version 5” on page 265 of this manual.

1 • Installation under the Linux operating system.

Under the Linux operating system, Arkeia can be installed in two different ways:

- For Red Hat, Mandrake and SuSE distributions (and other RPM-based distributions), Arkeia is supplied as an “.rpm” file, which can be installed with the utility of the same name.
- For Linux systems, and many other UNIX systems, that do not use or support the “.rpm” utility, Arkeia version 5 is also available as a standard “.tar.gz” file, which can be installed with the standard UNIX utilities “tar” and “gzip”, and the installation utility which is contained in the file.



Please note: Under most operating systems, Arkeia version 5 is supplied as a “.tar.gz” file. The installation procedure is therefore very close to the Linux procedure.

Nevertheless, since subtle differences may exist between the different UNIX flavors, each operating system will have its own section, and it is strongly advised to read each section, if necessary, to understand the differences between the different operating systems.

The rest of this section will study the two different installations, based on the type of files used.

1.1 Installation using the “.rpm” file

Installation of Arkeia version 5 under Linux, using the .rpm file format is straightforward. To complete this installation, please follow the steps outlined below.

- *Check the downloaded file*

If you have downloaded the Arkeia version 5 files from the Internet, you should now have one “.rpm” file, which contains the server, the client, the graphical user interface (GUI) and the command line interface. For instance, this file could be named:

```
arkeia-5.0.0-15.i386.rpm
```

For more information on the differences between the server, the client and the interfaces available in Arkeia version 5, refer to “5: GUI Concepts.” on page 62.

For more information on the “.rpm” utility and the different options it supports, please check its manual page, by entering the following command:

```
man rpm
```



Please note: If you have obtained Arkeia version 5 on a CD-ROM, the file indicated above should be available as well on your CD-ROM. To proceed with the installation, you should just mount your CD-ROM with the following command (as the “root” super-user):

```
su root
Password: <Enter the root password>
mount -t iso9660 /dev/cdrom /mnt/cdrom
```

If necessary, replace the device “/dev/cdrom” with the appropriate Linux device, and “/mnt/cdrom” with the directory where you want to mount the CD-ROM. For more information on the “mount” command, refer to the “man” pages of your Linux system.

On some Linux systems (for instance, Mandrake Linux), the “mount” operation may be automatically performed by a dedicated daemon. For more information, refer to the documentation of your Linux system.

Please note as well that the name of the file you have received may be different from the example given above.

- *Login as root and install the software*

Only the “root” super-user can install software on a Linux server using the “rpm” utility. To do this, enter the following command on a standard Linux shell, followed by the relevant password when prompted to do so by the operating system:

```
$ su root
Password: <Enter the root password>
```

Run the “rpm” utility, by entering the following command:

```
# rpm --install arkeia-5.0.0-15.i386.rpm
```



Please note: The Arkeia version 5 package installs the following packages on your machine:

- Client
- Server
- GUI
- Arkeia command line interface (arkc)

Please note that additional packages and options should be installed separately. Refer to the documentation that was supplied with the optional Arkeia package for more information.

If the installation has been completed successfully, you should be able to use Arkeia version 5 immediately. If not, try to solve the installation problems by following the suggestions outlined below:

- Make sure you are logged in as “`root`”. If necessary, log out of your current session and log again as “`root`”. If you cannot login as “`root`”, please check with your system administrator that you have the necessary authorizations.
- If necessary, use the “`-vv`” and “`--force`” options of the “`rpm`” utility to display more messages and use these to solve the installation issues. For more information on the options of the “`rpm`” utility, refer to its “`man`” page.

1.2 Installation using the “.tar.gz” file

Installation of Arkeia version 5 under Linux, using the “.tar.gz” file format is slightly more complicated than the “.RPM” installation. To complete this installation, please follow the steps outlined below.

- *Check the downloaded file*

If you have downloaded the Arkeia version 5 files from the Internet, you should now have one “.tar.gz” file, which contains, as the “.RPM” file, the complete Arkeia v5 software. This file could be named, for instance:

```
arkeia-5.0.0-15.i386.tar.gz
```

For more information on the differences between the different parts of Arkeia version 5, refer to chapter “5: GUI Concepts.” on page 62.

For more information on the “`tar`” and the “`gzip`” (GNU-ZIP) utilities and the different options they support, please check their manual pages, by entering the following commands:

```
man tar
man gzip
```



Please note: If you have obtained Arkeia version 5 on a CD-ROM, the file indicated above should be available as well on your CD-ROM. To proceed with the installation, you should just mount your CD-ROM with the following command (as the “root” super-user):

```
su root
Password: <Enter the root password>
mount -t iso9660 /dev/cdrom /mnt/cdrom
```

If necessary, replace the device “/dev/cdrom” with the appropriate Linux device, and “/mnt/cdrom” with the directory where you want to mount the CD-ROM. For more information on the “mount” command, refer to the “man” pages of your Linux system.

On some Linux systems (for instance, Mandrake Linux), the “mount” operation is automatically performed by a dedicated daemon. For more information, refer to the documentation of your Linux system.

Please note as well that the name of the file you have received may be different from the example given above.

- *Login as root and install the software*

Only the “root” super-user can install software on a Linux server using the “.tar.gz” package. To do this, enter the following command on a standard Linux shell, followed by the relevant password when prompted to do so by the operating system:

```
$ su root
Password: <Enter the root password>
```

Run the “tar” utility, by entering the following command:

```
# tar -xvzf arkeia-5.0.0-15.i386.tar.gz
```

The “tar” utility should create a directory named “arkeia-5.0.0-15”, which contains the Arkeia version 5 software. To complete the installation, enter this directory with the following command:

```
# cd arkeia-5.0.0-15
```



Please note: For the latest information on the Arkeia installation, please read carefully the “README” file that should be present in the installation directory.

This file may contain last-minute information that could not be integrated in this manual.

Run the supplied installation utility with the command:


```
# ./install
```

The installation procedure should then start and ask some questions before proceeding with the installation of Arkeia version 5 itself:

```
$ su root
Password:
<Enter root password, as required by su>
$ ./install
I1800020 main: ***** Start of installation *****
I1800020 distrbininstall: The path is
/home/jack/files/download/arkeia.rpm/arkeia-5.0.2
I1801200 mkplinstall: Installing "arkeia" Version "5.0.2" ("Arkeia")
I1801210 mkplinstall: OS: "Linux" Version "2.2.19"
Q1801002 *** Enter main ARKEIA_DIR directory [/opt/arkeia]? (q to quit)
<Press 'Enter' to accept default install directory>

I1801170 validateknox: Got from user $ARKEIA_DIR='/opt/arkeia'
I1801180 validateknox: Directory $ARKEIA_DIR='/opt/arkeia' OK
Q1801002 *** Enter the temporary directory [.]? (q to quit)
<Press 'Enter' to accept the default temporary directory>

I1801170 validatetmpdir: Got tmpdir = '.'
I1801180 validatetmpdir: Directory tmpdir='.' OK
E1801150 validatetmpdir: Normal end
I1801350 ai1chk: tmpspace needed: 50000 KB, available in '/home': 3422128 KB
I1801360 ai1chk: arkeiaspace needed: 25000 KB, available in '/': 3971916 KB
Q1800001 *** OK to start installation in '/opt/arkeia' (y/n) [y]? (q to quit)
<Press 'Enter' to launch installation>

I1802020 untz: Uncompressing and extracting 'arktar.z'...
I1802050 untz: 'arktar.z' extracted successfully
I1802510 arkbootstop: No arkeiad to stop (2: No such file or directory)
I1801080 mkdirknox [L1200]: Directory $ARKEIA_DIR='/opt/arkeia' created
I1801060 mkfsu: /etc/opt/arkeia/ARKEIA_DIR written ('/opt/arkeia')
I1804210 docopy: Installing files...
I1804190 cp1tree: Tree './arkeia/arkboot' successfully installed
I1804190 cp1tree: Tree './arkeia/arkc/arkc.param' successfully installed
I1804190 cp1tree: Tree './arkeia/arkc/arkc.prfl' successfully installed
I1804190 cp1tree: Tree './arkeia/arkc' successfully installed
I1804190 cp1tree: Tree './arkeia/arkeiad' successfully installed
I1804190 cp1tree: Tree './arkeia/lang' successfully installed
I1804190 cp1tree: Tree './arkeia/log' successfully installed
I1804190 cp1tree: Tree './arkeia/server/dbase/f3drv/drvtypes.lst' success-
fully installed
I1804190 cp1tree: Tree './arkeia/server/dbase/f3drv/scsi_detect_drvtypes.lst'
successfully installed
I1804190 cp1tree: Tree './arkeia/server/dbase/f3sec/defrole.lst' successfully
installed
I1804190 cp1tree: Tree './arkeia/server/dbase/f3lib/scsi_detect_libtypes.lst'
successfully installed
I1804190 cp1tree: Tree './arkeia/server/dbase/f3lib/libtypes' successfully
installed
I1804190 cp1tree: Tree './arkeia/server/dbase/f3tape/tptypes.lst' success-
fully installed
I1804190 cp1tree: Tree './arkeia/server/dbase' successfully installed
I1804190 cp1tree: Tree './arkeia/server/server.prfl' successfully installed
I1804190 cp1tree: Tree './arkeia/server/customer/defrole.lst' successfully
installed
I1804190 cp1tree: Tree './arkeia/server/customer' successfully installed
I1804190 cp1tree: Tree './arkeia/server' successfully installed
I1804190 cp1tree: Tree './arkeia/xarkeia' successfully installed
I1804190 cp1tree: Tree './arkeia/bin' successfully installed
I1804190 cp1tree: Tree './arkeia' successfully installed
I1802175 startarkeiad: cp /opt/arkeia/arkboot/sysvinit ->
/etc/rc.d/init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc0.d/K10arkeia -> ../init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc1.d/K10arkeia -> ../init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc2.d/S60arkeia -> ../init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc3.d/S60arkeia -> ../init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc4.d/S60arkeia -> ../init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc5.d/S60arkeia -> ../init.d/arkeia
I1802177 startarkeiad: ln /etc/rc.d/rc6.d/K10arkeia -> ../init.d/arkeia
I1802190 startarkeiad: Daemon '/usr/bin/arkeiad' will start automatically at
boot
I1801520 mk1usrbin: Created command 'xarkeia' in '/usr/bin'
I1801520 mk1usrbin: Created command 'arkc' in '/usr/bin'
arkeiad is starting...
I1800010 myexit: ***** End of installation (0) *****
```

The above installation procedure was taken from a Linux installation on a Slackware 8.0 machine. Installation on a different Linux distribution may be slightly different.



Please note: The Arkeia version 5 “.tar.gz” package installs on your machine the following packages:

- Client
- Server
- GUI
- Arkeia command line interface (`arkc`)

Please note that additional packages and options should be installed separately. refer to the documentation that was supplied with the optional Arkeia package for more information.


If the installation has been completed successfully, you should be able to use Arkeia version 5 immediately. If not, try to solve the installation issues by running the supplied installation procedure again.

1.3 Installation check

To check that your installation has been performed successfully, enter the following command in a standard Linux shell:

```
xarkeia &
```

This command should launch the Arkeia version 5 GUI:



2 • Post-installation root password

One very important post-installation configuration that must be done to ensure the security of the Arkeia server is to change the default password for the “root” user.

By default, there is only one user created on a newly-installed Arkeia server: “root”, and the password used by this user is blank. For instance, once the installation of the server has been completed, simply hitting the “Enter” key on the server login screen allows any user to connect to the Arkeia server.

Since the “root” user has access to all the functions of Arkeia, including the creation of the tapes, the Savepacks, etc. this could become a major security problem. Arkeia Corporation therefore strongly recommends to all system administrators to change this default password immediately after completing the installation of a new Arkeia server.



Please note: Changing the “root” password of an Arkeia server should not be considered an option in a production environment!!!

Changing the default “root” password of a complex system, such as Arkeia, should be treated as the most basic security measure to be applied.

To change the “root” password, please refer to “How to create users under Arkeia” on page 100 of chapter 6 (“Creating objects and users”).

3 • Installation of Arkeia for a client-only machine

3.1 Overview

By default, the installation procedures of Arkeia version 5 install all the components of the software on a given machine. These components are:

- The Arkeia server
- The Arkeia client
- The Arkeia Graphical User Interface (GUI) for X11

This standard installation is correct if the machine is designed to become an Arkeia server. However, running the server on a machine designed to be a client requires that some changes are entered in the Arkeia “`admin.cfg`” configuration files.

3.2 Configuration of an Arkeia client

The file to be modified to declare a client is:

```
/opt/arkeia/arkeiad/admin.cfg
```

To declare a client to a server, this file need to be modified to contain the name of the Arkeia server it should connect to.

For instance, let us assume that the client machine is named “`earl.snafu.com`” and the Arkeia server is named “`zeus.snafu.com`”. The file: “`/opt/arkeia/arkeiad/admin.cfg`” of “`earl.snafu.com`” should therefore contain the line:

```
zeus.snafu.com
```

If the Arkeia daemon, running on “`zeus.snafu.com`”, uses a port other than its default TCP port (617), then the line contained in the “`admin.cfg`” of the clients should be:

```
zeus.snafu.com:<Port_number>
```

In this line, “`<Port_number>`” should be replaced, of course, by the number of the port used, on “`zeus`”, by the Arkeia daemon. If the TCP port is different from the default value, it is up to the system administrator to make sure that the configured TCP port is not currently in use by another service.



Please note: At the time of this writing, it is not possible to declare more than one Arkeia server in the “`admin.cfg`” file of an Arkeia client.

Since the “`admin.cfg`” file was created, during the installation of Arkeia version 5, by the user “`root`”, it is necessary to have the proper rights and authorizations to be able to edit this file.

Once the “`admin.cfg`” file has been modified and saved, it is necessary to restart “`arkeiad`” to allow the client to declare itself to the server.

4 • Arkeia license management

4.1 Overview

Arkeia version 5 requires a license number to function correctly. This license number – which is supplied by Arkeia Corporation – should be entered on the Arkeia server by the system administrator, using the graphical user interface (GUI).

The procedure below explains how to enter the license after the installation of Arkeia version 5 and how to upgrade the license, if need be, to add other functionalities.

4.2 Entering the license number through the GUI

To enter the license number in the Arkeia GUI, follow the steps detailed below.

Login, as user “root”, on the server that should receive the license:



The screenshot shows a 'Server login' dialog box with the following fields and values:

Field	Value
Server:	voyager
Login:	root
Password:	

Buttons: OK, Cancel

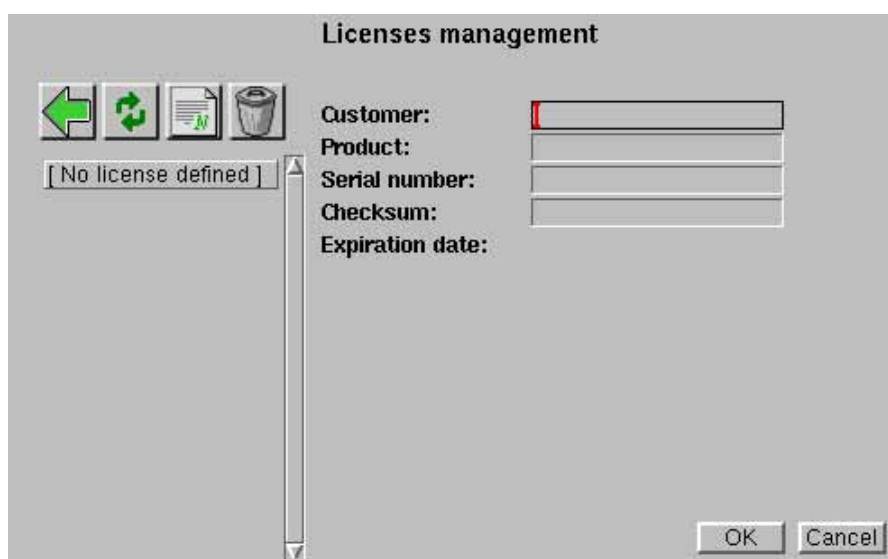
Within the “Main menu” of Arkeia version 5, click on the icon labeled “Administration”:



In the “Administration menu” of Arkeia version 5, click on the icon labeled “Licenses”:



This opens the Licenses Management screen within the main window of Arkeia version 5:



In this screen, the “Customer” ID, the “Product” ID, the “Serial number” and the “Checksum” should be entered.

Once the information has been entered, clicking on the “OK” button saves the license information and computes the “Expiration Date”. A license with the correct customer ID should then be displayed in the list of licenses, which is on the left in the screen printed above.

4.3 Additional functionalities that require a separate license

The following software products require a separate license. Please note that this is not an exhaustive list. For more information, contact the Sales Team at Arkeia Corporation.

- Tape libraries
- Encryption
- Compression

- Additional flows
- Additional clients
- Plug-in for most applications (for instance: Oracle)

For each of these products, a separate license number will be supplied by Arkeia Corporation and should be entered in the manner described above.

CHAPTER 3: HARDWARE INSTALLATION

- *Summary*

Installing hardware to work under Arkeia is a very easy process. However, caution and planning are absolutely crucial in order to create a stable, high-performance backup configuration. This chapter attempts to present the different points to consider before installing any hardware, as well as a simple installation procedure for a single tape drive.

1 • Overview of hardware management

1.1 Interface and host adapter supported

Arkeia version 5 uses SCSI backup devices exclusively. Cheaper IDE devices are not supported. This is due to the very high amount of data that must be transferred between the Arkeia server and the backup device itself.



Please note: Allow us to repeat what was said above: Arkeia version 5, at the time of this writing, supports SCSI backup devices only.

All other devices, using any other type of interface (USB, Firewire, IDE, parallel port, etc.) are unsupported at this time. Using these devices with Arkeia will either fail or will severely limit its performances and stability.

Arkeia Corporation therefore recommends that SCSI backup devices are connected to the Arkeia backup server, for use by Arkeia version 5.

As was noted in chapter “1: Installation pre-requisites” on page 6 of this manual, many problems that are encountered while using and testing Arkeia version 5 is due to SCSI configuration problems. These problems can be avoided through careful planning. Please make sure you read this chapter before starting your SCSI configuration.

1.2 Supported backup media

Please find below a short table with the most common backup media that are supported by Arkeia version 5. This table also tries to present the most important characteristics of these different medias and their possible applications.



Please note: If a backup media is not found in this table, try to find the closest media and try running some limited tests to see if the media itself can be used by Arkeia version 5.

Most UNIX (and UNIX-like) operating system offer many utilities to test the SCSI hardware connected to a server. These utilities are briefly described in this chapter. Arkeia version 5 contains many utilities that can also provide help in detecting and using SCSI hardware. Most are described in chapter “12: Arkeia Utilities” on page 178 of this manual.

(Please see the remarks below)

Backup Media supported by Arkeia.

Name:	Capacity (<u>without compression</u>):	Manufacturer:
8mm / “Exabyte”	from 2 to 7 GB	Exabyte
4mm / DAT / DDS 1, 2, 3, 4	from 1 to 20 GB	Hewlett-Packard and others
QIC / “Travan”	from 2.5 to 10 GB	Tandberg, Tecmar and others
DLT / DLT1 / SDLT	from 15 to 110 GB	Quantum Technology and others
AIT / AIT 1 / AIT 2 / AIT 3	from 256 to 100 GB	Sony
Mammoth 1 / Mammoth 2	from 20 to 60 GB	Exabyte
LTO (1st generation)	100 GB	Hewlett-Packard and others
9840 “Eagle”	20 GB	Storagetek
SD-3 “Redwood”	from 10 to 50 GB	Storagetek
3570 / 3590 B, E / “Magstar”	from 5 to 10 GB	IBM
VXA 1	12 to 33 GB	Ecix
DTF 1	from 12 to 42 GB	Sony
SLR 1, 2, 3, 4, 5, 50, 100	from 20 to 50 GB	Tandberg

Please note:



- The “manufacturer” column is for reference only, and may not be up-to-date.
- The capacities given above can be modified by the suppliers. As indicated, the capacity is without compression. Adding compression, through Arkeia or through the internal function of the hardware, may improve this capacity.
- While using the same backup media, backup hardware supplied by different manufacturers can have different characteristics, such as transfer speed and access time.
- Please note that many different formats are actually described in the table above. For instance, Tandberg’s SLR tape format is not just one, but 7 different formats (SLR 1, SLR 2, etc.) The capacity indicated is usually the maximum capacity of the latest format.
- Arkeia maintains up-to-date compatibility list for backup media and single tape drives at the following URL (click on the address shown below to open the Arkeia web site):

<http://www.arkeia.com>

For more information on the different backup media that exist, Arkeia Corporation recommends the invaluable (and exhaustive) resources of the following web site:

<http://www.storagemountain.com>

1.3 Supported tape drives

In general, even if you are using a single tape drive whose backup media is not listed in the table above, there is a good chance that this hardware can be supported if it strictly adheres to the SCSI standard.

Arkeia uses advanced SCSI functions to manage operations and/or SCSI events and to correctly interpret messages such as the end of a tape. If a backup hardware, regardless of the media it uses, complies with the SCSI II standard, Arkeia should therefore be able to manage it.

Unfortunately, while many hardware suppliers claim a strict adherence to the SCSI standard, this is not always the case. Many companies integrate advanced functions in their SCSI drivers, and not in the hardware itself. Therefore, as is often the case: “your mileage may vary”: some drives and/or media types that are not listed in the table of the previous page may function perfectly with Arkeia, while some others may not.



Please note: Again, Arkeia Corp. would like to strongly recommend to all prospective buyers of new hardware backup to check the compatibility lists of Arkeia version 5.

As noted above, these lists can be accessed from its web site, at the following addresses:

<http://www.arkeia.com/mass.html>

1.4 Handling and replacement of backup media

While this may seem obvious, backup media should be handled with care, especially if the data it contains is considered mission-critical.

Arkeia Corporation recommends to respect the following points when handling backup media:

- Always store the backup media on its side, in a cool dry place.
- Always store the backup media away from any magnetic field, especially computer monitors or advanced network equipment. For instance, storing backup media on top of the backup hardware itself is probably not a good idea.

- Always label backup media properly, indicating the content and the date of the backup. This is especially important for “Archive”, “Total” and “Disaster Recovery” backups. These are the most important backups that can be performed by Arkeia, and being able to quickly select the correct backup tape can be crucial, especially when restoring mission-critical systems.
- Always avoid dropping tapes and other backup media from a height. Some tapes or cartridges contain delicate mechanisms that can be badly damaged by shock. If re-used, a damaged media can in turn damage the mechanics of its drive. A damaged drive during an important restoration can make a bad situation even worse...
- Always respect the media life indications given by the supplier of the backup media itself. This is usually indicated in terms of “passes”. For instance, performing a backup and rewinding the tape after the backup count as two passes.
- Always replace backup media that shows signs of failure. Arkeia backup logs may sometimes indicate that a backup media has exhibited unusual behavior (such as read/write errors). Such a media should always be retired immediately, and the last backup operation it was used for performed again.
- Always store the backup media in a safe storage area that respects the security policies of your company. A stolen or “lost” backup tape is not just an inconvenience: it could represent an asset to your competitors! Arkeia offers advanced encryption functions that can be very helpful to secure the data on a tape. For more information, please refer to chapter “7: Interactive backup” on page 106.
- Finally, if the data that was saved on the backup media is considered as essential, copies of the media should be performed and these copies should be stored off-site. Arkeia offers a tape duplication utility which can speed up this process. For more information, please refer to chapter “12: Arkeia Utilities” on page 178 of this manual.



While the points indicated above may seem obvious, proper storage of backup media cannot be over-emphasized. A backup media that is handled improperly is, in effect, a media which may fail during restoration, negating all the hard work of a dedicated system administrator.

1.5 Advanced backup hardware

Apart from the single tape drive — which is useful for a small network, but impractical for large sites — there are three basic types of advanced backup hardware:

- **Stackers:** The original “stackers” were single tape drives to which a special “stack” device was added. These machines would simply insert a new tape in the drive every time a tape was ejected. As such, they were “sequential” machines, inserting each tape exactly in the order they were inserted in the stack. Today’s stackers are much more advanced and offer functions (such as a random access to any tape) that tend to blur the distinctions between them and the most simple libraries.
- **Library:** A library has several features that distinguish them from stackers. They usually offer several internal slots to contain tapes, several drives, and a robotic “arm” to move

the tapes from the slots to the drives and back to the slots. Some libraries now offer a “gate”, a small trapdoor, allowing a tape to be inserted or ejected from the library after an operation has been completed, without having to stop the entire library.

- Silos: These very high-end machines are usually connected to large servers (main-frame-class), that require the archival of several terabytes of data. Arkeia version 5 does not manage silos and they appear in this list of backup devices for the sake of completeness only.



Please note: This chapter will not cover the stackers and the libraries that are managed by Arkeia version 5. For more information on how to setup and use one of these advanced backup hardware, please refer to chapter: “10: Advanced Hardware” on page 158 of this manual.

2 • Basic SCSI configuration

This section will detail the different SCSI format available as well as some of the configuration tools that can be used to control these interfaces.

2.1 Overview of the SCSI interfaces

There are currently three broad types of SCSI interfaces available for computer systems. All of these are managed by Linux, and, consequently, by Arkeia version 5:

- **SCSI I:** The first standard of the SCSI interfaces, SCSI I was defined in 1986. This standard is now considered as obsolete, as it offers a limited bandwidth (5 MB per second). This first SCSI standard is not supported by Arkeia version 5.
- **SCSI II:** This evolution of SCSI I was defined in 1990. It introduced important new functions, such as a command queue and DMA access, while staying backward-compatible with the previous standard. SCSI II then offered an evolution of these features with the “Fast” SCSI II (with a throughput of 10 MB per second) and the “Wide” SCSI II (which can support up to 16 different devices) — of course, these two evolutions were soon combined into the “Fast Wide SCSI II”.
- **SCSI III:** This further evolution of the SCSI standard is not, at the time of this writing, completely finalized. It defines new and innovative type of interfaces, such as Fibre Channel and “Firewire” (which is itself defined in the IEEE 1394 standard), as well as many additional SCSI commands, covering functions such as auto configuration and multimedia devices. Many of the devices and adapters marketed today as “Ultra” SCSI really implement a subset of the proposed SCSI III standard. Some of these adapters actually offer speed of 80 MB per second or more.

2.2 SCSI device management under Linux

- *Standard devices*

Linux, like most other operating system in the UNIX category, offer several devices that are used to control the SCSI hardware connected to a server. These devices are the following:

- **“/dev/sgX”:** Under Linux, a device with a name such as “/dev/sgX” (where “X” is the device number) is a generic SCSI device. This type of device is used by Arkeia version 5 to manage the robotic arm of an SCSI library.
- **“/dev/stX”:** Under Linux, a device with a name such as “/dev/stX” (where “X” is the device number), is a rewindable SCSI device. This type of device is used by Arkeia version 5 to read and write to the tapes present in drives.
- **“/dev/nstX”:** Under Linux, a device with a name such as “/dev/nstX” (where “X” is the device number) is a non-rewindable SCSI device.

- *Vendor-supplied device drivers*

Many companies that market backup hardware supply their own device drivers for their hardware. These drivers may have different names, and may even be just symbolic links to existing devices.

For instance, IBM supplies the following drivers with its LTO library:

- “/dev/IBMtapeX”: for rewindable devices, where “X” is the device number.
- “/dev/IBMtapeXn”: for non-rewindable devices, where “X” is the device number.

For more information on these (possible) drivers, refer to the documentation supplied with your backup hardware.

2.3 Quick SCSI setup and troubleshooting

Whenever a problem has been encountered while trying to use SCSI devices, please use the points indicated below as a quick check-list before trying to contact the Arkeia technical support team.

- *Hardware checks*

- Make sure all devices are properly connected. If necessary, check the connecting cable on another, correctly working, SCSI device and replace the cable if it is found to be defective.
- Make sure that all SCSI devices connected to the same adapter are of the same type! Mixing SCSI II and SCSI III devices is a very good way to either block all communications between the devices and the adapter or to slow down the transfer of data from the devices to the server.
- Make sure that all devices are correctly configured within the SCSI chain, especially their device numbers.
- Do not use other SCSI devices, such as a hard disk or a CD-ROM, on the same SCSI adapter as the backup device.

- *Software checks*

There are several programs and utilities available, under Linux and most UNIX operating systems, that can help with the configuration and troubleshooting of a SCSI chain. Some may be available only from the Internet, but all are pretty valuable in their own right.

- The first step in a troubleshooting procedure is to make sure that the SCSI peripherals have been correctly identified and detected during the bootstrapping procedure of your machine. To do this, display the messages returned by the system by using the “dmesg” procedure or display the content of the “/var/log/boot.msg” file (or the equivalent

on your system). Printed below is a short section of these messages, showing the detection of the SCSI adapters.

```
(scsi0) <Adaptec AIC-7892 Ultra 160/m SCSI host adapter> found at PCI
0/9/0
(scsi0) Wide Channel, SCSI ID=7, 32/255 SCBs
(scsi0) Downloading sequencer code... 392 instructions downloaded
(scsi1) <Adaptec AHA-294X Ultra SCSI host adapter> found at PCI 0/11/0
(scsi1) Wide Channel, SCSI ID=7, 16/255 SCBs
(scsi1) Cables present (Int-50 NO, Int-68 NO, Ext-68 YES)
(scsi1) Downloading sequencer code... 422 instructions downloaded
scsi0: Adaptec AHA274x/284x/294x (EISA/VLB/PCI-Fast SCSI) 5.1.31/3.2.4
<Adaptec AIC-7892 Ultra 160/m SCSI host adapter>
scsi1: Adaptec AHA274x/284x/294x (EISA/VLB/PCI-Fast SCSI) 5.1.31/3.2.4
<Adaptec AHA-294X Ultra SCSI host adapter>
scsi: 2 hosts.

(scsi0:0:4:0) Synchronous at 160.0 Mbyte/sec., offset 63.
Vendor: IBM          Model: DDYS-T09170N      Rev: S80D
Type:   Direct-Access          ANSI SCSI revision: 03
Detected scsi disk sda at scsi0, channel 0, id 4, lun 0

Vendor: QUALSTAR    Model: TLS-4210          Rev: 2.09
Type:   Medium Changer        ANSI SCSI revision: 02
(scsi1:0:2:0) Synchronous at 20.0 Mbyte/sec., offset 14.

Vendor: ECRIX       Model: VXA-1              Rev: 2524
Type:   Sequential-Access     ANSI SCSI revision: 02
Detected scsi tape st0 at scsi1, channel 0, id 2, lun 0
(scsi1:0:3:0) Synchronous at 20.0 Mbyte/sec., offset 14.

Vendor: ECRIX       Model: VXA-1              Rev: 2524
Type:   Sequential-Access     ANSI SCSI revision: 02
Detected scsi tape st1 at scsi1, channel 0, id 3, lun 0

SCSI device sda: hdwr sector= 512 bytes. Sectors= 17916240 [8748 MB]
[8.7 GB]
sda: sda1
```

- The second step is usually to launch “`lspci`”, a utility which reports all the adapters found installed in a machine. If the utility “`lspci`” is not installed on your machine, it can be downloaded (and compiled) from the following URL:

<http://atrey.karlin.mff.cuni.cz/~mj/pciutils.html>

You'll find below an example of a normal output from “`lspci`” (Please note the two SCSI adapters detected in the machine):

```
00:00.0 Host bridge: VIA Technologies, Inc. VT8363/8365 [KT133/KM133]
(rev 03)
00:01.0 PCI bridge: VIA Technologies, Inc. VT8363/8365 [KT133/KM133
AGP]
00:07.0 ISA bridge: VIA Technologies, Inc. VT82C686 [Apollo Super
South] (rev 40)
00:07.1 IDE interface: VIA Technologies, Inc. Bus Master IDE (rev 06)
00:07.4 Host bridge: VIA Technologies, Inc. VT82C686 [Apollo Super
ACPI] (rev 40)
00:09.0 SCSI storage controller: Adaptec 7892A (rev 02)
00:0b.0 SCSI storage controller: Adaptec AIC-7881U (rev 01)
00:0d.0 Ethernet controller: VIA Technologies, Inc. Ethernet Control-
ler (rev 42)
01:00.0 VGA compatible controller: ATI Technologies Inc. 3D Rage P/M
Mobility AGP 2x (rev 64)
```

- The third step that allow a system administrator to check the configuration of a server is to use the auto-detection feature of Arkeia version 5. To do this, follow the steps outlined below:

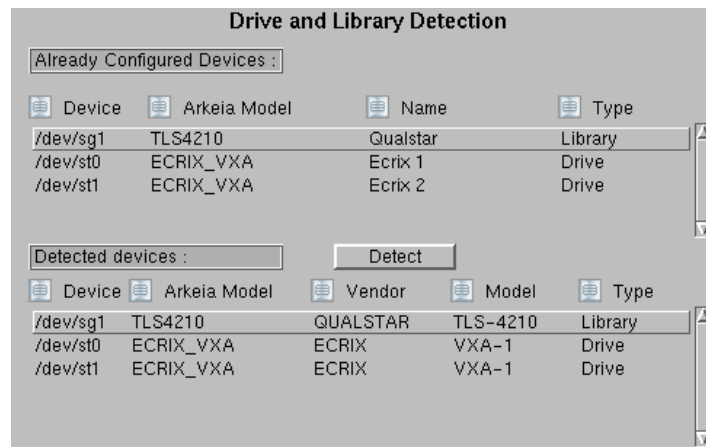


1. In the main menu of Arkeia version 5, click on the “Hardware” icon. This opens the main hardware menu in the window of Arkeia version 5.



2. In the “Hardware” menu, click on the icon labeled “Detect”.

This launches the Hardware auto-detection feature of Arkeia version 5, which searches for backup hardware on your server and reports its findings.



3. Once the auto-detection routine has been completed, it displays its results in the screen shown above. If no information appears in the “Detected devices” fields, clicking on the “Detect” button next to it will restart the hardware detection process.

- Another step that can be taken to check the configuration of an Arkeia version 5 server is the “arkdetect” utility. For more information on this utility, refer to the section named “Automatic hardware detection with arkdetect” on page 184 of this manual.
- One final utility that should be mentioned is the “mt” (which stands for “Magnetic Tape”) command, which is present on most UNIX systems. For more information on this command, enter the following command on a standard UNIX shell:

```
$ man mt
```

The “man” page for “mt” should contain all the relevant information to use this UNIX utility, which allows you to send commands to a tape drive from a shell.

If you are using the GNU version of “`mt`” (used on Linux), the following command will display the on-line “`info`” manual of this command:

```
$ info mt
```

3 • Basic installation procedure

This section will detail the installation of a single drive on an Arkeia server. To install more complex backup hardware, such as libraries and stackers, refer to chapter “10: Advanced Hardware” on page 158 of this manual.

To install a single tape drive on a machine running Arkeia version 5, follow the steps detailed below:

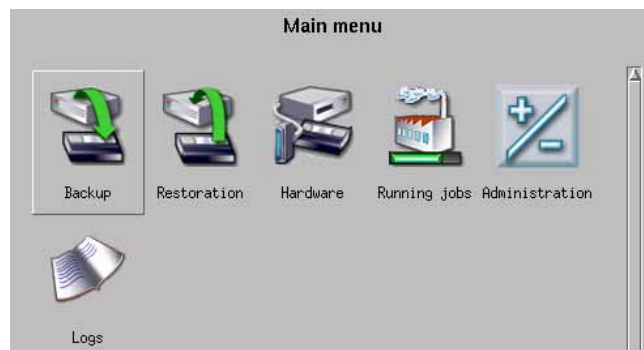
- *Auto-detection of the drive*

A good starting point for a successful installation is the auto-detection function, which is integrated into Arkeia version 5. This function is presented in more details in the previous section of this chapter.

The information supplied by the auto detection function will be very helpful in the next steps of the installation.

- *Creating a drive in the Arkeia GUI*

Once Arkeia version 5 is launched, it should display its main menu in its window.




Within this menu, clicking on the “Hardware” icon will open the relevant menu in Arkeia main window.

For more information on the other options of the menu shown below, refer to chapter “6: Objects and users creation” on page 71 of this manual.



Within the “Hardware” menu, clicking on the “Drives” icon will open the “Drives” screen within the main Arkeia window.

The “Drives management” menu, shown below, is divided into two parts. In the left-hand part, the drives already configured are displayed, while the right-hand side shows the different actions that can be applied to the drives. If no drive has been defined in Arkeia (as is the case below), Arkeia will simply display “[No drive]” in the list of drives.



Within the “Drives management” screen, clicking on the “New” icon (pictured on the left) will open the creation screen in the main window of Arkeia version 5.

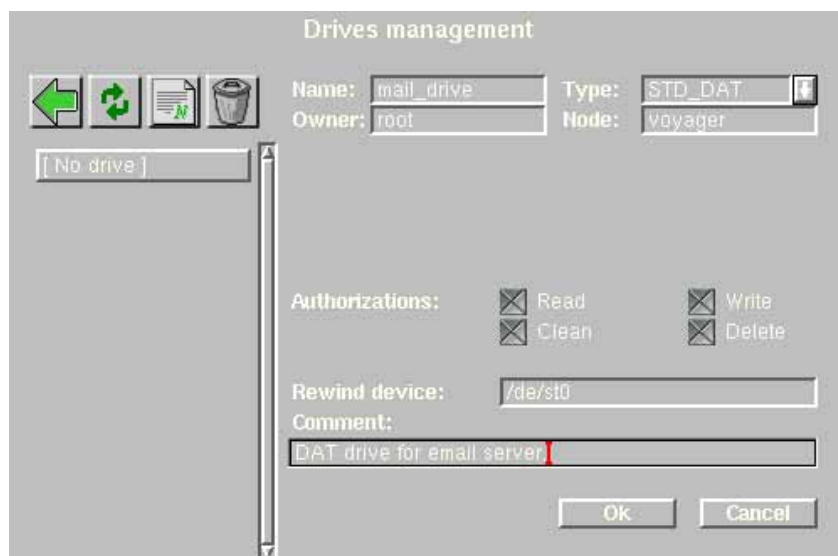
Within the creation screen, a number of fields have to be filled-up. These fields are detailed below:



The options that can be modified in this screen are the following:

- **“Name”**: This field is used to assign a name to the drive being created. This name will be used by Arkeia version 5 in its logs and messages.
- **“Type”**: This drop-down list is used to select one of the type of drives that can be managed by Arkeia version 5. For a complete list of supported tape drives, please refer to the Arkeia web site, at the following address: <http://www.arkeia.com/mass.html>.
- **“Owner” / “Node”**: These fields are filled-up automatically by Arkeia version 5, based on the owner and machine which has launched the creation process.
- **“Authorizations”**: These check boxes are used to set the authorizations that will be applied by Arkeia version 5 to every operation that will be performed on the drive. The four possible authorizations are **“Read”**, **“Write”**, **“Delete”** and **“Clean”** (they should be self-explanatory).
- **“Rewind device”**: This field should contain the UNIX device that will be used to write/read data to the tapes in the drive. For more information on these devices, refer to the rest of this chapter.
- **“Comment”** (optional): This field can contain an optional comment, to describe the drive or to supply more information on its location, type, supplier, etc.

In the screen shown below, the drive being created is a standard DAT drive (**“Type”** = **“STD_DAT”**), named **“mail_drive”**, which will be used for E-mail-related backup tasks, on a small E-mail server. All the authorizations have been checked, which gives Arkeia version 5 complete control over the drive:



Once the different fields have been selected, clicking on the “OK” button will create the drive within Arkeia version 5. The display then returns to the main “Drives” screen, updated to display the new drive:



Once this screen is displayed, the creation of a drive is completed. For more information on adding drives to a drivepack, please refer to chapter “6: Objects and users creation” on page 71 of this manual.

CHAPTER 4: BASIC CONCEPTS

- *Summary*

This chapter presents the basic concepts used by Arkeia version 5. It is recommended that new Arkeia users read and understand this chapter before attempting to use the software for the first time. Experienced users and system administrators may benefit as well, since some of the concepts presented in this chapter are new to version 5.

1 • Overview

Arkeia version 5 is a modern, multiple platforms client/server backup software. Its advanced functions include multiple simultaneous connections to different clients and highly customizable modes of operations.

Because of the flexibility and the complexity of Arkeia version 5, many of the concepts it uses to provide high-performance backup and restoration (such as Savepacks, Priorities, flows, etc.) are complex and deserve an explanation, which will be provided by this chapter.

Roughly speaking, the concepts that will be described and expanded upon in this chapter fall in two categories:

- The “*architecture*” concepts, such as the three-tiered architecture (client/server/interface), the platforms supported, etc.
- The concepts related to the internal working of the software. These are advanced functions aimed at simplifying the configuration and daily use of Arkeia version 5. Examples of these functions are the Savepacks, Scheduled backups, etc.

The rest of this chapter will be go through these two concept categories and explain each of them.

2 • Architectural concepts

This section will detail the most important concepts and characteristics of Arkeia version 5, presenting them in the context of the usage within a general computer LAN (Local Area Network).

2.1 Client/Server architecture

Arkeia uses what is called a “*Client/Server*” architecture. This type of architecture means that a single machine (the server) is in charge of all backup/restore operations of several other machines (the clients).

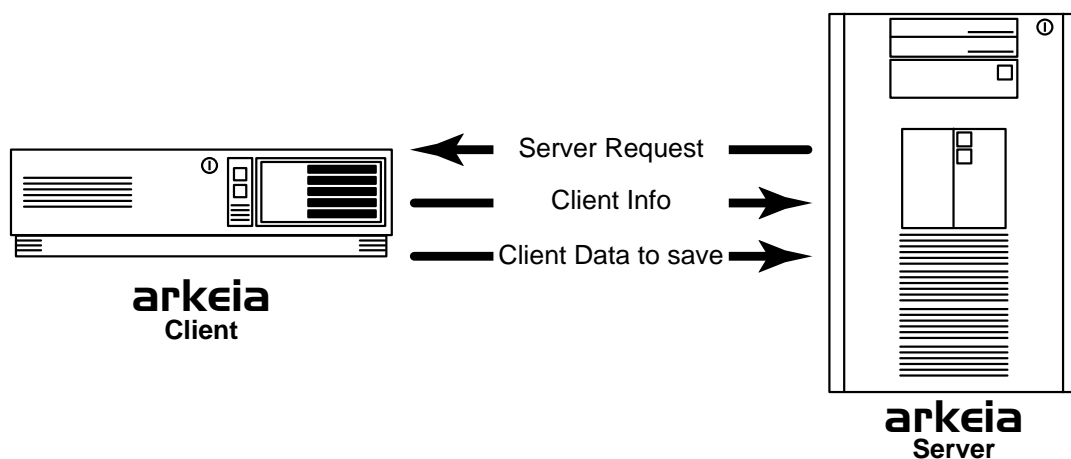
This architecture allows a company, using Arkeia version 5, to reduce its investment in hardware: instead of having to buy a backup device for every machine in a company, it is possible to “*focus*” the hardware investment on the backup server and realize substantial savings, while benefiting from high-end backup hardware.



Please note: It is highly recommended by Arkeia Corporation to dedicate a server to the backup tasks. While it is possible to install an Arkeia server on a machine that performs other tasks, this type of configuration will almost certainly create reliability, security and performance issues.

You will find below a graphical representation of a typical Client/Server architecture under Arkeia:

Arkeia’s Client/Server architecture



2.2 Three-tier architecture

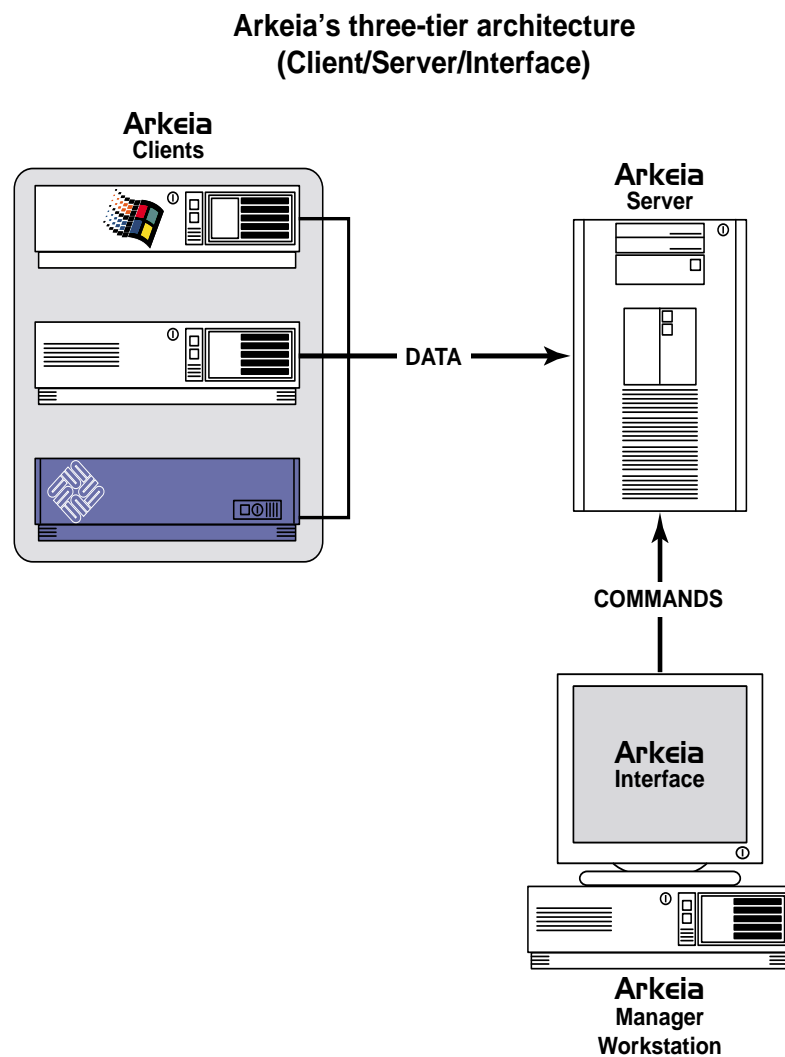
Arkeia adds another layer to the standard “*Client/Server*” architecture by adding a level called “*Interface*”. This layer is particularly important for large and complex networks, with several backup servers.

The “*Interface*” layer allows the system administrator(s) to install the managing interface on any machine of the network. From this machine he/she is then able to configure and control the different backup servers present on the network.

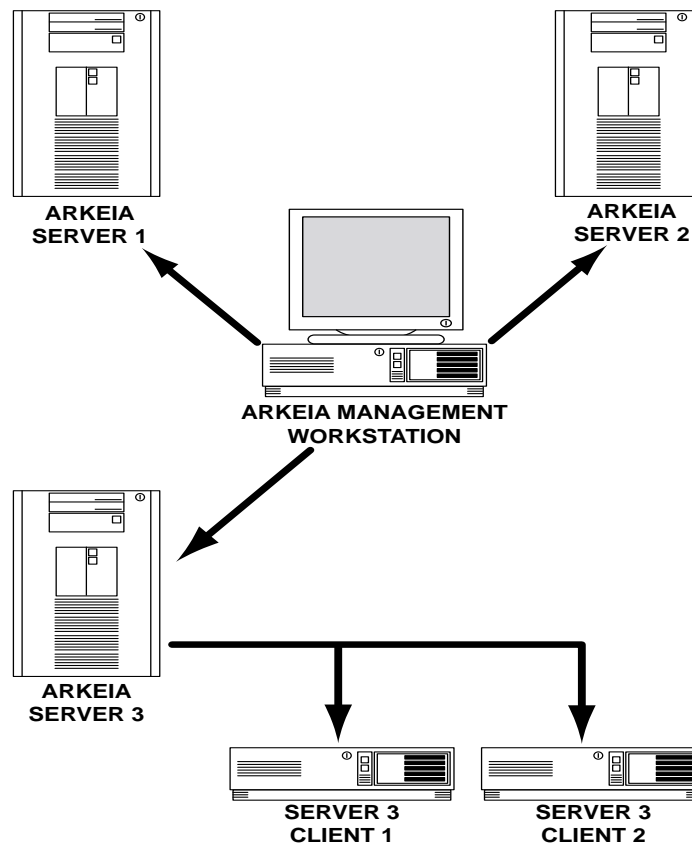
This function is especially important since it offers an unprecedented flexibility to system administrators who are now able to control all the backup and restoration operations performed on their network through a dedicated machine, located anywhere on the network.

With this three-tiered, *Client/Server* architecture, Arkeia version 5 provides system administrators more choices to manage their networked backup servers.

You will find below a graphical representation of this three-tiered architecture:



Another, more realistic representation of this architecture would be:



In the example above, a single administration workstation is used, with the Arkeia interface, to control several servers. Note that only two clients Arkeia clients are represented in the graph above. On a real network, each Arkeia server could backup dozens, or even hundreds, of different clients.

2.3 Heterogeneous networks.

Most companies make extensive use of Local Area Networks (LANs) in their organisation. These Local Area Networks bring together hundreds, and sometimes thousands, of diverse computers all communicating together through the use of the TCP/IP protocol.

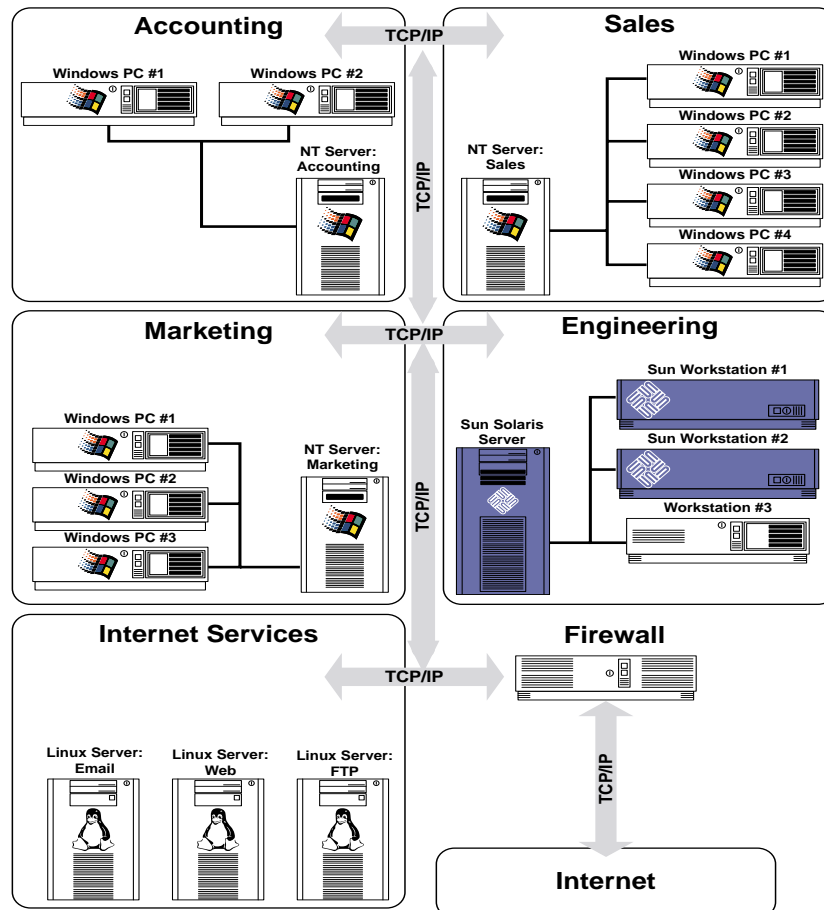
For instance, a typical company may use a large number of inexpensive PCs running several versions of the Microsoft Windows™ operating system, such as Windows 98 or Windows ME, for tasks ranging from simple accounting to desktop publishing (DTP). These microcomputers will, in turn, be connected to one or several local servers running the Windows NT or Windows 2000 operating systems and enterprise software, such as Microsoft Exchange for E-mail.

But the same company may also use UNIX servers and workstations for its engineering and CAD (Computer Assisted Design) departments, using the Sun Solaris™ operating system. Mission-critical 24/7 systems, such as the main accounting and company-wide data center can also be run on these machines which use a number of fail-safe protec-

tions (redundant hardware, RAID disk arrays, Uninterruptible Power Supplies, replication, etc.) to make sure these advanced systems are “always on”.

Finally, to maintain its World-Wide-Web and FTP sites, this company may use inexpensive Linux servers, which are known for their reliability and flexibility. Therefore, our example network could very well look like the one shown below.

A typical modern network architecture



Thanks to the quasi-universal adoption of computing standards, such as Ethernet networks and the TCP/IP protocol, all these computer platforms can communicate and function together in a coordinated manner. But communication is only one part of the equation: the other, much more difficult, part is saving the data these systems contain.

As one can easily deduce from the maze of different systems, saving the information they contain, in a timely and organized manner, is not a simple task. Beyond the communication, most of these computers use different file systems, different authorization procedures and different backup methods.

On the other hand, it is unrealistic to ask the network administrators to perform a separate backup operation for every computer used in a company. Time-consuming, these operations would simply be a duplication of effort and would multiply backup systems, with the associated costs and possible incompatibilities. Another possible problem is the lack of performance of one computer, compared to others: the slowest workstations

could tie up the backup of an entire network, including systems that are mission-critical.

If we return to the network example detailed above, it is easy to see why Arkeia is an ideal solution for an heterogeneous network: it is probably the only backup solution designed from the very beginning to function seamlessly on multiple platforms.

Thanks to Arkeia's Client/Server architecture, all that is required to backup a computer, whatever its type, is an Arkeia client installed on the machine. Currently, Arkeia servers or clients are available for more than 30 operating systems, including market leaders such as IBM's AIX, Sun Solaris, SGI's IRIX, Hewlett-Packard HP-UX, Linux (on several architectures) and Microsoft Windows 9x, Windows ME, NT, XP and Windows 2000.

An Arkeia client, regardless of the platform it is installed on, will be automatically recognized by an Arkeia server. Once recognized, it can be backed up by the server, even if they are not of the same type. For instance, a Sun Solaris or a Windows 98 client can be backed-up by a Linux Arkeia Server.

Because of this architecture, a company is free to use its existing systems as backup servers. With the example detailed above, a logical choice would be to use an Intel-based PC, running the Linux operating system. This would yield two benefits: the network administrators would not have to be trained on a new operating systems, since they supervised Linux systems before, and a cheaper PC would mean a larger budget allocation for the backup hardware itself (as we have seen in section "Client/Server architecture" on page 52).

Advanced benchmarks performed by Arkeia Corporation have shown that an average PC, running under Linux, and connected through a standard SCSI interface to a backup system, can be used to save the critical data of an entire network. This configuration provides an excellent alternative to more expensive backup systems, and a quick return on investment.

3 • Arkeia Internal Architecture

Arkeia's benefits are not limited to its “*general*” Client/Server architecture. Many of its functions are optimized to offer its users better performance and ease-of-use when backing up or restoring large networks.

3.1 Modular architecture

Arkeia version 5 introduces a new modular architecture, which allows the use of additional “plug-ins”. These plug-ins will allow very specialized options to be easily incorporated into Arkeia version 5.

For instance, a short list of the possible plug-ins that will be made possible by the new modular architecture include:

- Database back-up and restoration (MySQL, PostGreSQL, Oracle, etc....).
- Disaster Recovery of damaged machines.
- Application backup and restoration (Apache, MS Exchange, SAP, PeopleSoft, Lotus Notes, etc...).
- Advanced encryption features for data protection (AES, Twofish, etc...).
- “*Hot*” plug-ins, that allow the backup and the restoration of an application data while the application itself is in use.
- And many others!

Arkeia Corporation will announce in the near future several plug-ins designed specifically for Arkeia version 5, as well as their availability. These announcements will be made through the web site of Arkeia, at the following URL:

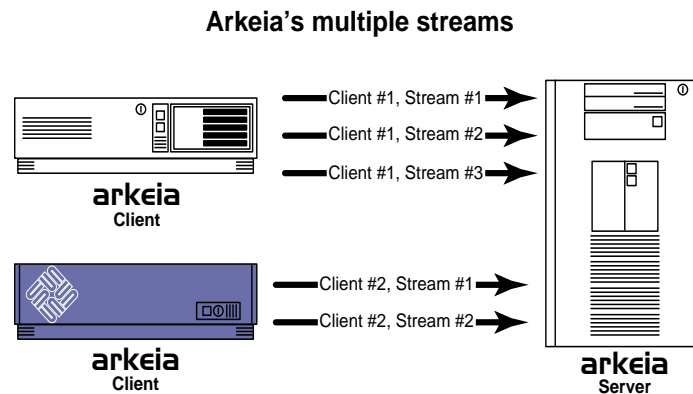
<http://www.arkeia.com>

3.2 Multiple data flows

To maintain high performance during a backup operations, an Arkeia server is able to process several clients at once, which is known as a multiple data flow architecture. This has two benefits: machines with very large amounts of data can open several “*flows*” of data to the server, to be backed up faster, and slower machines do not tie up the backup server.

By default, an Arkeia server uses only one flow per client machine. However, the system administrators using Arkeia version 5 can configure multiple data flows toward a single machine during the creation of a backup.

You will find below a graphical example of a multiple flow architecture:



In the example above, when the Arkeia server has to backup the first client, it opens 3 different data flows as this client may have more data to backup or may be considered more important than the second client. At the same time, it only opens two data flows to the second client, which may have less data to be backed-up than the first one.

3.3 Savepacks

An essential feature of Arkeia, the Savepacks allow system administrators to group the data they need to backup in one easily manipulated set. A Savepack can include files, disk partitions or even complete machines.

Savepacks allow a system administrator to define once, and only once, what they'd like to have backed-up. When the time comes to perform a backup, the only thing left to do is to select the Savepack that has been previously defined.

Since Savepacks can extend over several machines, disks and folders, complex backup tasks need to be defined once, and only once, making the overall backup tasks much easier, even on large and complex networks.

This manual will detail the creation and the modification of Savepacks in chapter “6: Objects and users creation” on page 71.


3.4 Tapes and tape pools

Within Arkeia version 5, each “*tape*” used for backup has to be created. The system therefore creates a crucial distinction between “*physical*” and “*logical*” tapes. A physical tape is a tape as it exists in reality. A logical tape is a tape that has been declared to Arkeia version 5 and that can be used by the software for all its operation.

It is therefore possible to assign some physical tape to Arkeia, by creating an equivalent as a logical tape, all the while reserving other physical tapes to applications other than Arkeia.

In the same way that Savepacks allow a system administrator to combine different elements into a unique element for ease of use, the “*Tape pools*” allow a system administrator to combine different tapes into one coherent whole for back-up and restoration.

For instance, if a system administrator has a backup system which can manage up to 16 tapes, as seen in the figure below:



The system administrator can choose to assign 10 tapes to a single tape pool. A backup operation and/or Savepack can then be configured to use this Tape pool, without using the other 6 tapes.




Please note: It is highly recommended not to mix tapes of different types within the same Tape Pool. Arkeia Corporation recommends that, within a given Tape Pool, all the tapes should be of the same type.

These other tapes can then be assigned to other backup tasks, without interfering with the main Tape pool. This allows system administrators to tailor very precisely the number of tapes that will be used by a backup, without having to create a separate configuration for each backup.

Arkeia manages an internal database, which is used to store all the system information on Savepacks and Tape pools. This allows Arkeia to retrieve the correct tape automatically, for instance if this tape is inserted in a library.

Arkeia uses a default Tape Pool, which is named the “*Scratch Pool*”. Every tape “*created*” (see the paragraph above for the distinction between logical and physical tapes) within Arkeia is first assigned to this pool. Once created and assigned to the “*Scratch Pool*”, a tape can then be moved to other pools. This is explained in the following figure:



- *Scratch Pool*

In the figure above, a new tape is created and automatically assigned to the “*Scratch Pool*” of Arkeia. From the “*Scratch Pool*”, it is moved by the user to the main Tape Pool, which contains all the tapes used for a particular backup operation. Therefore, an administrator could create a Tape Pool for database backup, another for E-mail backup, and so on.

If a Tape Pool is created “*empty*”, without any tape assigned to it, the first backup task that will use this Tape Pool will try to transfer tapes available in the Scratch Pool. This is also the case if, for instance, a Tape Pool has used all its tapes during a backup operation.

Tapes can be created from a Tape Pool management screen, and all tapes created in this way are automatically assigned to this Tape Pool.

Any tape used for a certain period of time can be erased and re-used for other backup tasks. This is called the “*recycling*” of a tape. Based on the configuration chosen by the user, a tape can either be recycled within its own Tape pool or moved back to the “*Scratch Pool*”. For more information on tape recycling and tape assignments, refer to chapter “6: Objects and users creation” on page 71.

3.5 Drivepacks

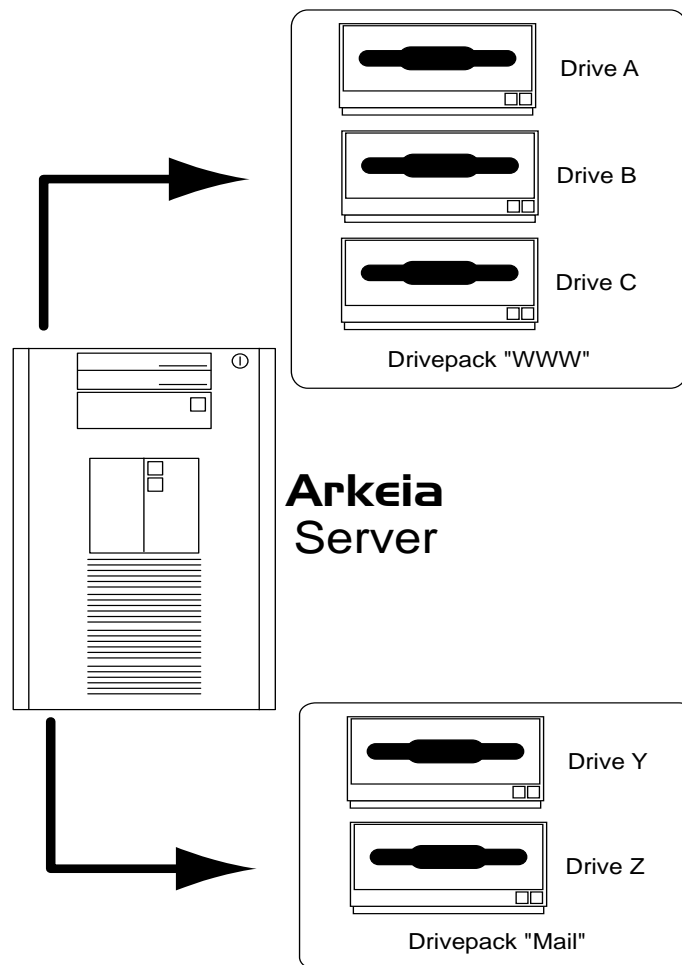
Drivepacks are, within the Arkeia architecture, the equivalent of Save Packs and Tape Pools. A Drive pack allows system administrators to combine several drives in a single unit, which can then be managed by Arkeia version 5.

The advantages are obvious: with a Drivepack, an Arkeia backup server can use several Drives at the same time and therefore complete its backup tasks faster. Through the use of Drivepacks, it is also possible to use larger capacity drives for the most important tasks, while using drives that are slower, or drives that have a smaller capacity, for less urgent or smaller backups.



Please note: Within a given Drivepack, it is highly recommended that all the drives should be of the same type. Using different types of drives within the same Drivepack, while possible, can create many unforeseen problems.

For instance, you will find below an example of the use of Drivepacks:



In the example above, the drives “A” through “c”, combined in the Drivepack “www” with a large capacity, are used by Arkeia to backup the web server, while drives “Y” and “z” are combined in the Drivepack “Mail” and used to backup the e-mail server.



Please note: It is possible, within Arkeia version 5, to create Drivepacks that “share” drives. For example, in the figure above, it would be possible to create a drive named “x”, that belongs to the “Mail” and to the “www” Drivepacks at the same time.

If that is the case, Arkeia version 5 will lock the “x” drive in such a way that it will be only be used by one Drivepack or another, but not by the two Drivepacks at the same time.

For instance, if the Drivepack “Mail” starts using the “x” drive before the “www” Drivepack, the “x” drive will be assigned to “Mail” and “www” won’t be able to access it to perform its own tasks.

CHAPTER 5: GUI CONCEPTS.

- *Summary*

Arkeia Corporation has created, for the version 5 of its flagship product, a very innovative and practical Graphical User Interface (GUI). This chapter is designed to introduce this interface, and make Arkeia users as productive as possible in a short amount of time.

1 • Overview

As has been noted in the chapter: “4: Basic Concepts” on page 50 of this manual, Arkeia has actually a three-tier architecture: client, server and interface. Having an interface which is separate from the client and the server allows the system administrators to control the backup and restoration functions of Arkeia version 5 from anywhere within a network.

Arkeia has actually two interfaces that can be used by system administrators: the Graphical User Interface (GUI) and “arkc” which is the command-line interface. This chapter will only detail the GUI, but a complete manual is available for “arkc” on the Arkeia web site, at the following address:

`http://www.arkeia.com/manuals.html`

This chapter will study three important components of the GUI: the window-related functions, the “list” functions and the main components of the main menu.

2 • The Window functions

The window functions can be divided into two large groups:

- The functions that control the main window of Arkeia version 5.
- The different fields that display information within the window.

These two groups of functions are going to be detailed in the rest of this section.

2.1 Window controls

These functions are grouped in the distinctive “wheel” of the Arkeia version 5, which is located in the upper-left corner of the window:



The following functions are available directly within this “wheel”:

- **“Close”**: Clicking on this button closes the Arkeia window and exits the software immediately (if the window being closed is the last one opened by the software).
- **“Help”**: Clicking on this button displays a short, context-sensitive help on the function or menu opened within the window at the time.
- **“Minimize”**: Clicking on this button will reduce the entire Arkeia window to an icon on your desktop (this may vary, depending on the window manager being used).
- **“Duplicate”**: Clicking on this button automatically duplicate the main window of Arkeia version 5. This allows a system administrator to accomplish two different tasks on the same server (for instance: watching over a running backup in one window while creating another task in the other window) or to work on two different servers at once.
- **“Maximize”**: Clicking on this window makes the Arkeia version 5 large enough to fill the whole screen.

For more control over the size of its window, Arkeia version 5 also provides a resize zone in the lower-right corner of its main window:



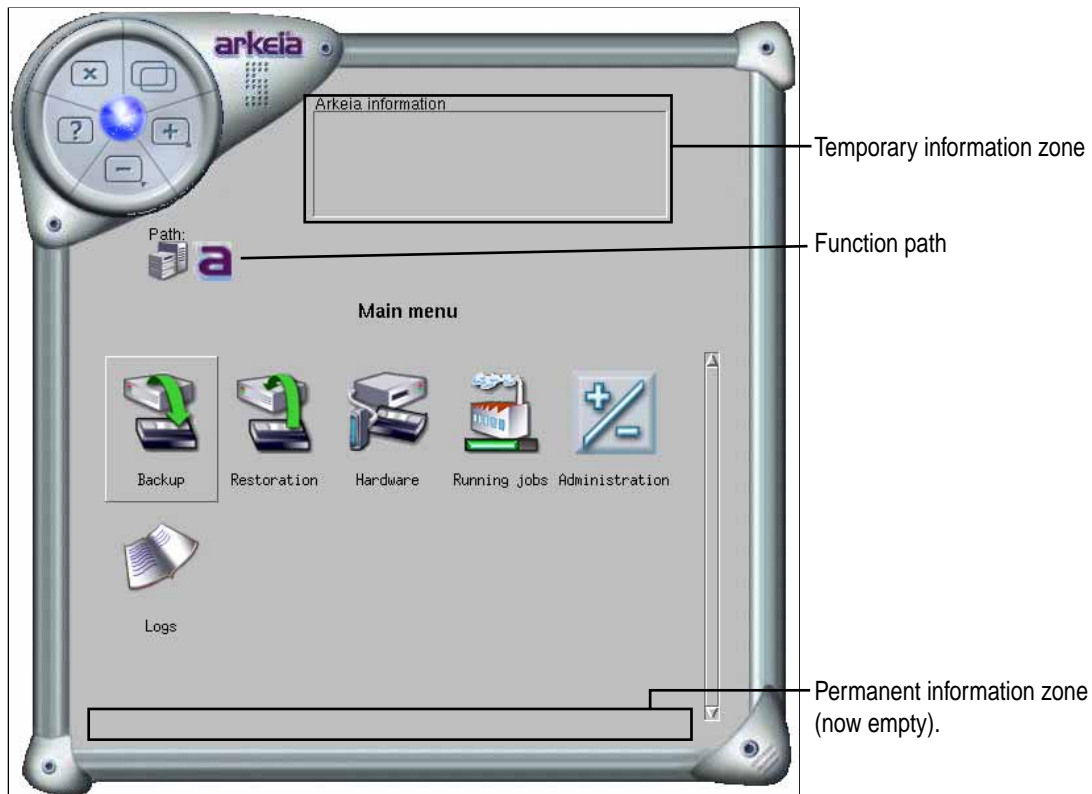
If the Arkeia window needs to be resized, for instance to make some selection easier, this can be achieved by clicking on this zone and dragging the mouse, while keeping the left mouse pressed.

2.2 Information fields

These information fields can be sub-divided into two categories:

- The information fields themselves.
- The “function path” which provides a quicker way to navigate through the different menus of Arkeia version 5.

These different fields are indicated below in the window of Arkeia version 5:



Within this window, the three fields indicated above can be found quickly:

- The “Temporary information zone” displays the Arkeia messages for a little while and is then cleared out.

- The “**F u n c t i o n p a t h**” displays the different menus that have been browsed to reach the current menu. Clicking on each icon allows the operator of Arkeia version 5 to quickly “jump back” to previous menus. These icons are also displayed next to the menu title.
- The “**P e r m a n e n t i n f o r m a t i o n z o n e**” displays the latest message and/or confirmation from Arkeia version 5 in a more permanent manner. This zone usually confirms that an operation has been performed correctly by the Arkeia server. Clicking on this zone will open a separate window, that contains all the messages about the operation being performed.

Using these fields can be a tremendous help in the day-to-day operation of Arkeia version 5. All the users of Arkeia version 5 should therefore try to learn and use these functions of the GUI.

3 • The list functions

Arkeia divides all its menus into two separate lists: the “object” list and the “function” list.

Objects are the Arkeia entities that are going to be created, modified, deleted, or configured through the list of functions. We are going to use the tape, which is one of the most basic entities of Arkeia, as an example in this section.

Here is the “Tapes Management” menu of Arkeia, as it is displayed in the main Arkeia version 5 window:



As can be seen in the example above, this menu is divided into two different zones: The zone on the left is the “object list”, which contains 8 different tapes, named “test01” to “test08”.

The zone on the right of the menu is the “function list”. All the functions that are listed there can affect all the objects in the object list, or just the object that have been selected by clicking on the grey check box next to their names. In the example above, it is possible to “recycle” a tape by first selecting it, then clicking on the “Recycle” function. Or, it is possible to recycle all the tapes by clicking on all the grey check boxes, then on the “Recycle” function.

Above the “object list” (the list of tapes of our example above), Arkeia offers four basic functions that are very useful for quick operations. These functions are, from left to right:

- “Back” (Green arrow): Clicking on this button takes Arkeia back to the previous menu.
- “Refresh” (Dual green arrows): Clicking on this button redraws the entire list.

- “**New**” (Blank document): Clicking on this button starts the creation process of a particular object. In our example above, this would start the tape creation process.
- “**Delete**” (Dust bin): Clicking on this button deletes the objects that has been selected in the list of objects. In the example above, this would delete one or several tape(s), depending on the check boxes that have been selected.

Simply leaving the mouse pointer on these buttons for a few seconds will display a short help “balloon”, which contains a definition of each button function.

4 • A short example: the main menu

As a final, short example, this section will detail the most important functions of the main menu of Arkeia version 5:

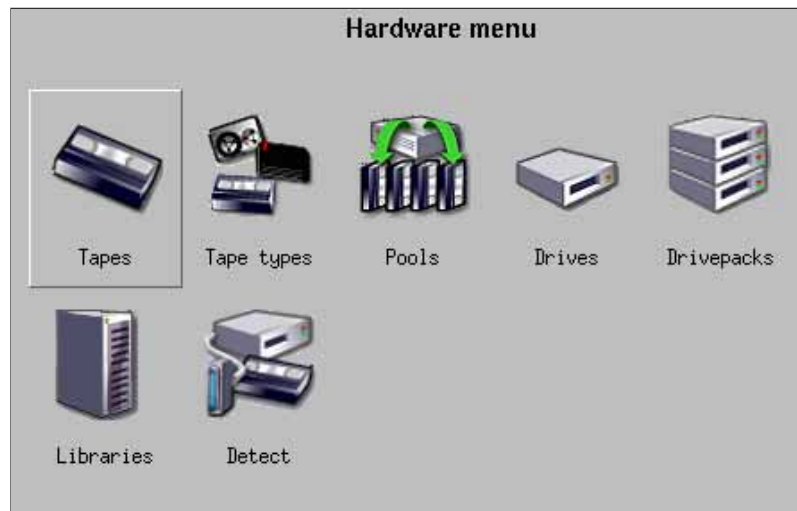


In the example above, it is possible to see the main functions of Arkeia version 5, as well as the functions we detailed in the rest of this chapter. The functions that can be accessed through a single click on one of the major buttons in the main menu are the following:

- **“Backup”**: Clicking on this button provides functions to define, configure and launch the backup function of Arkeia, including the Savepacks.
- **“Restoration”**: Clicking on this button provides functions to define, configure and launch the restoration function of Arkeia version 5.
- **“Hardware”**: Clicking on this button provides functions to define and configure the hardware (drives, tapes, libraries) attached to the Arkeia server.
- **“Running jobs”**: Clicking on this button displays all the tasks currently running on the Arkeia server.
- **“Administration”**: Clicking on this button provides functions to precisely configure Arkeia version 5.

- “Logs”: Clicking on this button provides functions to display and configure the logs of Arkeia version 5.

Some functions of this main menu may take you to a sub-menu. This is the case, for instance, of the “Hardware” button, which will take you to the “Hardware menu” of Arkeia version 5:



The rest of this manual will explain, in a detailed way, how to use the different menus and sub-menus of Arkeia version 5.

CHAPTER 6: OBJECTS AND USERS CREATION

- *Summary*

Before using Arkeia version 5, a system administrator should define the basic objects that will be used by the software itself. This chapter will detail the different objects that can be created within Arkeia version 5 and their usage.



Please note: This chapter only presents the creation of “logical” objects. To create “physical” objects, such as drive and libraries, refer to “3: Hardware installation” on page 35 of this manual, or to chapter “10: Advanced Hardware” on page 158. For a discussion on the different between logical and physical objects, refer to chapter: “4: Basic Concepts” on page 50 of this manual.


1 • Creating tapes

The first step any user of Arkeia version 5 should take is to create tapes within the system in order to be able to backup data.

Creating tapes within Arkeia is a simple task, which we are going to detail step by step below:

1.1 Log on the Arkeia server

Once the Arkeia graphical user interface (GUI) is running, the login screen should be the first one to appear:




In the screen displayed above, only one server is available, the machine named: “Voyager”. To log on a machine, click on its name. Arkeia version 5 will then request the user name and its password:



In the example above, we enter the user name “root”, followed by the password. Note that your installation of Arkeia version 5 may use different names, and that users with different privileges may create tapes. For more information on users and their different privileges, refer to “Creating users” on page 99. For more information on your local installation of Arkeia, contact your system administrator or IT Department.

Once you have logged on the requested server, the main menu appears:



If you need more information on the login procedure, refer to chapter “5: GUI Concepts.” on page 62 of this manual.

1.2 Create the tape(s)



To create a tape, click on the “Hardware” button of the main menu. This button is pictured on the left.

This will open the “Hardware menu” of Arkeia version 5.



Within the “Hardware” menu, click on the “Tapes” button.

Arkeia then opens the main “Tapes” menu, which is a list of all the tapes defined on the selected server.



Please note: If no tapes have been defined on the selected Arkeia server, a message will appear in the tape list, saying:

[No tapes]

Once the “Tapes management” screen has been displayed, it is possible to create tapes by using the buttons displayed above the main list of tapes. The options of this screen are detailed below:



- “Refresh”: This option refreshes the list of tapes and redraws the contents of the screen.
- “New”: This option creates one or more backup tapes or cleaning tapes.
- “Delete”: This option deletes one or more tapes. To delete several tapes at the same time, click on the check boxes next to their names.

- “**M o d i f y t a p e (s)**”: This option allow one or more tape(s) to be modified once created.
- “**P r o p e r t i e s**”: This options display the current tape properties.
- “**A d v a n c e d o p t i o n s**”: This option allows the properties of the tape to be modified once it has been created.
- “**W r i t e l a b e l**”: Allows the tape labels to be re-written manually.
- “**H i s t o r y**”: Displays information about the tape usage.
- “**R e c y c l e t a p e (s)**”: This function allows Arkeia to reuse a tape from the beginning. The tape contents are not erased until a new backup uses the recycled tape.
- “**S t a t i s t i c s**”: Displays relevant usage statistics for the selected tape.

Clicking anywhere in the “**T a p e s m a n a g e m e n t**” screen with the right mouse button displays a contextual menu, with the two following options:

- “**O K**”: Executes a command or goes back to the previous “**H a r d w a r e**” menu.
- “**C a n c e l**”: Cancels a command and go back to the “**H a r d w a r e M e n u**”.

If some of the above options are not activated, they will appear greyed-out.

To create a new tape, click on the “**C r e a t e t a p e (s)**” option. This will open the tape creation screen, which is pictured below.

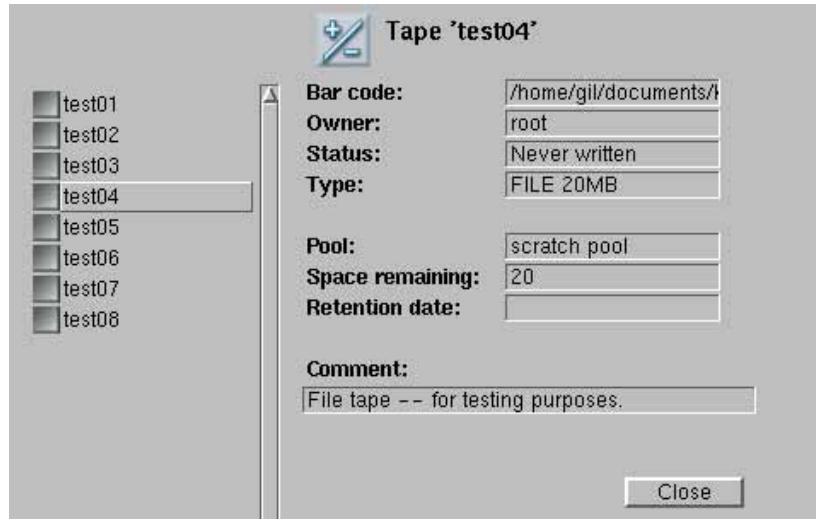
The following options are available in this screen:

- “**Backup/Cleaning tape**”: These check boxes determine the type of tape that will be created by Arkeia version 5.
- “**Tape name**”: Self-explanatory. Enter a descriptive name in this field. For instance: “**Web tape 001**”.
- “**Type**”: Click on the drop-down menu to select the type of tape to be created.
- “**Pool**”: Allows you to choose the pool the tape(s) will be assigned to. The default pool is the “**Scratch Pool**”. The “**Scratch Pool**” is used by Arkeia as a reserve of blank (unused) tapes. Whenever a backup operation needs tapes to be completed, Arkeia automatically assigns tapes from the “**Scratch Pool**” to the task being performed.
- “**Bar code**”: This field is used only if a “**File tape**” is used. In this case, the bar code is the path to be used when saving data to the tape. Otherwise, the bar code, is automatically assigned by the tape drive the first time the tape is inserted in the drive.
- “**First #**” and “**Last #**”: If it is necessary to create several tapes in one operation, enter the first number and the last in the corresponding fields. For instance, to create 5 tapes, enter “1” in the “**First number**” field and “5” in the “**Last number**” field. Please note: a digit will be appended to the name of the tape, to reflect the creation of several tapes at the same time.
- “**Owner**”: This field indicates the owner (or the “creator”) of the tapes. Entering a name other than “**root**” in this field will give control over the tapes to the user whose name has been entered. If unsure, leave this field as is.
- “**Comment**”: This field can contain a short descriptive comment about the tape.

Once these parameters have been set, clicking on the “OK” button at the bottom of the “**Create tape**” screen will automatically create the tape(s).

Once a tape has been created, it should appear in the list of tapes, which is displayed in the left-hand side of the tapes-related screens. Double-clicking on its name will then display the relevant information for this tape.

You’ll find below an example of the values that can be displayed for a tape.



The example shown above is for 20 MB “F i l e” tape, which is located in a directory on the hard disk of the Arkeia server.

The tapes you create on your system may have different values, instead of the ones shown above.

1.3 Tape types

The tapes supported by Arkeia can be displayed in a window of the main Arkeia interface.

To display the list of tapes supported by Arkeia, follow the steps detailed above:

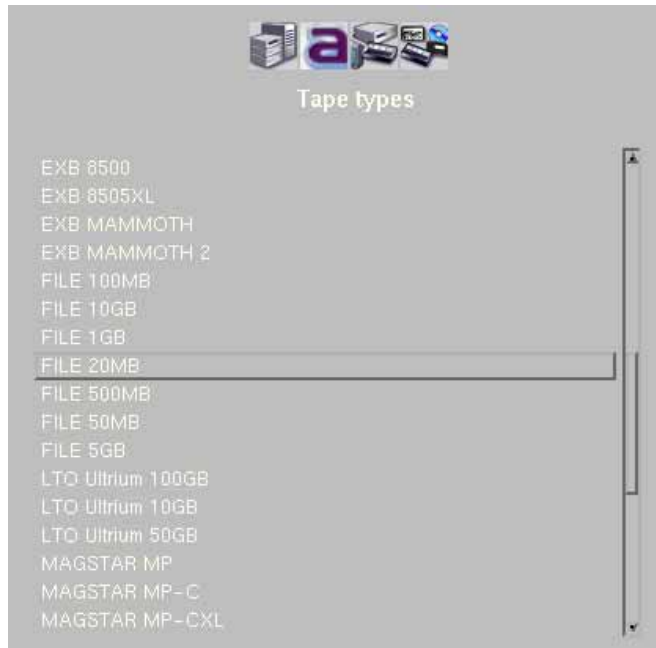


In the main menu of Arkeia, click on the “H a r d w a r e” icon pictured on the left.



In the “H a r d w a r e” menu of Arkeia, click on the “T a p e t y p e s” icon pictured on the left.

Arkeia then displays all the tapes that it can manage, in a list that you will find below.



To display more information on a given tape type, double-click on its name. This will display a small summary of the capacities of the tape.



As you can see above, this example screen applies to the tape type named "EXB MAMMOTH 2".

2 • Creating tape pools

Tape pools are groups of tapes that are created to perform a well-defined backup/restore task. For instance, an Arkeia administrator can create one tape pool for every server that he has to manage:

- Tape pool: “FTP” (for the FTP server backups and restorations).
- Tape pool: “www” (for the World wide web server)
- etc....

Or a tape pool that spans several servers of the same type.



As noted in the section named “Tapes and tape pools” on page 58 of chapter “6: Objects and users creation”, tape pools should only contain one type of tape. Tape pools containing more than one type of tapes may end up producing serious errors and incompatibilities.

Creating a tape pool is a very simple task, since a tape pool is an “empty” object. To add tapes to a tape pool means entering the name of the pool while creating the tape. For more information on how to create a tape, refer to section “Create the tape(s)” on page 74 of this chapter.

To create tape pools, follow the steps detailed below.

2.1 Log on the Arkeia server

To complete the logon procedure, refer to the section named “Log on the Arkeia server” on page 72 of this chapter.

If you are already logged on an Arkeia server, you do not need to log off and on again, unless you would like to create tape pools on another Arkeia server.

2.2 How to create a tape pool

Once logged on the server where a tape pool should be created, Arkeia’s main menu is displayed on the screen.



To create a tape pool, click on the “Hardware” button of the main menu. This button is pictured on the left.

This will open the “Hardware menu” of Arkeia version 5.



Within the “Hardware” menu, click on the “Pools” button.

Arkeia then opens the main “Pools management” menu, which is a list of all the tape pools defined on the selected server.

A list of the tape pools already defined on the server should then be displayed, in a window titled: “Pools management”, which you will find below. This window is divided into two list: the left-hand list contains all the pools that are defined on the Arkeia server, while the right-hand list contains the functions that can be applied to the different pools.



In the example above, there is only one tape pool defined on the Arkeia server, the “Scratch pool”, which is the default Arkeia pool. Note that the four icons (“Back”, “Refresh”, “New” and “Delete”) above the left-hand pool list are the same than the ones described in the “Tapes management” screen.

To create a new tape pool, click on the “New” button above the left-hand pool list to display the “Pool creation” screen. This screen, along with a description of its different options, is detailed below.



Please note: To perform the “Delete” and the “Statistics” operations offered in the pool management screen, a tape pool should be selected first with the help of the left mouse button.

Clicking on the “Create pool” option in the contextual menu will display the creation menu, that you’ll find below.

The information to be entered in this window is just the name of the tape pool (“Pool name” field), the “owner” of the tape pool (refer to the section titled “Arkeia users categories” for more information on users) and a short “comment” if necessary. By default, the “owner” of the tape pool is the user currently logged on the Arkeia server.

To go from one field to the next, use the “Tab” key on your keyboard. To create the tape pool, just press the “Enter” on your keyboard, or right-click anywhere on the screen to display the contextual menu and click on the “OK” option of this menu.

Once a tape pool has been created, it should appear within the list of tape pools, as can be seen below, where tape pools “Documents”, “foobar01” and “foobar02” have been created:





Please note:

It is possible to create tapes from within a tape pool. To do this, double-click on the name of a tape pool in the “**Pools Management**” screen, and click on the “**New**” button above the tape list. Refer, in this chapter, to section “**Create the tape(s)**” on page 74 for more information on the tape creation procedure.

It is possible to create tape pools that do not contain any tape. If a tape pool is created empty, it will request “free” tapes from the “**Scratch pool**” when tapes are needed to perform a backup operation. If that is the case, it is up to the Arkeia administrator to make sure the “**Scratch pool**” contain enough tapes to supply them when needed.

2.3 Checking a tape pool statistics

After creating a tape pool, it may be useful to check its statistics, to make sure it was created with correct characteristics. To do this, click, within the “**Pools management**” screen, on the “**Statistics**” option.

Arkeia should then display a screen such as the one you’ll find below, that displays the statistics for the tape pool named “**Documents**”.

Statistics of pool 'Documents'	
Backup tapes	
Total:	2
Free:	2
In use:	0
Full:	0
Worn-out:	0
Pool space (in MB)	
Total:	40
Used:	0
Available:	40

On this screen, you can notice that tape pool “**Documents**” contains two backup tapes, that have not been used so far. Displaying tape pool statistics is therefore an easy way to determine if a given tape pool has been created correctly or needs more tapes to perform its functions.

3 • Creating savepacks

3.1 Overview

Within Arkeia version 5, a Savepack can be defined as a group of objects that should be saved to tape by a backup operation. The objects that Arkeia can backup can be machines, entire disks, directories, individual files or plug-ins. Within the Arkeia architecture, a plug-in is an additional program that is treated as a file system by Arkeia.

For instance, within a Savepack named “**L a r g e B / U**”, it is possible to add:

- Six entire disks from a machine named “**f o o b a r**”.
- Several directories from a machine named “**s n a f u**”.
- One plug-in from a machine named “**r t f m**”.

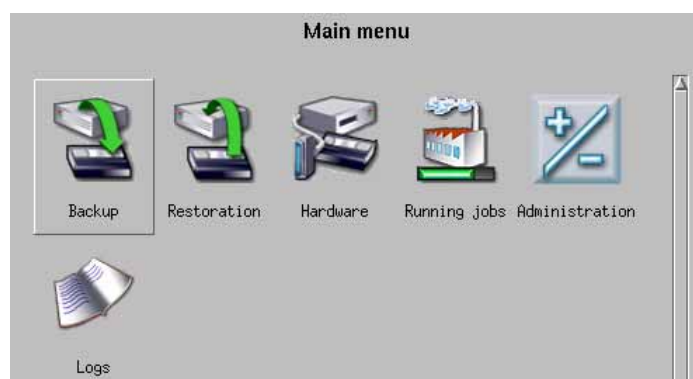
Whenever the “**L a r g e B / U**” Savepack is chosen as the target of a backup operation, all these objects, from the three different machines, all these object will be saved together, from the three different Arkeia client to the Arkeia server.

Furthermore, all the objects (or “trees”) can have different options applied to them. For instance, the directories from the machine named “**s n a f u**” can be compressed if they are quite large, while the data from the plug-in should be encrypted to protect its confidentiality. All these operations are made possible through the definition of the Savepack.

A Savepack is therefore a very powerful tool in the hands of the person responsible for protecting the data of a network. This section will study the creation of a Savepack, as well as the different options that can be applied to a Savepack.

3.2 Creating an empty Savepack

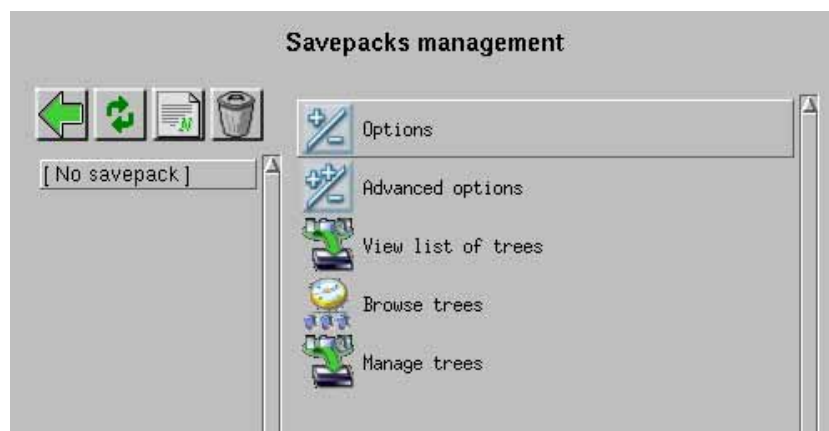
To create an empty Savepack, that will receive trees for backup, click on the “**B a c k u p**” icon in Arkeia’s main menu:



Then, click on the “**savepack**” icon in the “**Backup**” menu:



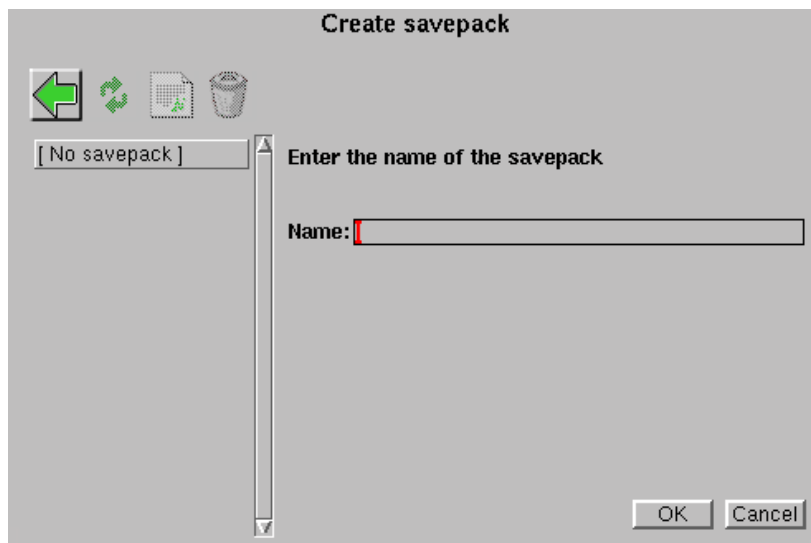
The “**savepacks management**” screen should then be displayed within the main window of Arkeia version 5:



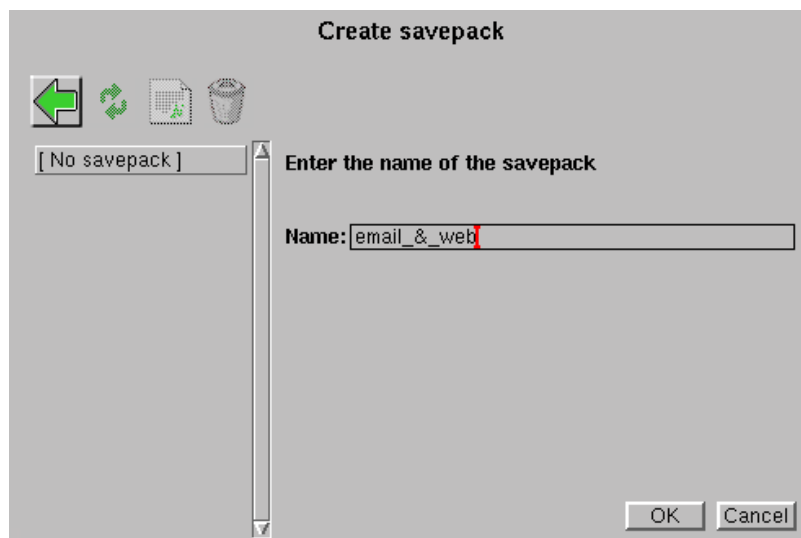
If there are no Savepacks defined in Arkeia (as is the case in the screen pictured above), Arkeia will simply display the following message: “[**No Savepack**]”. Within this screen, the following functions are available:

- “**Options**”: This function defines the basic Savepack options: compression, encryption, command to be executed before backup, command to be executed after backup and comment.
- “**Advanced options**”: This function defines additional options that can be globally applied to a Savepack: number of retries, file system management, access time management and include/exclude filters.
- “**View list of trees**”: This function displays the list of trees that are contained within the Savepack.
- “**Browse trees**”: This function lets the backup operator select which trees (machines, drives, directories, files and plug-ins) should be included in the Savepack. This selection is made through a file system navigator.
- “**Manage trees**”: This function lets the backup operator add and/or delete trees from the Savepack in a quicker manner. Within this list, it is possible to configure the options for a specific tree. Going back to our first example, it is possible to use this function to include/exclude certain files, or apply a compression strategy to a specific tree, which is different from the overall Savepack compression strategy.

Within the “Savepacks management” window, clicking on the “New” icon – above the list of Savepack – will open the “Create Savepack” screen:



After entering the name of the Savepack, clicking on the “OK” button will return to the Savepacks management screen. In the example below, a Savepack named “email_&_web” has been created.

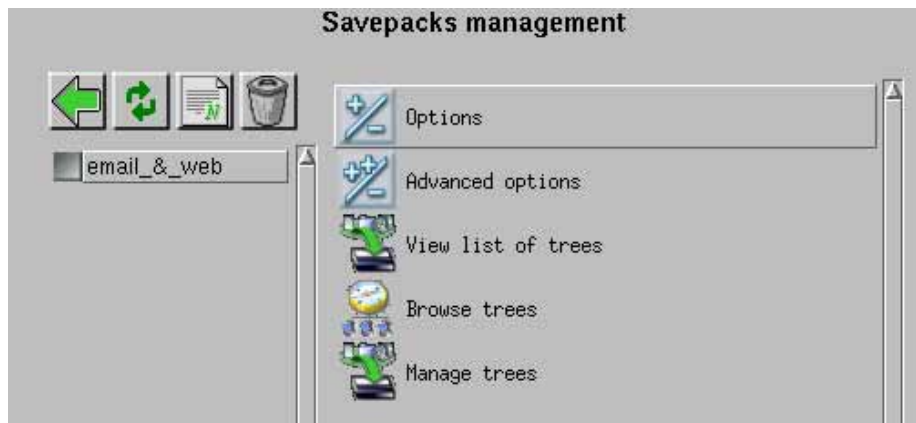


3.3 Adding trees to a Savepack

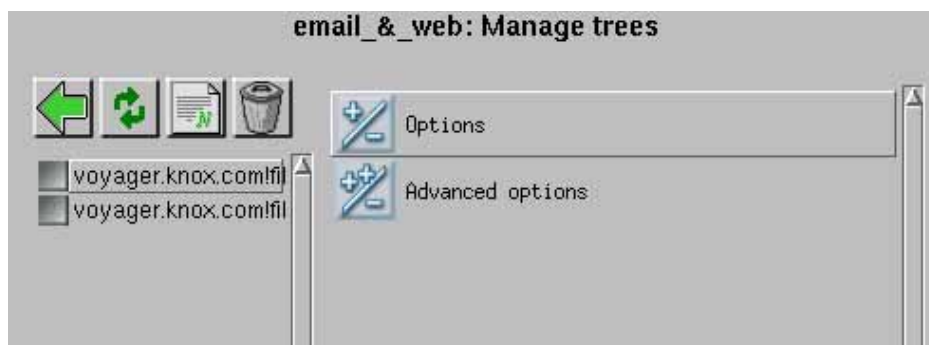
Adding trees to a Savepack can be done in two different ways: by entering the complete path to a given object or by choosing the object within the Arkeia browser. It is highly recommended to use the browser, but both methods will be described below for the sake of completeness.

- *Adding an object path to a Savepack*

To add an object to a save pack, click on the option named “Manage Trees” in the “Savepacks management” screen:



This will open the “Manage trees” menu within the main window of Arkeia version 5:



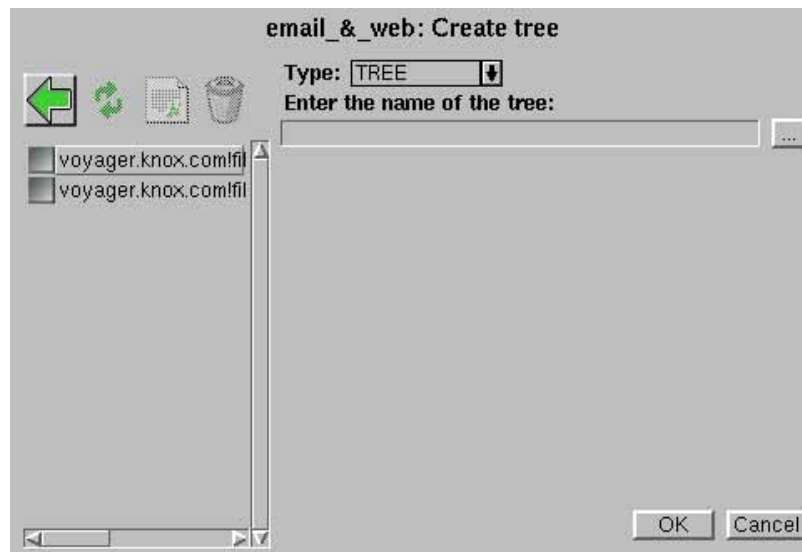
The following options are available through this screen:

- “Options”: This function defines the basic tree options: compression, encryption, command to be executed before backup, command to be executed after backup and comment.
- “Advanced options”: This function defines additional options that can be globally applied to a tree: number of retries, file system management, access time management and include/exclude filters.



Please note: The options described above apply only to the selected trees that appear in the left-hand tree list. They do not apply to the Savepack as a whole.

In the “Manage trees” screen, clicking on the “New” icon opens the “Create tree” window within the main window of Arkeia version 5. This screen allows a new tree to be created within the selected Savepack:



Within this screen, The following options are available:

- “Type”: This option determines the type of the tree. There are four different types of trees:
 - “Tree”: A tree can be the name of machine, or “hostname” (to backup an entire machine), or a directory to backup only this specific directory, and all the files and sub-directories it contains, or a single file. Please note that it is possible to choose the object through the Browser by clicking on the browser (“...”) button next to the field under “Enter the name of the tree”.
 - “Object”: An “Object” tree will back-up and restore the result of a command. This command should be entered in the “Path” field. Please note that the file created by the “Object” command will not reside on the Arkeia client, but an entry will be created in the Arkeia database to restore the data backed-up.
 - “Raw”: A “Raw” tree is either an entire disk or an entire partition. If a “Raw” tree is to be used, the Arkeia Browser should be used to select the disk or the partition present in the “/dev” directory of the UNIX machine.
 - “Savepack”: A “Savepack” tree appears only if the chosen tree is, itself, a Savepack.

In the most frequent case, the tree to be added to the Savepack should use the “Tree” format. In which case a complete case can be entered by respecting the following syntax:

```
machine_name!plugin_name:/path1/path2/.../file
```

For instance, the following paths are considered correct for Linux and UNIX machines:

```
foobar.acme.com!file:/home/jack/files/reports/pdf/dfgf2002.pdf
```

```
snafu.acme.com!file:/home/tina/Mail/
```

Please note that the full path always uses slash “/” characters, even if the machine to be backed-up is a Windows machine, which usually use backslash characters “\” in its path syntax. Windows machines are also a special case, since a path on a Windows machine must indicate which volume/hard disk is being used. A normal path for a Windows machine should therefore be:

```
machine_name!plugin_name:volume_letter:/path1/path2/.../file
```

For instance, the following paths are considered as correct for an Arkeia Windows client:

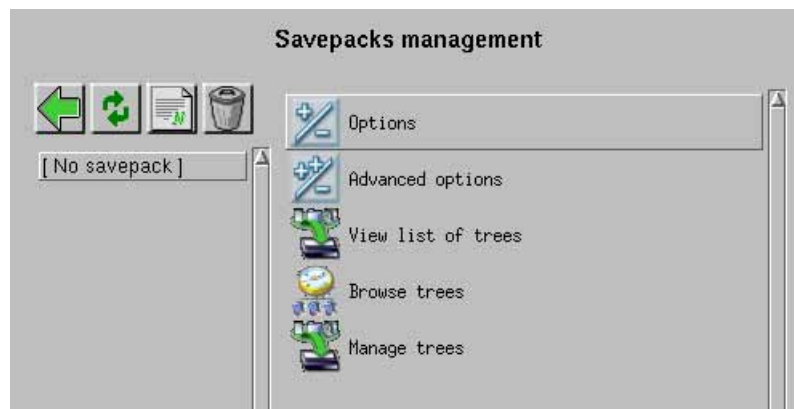
```
rtfm.acme.com!file:c:/my documents/reports/2002/sales_q3.xls  
jaguar.vector.com!file:e:/archive/intel/tech/PalladiumTechNotes.pdf
```

Once the path has been correctly defined, clicking on the “OK” button returns to the main “Manage trees” screen.

While entering a tree by its complete path/reference can be practical in some instances, (especially when creating “RAW” trees), creating complex Savepacks in this way can be very time-consuming. To make the creation of a Savepack much easier, Arkeia Corp. has created the “tree navigator”, which is an advanced function to rapidly select different trees to incorporate them in the same Savepack.

- *Adding a tree through the Navigator*

To add a tree to a Savepack by using the Navigator, click on the “Browse trees” option, within the main “Savepacks management” screen:



This opens the navigator in the main window of Arkeia version 5:



All the clients that have been declared on the Arkeia backup server appear in this list. In the example above, there is only one client, named “`voyager.knox.com`”. To browse the file systems contained in a client, double-click on its name, then on the “`file`” directory. The entire disk layout of the client is then displayed within the Navigator.



In the example represented above, the entire disk layout of a Slackware 8.0 machine is available for backup. To select a single directory (as well as all its subdirectories), the System Administrator using Arkeia can click on the grey check box next to its name. Clicking on the name of the directory allows the administrator to open the directory and select individual files or subdirectories.

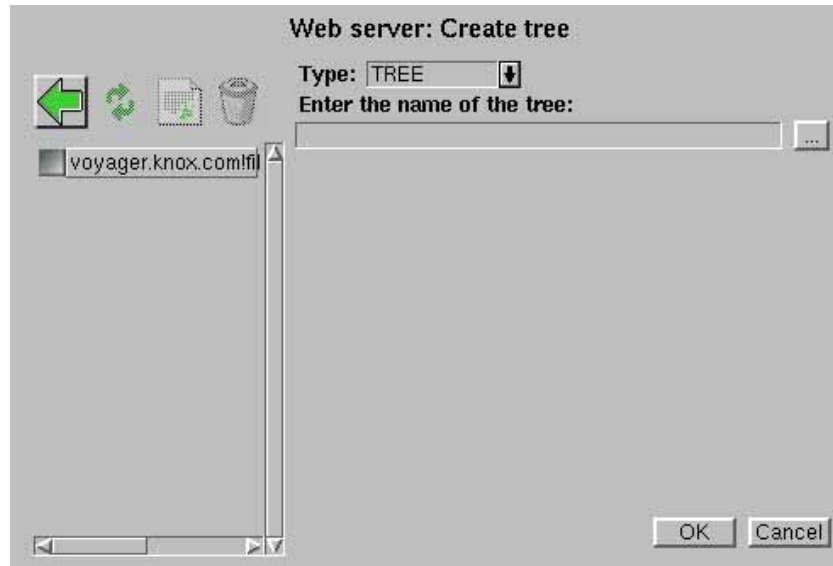
Once the different trees to be added to the Savepack are defined, right-click anywhere in the Navigator window and click on the “`Validate`” option to confirm the choices and return to the main Savepack menu.

- *Adding a Savepack to another Savepack*

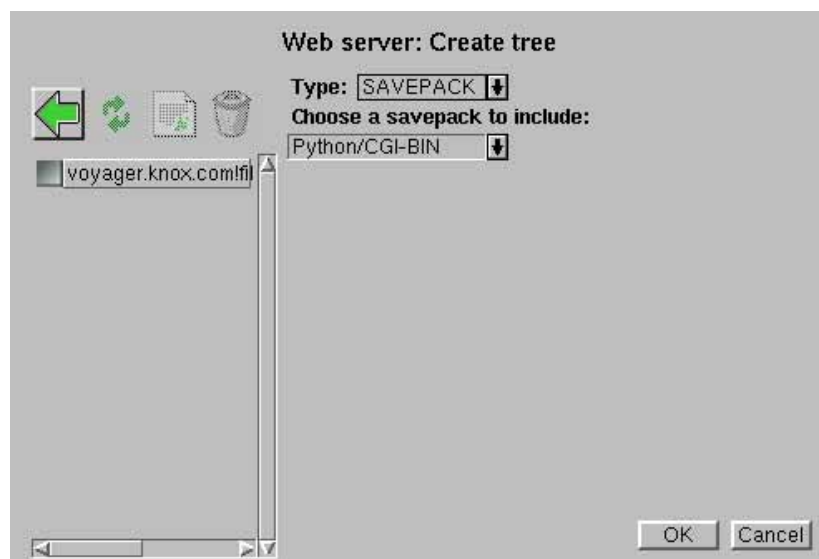
It is possible to incorporate a Savepack into another Savepack. This allows complex configuration to be realized in a more simple manner, since each Savepack can have its own configuration, and modifying the parameters of one does not mean modifying the entire “meta” Savepack.

To add a Savepack to another Savepack, follow the procedures described in the section titled “**Advanced options**”: This function defines additional options that can be globally applied to a Savepack: number of retries, file system management, access time management and include/exclude filters.” on page 84 of this manual, then proceed as described below.

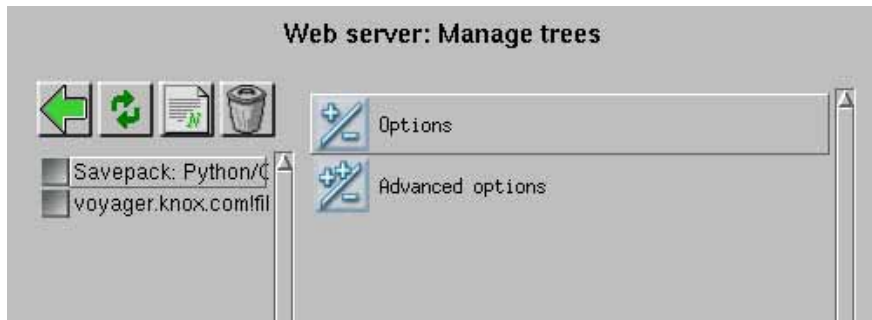
When Arkeia version 5, in the “**create tree**” screen, requires the operator to select a tree to be added, the “**savepack**” option should be chosen in the “**Type**” drop down list:



Once the “**savepack**” option has been selected, the operator should select the name of the Savepack to be added. In the example shown below, the Savepack named “**Web server**” will receive a secondary Savepack named “**Python/CGI-BIN**”:

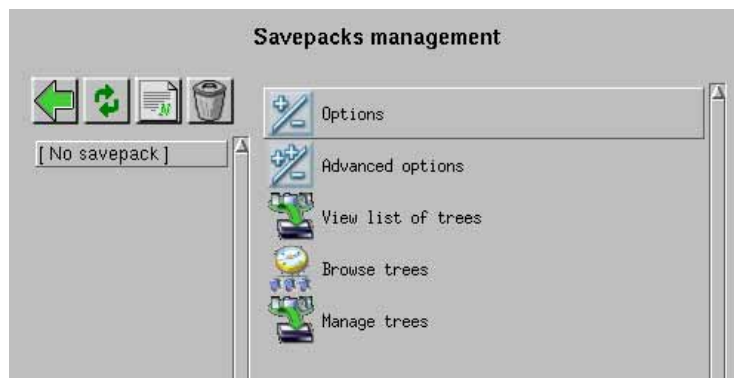


Once the Savepack to be added has been chosen, clicking on the “OK” button returns to the “Manage trees” menu. This menu displays, on its left, the list of trees to be backed-up, including the Savepack that has been added, as can be seen below.



3.4 Defining options for the Savepack

The options that can be applied to a Savepack are controlled through the “Options” and “Advanced options” options of the main “Savepacks management” screen:



Clicking on the “Options” function opens the corresponding screen within the main Arkeia version 5 window:



Note that, in the screen pictured above, the options chosen will apply globally to the entire Savepack (in this case: “email_&_web”). The following options are available in this screen:

- “**Compression type allowed**”: Click on this drop-down list to select the type of compression that will be applied to the files and directories included in the Savepack. Arkeia includes the “Lempel Ziv” compression scheme to obtain better performance on large files. The types of compression are: “**Only LZ1**” (fast compression, lower compression ratio), “**Only LZ3_3**” (slower than LZ1 but better compression ratio), “**LZ1 or LZ3_3**” (Arkeia determines which compression scheme is to be used based on a configuration file) and “**No compression**”.
- “**Encryption type allowed**”: Click on this drop-down list to select the type of encryption that will be applied to the files and directories included in the Savepack. The encryption schemes that are available are: “**Only DES**” (standard encryption scheme), “**Blowfish**” (more advanced than DES, but less standard), “**DES or Blowfish**” (Arkeia determines the encryption scheme to be used, based on a configuration file) and “**No encryption**”.
- “**Command before savepack**”: It is possible to enter a command in this field that must be executed on the client before the backup of the savepack.



Please note: Arkeia won't perform the backup if the command entered in the "Command before savepack" field fails and returns a standard error code. For instance, if a "tar" command should be performed on some files before the backup, and "tar" returns a standard error code (an error code larger than "0 -- Success"), then Arkeia will not perform the backup.

If the command did not perform as expected and returned an error code, but the backup should be performed nonetheless, click on the check box "Backup Savepack if command fails" above the entry field.

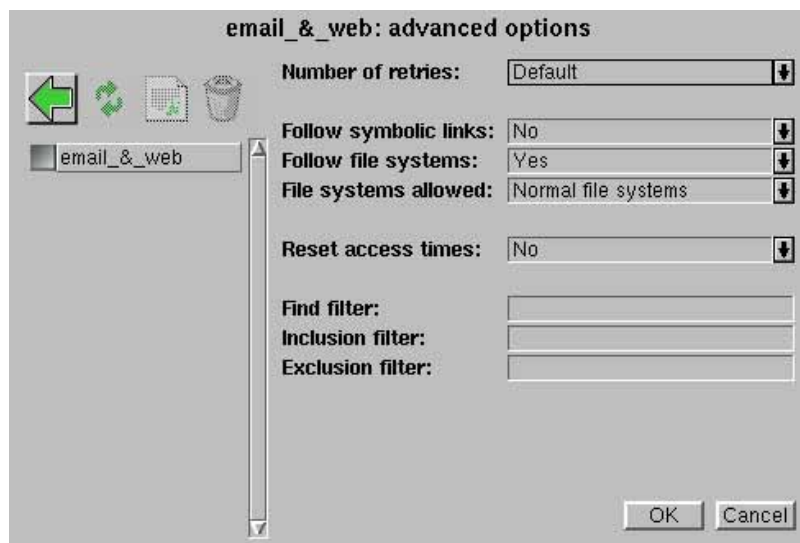
Please note that a command can be any standard UNIX command, a shell program, a batch file (on Microsoft Windows machines) or any other program that can be launched from a command line. Therefore the following commands are valid entries for a "Command before Savepack":

```
/home/john/scripts/dump.pl (Perl script)
./zcache -xvz --quiet (local executable)
/opt/tc/tcversion.py (Python script)
```

- "Command after savepack": In the same way that a command can be entered to be executed before the backup of the Savepack, another command can be entered to be executed after the backup itself (see above for more details). To allow the execution of this command even if the backup itself has failed, click on the "Execute if savepack backup fails" check box above the entry field.
- "Comment": This field allows a short descriptive comment to be included in every Savepack.

Once the above options have been configured, clicking on the "OK" button confirms the choices and returns to the "Savepacks management" screen.

Like the "Options" function, clicking on the "Advanced options" function in the "Savepacks management" screen opens the second options screen that can be applied to a Savepack:



- **“Number of retries”**: Click on this drop-down list to select the number of retries that will be applied globally to the Savepack. The default value is “3” and the maximum value is “10”. Entering a value of “0” will force Arkeia not to retry the backup of the Savepack.
- **“Follow symbolic links”**: Click on this drop-down list to choose if Arkeia should follow symbolic links or not during the backup of the Savepack. The default value is “No”.



Please note: Arkeia does not detect circular symbolic links. If this type of links exist in the directories backed-up - for instance if directory “A” contains a link to directory “B”, which contains, in turn, a link to “A” - the backup function may fail.

- **“Follow file systems”**: Click on this drop-down list to choose if Arkeia should backup different file systems or not during the backup. The possible values are “Yes” and “No” and the default value for this option is “Yes”.
- **“File systems allowed”**: Click on this drop-down list to choose which type of file systems will be accepted (and backed-up) by Arkeia.

The possible values are:

- **“Normal file systems”** (excluding CD-ROM file systems).
- **“All file systems”** (this includes NFS and CD-ROM file systems).
- **“All except NFS”** (this option includes CD-ROM but not NFS file systems).
- **“Reset access times”**: Click on this drop-down list to choose whether or not Arkeia should reset the access time of the files that are backed-up.

Whenever a backup is launched, Arkeia will modify the access time of the files it backups. If this access time should not be modified on the system being backed-up, this option should

be set to “Yes”.

The possible values are “Yes” and “No” and the default value is “No”.

- “**Find filter**”: This entry field restricts the backup to the files meeting a UNIX “**find**” criteria. All files that do not match this criteria are excluded from the backup. For more information on the “**find**” command, enter “**man find**” at any UNIX prompt.
- “**Inclusion filter**”: This entry field restricts the backup to the files matching a UNIX regular expression filter.

For instance, the regular expression “**tcM.*\.***” will match any files beginning with the letters “**tcM**”, followed by any characters, followed by a “**.**” and ending with any characters. All other files would be excluded. Given this example, the following file would be included in a backup:

```
tcM000006.dat
tcM000007.dat
tcM_.idx
tcM_27.report
```

While the following files, with the same inclusion filter, would be excluded from the backup:

```
AAAA.dat
M57887786.txt
beta0124554.log
tcM3696abc
```

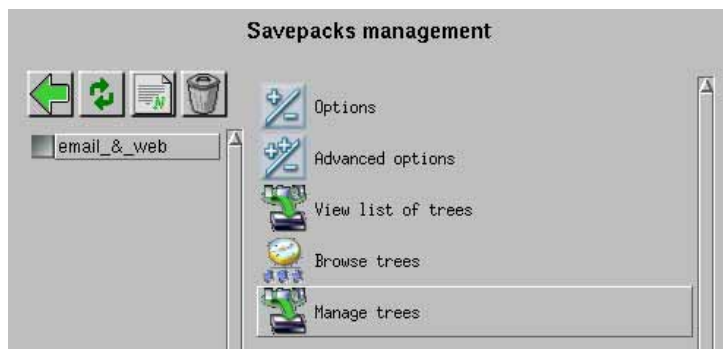
- “**Exclusion filter**”: The exclusion filter allows files to be excluded from the backup of the Savepack, by using the same UNIX regular expressions that are used by the inclusion filter.

For instance, an exclusion filter that contains the regular expression “***\.*jp***” would exclude all files ending with a suffix “**.jp**” such as “**.jpeg**” and “**.jpg**”.


As noted above, these advanced options apply globally to the selected Savepack. It is possible to apply different options to the trees contained in a Savepack, by selecting the “**Manage trees**” option in the main “**Savepacks management**” menu. Therefore, it is possible to define several parallel trees within the same Savepack, for instance one to backup HTML files with a high compression, and one to save the graphical elements of a web site, without compression applied.

3.5 Setting up the options of a tree.

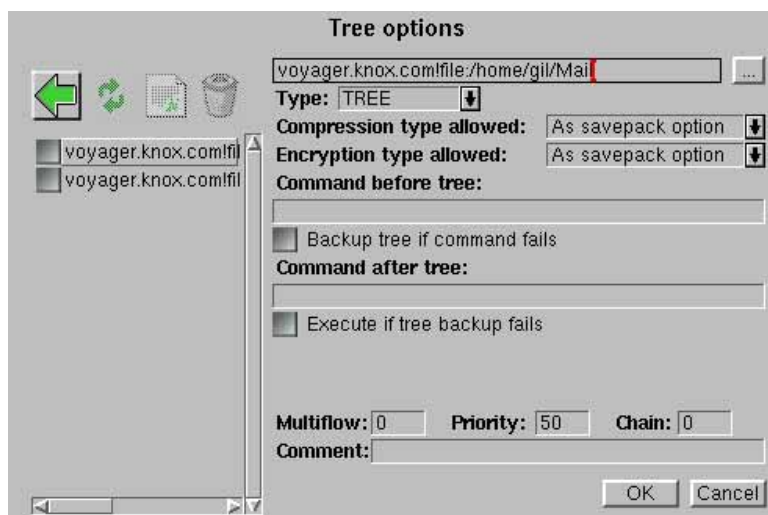
Within a Savepack, it is possible to set up backup options for each individual tree. To do this, in the main “**Savepacks management**” screen, click on the name of the Savepack, then on the “**Manage trees**” option. This will display the list of trees that are contained within the selected Savepack.



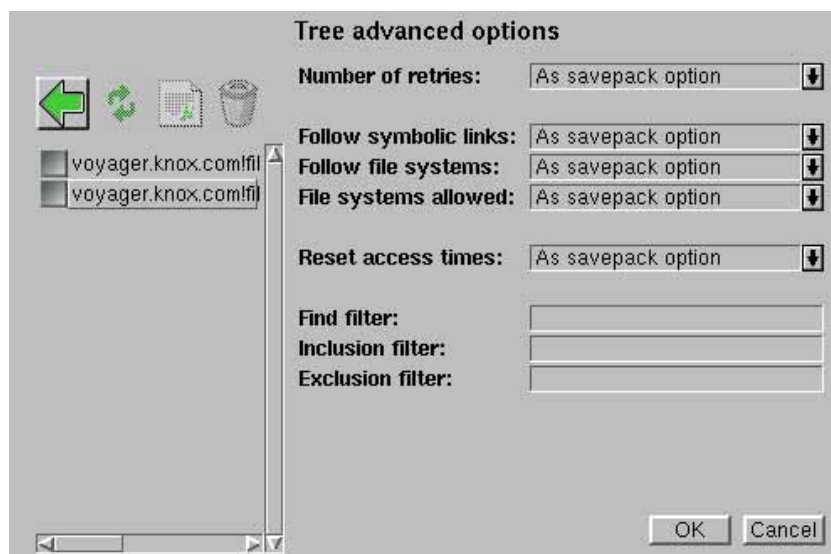
Clicking on the “**Manage trees**” option opens, logically, the menu of the same name within the main window of Arkeia version 5:



Within this menu, to modify the options applied to tree, click on the check box next to the tree that should be modified and select either the “**Options**” item (which opens the “**Tree options**” screen):



Or, click on the name of the tree to be modified, then on the “**Advanced options**” item to open the equivalent screen (see below):



Please note: Several options that are displayed in the “**T**ree **a**dvanced **o**ptions” screen (see above) are the same than the equivalent Savepack options. As such they won’t be detailed below. For more information on these options, refer to the section named “Defining options for the Savepack” above.

The following options are available in the tree options screen, and are specific to the tree configuration:

- “**T**ype”: This option determines the type of the tree. There are three different types of trees:
 - “**T**ree”: A tree can be the name of machine, or “**h**ost**n**ame” (to backup an entire machine), or a directory to backup only this specific directory, and all the files and sub-directories it contains, or a single file.
 - “**O**bject”: An “**O**bject” tree will back-up and restore the result of a command. This command should be entered in the “**P**ath” field. Please note that the file created by the “**O**bject” command will not reside on the Arkeia client, but an entry will be created in the Arkeia database to restore the data backed-up.
 - “**R**aw”: A “**R**aw” tree is either an entire disk or an entire partition. If a “**R**aw” tree is to be used, the Arkeia Navigator should be used to select the disk or the partition present in the “**/d**ev” directory of the UNIX machine.
 - “**S**avepack”: A “**S**avepack” tree appears only if the chosen tree is, itself, a Savepack.
- “**M**ultiflow”: The “**M**ultiflow” value determines if Arkeia should launch backup a tree with several concurrent data flows.

By default, the value of “**M**ultiflow” is “0” for all trees, meaning that Arkeia will use only one data flow between the server and the client, and will backup the trees in a sequen-

tial order.

To create more than one data flow (and, therefore, launch parallel backups), a value other than “0” should be entered in this field. For instance, if three trees are contained in a Savepack, the following values could be applied:

Tree/Multiflow example:

Tree:	Multiflow value:
A (sargon:/home/pete/)	1
B (sargon:/home/tina/)	2
C (sargon:/home/www/)	3

In the example above, three trees are created in a Savepack, and each of these trees has been assigned a different flow number. When the Arkeia server is going to backup the Savepack, it will open three different flows to the machine “sargon”, one for each of the trees, and backup all the trees at the same time.

If two of these trees were assigned the same Multiflow number, they would be backed-up sequentially, instead of at the same time.

- “**P r i o r i t y**”: This field determines the order in which trees should be backed-up by the Arkeia server. The possible values for “**P r i o r i t y**” go from “1” (highest priority) to “100” (lowest). Trees that have a higher “**P r i o r i t y**” value will be backed-up before the trees that have a lower priority. The default value for all trees is “50”. The priority level is only taken into account when trees have the same Multiflow number.
- “**C h a i n**”: This field determines the relationship/dependency between different trees. If two trees have the same “**C h a i n**” value, but different “**P r i o r i t y**” values, the tree with the smallest “**P r i o r i t y**” value will only be backed-up at the end of the other tree backup. For instance:

Example of chaining values:

Tree name:	Chain value (Priority value):
A (austin:/home/sal/)	12 (45)
B (abilene:/home/tc/)	12 (50)

In the example above tree “A” and tree “B” are chained together, since their “**C h a i n**” value is “12”. Since “A” has a higher priority (45) than “B”, tree “A” is going to be backed-up first. The backup of “B” will start once the backup of “A” has been completed.

Using the values of the “Multiflow”, “Priority” and “Chain” tree options allow the system administrator of an Arkeia server to control very precisely the order in which the different backups are going to be performed.

4 • Creating users

4.1 Overview

Since Arkeia is a network client/server backup software, it often has to backup/restore sensitive and/or confidential data. While it offers advanced encryption function to protect that data, another layer of security can be added by restricting access to certain data (and the backup operations of the data) to certain users.

To restrict this access, it is therefore necessary to create different users, who are in charge of different subsets of data on the network machines backed up by Arkeia. This is what this section of the manual will describe.



Please note: An important point to take into account is described in section “Post-installation root password” on page 29 of this manual.

Arkeia Corporation, as described in this section highly recommends that the default (blank) “root” password is changed immediately after the installation of a new Arkeia server.

4.2 Arkeia users categories

There are three main categories of users under Arkeia. The table below describes these three categories and the privileges assigned to each.

Arkeia User categories

Name	Privileges
ADMINISTRATOR	<ul style="list-style-type: none">• Can perform all operations• Can create or modify all objects• Can create or modify all users
OPERATOR	<ul style="list-style-type: none">• Can perform all operations• Cannot modify the Arkeia objects• Cannot create other users
USER	<ul style="list-style-type: none">• Can only restore the files of the machine he/she is currently logged in

As noted elsewhere in this documentation, the default Administrator is the user named “root”. This user is not the same as the “root” of a UNIX system.

4.3 How to create users under Arkeia

To create users under Arkeia, follow the procedure described below:



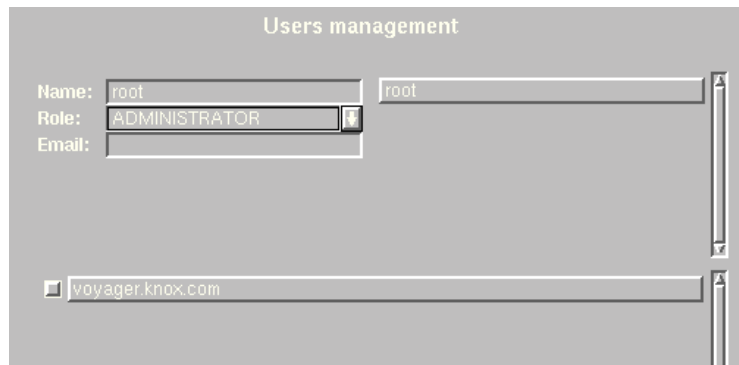
Within the main menu of Arkeia, click on the “Administration” icon.

Clicking on the “Administration” icon will open the “Administration menu” in the main window of Arkeia version 5.



Click on the “Users” icon to order the “User management” menu within the window of Arkeia version 5.

The main “Users management” screen displays a list of all the users, including the servers they can have access to:



To create an Arkeia user, right-click anywhere in the screen to display the “Users management” contextual menu.

This contextual menu offers several options to create, delete and modify users:



The following options are available in this contextual menu:

- “**C**reate **u**ser”: Clicking on this option starts the creation process of a new user within Arkeia version 5.
- “**D**el~~e~~te **u**ser”: If a user has been selected in the left-hand user list, clicking on this option deletes him or her. Arkeia asks the administrator to confirm the deletion.
- “**C**hange **p**assword”: If a user has been selected in the left-hand user list, clicking on this option will change the password it has been assigned.
- “**O**K”: Clicking on this option confirms the selected action.
- “**C**ancel”: Clicking on this option exit the “**U**ser **M**anagement” menu.

If the “**C**reate **u**ser” option has been chosen in the contextual menu, Arkeia version 5 then allows the operator to enter a name, choose a category and configure the E-mail for the user being created. The email entered will be used by Arkeia to send reports on backups to the owner of a backup:

Name:	Jack
Role:	ADMINISTRATOR
Email:	jack.smith@snafu.acme.com

In the example pictured above, the “**A**DMINISTRATOR” of Arkeia version 5 (by default, the user named “**r**oot”) is creating another Arkeia Administrator named “**J**ack”, whose E-mail address is “**j**ack.**s**m**i**th@**s**n**a**f**u**.**a**c**m**e.**c**o**m**”.

To create a user, simply enter the relevant information in the different field, right-click to display the contextual menu and click on the “**O**K” option to confirm the creation of a user. The newly-created user should then be displayed in the list of users, on the left of the “**U**sers **m**anagement” menu.



IMPORTANT. Please note: Whenever a new user is created, no password is assigned to this user. Arkeia Corporation recommends to assign a password to a new user immediately.

To assign a password to a user, select the name of the user in the list of users and right-click anywhere to display the User management contextual menu. In this menu, click on the “**C**hange **p**assword” option.

This displays the “**C**hange **p**assword” menu in the main window of Arkeia version 5:



Within the “**C h a n g e p a s s w o r d**” menu, the different password fields should be filled in order to change the password. The old user password should be entered in the relevant field, and the new password entered twice — once in the “**N e w p a s s w o r d**” field and a second time in the “**C o n f i r m p a s s w o r d**” field — to be defined.

If the user has just been created, the “**o l d p a s s w o r d**” field should be left blank, as the password of a new user is always blank.



Once the password fields have been filled, right-click anywhere in the Arkeia version 5 window to display the contextual menu and click on “**o k**” to confirm the password assigned to the user.

5 • Creating Drivepacks

5.1 Overview


Once drives have been physically attached to the Arkeia server and logically configured with the help of the graphical user interface, they should be logically grouped for an easier management.

This section will briefly describe the possibilities and the uses of the Arkeia version 5 Drivepacks.

5.2 How to create Drivepacks

To create Drivepacks, follow the steps detailed below.

Within the main menu of Arkeia version 5, click on the “Hardware” button:



Within the “Hardware menu”, click on the “Drivepacks” button:



Clicking on the “Drivepacks” button will open the following screen within the main window of Arkeia version 5:



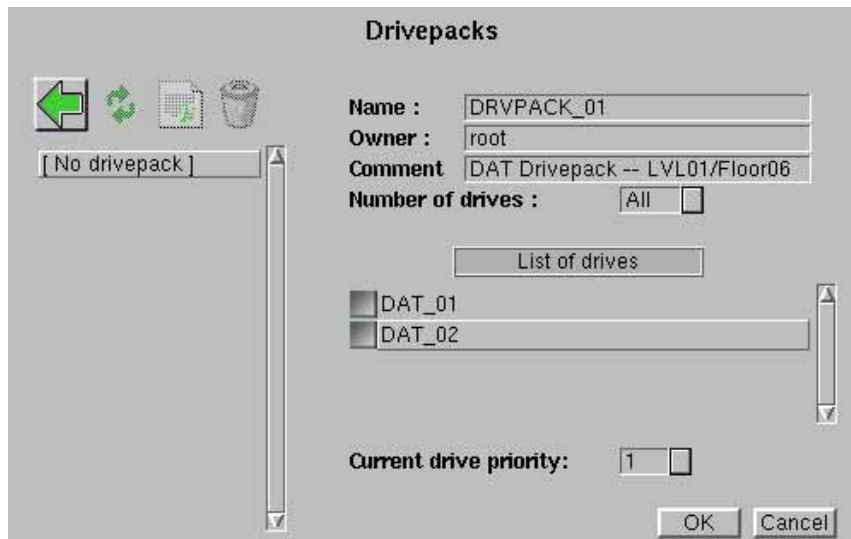
To create a new Drivepack, click on the “**New**” button, which is above the list of Drivepacks.

This button is pictured on the left of this paragraph.

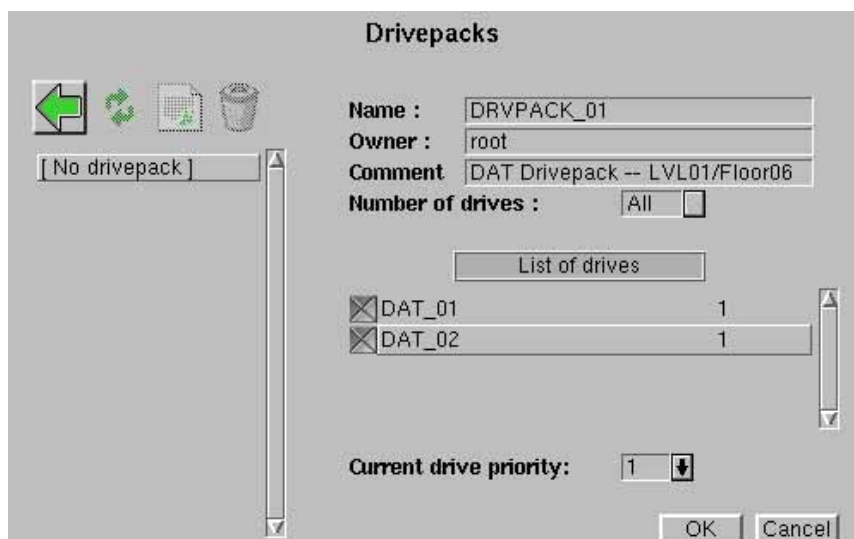
This screen contains the following options:

- “**Name**”: Enter the name of the Drivepack in this field.
- “**Owner**”: This is the name of the Arkeia user who is creating the Drivepack, usually “root”.
- “**Comment**”: If needed, it is possible to enter a descriptive comment in this field.
- “**Number of Drives**”: This drop-down list allows the system administrator to restrict the number of drives that can be used by the Drivepack at the same time. For instance, if a Drivepack contains 6 drives, it is possible to restrict the “**Number of Drives**” to “2”. In this way, the first backup to use the Drivepack will only use the first two drives (for example: “DAT_01” and “DAT_02” as in the example above), the second backup to use the Drivepack will use the next two drives, and so on until the total number of drives has been exhausted. The drives are chosen (or attributed) based on the “**Priority**” (see below) that has been assigned to it.
- “**List of Drives**”: This displays all the drives that have been configured (and are available) on the Arkeia version 5. To select one (or several) drive(s), click on the grey selection box.
- “**Current drive priority**”: If a drive has been selected in the “**List of Drives**”, it is possible to set its priority with this drop-down box. Setting the priority of a drive configures which drive will receive the data first. A priority of “1” is the highest priority that can be set. By default, all drives receive a priority of “1”. Once a drive has been selected, its “**Priority**” is displayed to the right of its name.

For instance, the screen below presents a typical Drivepack configuration:



To complete the Drivepack configuration shown above, all that is needed is to select the drives to be added to the Drivepack – by clicking on the grey, square, selection box next to their name – and clicking on the “OK” button at the bottom of the screen, as can be seen in the following screen:



As can be seen above, the Drivepack named “DRVPACK_01” is ready to be configured. It contains two drives (“DAT_01” and “DAT_02”), with the same “Priority” of “1”. Furthermore, this Drivepack is configured to use all the drives that have been attached to it.

Clicking on the “OK” button at this stage would create the Drivepack and display it in the list of Drivepacks, which is on the left of the screen pictured above.

CHAPTER 7: INTERACTIVE BACKUP

- *Summary*

This chapter introduces the interactive backup function of Arkeia. The interactive backup is the most simple backup provided by Arkeia. It introduces important concepts that will be used in the rest of this manual.

1 • Overview

Arkeia version 5 can start backup operations in two different modes, either “Interactive” or “Periodic” mode.

The difference between these two modes is that the “Interactive” mode launches the backup itself immediately, while the “Periodic” mode launches a backup only when such an operation has been scheduled.

While there is an overlap between these two functions, the “Interactive” mode is best used for backups that are launched and controlled by a system administrator, while “Periodic” backups should be used for regular unattended network backup tasks, that are performed on a regular basis.

This chapter will detail the functions of the “Interactive” mode, while the next one will detail the scheduling functions of the “Periodic” mode.



Please note: For more information on how to create logical objects within Arkeia version 5, such as Savepacks and Tape pools, please refer to the chapter “6: Objects and users creation” on page 71 of this manual.

An interactive backup requires several objects to be defined before the backup itself is launched.

2 • How to access the Interactive backup

When Arkeia version 5 is launched and the user is logged on an Arkeia server, the main menu of the Graphical User Interface (GUI) is displayed on the screen.



Within the main menu, click with the right mouse button on the “Backup” icon to access the related menu.

The “Backup menu” is then displayed within the GUI:



Within the “Backup menu” of Arkeia version 5, click on the “Interactive” icon to access the Interactive backup functions.

The main “Interactive backup” functions are then displayed on the screen. The rest of this chapter will detail these functions and their usage.

Interactive backup

Savepack: Email files

Drivepack: drive_pool

Pool: documents

Type: Total Backup Standard

Tape strategy: Complete existing tapes

Valid for: 30 Day(s)

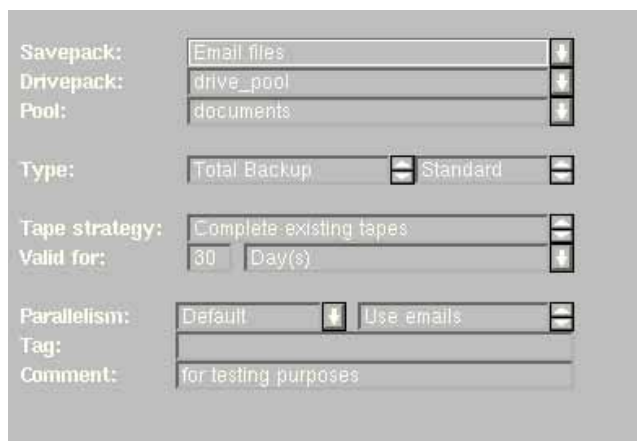
Parallelism: Default Use emails

Tag:

Comment: for testing purposes

3 • Main functions of the interactive backup

The main functions of the interactive backup can all be found on its main screen.



The screenshot shows a configuration window with the following fields and values:

- Savepack: Email files
- Drivepack: drive_pool
- Pool: documents
- Type: Total Backup, Standard
- Tape strategy: Complete existing tapes
- Valid for: 30 Day(s)
- Parallelism: Default, Use emails
- Tag: (empty)
- Comment: for testing purposes

The following functions can be configured through this screen:

- “**Savepack**”: Clicking on this drop-down list presents the different Savepacks that have been defined and can be selected to be backed-up.
- “**Drivepack**”: Clicking on this drop-down list presents the different Drivepacks that have been defined and can be used to perform the backup operation.
- “**Pool**”: Clicking on this drop-down list presents the different Tape Pools that have been defined and can be used to contain the backup data.
- “**Type**”: This field allows you to choose the type of backup that will be started, as well as the completion strategy for a backup.

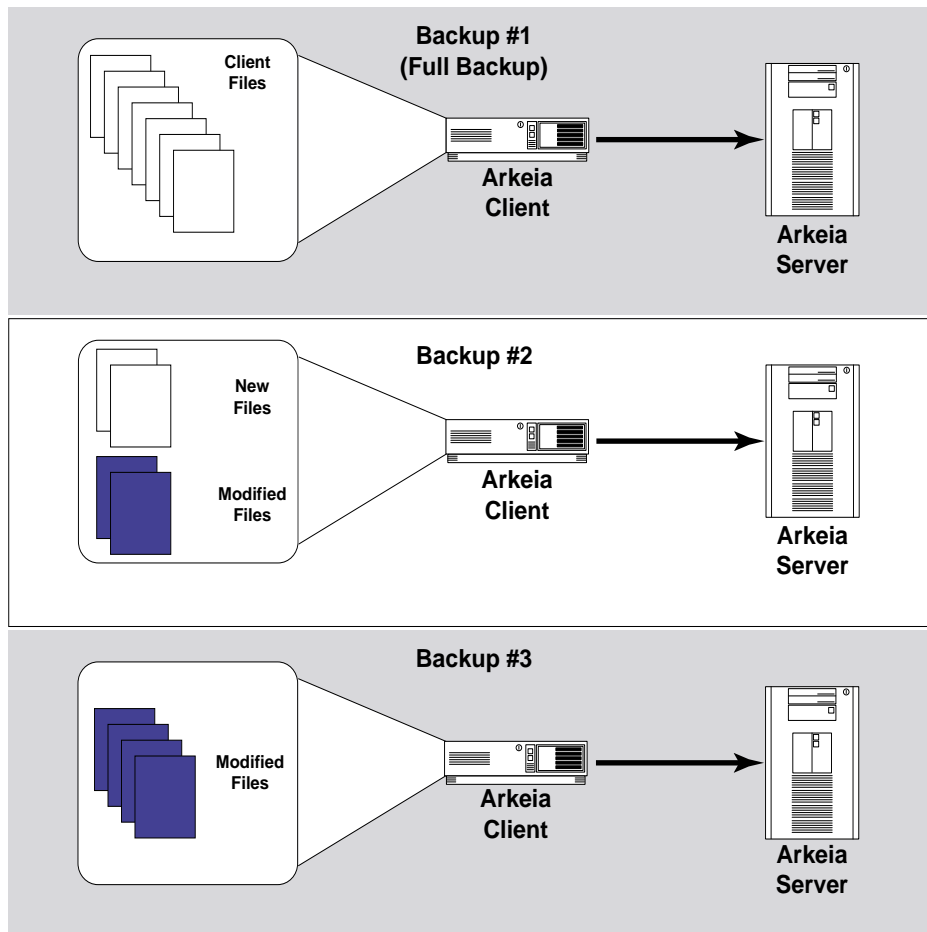
The three different types of backups are the following:

- “**Archive**” backup: This type of backup is the most simple backup Arkeia version 5 has to offer. As its name implies, an “**Archive**” backup, once it has been launched, will simply backup all the files and directories selected in the Savepack (see “**Total**” backup below) but will retain the data indefinitely. This prevent Arkeia from recycling the tape automatically, and ensures that the most crucial files will remain on tape as long as necessary. Refer to the “**Valid for**” parameter below for more information.

- “**Total**” backup: A “**Total**” backup will backup all files and directories defined in the Savepack, while allowing the data tape to be recycled once the “**Valid for**” threshold of the tape has expired. The most common use for this type of backup is to be performed periodically — every two weeks, for instance — to form the basis of a “**Incremental**” backup (see below).

- “**Incremental**” backup: An “**Incremental**” backup is the most complex backup function of Arkeia version 5, but it is also the one that allows the system administrators to define the most complete backup strategy. An “**Incremental**” backup can be

based either on a “Total” backup or on a previous “Incremental” backup. When this type of backup is launched, it will create a list of files that have been created and/or modified since the last backup, and it will then backup any and all files that are on this list. You will find below a graphical example of a incremental backup.



In the example above, the first backup (“Full backup”) is a “Total” backup, that transfers all the files of the Arkeia client to the Arkeia server, to be saved on tapes. The second backup is a incremental backup, based on the first backup, which saves only the few files that have been modified or created. The third backup is another incremental backup, that saves only the files that have been modified since the second backup.

As you can notice, this type of backup presents two great advantages: it reduces the time and the bandwidth necessary for the backup, since the amount of data to be sent to the server is smaller than on the first backup. The disadvantage of the “Incremental” backups is that the restoration of the data requires the system administrator to go through the entire tape set (from Backup #1 to Backup #3 in the example above) in order to restore the data in the exact state it was in when Backup #3 was launched. However, “Incremental” backups are the only way to backup large amounts of data without tying up the resources a network.

Each type of backup can be further configured with the following completion strategy:

- “Standard”: The “Standard” completion strategy completes the backup once all the data contained in the defined Savepack has been sent to the Arkeia server. The backup itself is considered as finished, logged as such in the Arkeia logs, and all hardware peripherals

used by the backup are released by the server and made available for other tasks. This is the default completion, and it is highly recommended to use this strategy most of the time.

- **“Continuous”**: A **“Continuous”** completion strategy means that, once the backup has finished saving the data contained in its Savepack, it stays “active” and is not considered as completed by Arkeia version 5. The hardware used for the backup operation is still considered as reserved by the Arkeia server, and cannot be used for other operations. It is possible to add new Savepacks to the backup, as long as the backup is not terminated by the user.



Please note: While a **“Continuous”** completion strategy may be interesting in some situations, it should be used with caution. It may be required by some Arkeia plug-in to perform its function correctly — please refer to the documentation delivered with the plug-in to determine if this is the case. In the same way, a continuous backup may be required if the system administrator needs to quickly add savepacks to an ongoing backup.

In any case, **“Continuous”** backups should never be used for unattended (**“Periodic”**) backups! For more information, refer to chapter **“8: Periodic backup”** on page 123 of this manual.

- **“Tape strategy”**: This field provides the different strategies that can be used to control the usage of tapes by a backup. The two different strategies are the following:
 - **“Use new tapes”**: This strategy forces Arkeia to use new tapes every time a backup is started. Use this strategy if you would like to keep the data of a backup on new tapes.
 - **“Complete existing tapes”**: This strategy will fill tapes that have already been used for previous backup.

The strategy used for the tape usage should take into account the amount of data to back up, the space remaining on a tape after a **“Total”** backup has been performed, the life expectancy and overall performances of the backup media, as well as the time necessary to perform a complete restoration of the data saved.

- **“Valid for”**: This field determines the duration of data retention. Once the data retention defined has expired, the tape is recycled, and the data it contains will be erased in the future by other backups. Note that this parameter is modified if the backup type is set to **“Archive”**: in this case, the retention data is overridden and the tape is never recycled.
- **“Parallelism”**: This field determines the number of parallel “flows” that should be open between the Arkeia server and the Arkeia client when performing the backup. For a more detailed explanation of the Arkeia data flows, refer to the section named **“Arkeia Internal Architecture”** on page 57 of this manual.
- **“Use E-mails”**: This field will determine if Arkeia version 5 is to send a confirmation E-mail to the current user once the backup has been completed.

- “**T a g**”: This field determines a “name” that can be used to identify the backup. Once this name has been entered, it can be used to identify the backup and all the data it contains when performing a restoration.
- “**C o m m e n t**”: This field can be used to enter a longer, more detailed description of the backup.




Please note: For more information on how to create logical objects within Arkeia version 5, such as Savepacks and Tape pools, please refer to the chapter “6: Objects and users creation” on page 71 of this manual.

An interactive backup requires several objects to be defined before the backup itself is launched.

Once the Interactive backup has been defined, it can be launched by clicking on the “OK” button at the bottom of the window.

Once the backup has been launched, Arkeia displays the main backup window and proceeds with the defined operation. This window, the information it contains and the modifications it allows, is described in the next section of this chapter: “Monitoring a running interactive backup” on page 116.



3.1 Main functions: Archive backup only

The “**A r c h i v e**” backup offers a simplified interface, given that it does not allow the tapes it uses to be recycled.

The main interface for the “**A r c h i v e**” backup can be found below.

Savepack:	Email files	↓
Drivepack:	drive_pool	↓
Pool:	documents	↓
Type:	Archive	Standard
Tape strategy:	Complete existing tapes	
Parallelism:	Default	Don't use emails
Tag:	backup-email-files	
Comment:	For testing purposes	

As noted above, the “Archive”-type backup is the most simple. All the fields available in this window have already been described. Please note the use of the “Tag” field (to identify the backup later on) in the screen shown above.

3.2 Main functions: Incremental backup only

The “Incremental” backup offers a field which is specific to the Incremental functions of Arkeia version 5.

Savepack:	Email files	↓
Drivepack:	drive_pool	↓
Pool:	documents	↓
Type:	Incremental Backup	Standard
From:	March 19, 2002 17:41 (Email files)	
Tape strategy:	Complete existing tapes	
Valid for:	30	Day(s)
Parallelism:	Default	Use emails
Tag:		
Comment:		

The “From:” drop-down list allows the system administrator to choose the previous backup which will be used as a reference by the current backup. Clicking on the arrow next to this list will display the list of the last 10 backups that have been performed on the same Savepack. This allows a quick selection of the reference backup.



Please note: As noted above, the reference backup should be a “Total” backup. It is possible to base an “Incremental” backup on a previous “Incremental” backup, but only if this previous backup has been based on a previous “Total” backup.


This point will be treated in more details in the chapter “8: Periodic backup” on page 123 of this manual.

The remaining sections of this chapter will deal with the monitoring and the modification of an interactive backup.

4 • Monitoring a running interactive backup

4.1 Monitoring a backup through the main screen

Once it has been launched, an interactive backup is fairly easy to monitor through the backup information displayed by Arkeia version 5.




This screen contains the following information on an Arkeia backup. On the left side of the screen:

- “E l a p s e d t i m e”: This is the first field displayed on the left of the running backup window. Its purpose is to display the time elapsed since the launch of the backup.
- “S e r v e r”: This is the name of the server the backup is running on.
- “S a v e p a c k”: This is the name of the Savepack which is being processed by Arkeia.
- “D r i v e p a c k”: This is the name of the Drivepack that has been chosen to perform the backup operation.
- “P o o l”: This is the tape pool used to save the data contained in the Savepack.
- “C r u i s e”: This slider allows the system administrator to control the amount of bandwidth consumed by the backup. If important network functions are sluggish while Arkeia is performing backups, use the “C r u i s e” control to reduce the amount of bandwidth used to perform a backup. Please note that this will slow down the backup operations.

In the middle of the screen, a graphical meter displays the amount of data being transferred between the client and the server.

On the right side of the screen, the following fields and drop-down boxes supply additional information on a backup:

- **“Global”**: Clicking on this drop-down list displays all the drives that are used by the backup. This allows the system administrator to restrict the display of information to a particular drive.
- **“Speed in KB/s -- Speed in MB/h”**: Clicking on this button changes the type of information that is displayed in the fields below it. Please note that this information is displayed in two separate fields: **“Instant”** indicates the current speed of the backup and **“Average”** the overall speed of the backup being performed.
- **“Backup Speed/Network throughput”**: Clicking on this button changes the type of information that is displayed in the fields below it, from the speed of the backup itself to the amount of information transferred by the network.
- **“>> (View)”**: Clicking on this button displays an additional information screen on the different flows that are launched by the Arkeia backup. Within this screen, press the **“Escape”** key on your keyboard to go back to the main backup screen.



Within this additional screen, the following information is supplied:

- **“Drive”**: This field displays the number of the drive being used by the flow.
- **“Compression ratio”**: This field displays the compression applied to the flow.
- **“Instant Speed”/“Average speed”**: See the description of **“Speed in ...”**.
- **“Current directory”**: This field displays the directory being processed.
- **“Message”**: This field contains the last message returned by Arkeia version 5.

4.2 Monitoring a backup through the list of operations

Another way to monitor a running backup is to click with the right mouse button anywhere in the window of a running backup. This displays the backup contextual menu, which you will find below.



In this menu, clicking on the fifth option “List of operations”, will display a screen that summarizes the operations and trees that are, or will be, performed by the backup.



In the example above, a backup has been launched that contains four different trees. The following fields offer additional information on tree number 1:

- “Tree number”: Self-explanatory. The trees are displayed (and processed) in the order they have been defined in the Savepack.
- “Type”: This field displays the type of tree that has been chosen by the user. The different types of trees are detailed in the paragraph “Type” on page 87 of this manual.

- **“State”**: This field displays the current state of the tree that has been selected by the user.

The three different states of a tree are the following:

- **“Init”**: This state indicates the tree is going through an initialization phase within the Arkeia backup subsystem.

- **“Running”**: This state indicates that the tree is being processed by the Arkeia backup subsystem.

- **“End of tree”**: This state indicates that the tree has been backed by Arkeia and is not processed any longer.

- **“Results”**: This field displays the results of the backup of the tree. These results should be self-explanatory.

- **“Retries left”**: This field displays the number of retries left in case a backup operation failed to perform — for instance, if a file contained in the tree was locked by the system.

Clicking on any tree, with the left mouse button, will display its information in the fields of the **“List of operations”**.

The next section of this manual will detail the other options of the backup contextual menu, that provides several functions to modify “on the fly” a running backup.

5 • Modifying a running backup

To modify a running backup, all that is needed is to click with the right mouse button anywhere in the backup window. This displays the following contextual menu, which provides options to quickly modify a backup.



The following functions are available in the backup contextual menu:

- “**Abort backup**”: Clicking on this option will immediately stop the current running backup.
- “**Set in continuous mode**”: Clicking on this option modifies the running backup and sets it into “**Continuous**” mode.
- “**Set in standard mode**”: Clicking on this option modifies the running backup and sets it into “**Standard**” mode.

For more information on “**Standard**” and “**Continuous**” mode, refer to the “**Type**” of backup, detailed on page 110 (and following) of this manual.

- “**Add savepack**”: Clicking on this option allows the backup operator to add a Savepack to the running backup. This operation is detailed below.
- “**List of operations**”: This option was detailed in the previous section of this chapter.
- “**Previous menu**”: Returns the user to the “**Backup menu**” of Arkeia version 5.

The rest of this section will detail some of the functions that are available in this contextual menu.

5.1 Adding a Savepack to a running backup

Adding a Savepack to a running backup is possible as long as the backup is running, regardless of the operating mode of the backup itself (see below “**Setting the operating mode of a running backup**” on page 121).

To do this, click on the “**Add savepack**” option detailed above. Arkeia version 5 will then ask which Savepack should be added to the running backup by displaying the screen below.



Clicking on the arrow on the right of the drop-down list will display the list of Savepacks, defined in Arkeia, that can be added to the running backup. Clicking on the name of the Savepack will select it.



Once the Savepack to be added has been selected, click anywhere in the Arkeia window with the right mouse button to display the contextual menu pictured on the left, and click, with on “OK” to confirm the operation.

Once the Savepack addition has been confirmed, Arkeia will return you to the main “Backup menu”.



Please note: Unless a backup is set in “Continuous” mode (see below), it is not possible to add a Savepack to a backup once the backup operation has been completed by Arkeia.

See below for a more in-depth discussion on “Continuous” versus “Standard” backup, or refer to the beginning of this chapter for more information.

5.2 Setting the operating mode of a running backup

Setting the operating mode of a running backup is possible through the two options (“Set in continuous mode” and “Set in standard mode”) of the contextual menu.

Clicking on “Set in continuous mode” in the contextual menu allows the system administrator to keep the backup task “open”, even after it has completed the backup of all the trees it was supposed to process. This allows the backup to accept additional Savepacks, even after it has completed all its assigned tasks, but it also ties up the hardware resources used to perform the backup.

Clicking on “Set in standard mode” allows the system administrator to change the “Type” of a “Continuous” backup back to “Standard”. Once the backup is back to this mode, it will be terminated by Arkeia as soon as its Savepacks have been pro-

cessed. This will free up (or “release”) the hardware resources that were used and make them available for other backups.



Please note: As indicated in this chapter, a backup should be run, by default, in “standard” mode. The “continuous” mode is mostly used for plug-in backups that require the special mode of operation of a “continuous” backup. For more information, refer to the documentation supplied with the plug-ins installed on your system.

CHAPTER 8: PERIODIC BACKUP

- *Summary:*

Periodic backups are the most frequently employed by system administrators. A periodic backup is a backup scheduled for execution at a given time and date.

The great advantage of periodic backups is that, once scheduled by the system administrator, they are executed automatically by the software itself, in a completely unattended manner. The most important data of a networked machine is therefore protected at regular, scheduled intervals to ensure the highest possible security.

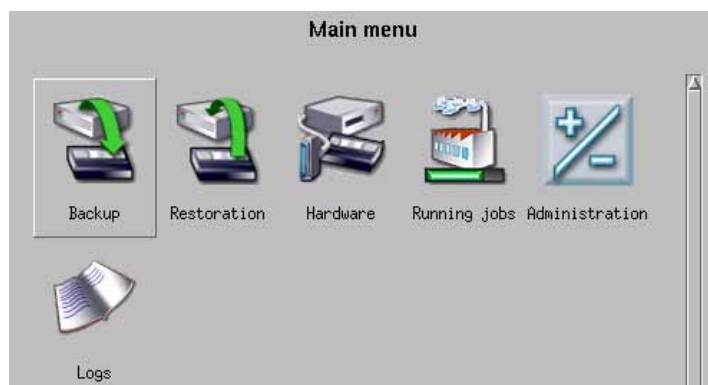


Please note: For more information on how to create logical objects within Arkeia version 5, such as Savepacks and Tape pools, please refer to the chapter “6: Objects and users creation” on page 71 of this manual.

A Periodic backup requires several objects to be defined before the backup itself is launched.

1 • How to access the Periodic backup menu

The main screen of Arkeia version 5 is usually displayed on the screen when the program is started:



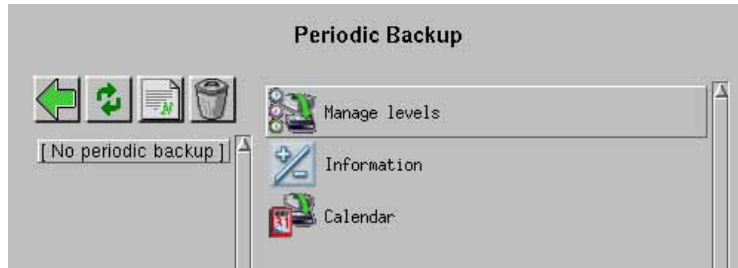
From this main menu, click on the “Backup” button (pictured on the left) to access the “Backup” menu.

The main “Backup” menu is then displayed in the Arkeia window.



From the main “Backup” menu, click on the “Periodic” button to access the “Periodic” menu.

The main “Periodic backup” menu is then displayed in the Arkeia window:



As many other screens in Arkeia version 5, the “**P e r i o d i c B a c k u p**” screen is divided into two parts: on the left side of the screen, a list of the periodic backups that have been defined (note the empty list in the screen pictured above) and, on the right of the screen, a list of functions that can be applied to the Periodic backups.

The functions that are available in the list of functions are the following:

- “**M a n a g e l e v e l s**”: Clicking on this icon allows the operator to create, configure or remove levels within the selected periodic backup.
- “**I n f o r m a t i o n**”: Clicking on this icon allows the operator to enable or disable the selected periodic backup and determine if the backup should use E-mail to report its success or failure.
- “**C a l e n d a r**”: Clicking on this icon displays the periodic backups that have been programmed during the current month.

Above the list of Periodic backups, four buttons offer the most common actions. These buttons are, from left to right: “**B a c k**” (to go back to the previous screen), “**R e f r e s h**” (to reload the defined backups), “**N e w**” (to create a Periodic backup) and “**D e l e t e**” (to erase an already-defined periodic backup).

The rest of this chapter will detail the different functions that are available from the “**P e r i o d i c b a c k u p**” menu.

2 • Overview: configuration of a periodic backup

Once a periodic backup has been created (refer to the next section for more information), it is possible to start its configuration process.



Please note: The configuration of a periodic backup is not a trivial task! It is highly recommended to read this chapter completely and to understand the concepts it introduces before attempting the configuration of a periodic backup for the first time.

Periodic backups are managed within Arkeia based on a “Level” architecture. Each Periodic backup can have up to three levels, numbered “1” to “3”. Within this hierarchy, the “Level 1” backup will replace or override the “Level 2” backup, which will in turn replace or override the “Level 3” backup. Therefore, it is important to determine from the very beginning of the configuration procedure when each level should be performed, and what will be the extent of the task to be performed by each level.

Typically, “Level 1” backups are complete backups, that are done monthly, while “Level 2” and “Level 3” backups are periodic backups that are more limited in scope, and which only backup the files modified since the last “Level 1” backup, or the latest “Level 2/3” backup.

For instance, a level 1 backup can be programmed to be executed on the first day of every month. Since this level 1 backup is a “Total” backup, the day after this backup can be excluded by the system administrator. A level 2 incremental backup, can be programmed to execute every Saturday of the month. Finally, a level 3 backup, designed to save frequently modified data, can be performed every Wednesday.

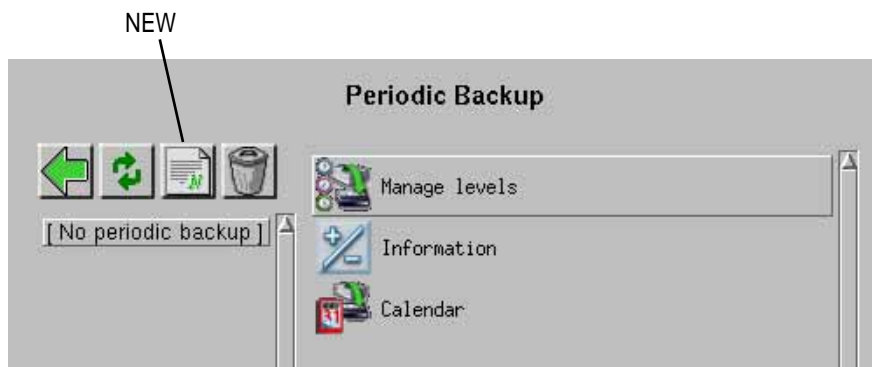
A policy such as the one presented above is easy to configure and put into place through the powerful scheduling functions of Arkeia v5. These functions are going to be detailed in the rest of this chapter.

3 • Creating a periodic backup

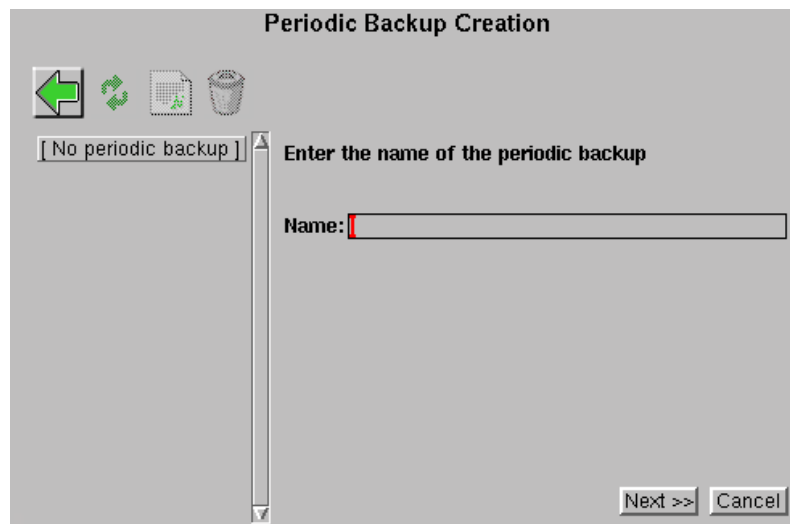
3.1 Step one – creating the backup itself

Creating a Periodic backup can be easily achieved by following the creation procedure supplied by Arkeia version 5.

To start this creation procedure, all is needed is to click on the “New” button in the “Periodic backup” window:



Arkeia version 5 then asks the operator to enter a name for the new periodic backup being created:



Once a name has been entered, clicking on the “Next >>” button at the bottom of the window displays the next step in the creation procedure: the definition of the first level of the periodic backup.

3.2 Step 2 – Defining the periodicity of the backup

The screenshot shows a dialog box titled "email_ &_ web level 1: Periodicity". On the left, there is a tree view with "Level 1" selected. The main area contains the following fields and options:

- Repeat at:** 17 Hr 45 Min
- every** 1
- Radio buttons for **day(s)** (selected), **week(s)**, **month(s)**, and **year(s)**.
- Starting on:** 2002/07/18
- Repeat until:**
 - End on 2002/07/18
 - backup has been done times
 - No ending date
- Buttons: **Next >>** and **Cancel**

In this screen, the first line, starting with the words “Repeat at:” defines the time at which the backup should be carried out. The hour is expressed in a 24-hour format.

The second (“every ____”) and third lines define the periodicity itself: for instance, it is possible to repeat the same backup every 3 days, every 5 days, etc., by entering the appropriate number in the field.

The next line defines the type of recurrence that will be applied to the backup: the choices are “Day(s)”, “Week(s)”, “Month(s)”, “Year(s)”. These options are defined below:

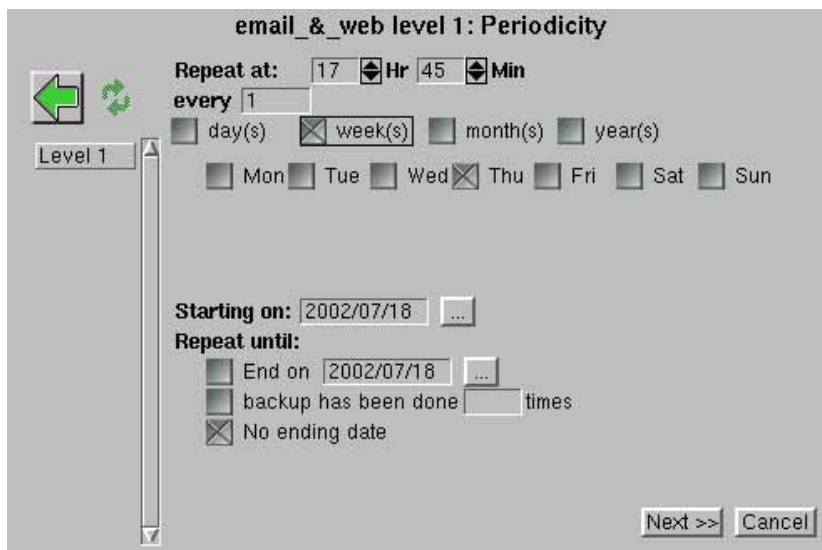
- *Periodicity Type: Daily*

Clicking on the “Day(s)” check-box selects the corresponding Periodicity for the level being configured. With a “Daily” periodicity, the backup is launched by Arkeia either every day, or every “x” days, depending on the value which is entered in the field “every ____”.

For instance, if the “Daily” check box is selected and “3” is entered in the “every ____” field, the current level of the Periodic backup will be launched every three days.

- *Periodicity Type: Weekly*

Clicking on the “Week(s)” check-box selects the corresponding Periodicity for the level being configured. The Arkeia screen then displays the corresponding options.

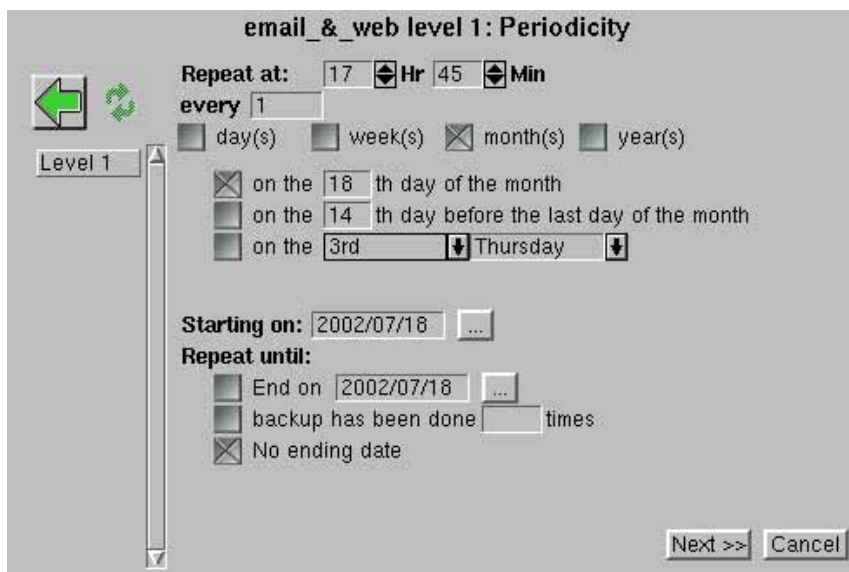


Within this screen, the backup can be configured to be launched on a given day of the week, by clicking on the corresponding check box. For instance, if backups are to be launched every Saturday, it is necessary to click on the “**S a t**” check box. In the screen pictured above, the backups will be launched every week, on Monday.

Within a week, the same level can be launched more than once by clicking on several days. Finally, it is possible to launch the backups every other week, by entering a number in the “**every _____**”. For instance, in the screen above, entering “**2**” in this field would launch the backup level every two weeks, on Thursday.

- *Periodicity Type: Monthly*

Clicking on the “**Month(s)**” check box selects the corresponding Periodicity for the level being configured. The Arkeia screen then displays the corresponding options.



The “Monthly” periodicity provides the following options:

- “every _____”: Entering a number in this field will determine the number of months that will separate two executions of the level. For instance: entering “1” will configure the execution of the backup every month. Entering “2” in this field will execute the backup once every two months, etc.
 - “on the _____ th day of the month”: Entering a number in this field determines on which day of a given month the level backup will be launched by Arkeia. In the screen pictured above, the backup will be launched on the 18th day of the month, regardless of the day itself (so the backup would be carried out on Tuesday, January 18th, Friday, February 18th, etc.).
 - “on the _____ th day before the last day of the month”: This field is the equivalent of the previous field, but it uses the last day of the month to determine when the backup should take place. For instance, entering “10” in this field would carry out the backup either on the 20th or the 21st, whether the current month has 30 or 31 days (February being, of course, a special case).
 - “on the _____ th _____ of the month”: This field and drop-down box provides a way to choose the day of the backup in a more precise manner. Clicking on one of the drop-down box provide the different choices. In the example pictured above, the user has chosen to execute the backup level every third Thursday of a month.
- *Periodicity Type: Yearly*

This check box allows a backup level to be configured for a yearly execution. Clicking on the check box will display all the relevant options in the main Arkeia window.

The screenshot shows a dialog box titled "email_ &_ web level 1: Periodicity". It contains several configuration options:

- Repeat at:** 17 Hr 45 Min
- every** 1
- Radio buttons for frequency: day(s), week(s), month(s), year(s)
- Radio buttons for specific day options: on the 199 th day of the year, on the 167 th day before the last day of the year, on July 18 th
- Starting on:** 2002/07/18
- Repeat until:**
 - End on 2002/07/18
 - backup has been done times
 - No ending date
- Buttons: Next >>, Cancel

This Recurrence type offers the following options:

- “on the _____ th day of the year”: Entering a number in this field defines the day on which Arkeia will perform the backup level. In the example above, the backup will be run on the “199th” day of the year.
- “on the _____ th day before the end of the year”: Entering a number in this field defines the day on which Arkeia will perform the backup level, starting from the end of the year. This field is the inverse of the previous field, as it counts from the end (as opposed as counting from the beginning) of the year.
- “on <month> <date> th”: Clicking on these drop-down boxes provides a way to select a precise date for the execution of the backup level.



Please note: Whenever Arkeia encounters an incorrect date within the planning of a Periodic backup, the operation scheduled is not carried out.

For instance, if a backup has been planned for the 5th Sunday of a month that only has 4 Sundays, the operation will not be executed.

This limitation is valid for all the Recurrence types that are available within Arkeia version 5.

- *Periodicity range configuration*

The “Periodicity Range” is the number of times a backup level should be repeated before it is disabled automatically by Arkeia version 5. The following options are available in the main “Periodicity” screen of Arkeia version 5.

- “Starting on: _____”: This field usually contains the current date, which is selected as the start of the backup level. To choose another date for the start of the backup level, click on the “...” button next to the field. You can also enter a date in the following

format: “YYYY/MM/DD” (Year/Month/Date).

- “End on: _____”: This field contains the date at which the current backup level will end. To choose another date for the end of the backup level, click on the “...” button next to the field. You can also enter a date in the following format: “YYYY/MM/DD” (Year/Month/Date).
- “Repeat until: backup has been done ____ times”: This field provides a way to stop a backup level after it has been carried out “X” times by Arkeia version 5.
- “No ending date”: Clicking on this button will make the backup level permanent. As long as the Periodic backup exists, this level will be carried out by Arkeia at the defined time and date.



Please note: Use the “No ending date” option with caution, especially when applying it to large Savepacks and/or backup level which have a high periodicity.

3.3 Step 3 – Configuration of the backup

The next screen after the periodicity configuration provides several fields to quickly configure the periodic backup:

The screenshot shows a configuration window titled "email_&_web level 1: Configuration". On the left, there is a tree view with "Level 1" selected. The main area contains several configuration fields, each with a dropdown arrow: "Savepack:" is set to "email_&_web", "Drivepack:" is set to "email", "Pool:" is set to "email", "Type:" is set to "Archive", "Tape strategy:" is set to "Complete existing tapes", and "Parallelism:" is set to "Default". At the bottom right, there are "OK" and "Cancel" buttons.

The following options are available in this screen:

- “Type”: This field allows you to choose the type of backup that will be started, as well as the completion strategy for a backup.

The four different types of backups are the following:

- “Archive” backup: This type of backup is the most simple backup Arkeia version 5 has to offer. As its name implies, an “Archive” backup, once it has been

launched, will simply backup the Savepack but will retain the data indefinitely.

- **“Total” backup:** A **“Total”** backup will backup the Savepack, while allowing the tape to be recycled.

- **“Differential” backup:** A **“Differential”** backup will save all the data that has been changed since the last **“Total”** backup. When this type of backup is chosen, a **Periodic** backup should be chosen in the **“Based on”** drop-down box. The **“Differential”** is a more restricted form of the **“Incremental”** backup (see below).

- **“Incremental” backup:** An **“Incremental”** backup can be based on a **“Total”** backup, on a previous **“Incremental”** backup, on an **“Archive”** backup or on a **“Differential”** backup. Like a **“Differential”** backup, an **“Incremental”** backup will only save the data that has been modified since the last backup, but it is more flexible, since it can be based either on a **“Total”**, a **“Differential”**, or on another **“Incremental”** backup.

For a more detailed discussion of **“Incremental”** backups, refer to page 118 of chapter **“7: Interactive backup”** of this manual.

Each type of backup can be further configured with the following completion strategy:

- **“Standard”:** The **“Standard”** completion strategy completes the backup once all the data contained in the defined Savepack has been sent to the Arkeia server. The backup itself is considered as finished, logged as such in the Arkeia logs, and all hardware peripherals used by the backup are released by the server and made available for other tasks. This is the default completion, and it is highly recommended to use this strategy most of the time.

- **“Continuous”:** A **“Continuous”** completion strategy means that, once the backup has finished saving the data contained in its Savepack, it stays **“active”** and is not considered as completed by Arkeia version 5. The hardware used for the backup operation is still considered as reserved by the Arkeia server, and cannot be used for other operations. It is possible to add new Savepacks to the backup, as long as the backup is not terminated by the user.



Please note: While a **“Continuous”** completion strategy may be interesting in some situations, it should be used with caution. It may be required by some Arkeia plug-in to perform its function correctly — please refer to the documentation delivered with the plug-in to determine if this is the case. In the same way, a continuous backup may be required if the system administrator needs to quickly add savepacks to an ongoing backup.

In any case, **“Continuous”** backups should never be used for unattended (**“Periodic”**) backups! For more information, refer to chapter **“8: Periodic backup”** on page 123 of this manual.

- **“Savepack”:** Clicking on this drop-down presents the different Savepacks that have been defined and can be selected to be backed-up.

- **“Drivepack”**: Clicking on this drop-down list presents the different Drivepacks that have been defined and can be used to perform the backup operation.
- **“Pool”**: Clicking on this drop-down list presents the different Tape Pools that have been defined and can be used to contain the backup data.
- **“Tape strategy”**: This field provides the different strategies that can be used to control the usage of tapes by a backup. The three different strategies are the following:
 - **“Use new tapes”**: This strategy forces Arkeia to use new tapes every time a backup is started. Use this strategy if you would like to keep the data of a backup on new tapes. Please note that the **“Complete existing tapes”** strategy will be used if there are no tapes available.
 - **“Complete existing tapes”**: This strategy will fill tapes that have already been used for previous backup.
 - **“New tape at level break”**: This strategy fills up the tapes if the backups performed are of the same level. Otherwise, it uses a new tape each time the level is different. For instance, if there are two levels “1” and “2”, and each backup is performed twice. Assuming the level “1” backup is performed first, Arkeia will use a new tape for this backup and complete it twice, using the same tape. However, when the level “2” backup is performed, Arkeia will use a new tape, specifically for the backups of that level. Every level “2” backup will then be stored on the same tape, until the tape has been filled-up and a new one has to be used.

The strategy used for the tape usage should take into account the amount of data to back up, the space remaining on a tape after a **“Total”** backup has been performed, the life expectancy and overall performances of the backup media, as well as the time necessary to perform a complete restoration of the data saved.

- **“Validity”**: This field determines the duration of data retention. Once the data retention defined has expired, the tape is recycled, and the data it contains will be erased by future backups. Note that this parameter is modified if the backup type is set to **“Archive”**: in this case, the retention data is overridden and the tape is never recycled.
- **“Parallelism”**: This field determines the number of parallel “flows” that should be open between the Arkeia server and the Arkeia client when performing the backup. For a more detailed explanation of the Arkeia data flows, refer to the section named “Arkeia Internal Architecture” on page 57 of this manual.



Please note: For more information on how to create logical objects within Arkeia version 5, such as Savepacks and Tape pools, please refer to the chapter “6: Objects and users creation” on page 71 of this manual.

A Periodic backup requires several objects to be defined before the backup itself is launched.

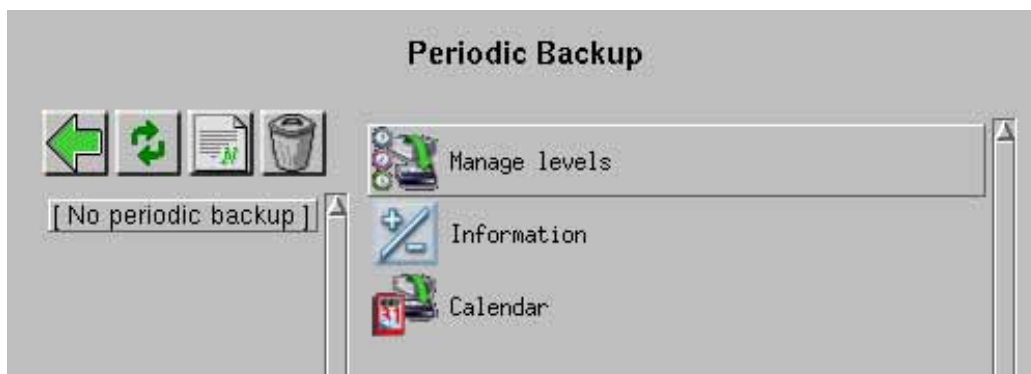
Once this step is completed, clicking on the “OK” button in the lower part of the window will end the creation of the Periodic backup and return to the main “Periodic backup” menu.

3.4 Step 4 – Defining exceptions and advanced options

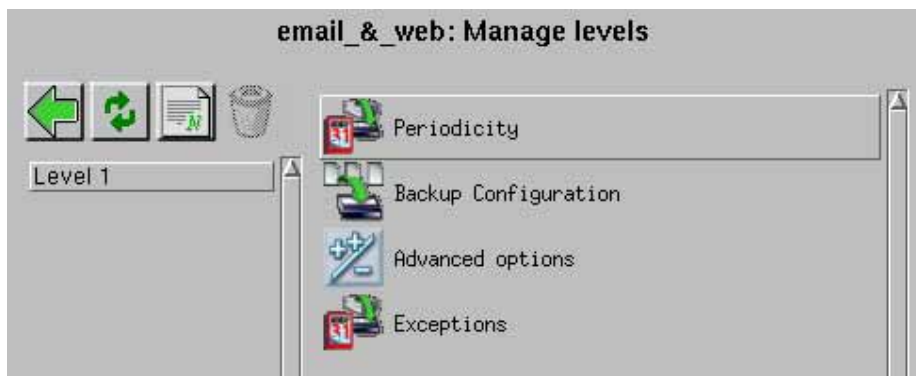
The last, optional, step is to define exceptions to the periodicity rules and commands that should be launched by Arkeia before performing the backup. To do this, follow the instructions detailed below:

- *Defining exceptions:*

To define exceptions, click on the “Manage Levels” icon in the main “Periodic Backup” menu:



This opens the “Level management” menu within the main Arkeia version 5 window:



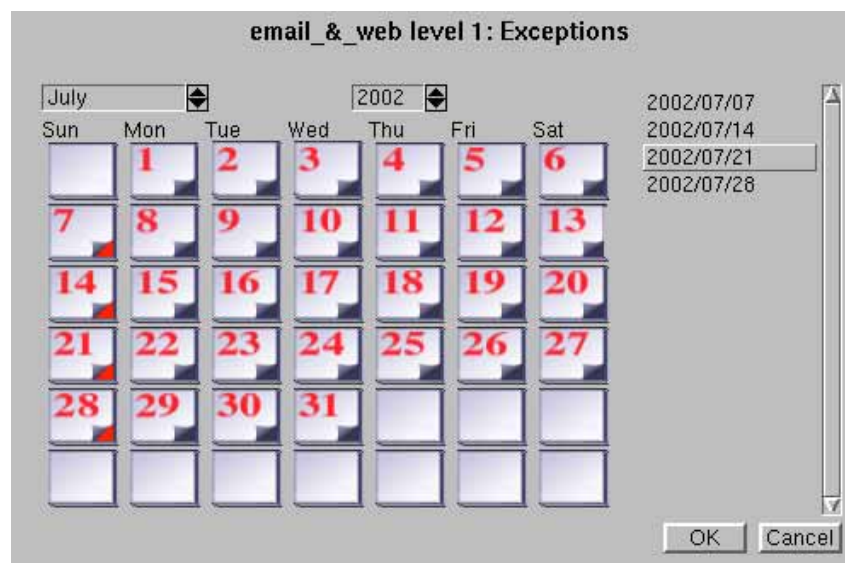
In this menu, the following options are available:

- “**Periodicity**”: Clicking on this icon allows the operator to re-configure the periodicity of the selected level. For more information on the configuration of the periodicity, refer to the previous sections of this chapter.
- “**Backup configuration**”: Clicking on this icon allows the operator to re-configure the backup of the selected level, by changing its configured Savepack or Tape Pool, for instance. For more information on the configuration of the backup, refer to the previous sections of this chapter.

- “**Advanced options**”: Clicking on this option allows the operator to re-configure the “**Command before**” and “**Command after**” that should be performed before launching the selected backup level. For more information on this option, refer to “**Creating savepacks**” in chapter “**6: Objects and users creation**” on page 71 of this manual.
- “**Exceptions**”: Clicking on this option allows the operator to define exceptions to the periodicity of the selected level.

As noted above, “**Exceptions**” defines days during which the backup level will not be executed.

To define Periodicity exceptions, click on the “**Exceptions**” button, to display the following screen:



In this screen, clicking on a day within the central planning will select the day as an exception. Exception days that are already defined appear in the short list on the right of the planning, and their corner appear in a red color.

As can be seen in the example above, four days have been selected as exceptions in the central planning – their dates appear as well on the right of the planning.

Clicking on the “**OK**” button returns to the “**Level management**” menu and confirms the exception configuration.



Please note: All exceptions must be explicitly defined for each level. This means that if, for instance, Tuesdays are considered permanent exceptions to backups, every Tuesday in every month should be selected in such a way.

- *Excluding a directory from all backup operations*

If, for any reason, a directory should be excluded from a backup operation, a special (empty) control file should be created in it.

Under a UNIX system, a control file can be created very easily, by entering the command:

```
touch <name of the control file>
```

With other operating systems, the most simple solution is to create an empty file with a standard text editor (such as the Windows Notepad) and to save this file in the correct directory.

The short list detailed below explains the control file that can be used to exclude a directory from a backup operation:

- If Arkeia version 4.x UNIX clients are used with an Arkeia version 5 server, use an empty file named “.OPB_NOBACKUP”. This file should be created on the client, in the directory to be excluded from all backup operations.
- If Arkeia version 5 UNIX clients are used with an Arkeia version 5 server, use an empty file named “.ARK_NOBACKUP”. This file should be created on the client, in the directory to be excluded from all backup operations. Please note that the Arkeia version 5 server can accept either name for the control file.
- If Arkeia version 5 Windows or Netware clients are used, use an empty file named “NOBACKUP.ARK”. This file should be created on the client, in the directory to be excluded from all backup operations.



Please note: This functionality should be used with caution, as putting a control file in a directory excludes all the content of this directory from all Arkeia backup operations – including all subdirectories.

If a directory that contains a control file is encountered by Arkeia version 5 during a backup operation, an entry is made in the Arkeia backup logs.

CHAPTER 9: DATA RESTORATION

- *Summary:*

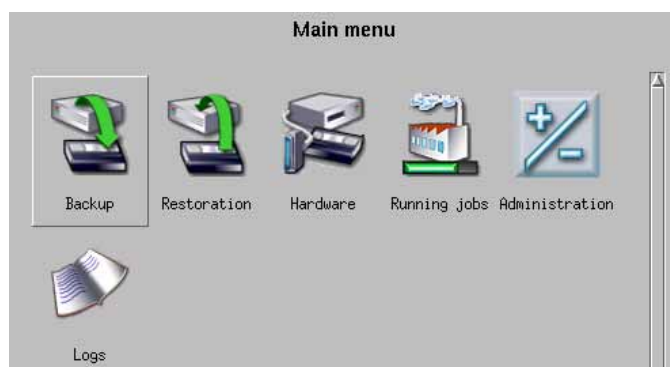
Arkeia version 5, like all backup programs, offers an advanced Restoration function. Through the use of this function, the data saved by one of the backup functions can be restored from tape and copied back to hard disk of the Arkeia client.

This chapter will present the normal procedure to perform a Restoration, as well as the different options that can be applied to a Restoration in Arkeia version 5.

1 • Accessing the Restoration function

To access the restoration function of Arkeia version 5, please follow the steps outlined below.

The Arkeia main menu is usually displayed when the program is first opened:



Within the main menu of Arkeia version 5, click on the “Restoration” icon to access the Restoration function.

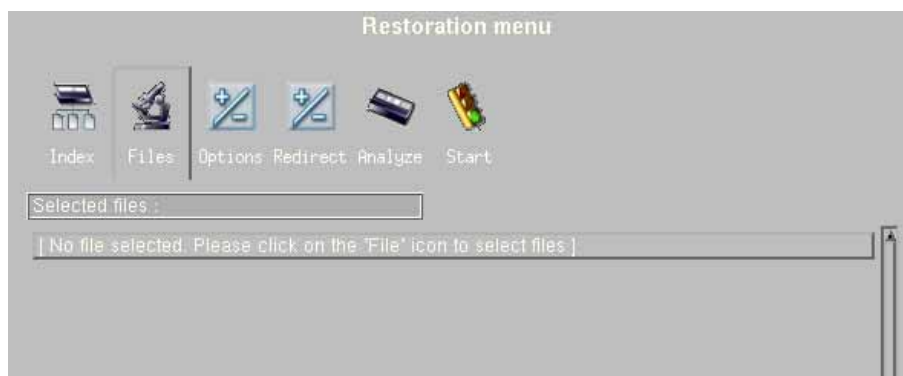
The “Restoration menu” is then displayed within the main Arkeia window.



The rest of this chapter will detail the different functions available in this menu, and the correct procedure to restore files that have already been backed up by Arkeia version 5.

2 • The Restoration menu

The “`Restoration menu`” contains the following functions.



- “`Index`”: Clicking on this button opens the “`Database Navigator`” window, a small browser that allows you to select files based on the time and date they were backed up by Arkeia version 5.
- “`Files`”: Clicking on this option allows the selection of the files that should be restored, based either on the time and date they were backed up (see the “`Index`” command above), or on the path they occupied at the time the backup was performed.
- “`Options`”: Clicking on this option allows the selection of the different options that should be applied to the Restoration operation being defined.
- “`Redirect`”: Clicking on this option allows the files that are going to be restored to be “redirected” to another path or machine. This, in turn, allows the restored files to be compared to their copies, in case you suspect either the backup or the existing files to be corrupted.
- “`Analyze`”: Clicking on this option will force Arkeia version 5 to open its database, search for the files that have been selected using the other options, and display the tapes these files were backed-up on. This option can save a lot of time if, for instance, some backup tapes are kept “off-site” in a storage facility.

The list under the main options of the Restoration menu displays all the files that have been selected, or (as is the case in the screen pictured above), a short message indicating that no files have been selected.

- *Overview of the restoration*

Regardless of its size, a Restoration is always performed, within Arkeia version 5, in the same manner:

- File selection.

- Redirection/Options settings
- Tape analysis (optional)
- Launch of restoration.

Within this procedure, the file selection and the Redirection/Option settings will be detailed in this chapter, the last two steps of the Restoration procedure (Tape analysis and launch) are self-explanatory.

3 • File selection

Selecting files for Restoration can be accomplished in three different ways within Arkeia version 5:

- The “**D a t a b a s e N a v i g a t o r**” (time-based file selection)
- The Database search engine (file name selection)
- The complete path and name of the file

All these methods have their advantages and their disadvantages. This section will detail the different methods, their pros and cons, and their general usage.

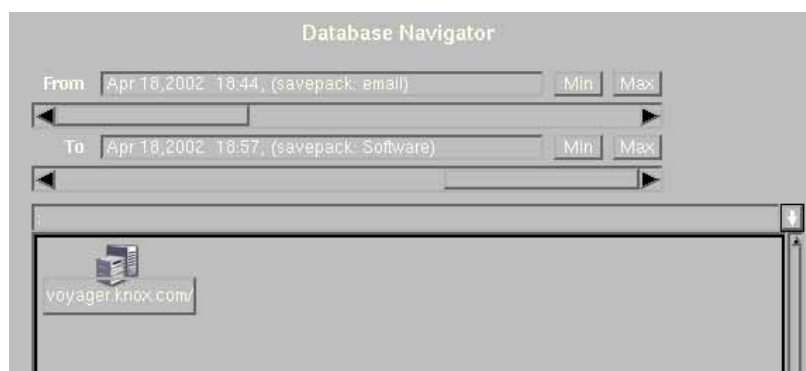
3.1 File selection with the Database Navigator

This function provides an easy way to explore the Arkeia database and select the files that should be restored. This function should be used when you know which files should be restored, and at which date and time they have been backed-up.



To access the Database navigator either click on the “**I n d e x**” button in the main Restoration menu, or click on the “**F i l e s**” index and then click on the “**N a v i g a t o r**” icon within the “**S e l e c t f i l e s**” menu.

The database navigator window is then displayed within the main Arkeia window.





Please note: if several backups have been performed on the same machine (for example a “Total” backup followed by several “Incremental” backups), it is possible to select the date of the “Total” backup with the “From” slider and the date of the last “Incremental” backup in the “To” slider.

Selecting the client to be restored with the sliders in the correct positions means that the client will be completely restored, as the data from the first, “Total”, backup will be restored first, and the different “Incremental” backups will then be restored.

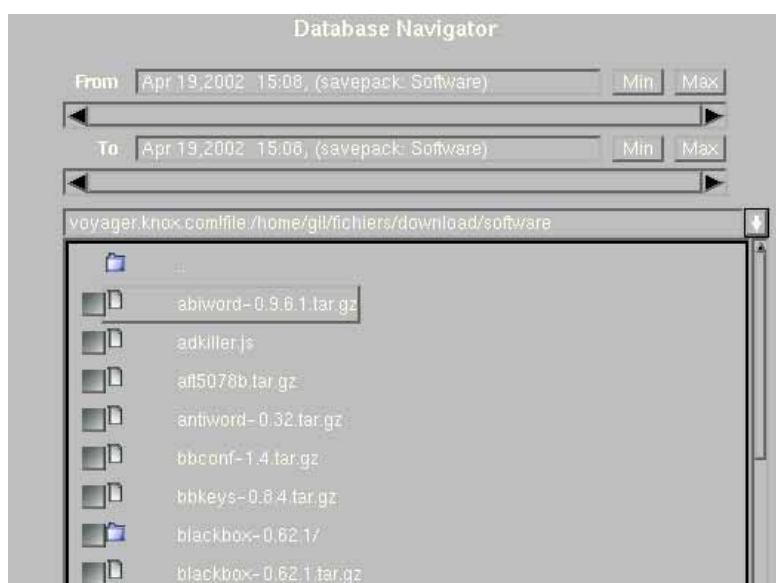
Within this screen, the following functions are available:

- “From/To”: Clicking on these sliders will limit the restoration to the files that have been backed up by Arkeia between the two dates they define (See the important note above).


In the screen shown above, the only entries in Arkeia database are two backups that were performed on April 18th, 2002, with different Savepacks and hours. With an Arkeia server that has been in use for a longer time, it’s recommended to limit the amount of data displayed by using the sliders.

- The “Navigator” window itself. Within this window, Arkeia displays the files, directories, machines (etc.) that have been backed up between the two dates defined with the sliders. Only the files backed up will be displayed. If a machine, or a directory does not appear in this Navigator window, it means it has not been backed up between the two chosen dates.

Using the “Navigator” window, it is possible to browse the Arkeia database and search for the files (or directories, etc.) that should be restored.

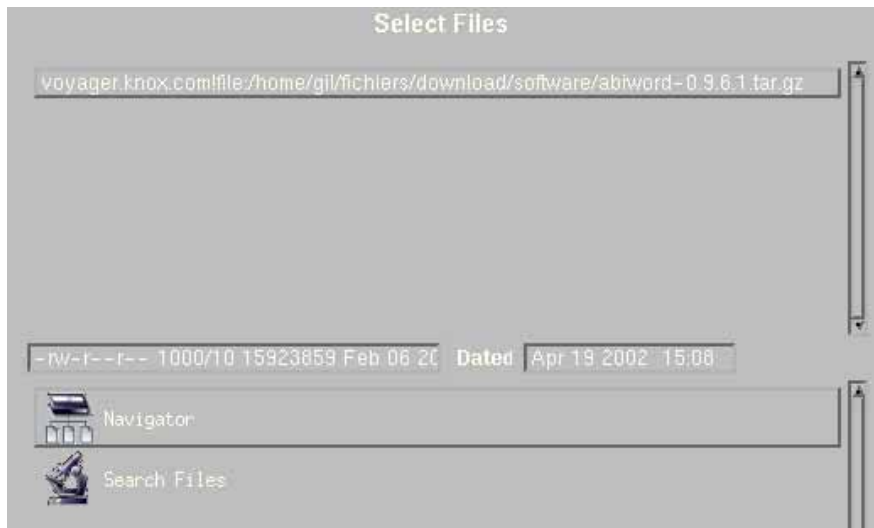


To select a file or an entire directory in the “Navigator” window, click on the check box next to its name.




Once the backup dates have been defined and the files to be restored have been chosen in the “Navigator” window, right-clicking anywhere in the window and selecting the option “Validate (selection)” returns to the main “Select Files” window, with the selection displayed.

In the screen shown below, a single file has been selected to be restored. This file is named “abiword-0.9.6.1.tar.gz” and has been found in the directory: “/download/software/” of the machine named “voyager”.



Once the files to be restored have all been selected, it is possible to right-click anywhere in the “Select files” window to display the relevant contextual menu.



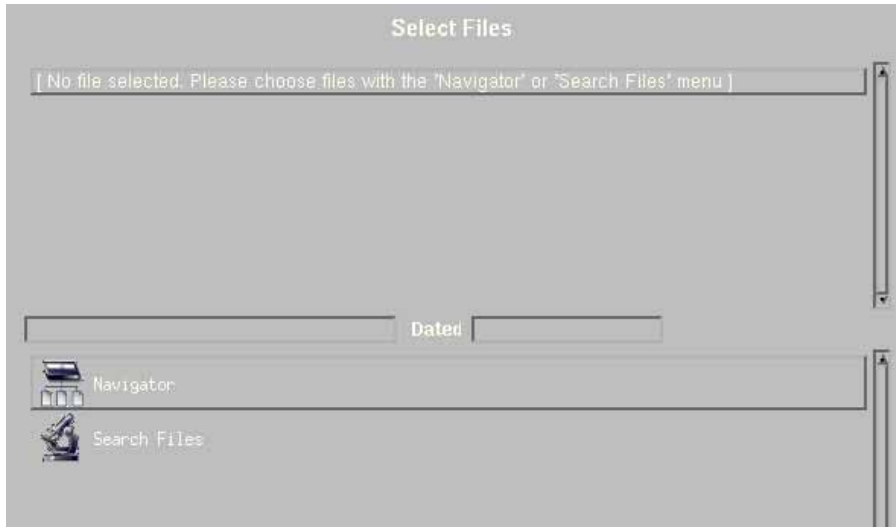
Within this contextual menu, the following options are available:

- “**A d d f i l e**”: Clicking on this option displays a request to enter the complete name and path of a file to be restored by Arkeia. This option is only valid if the user remembers precisely which file is supposed to be restored.
- “**R e m o v e f i l e**”: Clicking on this option removes the selected file from the list of files to be restored.
- “**M o d i f y f i l e**”: Clicking on this option displays a request to modify the path/name of a file to be restored. Again, this option is only useful if users remember precisely which file is to be restored.
- “**C l e a n l i s t**”: Clicking on this option removes all the selected files from the list of files to be restored (the files themselves are not deleted, nor their entries in the Arkeia database, only their selection).
- “**P r e v i o u s m e n u**”: Clicking on this option closes the “**S e l e c t f i l e s**” window and returns to the main “**R e s t o r a t i o n**” menu.

3.2 File selection through the file search

While the Database Navigator of Arkeia version 5 is a very practical way to select files to be restored, it can be tedious to browse the database tree to restore only one file or one directory. In many other instances, it is actually better to be able to select several files or directories at once.

To do this, Arkeia offers a powerful “**S e a r c h f i l e s**” function, which can be accessed from the “**S e l e c t f i l e s**” screen.



To access the “**Search files**” function, click on its icon in the bottom half of the “**Select files**” window. The main search windows is then displayed on the screen.



Within the “**Search files**” window, the following functions are available:

- “**Start of search**”: This field contains the path which is going to be used as the basis for the search. By default, this path is the “**root**” path of the Arkeia database of the machine currently selected. The fields detailed below contain the different search criterion that will be used by Arkeia to select the files to be restored from its database.
- “**filename matching exactly**”: Enter the filename to be searched exactly “as is” in the database. Click on the field check box to apply this criteria to the search.
- “**filename containing**”: Enter the portion of the file name to be searched for by Arkeia in its database. Click on the field check box to apply this criteria to the search.

- “filename ending with”: Enter the end of the file name to be searched for by Arkeia in its database. Click on the field check box to apply this criteria to the search.
- “filename starting with”: Enter the start of the file name to be searched for by Arkeia in its database. Click on the field check box to apply this criteria to the search.

The last field displays the concatenated search criterion.



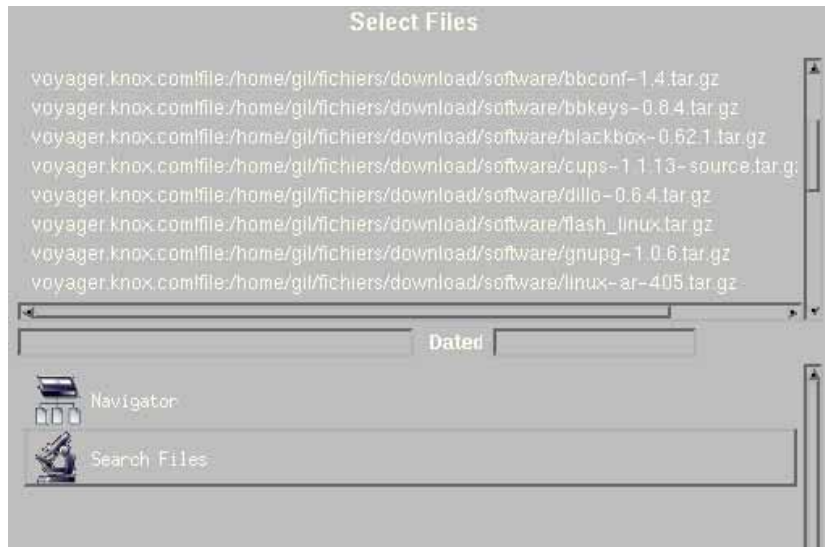
Once all the criterion have been selected, right-clicking anywhere on the “Search files” window displays the contextual menu pictured on the left.

In this menu, clicking on the “OK” option starts the search and adds the files found to the selected files list, while clicking on the “Cancel” option stops the searches and returns to the “Select files” window.

For instance, if the files to be restored all end with the “.tar.gz” extension, it is possible to enter the following information in the “Search files” window:



Clicking on the “OK” option of the contextual menu, then on the “Cancel” option of the contextual menu then returns to the “Select files” window, where all the files ending with the “.tar.gz” extension are displayed.



Within this window, clicking on the “Previous menu” option of the contextual menu returns to the main “Restoration” menu.

The rest of this chapter will describe the different options that can be applied to the Restoration.

4 • Setting the options of a Restoration

The “Restoration menu” of Arkeia version 5 offers many options. One of them is the “Redirection” option, which is detailed below, in its own separate sections. All other options will be detailed within this section.

4.1 Accessing the options

To access the options of the “Restoration” menu, click on the “Options” icon in the main “Restoration menu” window:



The main “Restoration options” screen is then displayed within the main Arkeia window:



As can be seen in the screen pictured above, all the restoration options can be selected very quickly, simply by clicking on the check box next to their title.

4.2 Configuring the options

The first group of options available in the “Restoration options” screen deal mainly with the characteristics of the files being restored:

- “**A c c e s s r i g h t s**” : Clicking on this option will restore all the original access rights of the files. If left unchecked, this option will simply restore the current access rights.
- “**M o d i f i c a t i o n d a t e**” : Clicking on this option will restore the original date of the file. If this option is left unchecked, Arkeia will replace the original date and time of the file with the date and time at which the Restoration is taking place.
- “**F i l e s m o d i f i e d s i n c e b a c k u p d a t e**” : Clicking on this option will replace all existing files by the restored files. This means a newer version of the restored files will be deleted and replaced by the older, backed-up, version.

The files ownership options are the following:

- “**B y u s e r n a m e**” : Clicking on this option will force Arkeia to search the “/ e t c / p a s s w d” file for the user ID that matches the owner of the files being restored.
- “**B y u s e r I D**” : Clicking on this option will force Arkeia to use the original user ID of the files being restored, regardless of the entries in “/ e t c / p a s s w d”.

4.3 Applying the options to the Restoration

Once the options of the Restoration have been correctly configured, clicking anywhere on the “**R e s t o r a t i o n o p t i o n s**” window displays the contextual menu.



In this menu, clicking on the “**OK**” menu item will confirm the options and return to the main “**R e s t o r a t i o n m e n u**”, while clicking on the “**Cancel**” menu item will cancel the option selection and return to the main menu.

5 • Redirecting a Restoration

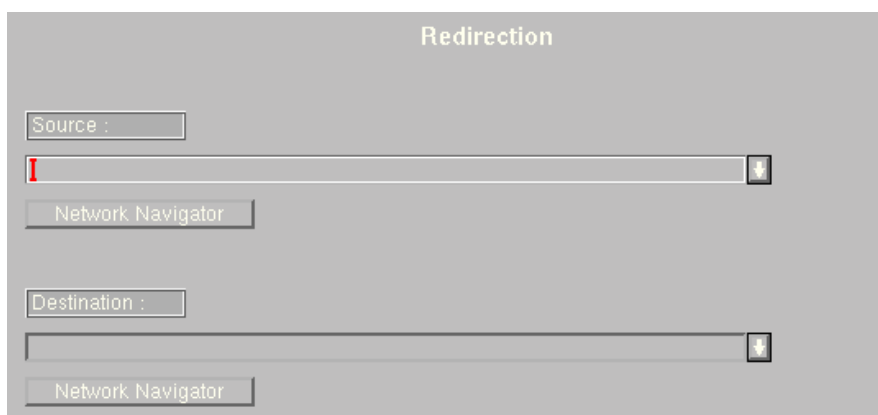
An important function of the Restoration in Arkeia is the ability to restore backed-up files from an original path to another path, or even another machine. This function even allows files to be restored on a different machine than the one which has been backed-up, which can be very practical if the original machine has been lost in a disaster, for instance.

5.1 Accessing the redirection function

To access the redirection, click on the “`Redirect`” icon in the main “`Restoration` menu”:



The “`Redirection`” menu is then displayed in the main Arkeia window:



The rest of this section will describe the different options of this function.

5.2 Options of the redirection

The following options are available within the “`Redirection`” menu.

- “**Source**”: Clicking on this drop-down list displays files that have been selected with one of the different selection function of the Restoration.
- “**Destination**”: Clicking on this drop-down list allows the modification of the path that will be used to restore the file chosen in the previous list.
- “**Network Navigator**”: Clicking on either of these buttons allows the original/redirection paths to be chosen within a network browser. This solution may sometimes be more convenient than entering the path/file names “by hand”.

For instance, if a file named “`abiword-0.9.6.1.tar.gz`”, which was originally contained in the directory “`/download/software/`” should be restored to the path “`/download/rst/`”, the following information should be entered in the “**Redirection**” screen:

The screenshot shows a graphical user interface for file restoration. It features two main sections: 'Source' and 'Destination'. Each section has a text input field containing a path, a 'Network Navigator' button, and a small square icon with a downward arrow. The 'Source' field contains the path `ager.knox.comfile:/home/gil/fichiers/download/software/abiword-0.9.6.1.tar.gz`. The 'Destination' field contains the path `voyager.knox.comfile:/home/gil/fichiers/download/rst/abiword-0.9.6.1.tar.gz`.

Please note that the paths above respect the standard Arkeia syntax, which is different from the standard Linux/UNIX hierarchical syntax. The Arkeia syntax will be explained in the following section.

The redirection works by replacing the “**Source**” path by the “**Destination**” path whenever restoring a file. For instance, if the original file is “`/home/jack/doc-31.txt`”, and the file “`doc-31.txt`” should be restored instead as “`/home/jill/restore/doc-31.txt`”, the following paths should be entered in the fields detailed above:

- “**Source**”: “`/home/jack/`”
- “**Destination**”: “`/home/jill/restore/`”

The rest of the syntax is detailed below.

5.3 Arkeia path syntax

When Arkeia is requesting a path to be entered in an entry field, in the “**Redirection**” menu or anywhere else in the program, the path and file name entered have to respect a special syntax. This syntax is required by Arkeia since it is able to backup different client machines through a local area network.

The following is the standard Arkeia version 5 syntax:

```
FQDN_machine_name[!plugin_name]:/path01/.../pathX/filename
```

Here is a detailed explanation for all the parts of this syntax:

- “**FQDN_machine_name**”: This is the “Fully Qualified Domain Name” (FQDN) machine name. This means that Arkeia version 5 accepts names such as “**zap.foo.com**”, where “**zap**” is the name of the client, and “**foo.com**” the name of the domain which contains the machine named “**zap**”.

Please note that an exclamation point “**!**” (also called a “**bang!**” in the UNIX world) should separate the “**FQDN_machine_name**” from the plug-in name (see below).

- “**!plugin_name**”: This is the name of the plug-in which is used by the Arkeia server to perform the backup and restoration operation. In the example screen above, the plug-in name is the default plug-in, “**file**”, which is used for all normal file system operations. Future plug-ins may be named, for instance “**oracle**”, “**mysql**” for plug-ins specialized in database backup and restoration. Please note that the plug-in name should be separated from the path by “**:**”.
- “**/path/.../filename**”: This is the path and the name of the file on which the operation should be performed.

Here are a few examples of correct paths/file names, that will be accepted by Arkeia as valid:

```
zap.foo.com!file:/home/jack/cad-files/DCCF_0124122589.dat  
zip.foo.com!file:/home/tina/docs/pdf/arkeia/man_arkc.pdf  
email01.neon.com!file:/etc/mail/statistics/
```



Please note: This syntax is valid whether the machine is a Windows or a UNIX machine. Arkeia uses slash characters “/” for paths on all operating systems.

The only difference in syntax between the Windows and the UNIX Arkeia client is that Windows machines require the volume letter to be inserted before the path and file name.

Therefore the following paths are correct only for Windows machines:

```
bilbo.shire.org!file:c:/files/docs/linux/pdf/LinuxFromScratch.pdf  
gandalf.mordor.com!file:d:/patrick/meeting/Conference2002.ppt
```


5.4 Applying the redirection

Once the redirection has been properly configured, it needs to be confirmed by right-clicking anywhere in the window to display the contextual menu.



In this menu, clicking on the “Ok” menu item will confirm the redirection and return to the main “Restoration menu”, while clicking on the “Cancel” menu item will cancel the redirection and return to the main menu.

6 • Tape analysis and launching a Restoration

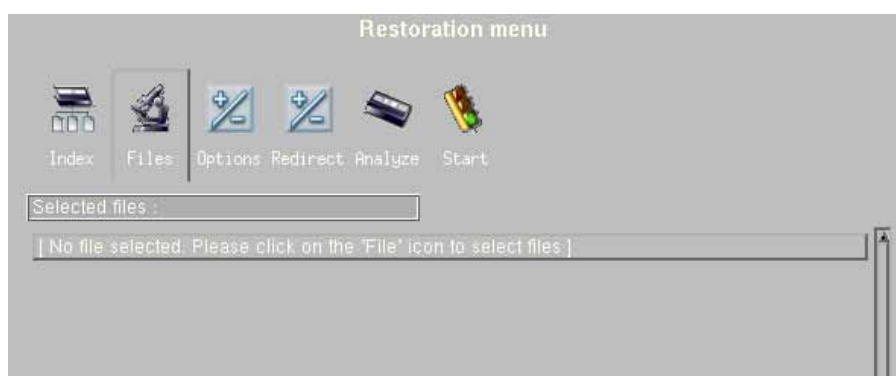
6.1 Tape analysis overview

Finding the right tape to perform a Restoration can sometimes be difficult, especially if the tapes have been moved from the backup device to a storage facility.


Arkeia version 5 offers a small utility that will search the database for the selected files and displays the tapes that will be required by Arkeia to perform the Restoration correctly.

6.2 Using the Tape analysis

To access the “Tape Analysis” feature of Arkeia version 5, click on the “Analyze” icon in the “Restoration menu”.



Arkeia version 5 then performs its database analysis and displays immediately the list of tapes that will be needed for the Restoration:



In the example above, only one tape will be needed by Arkeia version 5, called “email02”. If necessary, the user should then retrieve the tape in a storage facility for instance and make sure it is inserted in the backup device that will be used for the Restoration of the selected files.

Needless to say, if more files had been selected, the list would contain more tapes and not just one.

6.3 Launching a Restoration

To recapitulate the Restoration procedure that was described in the second section of this chapter, here are the steps to be followed to perform a Restoration:

- File selection.
- Redirection/Options settings
- Tape analysis (optional)
- Launch of restoration.

Once these steps have been followed as described in the rest of this chapter, the Restoration can be launched and performed on the selected files.

To launch a Restoration, click on the “Start” icon in the main “Restoration menu”:



If the Arkeia server uses only a stand-alone tape drive, Arkeia version 5 will display a message asking the system administrator to select the drive that should be used by the restoration, and to confirm that choice by clicking on “OK”.

Arkeia then opens the main “Restore” screen and performs the Restoration as configured by the user:



The following items are available in this window to monitor a running restoration:

- **“Elapsed time”**: This is the first field displayed on the left of the running restoration window. Its purpose is to display the time elapsed since the launch of the restoration.
- **“Server”**: This is the name of the server the restoration is running on.
- **“Cruise”**: This slider allows the system administrator to control the amount of bandwidth consumed by the Restoration. If important network functions are sluggish while Arkeia is performing restorations, use the **“Cruise”** control to reduce the amount of bandwidth used to perform a restoration. Note that this will also slow down the operation.

In the middle of the screen, a graphical meter displays the amount of data being transferred between the client and the server.

On the right side of the screen, the following fields and drop-down boxes supply additional information on a restoration:

- **“Global”**: Clicking on this drop-down list displays all the drives that are used by the restoration. This allows the system administrator to restrict the display of information to a particular drive.
- **“Speed in KB/s -- Speed in MB/mn”**: Clicking on this button changes the type of information that is displayed in the fields below it. Please note that this information is displayed in two separate fields: **“Instant”** indicates the current speed of the restoration and **“Average”** the overall speed of the restoration.
- **“Restore Speed/Network throughput”**: Clicking on this button changes the type of information that is displayed in the fields below it, from the speed of the restoration itself to the amount of information transferred by the network.

Once a restoration has been completed, Arkeia displays a message indicating the overall results of the operation. Checking this message is important, as it may provide crucial information in case the restoration failed.

CHAPTER 10: ADVANCED HARDWARE

- *Summary*

This chapter of the documentation will present the management and configuration of tape libraries under Arkeia version 5. Tape libraries are advanced backup devices that are ideal for saving and restoring large amounts of data to/from tape.

1 • Overview

1.1 Differences between libraries and other backup devices

Libraries usually contain one or more of the following components:

- **Robotic arm:** These devices are used to transfer tapes to/from the tape slots to the tape drives. They can also be called “medium changer” or “robot” and, in the rest of this chapter, these terms will be user interchangeably.
- **Tape slots:** Tape slots contain tapes that will be used for backup and restoration operations. A typical library can hold dozens of tapes, each in its own separate slot.
- **Drives:** These devices are, of course, used to read and write data to the tapes that are inserted in them by the robotic arm. A typical library contains usually two (or more) tape drives.
- **Gates:** A library “gate” can be used to eject a tape without having to open the entire library. Tapes inserted in the gate are then moved to free tape slots by the robotic arm. In the same way, tapes can be moved from their slots to the gate, for instance in order to be archived away from the server room. Most libraries have to stop all operations when they are opened completely and gates can therefore be very important to a busy library.

Of course, the precise configuration of a library can vary from model to model and from supplier to supplier. Refer to the documentation supplied with the library for more information.



Please note: There is a much longer discussion on backup devices in chapter “3: Hardware installation” on page 35 of this manual. Also, chapter “1: Installation pre-requisites” on page 6 contains additional information which may be useful during the configuration of a library.

1.2 Libraries managed by Arkeia

The list of libraries managed by Arkeia is quite long and contain, at the time of this writing, more than 300 devices. Due to the length of this list, it is not incorporated in this manual, but is available online, at the following URL:

`http://www.arkeia.com/mass.html`

It is highly possible that a device which is not listed at the web page indicated above can (and will) be supported. For more information, please contact:

`sales@arkeia.com`

1.3 Advantages of the library over other backup devices

The advantages of a library should be evident from the list of its components: libraries are powerful device, that allow completely automated tasks to run unattended on the Arkeia server.

Thanks to their slots and multiple drives, libraries offer generous amounts of storage space to medium-to-large networks or to servers needing a high-speed, high-volume storage solution.

Since Arkeia version 5 manages a lot of libraries “out of the box”, this solution is ideal for a lot of networks and for large volumes of data.



Please note: Using libraries under Arkeia version 5 requires a separate licence to be acquired from Arkeia Corp.

For any and all information, please contact:

`sales@arkeia.com`

2 • Creation of a library under Arkeia version 5

2.1 Creation of the library

- *Overview*

Creating a library under Arkeia version 5 is simple process, which is helped by the powerful functions offered by the GUI. However, it is important to respect a certain order to this process, which is detailed below:

- Create the drive(s) that will be added to the library.
- Create the tape(s) that will be added to the library.
- Create the library itself, including its robotic arm and its slots.
- Add the drive(s) to the library.
- Assign the tapes to the slots of the library.

Creating tapes and drives has already been covered in this manual, in the chapter “6: Objects and users creation”, and, therefore, it won’t be covered in this chapter. Refer to chapter “6: Objects and users creation” for more information on the creation and configuration process.

As noted in the previous section, chapter “3: Hardware installation” contains many information on how to detect, install and manage backup devices, either from the GUI or from the UNIX command line.

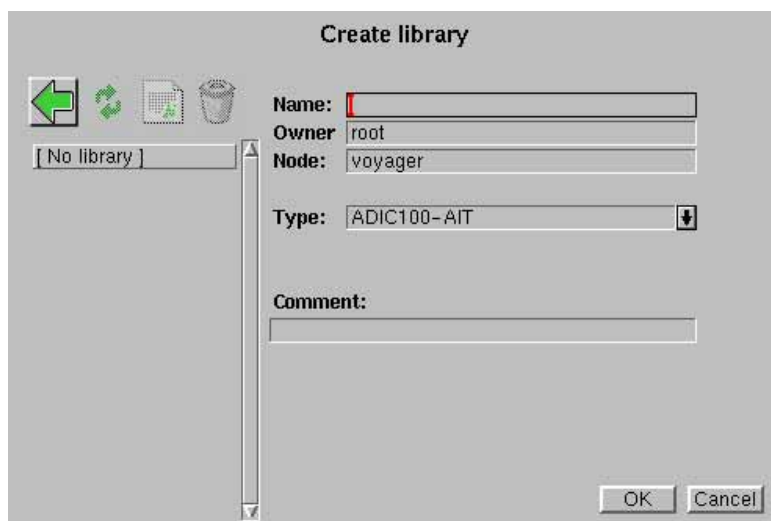


Please note: If you haven’t read the chapters listed above, please do so before reading the rest of this section! Doing this will probably save a lot of time, and allow the library installation to proceed more smoothly.

- *Creating a library*

To create a library under Arkeia version 5, follow the steps outlined below:

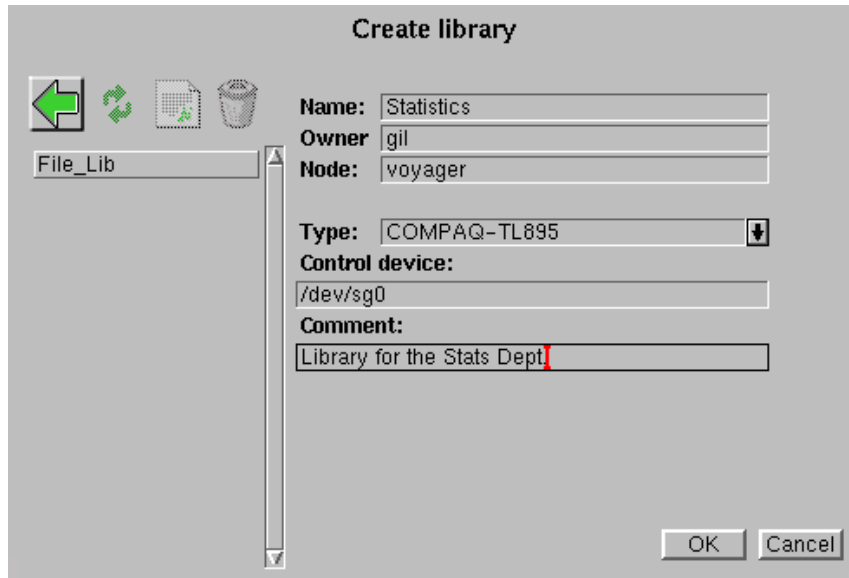
In the main menu of Arkeia version 5, click on “Hardware”, then on “Library”. In the “Library” main menu, click on the “New” icon (above the list of libraries) to start the library creation process. The following menu should then be opened in the window of Arkeia version 5:



In this menu, the following functions are available:

- “**Name**”: This field can be used to assign a name to the library. This name will then be used in the different logs of Arkeia version 5.
- “**Owner**”: The name of the owner of this library can be entered in this field. By default, the user who opened the GUI interface of Arkeia version 5 and created the library is the owner of the hardware device. This name can be changed to restrict the usage of the library to a specific user, which may be appropriate if several libraries are connected to a single Arkeia backup server.
- “**Type**”: This drop-down menu is a list of all supported libraries. To navigate more quickly in this list, it is possible to type the first letters of the library type: for instance, typing “**F**” while this menu is active will move automatically to “**FALCON-1530**”, which is the first library type whose name begins by “**FAL**”. Please note that most libraries also require an additional licence to be managed by Arkeia version 5.
- “**Control device**”: This field should contain the device used to control the library medium changer. For instance, under Linux, such a device could be “**/dev/sg0**”. Please note that this field will only appear if it is required by the given library type.
- “**Comment**”: This field can contain a descriptive comment, such as the precise location of the library or some other information.

When the library configuration is finished, clicking on the “**OK**” button at the bottom of the menu creates the library and returns to the main “**Libraries management**” menu. In this menu, the library that was created should appear in the left-hand list of libraries.



In the example above, we have created a “Compaq-TL895” library, named “Statistics”. This library is controlled, under Linux, by the device “/dev/sg0”. The exact control device may vary, based on the operating system used.



Please note: Under Arkeia, a special type of library exist, which is called a “File” library. These libraries function in the same way as normal libraries.

File libraries are an easy way to test the library configuration, and can be used to make backup to hard disks, instead of making backups to tape. File libraries are going to be detailed in section 4 of this chapter.

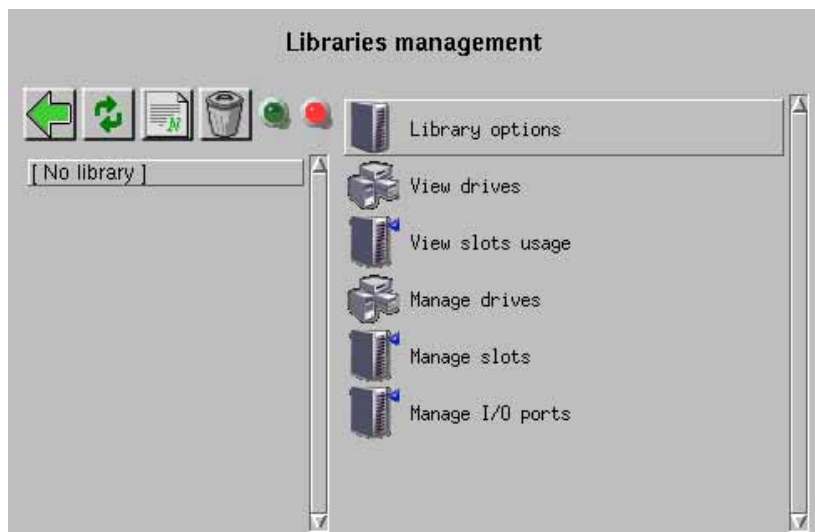
2.2 Adding and cleaning drives into a library



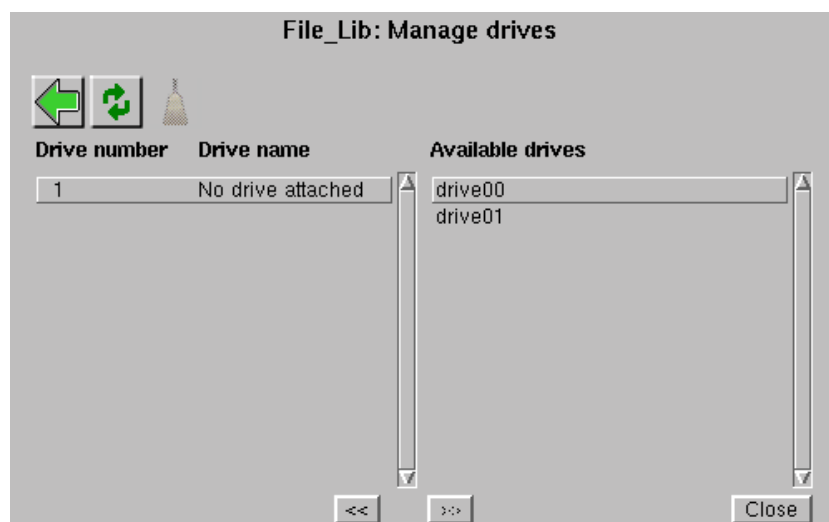
Please note: Allow us to repeat that drives should be created before being added to a library. Refer to chapter “6: Objects and users creation” for more information.

To add existing drives to a library, follow the steps outlined below:

Within the main “Libraries management” menu, click on the option “Manage drives” (refer to the screen pictured below):



The “**Manage drives**” menu is then opened within the main Arkeia window:



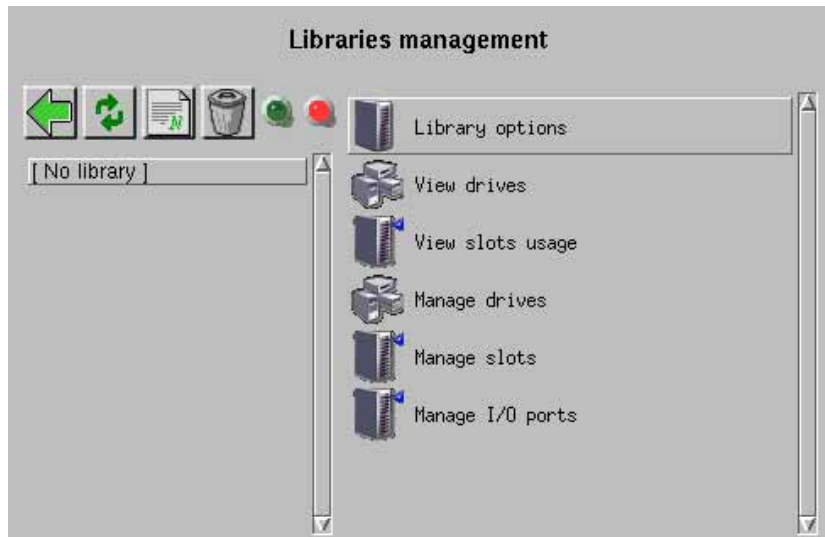
This menu is divided into two lists: on the left is the list of drives positions within a library, and, on the right, the list of available drives. Using the “<<” and the “>>” button allow you to attach and detach drives in a quick and easy manner. In the example above, two drives “drive00” and “drive01” can be assigned to the current library.

Select a drive position, then an available drive, and click on the “<<” button to attach the drive to the library. Or, select a drive position, then a drive, and click on the “>>” button to detach the drive from the library. Once the library configuration is finished, clicking on the “Close” button takes into account the drive assignments and returns to the main “Libraries management” menu.

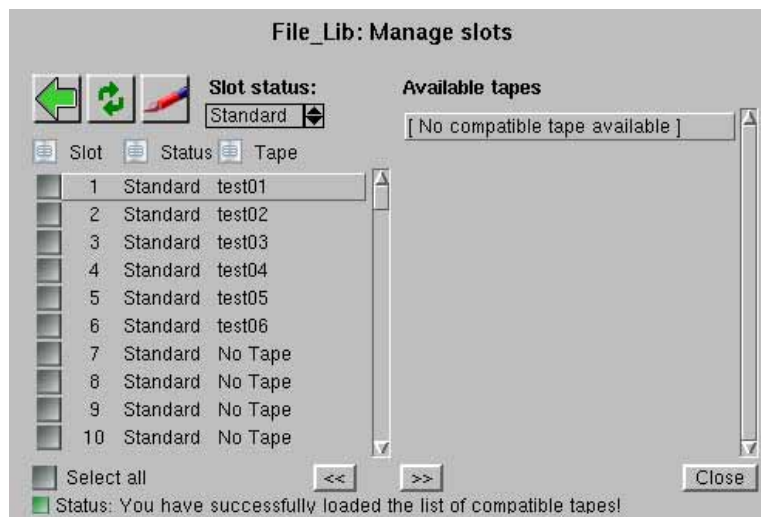
Also of interest in the “Manage drives” is the “broom” (Clean drive) button. To clean a drive, simply select the drive in the library, and click on the “broom” button. This will launch the cleaning process. During this process, Arkeia will transfer a cleaning tape from its slot (see below) into the selected drive and clean it automatically.

2.3 Adding tapes to a library (slot assignment)

To add (existing) tapes to a library, you have to assign these tapes to one of the internal slot of the library. To do this, follow the steps outlined below:



Within the main “**Libraries management**” screen pictured above, click on the option named “**Manage slots**”. The slot management screen should be displayed within the main Arkeia version 5 window:



This screen, like the “**Manage drives**” described above, is divided into two parts: the left-hand part is the list of slots within the library, while the right-hand is the list of available tapes.


To assign a given tape to a library slot, select both the tape and the slot by clicking on the grey check box next to their names, and click on the “<<” (assign) button.

On the other hand, to un-assign a tape from a library slot and make it available again, select the library slot where the tape is stored and click on the “>>” (un-assign) button.

When the configuration is complete, clicking on the “Close” button saves the slot status and returns to the main “Libraries management” screen.

Two other interesting items in the “Manage slots” menu are the “Slot status” and the “Write label” button:

- “Slot status”: This spin box provides a quick way to change the status of the selected slot. The status can be:
 - “Normal”: the slot can be used by Arkeia for backup and restore operations.
 - “Reserved”: the slot cannot be used as it is saved for another application.
 - “Cleaning”: the slot must contain a “Cleaning” tape. Self-explanatory.



Clicking on the “Write label” button will force Arkeia version 5 to write a new label on the selected tapes. While this may be useful in certain cases, Arkeia automatically writes a label to the tape during the first backup, so this operation should be considered as optional, especially if the library already has a bar code reader, which is recognized and used by Arkeia version 5.



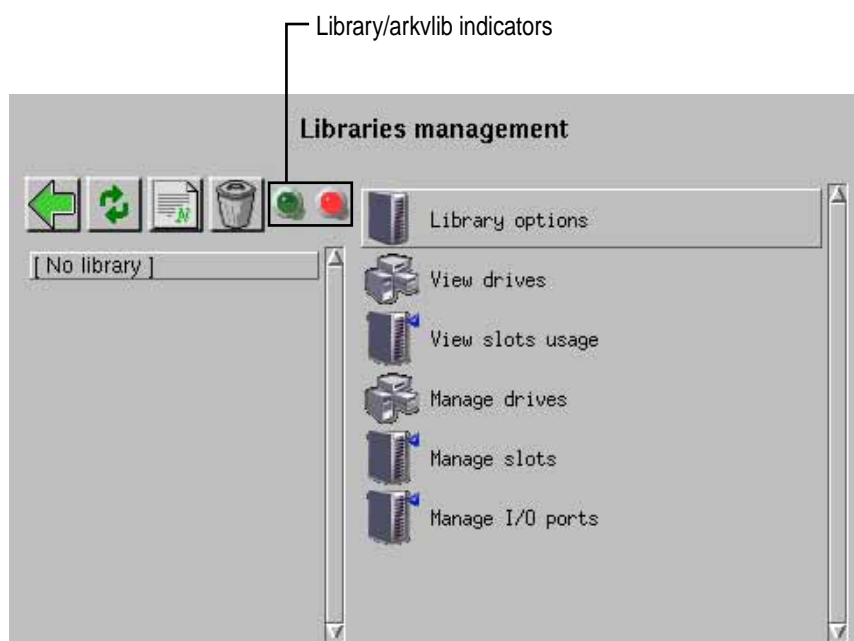
Please note: Since writing a label to tape assigns a new “name” to the tape, Arkeia is then unable to match the tape to its database records. Use this function with extreme caution, on tapes that do not contain important data, and only when directed to do so by a Arkeia Technical Support Specialist.

3 • Managing a library under Arkeia version 5

3.1 Starting and stopping a library

The very first task that should be undertaken, before using a library for any task, is to make sure that “arkvlib”, the Library manager of Arkeia version 5 is running correctly.

To do this, check the indicators that are between the library list buttons and the “Library options” item:



To start the library manager (or make sure it is already running), click on the green indicator. To stop the “arkvlib” library manager, click on the red indicator. Simply putting the mouse pointer over the indicators displays the current status, as returned by Arkeia version 5.

3.2 Using a library for a backup

There is no need to declare a library for a backup operation. To use an existing, configured, library in a backup operation, the system administrator should simply define a Drivepack (refer to chapter “6: Objects and users creation” for more information) that contains the drives contained in the library.

Whenever a backup is launched, the Drivepack that contains the drives of the library should be selected. Arkeia version 5 will then manage transparently the backup, using the library drives and all the advanced functions the library has to offer.

3.3 Tape operations under Arkeia version 5

To use the tapes contained in a library for a backup operation, all is needed is to create a Tape Pool that contains all the tapes present in a library, before the backup is launched. Of course, this presupposes that all the tapes have been created and assigned to the library slots.

When the backup operation itself should be executed, the system administrator should select, as the target media, the Tape Pool that contains the tapes of the library. Arkeia version 5 will then transparently use the library tapes, using the different SCSI device to change, when needed, the tapes in the library drives.

3.4 A special case: the file libraries.

File libraries are a special case among libraries: they do not use real “tapes”, but save the data they are supposed to contain to disk. This can be very useful for testing purposes, and to get a better sense of the inner working of Arkeia libraries in general.

File libraries are configured exactly the same as normal libraries, except that file tapes should be assigned a path during their creation. Also, File libraries cannot contain more than one drive, and should not be assigned any controlling device.

Arkeia Corp. recommends to its users to configure first a library to test the different possibilities offered by Arkeia before trying to install a real library.

The behavior of a file library is exactly the same as a normal library.

CHAPTER 11: OPTIMIZATION

- *Summary*

While the default performance of Arkeia version 5 can be very satisfying “out of the box”, there are cases in which the default configuration needs to be modified to produce better results.

This chapter, while it covers many subjects already detailed in other sections of this manual, attempts to present as much information as possible, with the specific aim of optimizing Arkeia’s performance.

1 • Overview

As a general rule, the best way to maximize performance for a network backup software is to send as much data, as fast as possible, from the disk drives to the tape drive.

From this point of view, the weak spot in a backup chain is often the tape drive: if not correctly fed with data, it will have to rewind very often and the throughput will plummet. Thus, the idea is to “saturate” the SCSI bus, to make sure that the tape drive gets all the data it needs to write continuously.

Even with a Fast-Ethernet network, with 100Mb/s bandwidth, meaning 12,5MB/s, you can't expect a single flow to saturate a 20MB/s SCSI bus. Moreover, latencies on the network and eventual delays on the clients can cut the throughput even more, resulting in even lower performance.

The idea behind Arkeia's technology is to allow the simultaneous backup of multiple sources, using several interlaced flows of data to make sure the SCSI bus will be saturated as often as possible. Indeed, the total performance of the backup system can be higher if you backup simultaneously three clients instead of one.

The idea of this section is to show you how to use this multi-flow feature to maximize the performance of your backup system.

1.1 Average network performance

On a 100 Mb network, you can expect backup speed anywhere from 300 – 1000 MB per minute. Backup speed will depend on the client, server, network and tape drive(s) or library configurations as well as data compression. 10Mb networks are proportionally slower. Faster networks are proportionally faster.

If the actual performance do not reach this level, it might be necessary to check if your network is setup properly.

The average backup on the client side is 0.83 MB per second, with a 100 Mb per second Ethernet network, it gives 12.5 MB per second of theoretical traffic, meaning 7 to 8 MB per second of real throughput.

2 • Use of multiple flows

2.1 Overview

Arkeia's performance is also based, in part, on the number of simultaneous backup flows that are active at anytime. Each flow represents a client machine or a disk drive of a client machines.

Parallel backup, or multi flow, increases backup speed and reduces the overall time required to backup a group of networked computers by interleaving the data from several clients/disks at the same time.

This allows for optimum network and tape drive usage even when the client machines are on different network loops and have different speed disk drives.

The backup can be configured to use one flow per disk drive in the file servers and 1, 2 or "n" flows for the entire group of desktop machines. This will backup the file servers very quickly and also backup the desktop pool in a reasonable period of time. When there are more clients/disks than flows, Arkeia uses a round robin strategy, which can be modified, to complete the backup. As one client/disk completes its backup the next available client/disk is started.

2.2 Parallelism of a single machine and of multiple machines

By default, Arkeia backs up multiple machines using one flow per machine. The parallelism is automatically configured and performance is optimum.

Also by default, Arkeia uses a sequential procedure to back up the trees selected on a single machine (same flow: 0). To run parallel backups, the trees from a single machine are separated into several groups via number changes made in the "Multiflow" field. In other words, giving each tree a different number will ensure all trees will be backed up simultaneously (provided you have enough available flows). Each number is, roughly speaking, a flow number.

Then trees with the same "Multiflow" number are backed up sequentially.

3 • How to limit backup speed

3.1 Why limit the backup speed?

By default, Arkeia tries to use as much bandwidth as possible to run its backup. However, you may want to save some network bandwidth for others jobs. You may also have some network problems as some parts of your installation may be slower, or have more issues, than others.

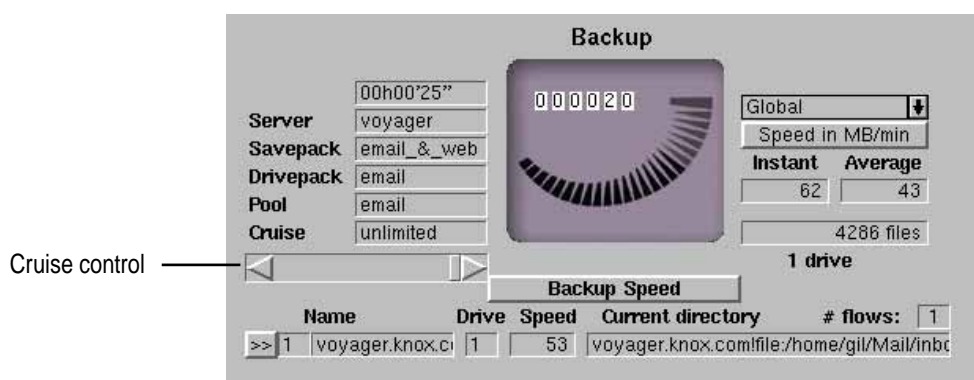
In this case, it can be useful to control the backup speed, thus making sure that your network won't be completely saturated.

3.2 How to limit backup speed

- *The graphical “cruise control”*

Within the Arkeia version 5 graphical interface, you can limit the backup speed by using the cruise control slider.

In the “Backup” window, use the slider to change the backup speed.



- *Default Backup speed limitation*

Usually periodic backups are running during night time. It may not be convenient to manually modify the cruise control. The easiest way to do this is then to set the “SOCK_SPEED” preference.

Add the following preference in the file:

```
/opt/arkeia/arkeiad/arkeia.prf  
  
SOCK_SPEED "[value]"
```

The “[value]” is the number of KB/s per drive. So, if “n” drives are available, this value should be divided by “n”.

Once this value has been defined, all the backups will use this preference.

You can still use the cruise control from the GUI to change the current speed of a specific backup.

For instance:

```
SOCK_SPEED "1024"
```

means a throughput limit of 1024 KB/s, (+/- 60 MB/minute).

4 • Priority

4.1 Overview

Some administrators may want to backup their various trees in a specific order. For instance, they may want to assign a higher priority to the backup of an heavily-used server to make sure it will be backed up during the night, and to let the other machines be backed up later, during the day if necessary.

The “`Priority`” feature of Savepacks is used for this, combined eventually with the “`Chaining`” command.

Priority is set between trees saved by the same flow. Trees backed up by different flows are saved simultaneously.

4.2 Usage

The use of Priority is extremely simple:

In the “`Tree`” Options, for each tree of the Savepack, it is possible to add a Priority field according to needs.

By default, all trees have a Priority of “50”. To set a higher priority, lower the number (the highest priority is “1”). To lower the priority, increase the number (the lowest priority is “100”)

Trees will then be processed from the highest priority down to the lowest.

The value of “`Priority`” is only used when considering trees saved by the same flow. It can only influence backups of various trees of the same machine or for machines backed up by the same flow, which is allowed if backup is restricted to the use of one flow. As this is not an optimum solution, the concept of Chaining comes into play.

5 • Chaining

5.1 Overview

While setting Priority to organize backups on a specific machine can be interesting, it is a lot more important to be able to set an order for the backups of various machines, while still using multiple flows.

This is the use of chaining: it creates a “link” between several trees of various machines. All the trees on the same Chain number will be processed in order of Priority, effectively allowing order control while keeping multiflow.

5.2 Usage

The use of Chaining is very simple.

In the “`Tree options`”, for each tree of the Savepack, it is possible to set a “`chain`” value (“0” by default, meaning no chaining).

Trees with the same value of chaining will be linked and processed according to their Priority.

Chaining has to be used with the “`Priority`” function. These two functions are complementary and are actually close to useless when set alone.

6 • Configuring compression in Arkeia

6.1 Overview

By default, Arkeia always compresses the data it backs up. This is done on the client itself and allows an optimal network throughput. By default, the compression algorithm is LZ1 (Lempel – Ziv Algorithm).

You may want to increase compression, or adjust it according to the type of file being backed up.

6.2 Compression settings

- *Savepacks*

You can set a specific compression method by using the graphical interface and modifying the option in the “Savepack Option” screen or in the “Tree Option” screen.

Go to the “Savepack Option” or “Tree Option” screen and set the desired compression type.

- *Setting a default compression method for a particular client.*

It is possible to set a specific configuration for each tree that will replace the other settings. Arkeia will then apply this configuration instead of the Savepack settings.

This specific configuration is created by adding the following line to the file:

```
/opt/arkeia/arkeiad/arkfs.cfg  
"PREFERRED_COMPRESSION" "[value]"
```

Where “[value]” can be: “0”, “1” or “2”. This code correspond to the following compression:

Code Compression algorithm

Value:	Compression;
0	No compression
1	LZ1, fast compression
2	LZ3, maximum compression ratio (slower)

CHAPTER 12: ARKEIA UTILITIES

- *Summary*

Arkeia version 5 is supplied with a large number of utilities. These utilities allow the user to complete the functions of Arkeia in specific situations. This chapter will detail these utilities and the extra functions they supply.



Please note: At the time of this writing, these utilities exist only for the architecture supporting a server version of Arkeia. If you are unsure if your architecture is a server architecture, refer to the section named “Available clients and servers” on page 7 of this manual.

1 • Extracting data from tape with readarkeia

1.1 Overview

The “`readarkeia`” program is a small utility that can be used to extract files from a tape created by Arkeia, without the need of licenses or even the need to install Arkeia. As Arkeia uses a specific format for its backup, there may be cases when this extracting utility will be your last chance to get a file back, after a crash of your database or in case of errors while writing the tape.

Of course, “`readarkeia`” is a very simple tool: it doesn’t deal with history and rights. It is an emergency tool only.

1.2 Obtaining readarkeia

A version of “`readarkeia`” is installed in the directory: “`/opt/arkeia/bin`”, which is the default Arkeia installation directory.

Alternatively, a small C source code is written at the beginning of each tape used by Arkeia. This is the source code of “`readarkeia`” that can be extracted and compiled with any standard C compiler. You can also find the complete C source code in the following directory:

```
/opt/arkeia/server/readarkeia.c
```

To extract the C code, first rewind the tape by using the “`mt`” command which is applicable to the rewindable device defined in chapter “2: Installation procedure” on page 21 of this manual.

For example, for Linux:

```
mt -f /dev/st0 rewind
```

Then extract the C source using the following “`dd`” command:

```
dd if=/dev/st0 of=readarkeia.c bs=64k
```

You should then edit the file named “`readarkeia.c`” file and delete the first line (it is actually the tape label). Replace the very first space character of the file by a “`/`” character.

Then compile the file “`readarkeia.c`”, using the following command:

```
cc readarkeia.c -o readarkeia
```

1.3 Using readarkeia

The “`readarkeia`” utility is used like a limited “`tar`” command. It uses the same basic syntax, though fewer options are actually implemented:

```
readarkeia [-L<loglevel>] x|t|i [mfv] [device|-] [filename1  
filename2...]
```

- *Options available for readarkeia*

The following options are available:

- “`-l`”: Change the verbosity level. Use it to obtain detailed information on what is really on the tape. Value can be from 10 to 80. Please note: this character is the lower case letter “`L`”.
- “`-x`”: Extract the named files from the tape. If no files are specified, all the files on the tape are extracted.
- “`-t`”: List the table of contents of the tape.
- “`-i`”: Get information about the tape.
- “`-m`”: This option tells “`readarkeia`” that we are in the middle of the tape. (There is no label to read).
- “`-f`”: Use the next argument as the name of the device. If device name is given as “`-`”, “`readarkeia`” reads the tape data from the standard input.
- “`-u`”: This option “updates” the file being restored to the newest version found on a tape. Without using this option, if several versions of the same file are found on the same tape, all versions are restored with a different numerical extension — for instance: “`foo.1`”, “`foo.2`”, “`foo.3`”, etc. With the “`-u`” option, only the newest version of “`foo`” would be restored.
- “`-v`”: Verbose.

1.4 List the content of a given tape

To list the files present on a tape, type:

```
./readarkeia tvf /dev/st0
```

1.5 Extract all the files on a tape

To extract all the files present on a tape, type:

```
./readarkeia xvf /dev/st0
```



Please note: the “`readarkeia`” utility extracts data locally and creates a subdirectory for each host backed up.

Please note also that the “`readarkeia`” utility scans a tape completely. This operation can take some time, especially if the tape to be read has a large capacity.

1.6 Extract a specific file

To extract a specific file available on a tape, type:

```
./readarkeia xvf /dev/st0 [machine_name]/[path]/[file_name]
```



Please note: “`readarkeia`” scans a tape completely, even if the requested file is at the beginning of the tape itself. Therefore, this operation can take some time, especially if the tape has a large capacity.

- *Example:*

If you want to extract a file named: “`/usr/src/kernel/make`” from a machine named “`ariane`”, you have to type (under Linux):

```
./readarkeia xfv /dev/st0 ariane/usr/src/kernel/make
```

2 • Recovering the Arkeia database with arkrstdb

2.1 Overview

On rare occasions, your backup server will crash and the Arkeia database may be corrupted. In order to be able to use your backup history again, you may have to rebuild the index, based on the files and directories that have been backed-up on tape(s). This is the purpose of the “`arkrstdb`” utility.



IMPORTANT NOTE: You need to be aware that this restoration is a long process, since this utility needs to read all the tapes contained in the tape pools to completely rebuild the Arkeia database.

2.2 General usage and syntax

The following is the correct syntax of “`arkrstdb`”. Type “`arkrstdb -usage`” for more information.

```
arkrstdb -d{rewindable device} -t{tape type} -r{drive type}
[-D {database}] [-p(pool)] [-v(voltag)] [-s(silent)] [-l <LogLevel>]
[-j(journal)]
```

2.3 Options

The following are the available options of “`arkrstdb`”:

- “`-d`”: This option specifies the rewindable device of the drive as you entered it in drive management.
- “`-t`”: This option specifies the tape type.
- “`-r`”: This option specifies the drive type.
- “`-D`”: This option specifies the pool name (optional).
- “`-p`”: This option specifies the path to your database (optional).
- “`-v`”: This option specifies the “`voltag`” (tape ID) of your tape (optional).
- “`-s`”: This option runs `arkrstdb` in silent mode (optional).
- “`-l`”: This option specifies the log level. Increase the log level to obtain more information on the operation of “`arkrstdb`”. The possible values of this option are: 10, 20, 30, 40, 50, 60, 70, 80, with “10” being the smallest log level (terse) and “80” being the highest log level (verbose). (optional).
- “`-j`”: This option specifies the log file in which “`arkrstdb`” will save the information on the current operation.

- *Example:*

You plan to restore the content of a DLT tape, from a DLT drive, to the Arkeia database.

To launch this operation, enter the following command:

```
/usr/knox/bin/arkrstdb -d/dev/st0 -t "DLT 4000" -rSTD_DLT
```

3 • Automatic hardware detection with arkdetect

3.1 Overview

The program “`arkdetect`” is a small utility designed to quickly retrieve as much information as possible on the back-up hardware attached to the Arkeia server. This allows the administrator of a given server to quickly select the correct backup hardware attached to his Arkeia server, check that the configuration that has been chosen within the Arkeia GUI is the correct one or check that the hardware attached to the server works properly.

By default, “`arkdetect`” is installed in the “`/opt/arkeia/bin/`” directory, with the rest of the programs of Arkeia version 5.

3.2 Running arkdetect

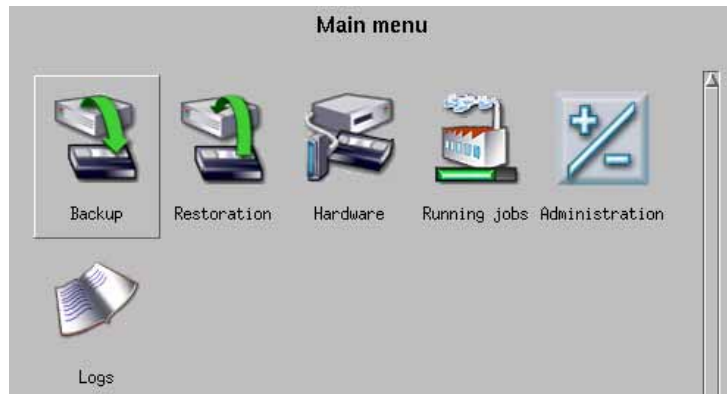
Running the utility “`arkdetect`” is a simple matter of entering the command at any command line.

However, please note the following important information:

- The utility “`arkdetect`” should always be run as “`root`” (the super-user of the Arkeia server).
- The back-up hardware should not be in use (by Arkeia or by any other program) while “`arkdetect`” is running.
- Finally, the “`arkdetect`” utility disregards any parameter entered on the command line.
- *Stopping a library within the Arkeia GUI*

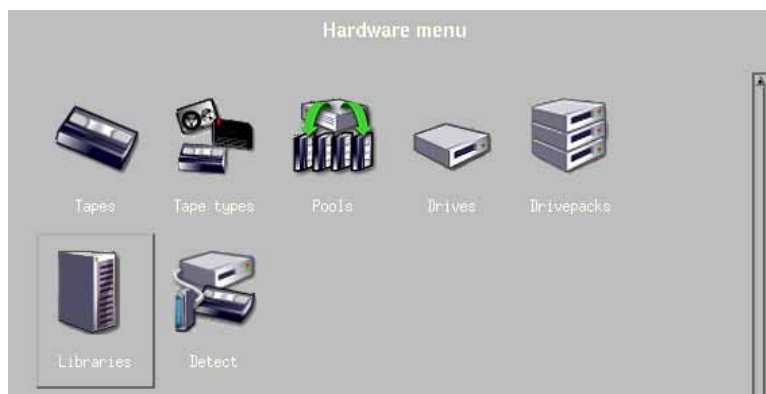
To make sure that Arkeia and its different components cannot use the library before launching “`arkdetect`”, follow the procedure outlined below:

When Arkeia version 5 is launched, its main menu is displayed in its window:



Within this menu, click on the “Hardware” icon, pictured on the left.

This will open the main “Hardware” menu in the window of Arkeia version 5 (see next page).



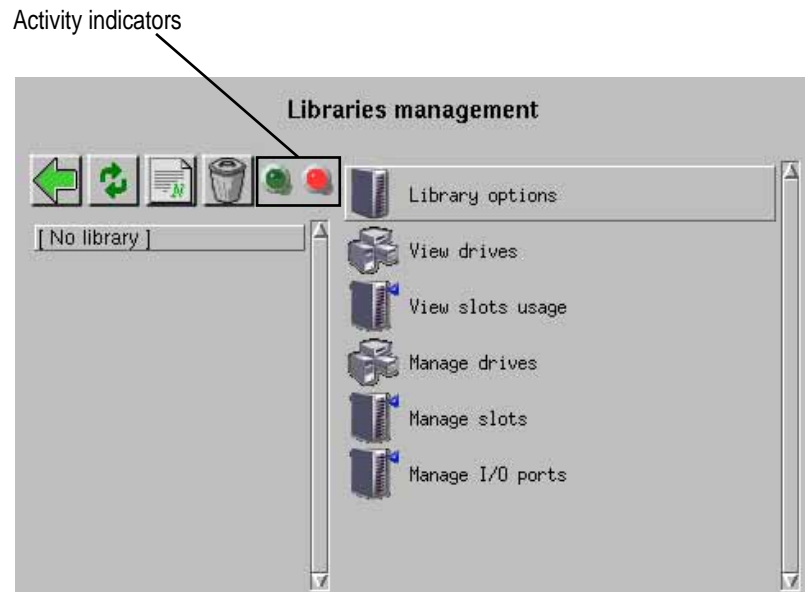
Within the “Hardware” menu, clicking on the “Libraries” icon (pictured on the left) will display the “Libraries management” screen of Arkeia version 5.



Please note: Even though the “Hardware” menu of Arkeia version 5 contains an option named “Detect”, the “`arkdetect`” utility displays much more information than the function integrated in the GUI.

For this reason, if a problem is encountered with the backup hardware used by Arkeia, it is recommended to use “`arkdetect`” rather than the detection functions of the Arkeia GUI.

Within the “`Libraries management`” screen, clicking on the red and green indicators allow you to stop or start the library selected in the left-hand list of libraries.



In this menu, double-clicking on the “`Stop Library`” command will force Arkeia to stop the relevant services and consider the library as inoperative. Arkeia should then return the following message (in this example, the library is named “`BigFileLib`”:

```
■ You have succesfully stopped library BigFileLib
```

Once every library has been stopped in the manner described above, “`arkdetect`” can be run to analyze and detect the hardware attached to the Arkeia server.

3.3 Analyzing the results of `arkdetect`

The best way to analyze the results returned by “`arkdetect`” is to save its messages into a file, by entering the following command at any prompt:

```
./arkdetect > results.txt
```

This will save all the “`arkdetect`” messages into the file named “`results.txt`”. It is then possible to look at the results by opening the file in a text editor or by entering the command:

```
less results.txt
```

- *Empty results: a failed hardware detection*



Please note: For reasons of space, this section will only survey the different messages returned by “`arkdetect`”. However, this summary should be enough to understand most of the output of the utility.

The “`arkdetect`” utility starts its hardware detection procedure by listing the possible backup devices contained in the “`/dev`” directory of the Arkeia server it is installed on. This is indicated by the following lines in the “`arkdetect`” output:

```
Linux Temple Of Hate 2.2.19 #93 Thu Jun 21 01:09:03 PDT 2002

-----
|                               DRIVE AUTODETECTION                               |
-----

arkdetect: Find tape device files...

Entering...
Scanning for tape devices...
File '/dev/st0' is good!
File '/dev/st1' is good!
File '/dev/stderr' is good!
```

The first line of this output displays the operating system used by the Arkeia server where “`arkdetect`” has been launched. It then launches its single drive auto detection procedure.

To do this, the utility first checks the presence of several devices and reports a cheerful message (“`File '/dev/stderr' is good!`”) if it has found the expected devices.

After this potential device detection, “`arkdetect`” will then try to connect to each device and query it to see if there is any backup hardware attached:

```
arkdetect: detect devices by mtio...
Entering for device: /dev/st1
error opening /dev/st1
bad end

arkdetect: detect devices by mtio...
Entering for device: /dev/st0
error opening /dev/st0
bad end
```

In the example above, “`arkdetect`” tried to connect to “`/dev/st1`” and “`/dev/st0`”, only to discover there were no hardware present, which it reports with a couple of error messages: “`error opening /dev/stX`” and “`bad end`”.

After trying to detect the single devices attached to the Arkeia server, “`arkdetect`” then tries to detect libraries, which is shown by the following lines:

```
-----
```

```

| _____ LIBRARY AUTODETECTION _____ |
|-----|

arkdetect: Find library device files...

Entering...
Scanning for library devices...
File '/dev/sga' is good!
File '/dev/sgb' is good!
File '/dev/sgc' is good!
File '/dev/sgd' is good!
[etc...]
```

After checking the existence of the potential devices, and indicating what it has found, “arkdetect” then proceeds to scan the devices themselves and reports its discoveries with the same messages used for the single drive detection:

```

arkdetect: detect devices by inquiry...
Entering
Trying to open device '/dev/sg2's_open('/dev/sg2') ERROR

bad end

arkdetect: detect devices by inquiry...
Entering
Trying to open device '/dev/sg1's_open('/dev/sg1') ERROR

bad end

arkdetect: detect devices by inquiry...
Entering
Trying to open device '/dev/sg0's_open('/dev/sg0') ERROR

bad end
```

As can be seen above, “arkdetect” will simply report an error if no device is present to answer its queries.

Finally, “arkdetect” will print a summary of its detection operation. In this case, of course, the summary will be empty:

```

| _____ DRIVES DETECTED _____ |
|-----|

none...

| _____ LIBRARIES DETECTED _____ |
|-----|

none...

arkdetect: normal end
```

After this somewhat discouraging example, the rest of this section will detail what are the messages returned by “`arkdetect`” if it detects backup hardware connected to the Arkeia server.

- *Results returned if detection has been successful*

The results shown below were obtained on an Arkeia server, to which an “`Autopak`” library has been attached. That library contained two “`VXA`” tape drives.

```
Linux gabriel 2.4.3-20mdk #1 Sun Apr 15 23:03:10 CEST 2002
```

```
-----  
|                               DRIVES DETECTED                               |  
-----
```

```
New elem: -----  
PROP: 'COMMENT           'VAL: 'ECRIX VXA-1'  
PROP: 'ARK_MODEL         'VAL: 'ECRIX_VXA'  
PROP: 'MODEL             'VAL: 'VXA-1'  
PROP: 'VENDOR            'VAL: 'ECRIX'  
PROP: 'DEVICE            'VAL: '/dev/st0'
```

```
New elem: -----  
PROP: 'COMMENT           'VAL: 'ECRIX VXA-1'  
PROP: 'ARK_MODEL         'VAL: 'ECRIX_VXA'  
PROP: 'MODEL             'VAL: 'VXA-1'  
PROP: 'VENDOR            'VAL: 'ECRIX'  
PROP: 'DEVICE            'VAL: '/dev/st1'
```

In the lines above (which contain only the “`arkdetect`” summary report), we can see a detailed report on the two “`VXA`” drives that have been found. The following values are reported by “`arkdetect`”:

- “`COMMENT`”: This line is an informative string, returned by “`arkdetect`”, based on the information returned by the backup hardware when it has been queried (see “`MODEL`”/“`VENDOR`” below).
- “`MODEL`”/“`VENDOR`”: These lines contain the information returned by the backup hardware, when queried by “`arkdetect`”.
- “`ARK_MODEL`”: This line indicates the configuration that should be selected, within the Arkeia GUI, to correctly manage the backup device.
- “`DEVICE`”: This line contains the UNIX device to be used to access the tape drive. For instance, the first “`VXA`” drive in the example shown above is accessed through the “`/dev/st0`” device.

The following information is returned by “`arkdetect`” when it has found library-class devices connected to the Arkeia server:

```
-----
```

```

| LIBRARIES DETECTED |
-----

New elem: -----

PROP: 'COMMENT          'VAL:'Spectra Logic SL 2000'
PROP: 'ARK_SUGGEST     'VAL:'SL2000-15'
PROP: 'ARK_SUGGEST     'VAL:'SL2000-30'
PROP: 'ARK_SUGGEST     'VAL:'AUTOPAK'
PROP: 'MODEL           'VAL:'215'
PROP: 'VENDOR          'VAL:'SPECTRA'
PROP: 'DEVICE          'VAL:'/dev/sga'

arkdetect: normal end

```

Each library is separated from the other by the “New elem[ent]: -----” line.

The fields “COMMENT”, “MODEL”, “VENDOR” and “DEVICE” have all been documented above, but this output introduces a new value field:

- “ARK_SUGGEST”: This field is a suggestion by “arkdetect”, in the case it has not been able to determine precisely the model of library it has detected. If this is the case, “arkdetect” will return the closest models it can find in its internal database.

In the rare case “arkdetect” cannot find any device that matches the information returned by the hardware, it will print “ARK_MODEL : unknown”.

As can be noted above, “arkdetect” was not able to guess the model correctly and returned several suggestions for the library named “Spectra Logic SL 2000”.

3.4 Additional options of “arkdetect”

__??Enter the rest of this section here??__

4 • Tape duplication with arkc

4.1 Overview

The program “`arkc`” is not a “utility” in the strict sense of the word: it is rather a complete interface to the Arkeia server. While most users will only use the GUI interface, “`arkc`” offers advanced user the power and the convenience of a command-line interface, which has many advantages over the GUI.



Please note: This chapter will not detail all the functions, options and commands of “`arkc`”. If a complete reference for this program is required, a separate manual is available.

This manual can be obtained on the Arkeia web site, at the following URL:

`http://www.arkeia.com`

This section of the Arkeia version 5, as its title implies, will only cover one function of the “`arkc`” interface: the tape duplication.

Duplicating backup tapes is a very useful function that allows a system administrator to keep perfectly identical copies of mission-critical tapes in separate locations.

For instance, it is possible to have two (or more) copies of an E-mail system, one which is kept on-site for quick restoration if necessary, and the other which is kept off-site in case a major problem makes the main site inoperative.



Please note: The tape duplication requires the following hardware

- Two tape drives and two tapes of the exact same type.
- Or a library that contains at least two drives and has two tape slots (one for source and one for destination).
- If a library is used, the destination slot must be assigned as “`reserved`” in the “`Slots in library`” screen.
- The destination tape should not have an Arkeia label, unless it is the exact same label as the source tape.

Attempting to copy a tape of a certain type (for instance DAT) to a tape of a different type (for instance DLT) will only result in an error message sent back by “`arkc`”.

4.2 Overview of the arkc syntax

The standard “arkc” syntax is as follows:

```
arkc -duplication -start -D <parameters...>
```

The “-D” parameter indicates to “arkc” that its arguments and parameters are entered on the command line.

The following parameters must be entered on the command line, when duplicating a tape from one drive to another:

- “src_tape_name” or “src_tapeid”:

These parameters are required by “arkc”.

These parameters should contain either the name or the ID of the source tape. For instance, the following names and IDs are considered valid, and will be accepted by arkc for a tape duplication command:

```
src_tape_name=DLT004  
src_tapeid=3bc1d2e2
```

As can be seen above, the “src_tape_name” may be easier to remember, but both these values are valid, and accepted, by arkc. To display a list of available source tapes, enter the command: “arkc -tape -list” (without the quotes). The tape name is assigned when a tape is created by Arkeia.

- “src_drv_name” or “src_drvid”:

These parameters are required by “arkc”.

These parameters should contain either the name or the ID of the source drive. For instance, the following names and IDs are considered valid, and will be accepted by arkc for a tape duplication command:

```
src_drv_name=RD_LTO  
src_drvid=3a530bcf
```

As can be seen above, the “src_drv_name” may be easier to remember, but both these values are valid, and accepted, by arkc. To display a list of available drives, enter the command: “arkc -drive -list” (without the quotes).

- “dst_drv_name” or “dst_drvid”:

These parameters are required by “arkc”.

These parameters should contain either the name or the ID of the destination drive. This drive must be compatible with the source drive! For instance, the following names and IDs are considered valid, and will be accepted by arkc for a tape duplication command:

```
dst_drv_name=RD_LTO  
dst_drvid=3a530bcf
```

As can be seen above, the “`dst_drv_name`” may be easier to remember, but both these values are valid, and accepted, by `arkc`. To display a list of available drives, enter the command: “`arkc -drive -list`” (without the quotes).

Additionally, the following parameters can be entered, when duplicating a tape that is contained in a tape library:

- “`dst_lib_name`” or “`dst_libid`”:

These parameters are required when using “`arkc`”, if duplicating in a library.

These parameters should contain either the name or the ID of the destination library. This library must contain the destination drive defined above! For instance, the following names and IDs are considered valid, and will be accepted by `arkc` for a tape duplication command:

```
dst_lib_name=RD_LT0
dst_libid=3a530bcf
```

As can be seen above, the “`dst_lib_name`” may be easier to remember, but both these values are valid, and accepted, by `arkc`. To display a list of available libraries, enter the command: “`arkc -library -list`” (without the quotes). Please note that defining “`dst_lib_name`” or “`dst_libid`” means that “`dst_slot`” (see below) must be defined as well.

- “`dst_slot`”:

This parameter is required when using “`arkc`”, but only if duplicating in a library.

This parameter should contain the number of the slot that contains the destination tape. This parameter must be defined if duplicating to a library! For instance, the following is considered as valid value, and will be accepted by `arkc` for a tape duplication command:

```
dst_slot=6
```

This slot must contain the tape that will be used as a target by the duplication operation! Please note that, when duplicating from/to a library, it is not necessary to define a source slot for the source tape, as Arkeia is able to retrieve the correct tape based on its ID and/or name. Also, please note that defining “`dst_slot`” means that “`dst_lib_name`” or “`dst_libid`” (see above) must be defined as well.

5 • Testing the Arkeia server connectivity with chknlp

5.1 Overview

Networking issues are the most frequently encountered problems when first configuring Arkeia. Most of these networking issues come from a badly configured network, failing DNS or misspelled machines names. Those issues are not always trivial to solve.

For this class of issues, Arkeia provides “`chknlp`”, a utility to check the correct connectivity between the backup server and its clients. It tests all connections, the presence of all needed services and all of Arkeia’s parameters. The “`chknlp`” utility can therefore help you find the weak spots of your network.

5.2 Usage

The “`chknlp`” program is a binary provided with client and server to check network connectivity between client machine and backup server. You must have “`root`” (Super User) access to use this command.

To start “`chknlp`”, just type the following command:

```
/opt/arkeia/bin/chknlp -pause
```

The “`chknlp`” utility then runs four or five tests, the last test being run only on the backup server.

5.3 chknlp test description

- *chknlp Test 0*

The “`chknlp`” utility first tries to locate the Arkeia directory. The directory itself is stored in the system variable named “`$ARKEIA_DIR`”.

If this test fails, an installation problem has occurred: try to reinstall the software.

An example of a correct output is shown below:

```
>>> chknlp 4.3.20 <<<
>>> TEST FOR CERTIFICATION LEVEL ONE <<<
*** TEST 0 IN PROGRESS...
>>> Getting ARKEIA variables:
  ARKEIA_DIR = /opt/arkeia
  NLPDIR = /opt/arkeia/arkeiad
  DIRAPPLIC = /opt/arkeia/chknlp
  PATHAPPLIC = /home/gil/.chknlp:/opt/arkeia/chknlp/chknlpctx:/opt/
arkeia/chknlp:/opt/arkeia/gui:/opt/arkeia/arkeiad
*** TEST 0 OK
```

- *chknlp Test 1*

During Test 1, the “chknlp” utility will attempt to retrieve network preferences like:

- The port number used for communications (`PORT_NUMBER`). The port number is not displayed.
- The backup server name (`ADMINSERVER`)
- The Internet address of the backup server (`Internet address`)
- The local host name (`HOSTNAME`)
- The local address (`Internet address for local HOST`)

If the “chknlp” utility cannot reach the “ADMINSERVER” during this test, check that the file “`/opt/arkeia/arkeiad/admin.cfg`” contains the name of the backup server.

If the “chknlp” utility cannot obtain the Internet address, a network configuration problem has occurred. You should add the backup server in your host file. If you are using a Domain Name Server, check that the backup server is listed and check that the local machine access to the DNS. To be sure, “ping” the backup server with its host name

If “chknlp” cannot reach your local host name or local address, check the network configuration.

An example of a correct output is shown below:

```
*** TEST 1 IN PROGRESS...
>>> Testing variables in configuration file
   Internet address for Licence server `voyager' = 192.168.7.139
   Internet address for local HOST `voyager.knox.com' = 192.168.7.139
*** TEST 1 OK
```

- *chknlp: Test 2*

This test checks that the local daemon or service is running correctly.

If this test fails, the local daemon or service is not running correctly, or does not accept connection to port 617. You should then restart the daemon with the “arkboot” command (see below) or restart the “arkeiad” daemon.

An example of a correct output is shown below:

```
*** TEST 2 IN PROGRESS...
>>> Testing local (voyager.knox.com) ARKEIAD daemon
      Message from arkeiad daemon:      Test OK
*** TEST 2 OK
```

- *chknlp: Test 3*

This test checks the network addresses of the client and the backup server as seen by each other and compare both.

If this test fails, it is quite common to configure the “ADMINSERVER” as the backup server, although this can be configured differently.

To do this configuration, the host name of the “ADMINSERVER” on the local machine (in the “admin.cfg” file) must be the name of your Arkeia “ADMINSERVER”. This machine is typically the backup server.

If the “ADMINSERVER” IP address on the local machine is different than the “ADMINSERVER”’s true IP address, check the IP address of the “ADMINSERVER” on both machines. If you ping the “ADMINSERVER” when logged onto the local machine you should get the same IP address as you get when pinging the “ADMINSERVER” while logged onto the ADMINSERVER.

If the local host name is different on the “ADMINSERVER”, check the IP addresses of the local host on both machine.

If the local host address is different on “ADMINSERVER”, check the IP addresses on both machines.

If you got the message: “name & address of local host are not on the licensed server list”, the local host has not successfully been declared to the “ADMINSERVER”.

An example of a correct output is shown below:

```
*** TEST 3 IN PROGRESS...
>>> Connection test between LOCALHOST and ADMINSERVER
>>> Accessing to ADMINSERVER services
Remote host 'jupiter.knox.com' is Admin Server
Admin Server name on Local Host is jupiter.knox.com
on Admin Server is jupiter.knox.com
Admin Server address on Local Host is 192.168.8.22
on Admin Server is 192.168.8.22
Local host name on Local Host is jupiter
on Admin Server is jupiter
Local host address on Local Host is 192.168.8.22
on Admin Server is 192.168.8.22
>>> Name & Address of Local Host are already in License Server List
*** TEST 3 OK
```

- *chknlp: Test 4*

This test only runs on the server side. It checks the network addresses of all the clients listed in the “/opt/arkeia/arkeiad/rhosts.lst” file, and of the backup server as seen by each other.

If this test fails, the parameters to be checked are basically the same as those in Test 3, (see above) though you should have a complete set of information about your network in this way.

As “chknlp” stops at each error, it should be executed again after resolving an issue, to make sure there are no other problems with another machine.

An example of a correct output is shown below:

```
>>> TEST FOR CERTIFICATION LEVEL TWO <<<
*** TEST 4 IN PROGRESS...
>>> Testing connection between the local host and all known clients
>>> Connecting to: Host = c4 Internet address = 89.0.0.5
Connection to host is OK
>>> Connecting to: Host = c6 Internet address = 89.0.0.8
Connection to host is OK
>>> Connecting to: Host = bindernt2 Internet address = 89.0.0.20
Connection to host is OK
>>> Connecting to: Host = c1 Internet address = 89.0.0.1
>>> Name & Address of Local Host are already in License Server List
*** TEST 4 OK
```

6 • Renaming servers/clients with arkmvhost

6.1 Overview

Arkeia version 5, as this manual has described, is able to manage a large number of machines, both clients and servers.

However, it is possible that, within a network, some machines will change names. This usually creates configuration problems ranging from the easy-to-solve to the devious, particularly if additional information has been lost along the way.

Arkeia version 5 manages this issue by assigning a unique machine ID to each computer it is installed on, once this machine attempts to declare itself to the Arkeia server.

On an Arkeia client, the update is then automatic: if the machine name is changed, for instance from “foo.acme.net” to “bar.acme.net”, this change is automatically taken into account and the new name is communicated accordingly to the Arkeia server indicated in the file “/opt/arkeia/arkeiad/admin.cfg”.

However, the server database then needs to be updated so that the server Savepacks (and other items) can be automatically updated. This is done through the “arkmvhost” utility.

6.2 Syntax of arkmvhost

The following syntax is accepted by “arkmvhost”:

```
arkmvhost --usage
arkmvhost --list
arkmvhost --allhosts [--no_confirm|--keep_info] [options]
arkmvhost --host <old_name> --newhost <new_name> [options]
```

To use “arkmvhost” to update the server database, you must be logged on the server as “root” (Super User). To start the update process, it is necessary to enter the following command:

```
#!/opt/arkeia/bin/arkmvhost
```

To have a brief reminder of the options detailed below, the following command can be entered:

```
arkmvhost --usage
```

6.3 Listing all the (modified) clients in the Arkeia database

To list all the clients that are present in the database of an Arkeia version 5, that have a modified name, use the following option:

```
arkmvhost --list
```

This options lists all the clients which have been detected as modified and exits immediately. This is useful to force the Arkeia server to display which clients have a changed name and, therefore, may require an update.

6.4 Complete update of the Arkeia database with arkmvhost

To do a complete update of the Arkeia database, use the following “arkmvhost” syntax:

```
arkmvhost --allhosts [--no_confirm|--keep_info]
```

This mode updates the entire Arkeia v5 client database. Only the machines detected by the server as having a modified name are going to be updated. To obtain a list of these machines, see the option “--list”, detailed above. All the machines which do not have a changed name are left intact in the database.

In this mode of operation, the following options are accepted:

- “--no_confirm”: Update all the clients of the Arkeia v5 server without confirmation.
- “--keep_info”: A copy of the client-related information is kept intact in the database. This allows a modification to be performed in several separate stages (for instance, modifying only the backup index, then the Savepack, etc).

For instance:

```
arkmvhost --allhosts --no_confirm
```

The command above would update the entire Arkeia v5 database, without asking the user to confirm every modification.

6.5 Renaming a single client in the Arkeia database

It is possible, instead of updating the entire Arkeia database, to update one client. To do this, use the following syntax of “arkmvhost”:

```
arkmvhost --host <old name> --newhost <new name>
```

This mode allows the system administrator to specify which client should be renamed in the Arkeia v5 database.

For instance:

```
arkmvhost --host foo.acme.com --newhost bar.acme.com
```

The command above would rename the client named “foo.acme.com” into “bar.acme.com” in the Arkeia v5 database.

6.6 Common options of arkmvhost

The following options can be used in all modes of operation (except the “`--list`” mode detailed above):

- “`--index`”: This option updates the backup index.
- “`--savepack`”: This option updates the list of savepacks.
- “`--device`”: This option updates the list of devices (libraries, drives).
- “`--user`”: This option updates the list of users.

If the options above are not specified on the command line, the default behavior of “`arkmvhost`” is to update all the items listed above.



Please note: When running the utility “`arkmvhost`”, the only Arkeia process running should be the “`arkeiad`” process, which will update the Arkeia database with the new name.

To list all the running Arkeia processes, use the “`ps`” or the “`top`” commands as “`root`”. To stop and restart Arkeia processes, please see the “`arkboot`” utility in this chapter.

7 • Removing a host with arkrmhost

7.1 Overview

It may, sometimes, be useful to quickly delete a client (or “host”) from the Arkeia backup database. This functionality is provided by the “`arkrmhost`” utility, which is part of the standard Arkeia version 5.0.12 package.

If “`arkrmhost`” is not available in your installation or Arkeia, please upgrade to version 5.0.12 or higher. For more information about the upgrade possibilities, please contact Arkeia through its web site, at the following address:

`http://www.arkeia.com`



Please note: Use the utility “`arkrmhost`” with extreme caution! If a client has been deleted from the Arkeia backup database, this has two very serious consequences:

- The client will be unreachable from the Arkeia server and will not appear in the list of possible clients within the Arkeia Navigator.
- Since all reference to the client are deleted from the Arkeia database, this prevent Arkeia from restoring the client data from the tapes! The only solution to restore previously saved data is to use the utility “`arkrstdb`”, also described in this chapter – but this may create additional problems if the client name has been changed before its final deletion.

7.2 Syntax of “arkrmhost”

The syntax of “`arkrmhost`” is quite simple:

```
arkrmhost <host name>
```

Where “`host name`” is the name of the host to be deleted from the Arkeia backup database.

For instance, the following command will delete the Arkeia client named “`fubar.acme.com`” from the Arkeia database:

```
arkrmhost fubar.acme.com
```



Please note: as indicated above, “`arkrmhost`” must be able to communicate with “`arkeiad`”. Therefore, “`arkeiad`” must be running in order to carry out the deletion of a client machine with “`arkrmhost`”.

8 • Checking the database integrity with arkdbchk

8.1 Overview

It is possible that the Arkeia database index may become damaged in the event of a system crash. The Arkeia index database is a list of all the operations performed by Arkeia. This database is used by Arkeia to rapidly retrieve the information on a specific backup operation.

The database integrity checks are performed by the “`arkdbchk`” utility.



Please note: The usage of “`arkdbchk`” may have serious consequences on the database index file. Use of this utility should only be attempted after thoroughly reading this documentation and consultation with the Arkeia Technical Support.

8.2 Usage and syntax

To run the “`arkdbchk`” utility, the following command should be entered on a command, in “`root`” (Super User) mode:

```
#!/opt/arkeia/bin/arkdbchk
```

The standard syntax for the “`arkdbchk`” is the following one:

```
arkdbchk [-p|r|m|b|D<database>] [-s] [-l<log level>] [-j<file>]
```

The options indicated above can be entered on a command line to modify the “`arkdbchk`” execution:

- “`-m`”: This option forces “`arkdbchk`” to check the master Arkeia index. It also lists its contents. The master database index is usually “`/opt/arkeia/server/database/f3bkp/bkmaster`”.
- “`-r`”: This option forces “`arkdbchk`” to remove corrupted entries (also called “`o3_cpnt`”) that can be found within the Arkeia database index. This operation will usually recurse the directories or the Arkeia database and/or Arkeia index. The index can be specified with the use of the “`-D`” option (see below).

- “-b”: This option forces “arkdbchk” to rebuild the Arkeia database index, based on the information that can be found in the Arkeia database. Unless the “-D” option is used to specify a path, the master index is used.
- “-p”: This option forces “arkdbchk” will purge the database index of the unnecessary information it contains. This option will reduce the number of inodes used by the index itself. A path can be specified with the “-D” option.
- “-D<database>”: This option indicates which part of the Arkeia database index will be the target of “arkdbchk”. Its value, which should be a path pointing to the desired part of the Arkeia database index, will be used by the three options detailed above (“-r”, “-b” and “-p”).
- “-s”: This option forces “arkdbchk” to be executed in “silent” mode (no messages).
- “-l<log level>”: This option accepts value ranging from “10” to “90”. It determines the number of messages returned by “arkdbchk” during its execution, with “10” being the tersest level (few messages) and “90” being the most verbose. The default value is “10”.
- “-j<journal>”: This option will force “arkdbchk” to store its messages in a separate log/journal file.

8.3 Examples

Displaying the content of the Arkeia database index, with a high level of verbosity:

```

/opt/arkeia/bin# ./arkdbchk -m -l90

I1132150 main [L114]: ***** Start of database checking *****
D2702040 df3_init: entering
V2709010 db3_gendirs: Entering (/opt/arkeia/server/dbase/f3tmp,-1)
(2: No such file or directory)
V2709060 db3_gendirs: Normal end (/opt/arkeia/server/dbase/f3tmp)
D2702040 df3_set_dbtree: entering
V2709010 db3_gendirs: Entering (/opt/arkeia/server/dbase/o3dbtree,-1)
V2709060 db3_gendirs: Normal end (/opt/arkeia/server/dbase/o3dbtree)
D2702050 df3_set_dbtree: return "/opt/arkeia/server/dbase/o3dbtree"
D2702050 df3_init: return "/opt/arkeia/server/dbase"
D2804010 dc3_load_bkm: Entering, mode=00
V2801010 dc3_open: Entering, name=/opt/arkeia/server/dbase/f3bkp/bkmaster, dc3_mode=1
V2801010 dc3_open: Continuing with name=/opt/arkeia/server/dbase/f3bkp/bkmaster
V2801020 dc3_open: Normal end
V2804020 dc3_load_bkm: reading `2' records
D2804030 dc3_load_bkm: Normal end
D2804160 dc3_list_bkp: Normal end
V2301180 display_list: Entering
I2301190 display_list: ----- Content of list ': Bkmaster
-----
V2301120 display_local_props: Entering
V2301140 display_elem: Entering
I2301150 display_elem: ITEM {
I2301160 display_elem: PROPERTY 'DATE' ALFA '1019221687'
I2301160 display_elem: PROPERTY 'BKSID' ALFA '3cc016b7'
I2301170 display_elem: }
D1132180 main: normal end

```

```
D1132130 myexit: entering, exit code is `0`
I1132140 myexit: ***** End of database checking *****
```

As can be noted in the example above, the Arkeia index only contains two items (indicated above in the “{ITEM... }” lines, that are related to a single backup performed with Arkeia version 5. Please note as well that there is no space between the “-l” option and its value.

Removing corrupted entries in the Arkeia database, with a low verbosity:

```
# ./arkdbchk -r -l20
I1132150 main [L114]: ***** Start of database checking *****
I1132050 check_o3cpnt: Component file of "/opt/arkeia/server/dbase/
o3dbtree/omega.vertex.net/file/home/jack/Mail/trash" is empty
I1132050 check_o3cpnt: Component file of "/opt/arkeia/server/dbase/
o3dbtree/omega.vertex.net/file/home/jack/Mail/draft" is empty
I1132050 check_o3cpnt: Component file of "/opt/arkeia/server/dbase/
o3dbtree/omega.vertex.net/file/home/jack/Mail/queue" is empty
I1132050 check_o3cpnt: Component file of "/opt/arkeia/server/dbase/
o3dbtree/omega.vertex.net/file/home/jack/Mail/outbox" is empty
I1132050 check_o3cpnt: Component file of "/opt/arkeia/server/dbase/
o3dbtree/omega.vertex.net/file/home/jack/Mail/inbox/Ref" is empty
I1132050 check_o3cpnt: Component file of "/opt/arkeia/server/dbase/
o3dbtree/omega" is empty
I1132140 myexit: ***** End of database checking *****
```

Please note that, in the example given above, the database did not contain any corrupted entries.

9 • Stopping and starting Arkeia with arkboot

9.1 Overview

Like most advanced software systems, Arkeia version 5 has to perform many operations before it can be safely launched or stopped. Since Arkeia version 5 also has a modular architecture, this means that several Arkeia modules can be stopped and started in a completely independent manner.

These operations are controlled by the “arkboot” utility, which gives the system administrator the capacity to control precisely which Arkeia services are launched and/or executed at any given time.

Please note that the “arkboot” utility can only be launched by the “root” (the super user) of the Arkeia server.

9.2 Usage, configuration and syntax

The utility “arkboot” is configured through the use of “arkboot.cf”, an XML file, which is contained in the directory “/opt/arkeia/arkboot”. The name and the location of this file can be modified by editing the relevant entry (the variable named “ARKBOOTFILE”) in the file “arkeia.prf”.

The content of this file are as follows:

```
<?xml?>

<arkml:boottable>

  <arkml:bootentry id="617" owner="arkeiad" startprio="-1000"
stopprio="1000" startcmd="/opt/arkeia/bin/arkeiad" stopcmd="/opt/
arkeia/bin/arkeiad stop"/>

  <arkml:bootentry startcmd="/opt/arkeia/bin/arkvlib -i 3cbef720"
id="3cbef720" owner="arkvlib" startprio="50" stopprio="50"
stopcmd="kill 5864"/>

</arkml:boottable>
```

As can be seen in the file above, every program that can be launched has its own line in the configuration file, enclosed within the XML tag: “<arkml:bootentry... />”. Each program is then assigned a unique ID and a name, as well as a “start” command to be executed.

As can be noted, the file above contains only two entries, one for “arkeiad”, which is the main Arkeia UNIX daemon and one for “arkvlib”, which is the Arkeia Library Manager daemon. These two daemons are able to control most of Arkeia operations.



Please note: Based on the configuration of the Arkeia server, there can be several more entries in the “`arkboot.cf`” file. For more information on these entries, please refer to the documentations of the different Arkeia modules.

The standard syntax for “`arkboot`” is straightforward:

```
# arkboot [-start|-stop]
```

The options “`-start`” and “`-stop`” are self-explanatory, and, therefore, won’t be detailed in this manual.

For instance, the following command will request “`arkboot`” to stop the entire Arkeia server (once all pending operations have been completed):

```
#arkboot -stop
```

CHAPTER 13: DISASTER RECOVERY

- *Summary*

One of the most important features of backup software is the ability to quickly recover from a disastrous event. In the 24/7 connected world we live in, even a few hours of interrupted service can translate into huge losses for companies. Arkeia version 5 provides the ideal way to quickly recover both essential servers and mission-critical data: a complete, highly-automated Disaster Recovery procedure.

This chapter will detail this procedure and the Recovery functions, as well as the best way to prepare for emergencies.



Please note: Using the Disaster Recovery procedure for an Arkeia server machine may require an additional license to be obtained from Arkeia Corporation.

Please make sure that you have the correct license, by checking with Arkeia Corp., before attempting to start a Disaster Recovery on an Arkeia server or client. Saving the information required by the Disaster Recovery procedure, on the other hand, does not require a special license.

IMPORTANT: Since the Arkeia Disaster Recovery procedure is a complex process, that requires many steps to be performed in a very precise order, it is highly recommended that system administrators who want to use this feature fully read and understand this chapter. Attempting to use the Disaster Recovery feature of Arkeia version 5 without reading this chapter will almost certainly make a bad situation even worse!

Also important is to test the Disaster Recovery procedure to make sure it supports the hardware that will be restored, especially the drivers that are present on the Arkeia Disaster Recovery CD.

1 • Introduction to Disaster Recovery

1.1 Overview

The Disaster Recovery procedure of Arkeia version 5 offers advanced system recovery and data restoration functions. It allows system administrators to quickly recover from critical storage failure or common hardware problems, but it also allows systems to be quickly put back online in the case of more important events, such as fire or floods, as long as backup tapes are available.



VERY IMPORTANT. Please note: As mentioned above, the Disaster Recovery procedure is entirely based on the availability of backup tapes/media.

If the backup media is damaged or missing, there is, obviously, nothing that Arkeia version 5 (or any other backup software) can do. For this reason, Arkeia Corp. **highly recommends** to make copies of all backup media.

Arkeia version 5 provides a simple utility to quickly duplicate backup media in its standard distribution. For more information on this utility, refer to the section titled “Tape duplication with arkc” on page 191 of the chapter “12: Arkeia Utilities” of this manual.

Please note: At the time of this writing, Disaster Recovery works correctly with Linux clients and servers only.

The Disaster Recovery itself is a three-stage process:

- First of all, some crucial system information must be backed-up by Arkeia. This information will allow the Disaster Recovery procedure rebuild the network and the disk configuration of the damaged computer.
- Second, a complete backup should be completed by Arkeia, in order to have a “starting point” for the recovery procedure.
- Finally, it is crucial to set up a complete backup procedure. Based on the importance of a given machine, the size of its mission-critical data, and the rate at which this data is updated, it may be necessary to create a backup schedule, using the Periodic backup function of Arkeia version 5. For more information on this function, refer to chapter “8: Periodic backup” of this manual.

Based on these three stages, Arkeia version 5 is able to restore any Linux machine back to the state it was before disaster struck. The rest of this chapter will present the different steps listed above.

1.2 Disaster Recovery issues

The Disaster Recovery process attempts to restore a machine in the exact state it was right before it was damaged. To do this, it has to take into account the following issues, whether the machine is an Arkeia client or an Arkeia server.

- *Hard disk drive parameters recovery*

To restore a Linux machine, it is necessary, first of all, to restore the parameters of its hard disk drive (HDD). Of course, to be able to restore these disk parameters, they need to be backed up before the disaster in the first place! These parameters are the followings:

- **Size:** The overall size of the hard disk(s) that are installed within a Linux machine. For instance, a personal workstation may have a single 20 GB IDE hard disk, while a Linux server may have several 60 to 80 GB SCSI hard disks (or more!).
- **Partition information:** Linux divides all the hard disks present in a machine into several partitions. These partitions can then be used to separate the file system into different, independent parts, such as “/home” or “/usr”.
- **File system:** Linux offers a plethora of journaled, and non-journaled, file systems to choose from. While this allows system administrators to choose a given file system, based on the needs of the server, or the requirements of a given application, it also forces the Disaster Recovery procedure to take several additional parameters into account.



Please note: At the time of this writing, the supported file systems are the following:

- Linux “**ext2fs**”
- Linux “**ext3fs**” (Experimental support only, due to kernel limitations)
- Linux “**reiserfs**”
- SGI “**xf s**”
- IBM journaled file system “**jfs**”
- MS-DOS “**fat**” and “**vf at**” (or: “**fat32**”)
- ISO9660 file systems (including the “**Joliet**” extensions).
- Minix file system

This chapter will cover the methods that should be used to save these crucial parameters, from a client machine to an Arkeia server.

- *Network parameters recovery*

Another set of parameters that are absolutely essential for the Disaster Recovery procedure are the network parameters. These parameters include:

- **Network card:** Since Arkeia is a client/server software, the Disaster Recovery procedure of a “client” machine will attempt to connect to an Arkeia server. In order to do this, a network card (and network access) is, of course, mandatory. The Disaster Recovery procedure includes advanced auto-detection routines in order to discover and configure automatically the network card of the client machine, but still asks the system administrator to confirm its results.
- **TCP/IP parameters:** these parameters, which are required for a complete TCP/IP connection to the Arkeia server, include: the name of the machine (for instance: “`sna fu . usa . acme . com`”), its IP address, its netmask, gateway and DNS server.
- **Arkeia server parameters:** these parameters, that are required to properly access the Arkeia backup server and restore the client data, are the followings: the Arkeia server name (for instance: “`arkeia . usa . acme . com`”), the TCP/IP address of the Arkeia server, the port that it uses (and the port that should be used on the client), as well as the login/password combination that should be used to access the server.

This chapter will cover the methods that should be used to save these crucial parameters, from a client machine to an Arkeia server.

- *Server database/index recovery*

If the machine that suffered a crash and/or major failure is an Arkeia server, the parameters that should be restored are the ones indicated above (Hard disk drive parameters and network parameters), but should be completed by a complete restoration of the server data contained in the Arkeia index and database.

Whenever an Arkeia server performs an operation (backup or restore), it stores information about this operation in its database and in its index. While the Arkeia database is global, the index is divided by machine. Therefore, if an Arkeia server has three clients: “`sna fu`”, “`foobar`” and “`rtfm`”, its index will contain three branches, one for each client. If an Arkeia server is used to backup itself, its index will also contain a branch with its own name. To recover an Arkeia server, it is therefore necessary to restore these important files.

Since an Arkeia server can backup itself, recovering its index and its database is the only way to restore Arkeia to its normal (meaning, before failure) state.

However, since the database and the index can become quite large on an Arkeia server providing backup and restoration services to many client, it is strongly suggested to perform the Disaster Recovery of an Arkeia server in three steps: (1) recovery of the index and database information related only to the server itself, which allows (2) the restoration of the entire server and (3) the recovery of the clients-related information contained in the Arkeia database and index. This procedure is detailed in the rest of this section.

Once these three steps have been performed, the Arkeia server has been completely recovered and it can then proceed, if needed, with the Disaster Recovery procedures of its different clients.

2 • Preparing for a Disaster Recovery

As was noted above, the Disaster Recovery feature of Arkeia version 5 cannot work in a “vacuum”, but requires precise information to be saved. This section contains a detailed list of tasks to be completed, and system administrators who want the additional security of Arkeia version 5 Disaster Recovery should complete all these steps.

2.1 Creating a specific Disaster Recovery configuration

Depending on the importance of a computer and the data it contains, when a Disaster Recovery procedure becomes necessary to restore the system, it may well be an emergency operation.

Mission-critical computers, and the data they contain, can cost a company thousands of dollars for every hour of downtime. Therefore, if a Disaster Recovery procedure is put into place at any site, it is worthwhile to prepare a specific Disaster Recovery configuration, in order to optimize the time needed to recover a machine.

System administrators are therefore strongly advised to follow the suggestions below, in order to recover their systems as quickly as possible:

- *Dedicated tapes and tape pools*

When the Disaster Recovery will be applied to an Arkeia server, it is strongly suggested to create dedicated “Disaster Recovery” tapes and tape pools, which are used to backup the information essential to a Disaster Recovery procedure. Such tapes should be kept in a safe place, either off-site – such as the safe in a bank – or on-site, for instance, in a fire-proof safe, away from the main computer/server room.

It is possible to duplicate the Disaster Recovery tapes, with the help of the tape duplication utility of Arkeia version 5. For more information on this utility, refer to the section named “Tape duplication with arkc” on page 191 of this manual.



Please note: It is highly recommended to properly label the tapes used for the Disaster Recovery procedure.

Since Arkeia version 5, during the Disaster Recovery itself, will ask the operator for the different tapes by name, labeling the tapes can therefore substantially reduce the time necessary for the operation.

Proper handling and storage of the Disaster Recovery tapes is also strongly recommended. For more information on this subject, refer to the documentation supplied by the maker of the tape/backup media.

For more information on the creation of tapes and tape pools, refer to chapter “6: Objects and users creation” on page 71 of this manual.

- *Creating a recovery diskette for a specific system*

The Disaster Recovery procedure of Arkeia version 5 also allows a “Recovery diskette” to be created for a specific system. This diskette, once created, contains the configuration information needed to quickly recover the network parameters of the “reference” computer (the undamaged system).

If a Recovery diskette has been created, for instance for a system named “`snafu.acme.com`”, the Disaster Recovery procedure will be able to use it to shorten the lengthy question/answer session that starts the entire procedure.

To create a recovery diskette for a given system, follow the steps detailed below:

- Use the instructions below (in the section titled “How to backup a client/server parameters”) to perform a complete backup of the system configuration (with the “`sys-info:/config`” plug-in described below).
- Once the backup is complete, restore the system configuration, contained in “`sys-info:/config`”, to disk, by using the redirection function of Arkeia version 5. For more information on the redirection of a restoration, refer to the section named “Redirecting a Restoration” on page 151 of this manual.
- Copy the file created by the restoration to a diskette. The Disaster Recovery procedure is now able to use this diskette to automatically offer correct default values to the questions asked in the first step of the Disaster Recovery procedure.



Please note: The diskette used in the Disaster Recovery procedure should be formatted either with the MS-DOS “`fat`” or with the Linux “`ext2`” file systems. For more information on the formatting commands, refer to the documentation and to the “`man`” pages of your system.

Since the Recovery file stored on the diskette is quite small, it is recommended to store several of these files on a single diskette, by giving different names to each file. For instance, if a Disaster Recovery diskette should be created for the three following systems:

- Machine 1: “`snafu.usa.acme.com`”
- Machine 2: “`fubar.usa.acme.com`”
- Machine 3: “`rtfm.usa.acme.com`”

It is possible to create three separate files, named: “`snafu.cfg`”, “`fubar.cfg`” and “`rtfm.cfg`” that can then be stored on the same diskette. When the Disaster Recovery procedure is launched, it will ask first for the Recovery diskette, then, if several files are found, it will ask which one should be used for the current target machine.

2.2 How to backup a client/server parameters

Making a backup of the crucial system parameters is essential to the Disaster Recovery procedure. The system parameters are divided, as was detailed above, into three different classes:

- The network configuration information.
- The hard disk drive parameters.
- The Arkeia index and database (for servers only).

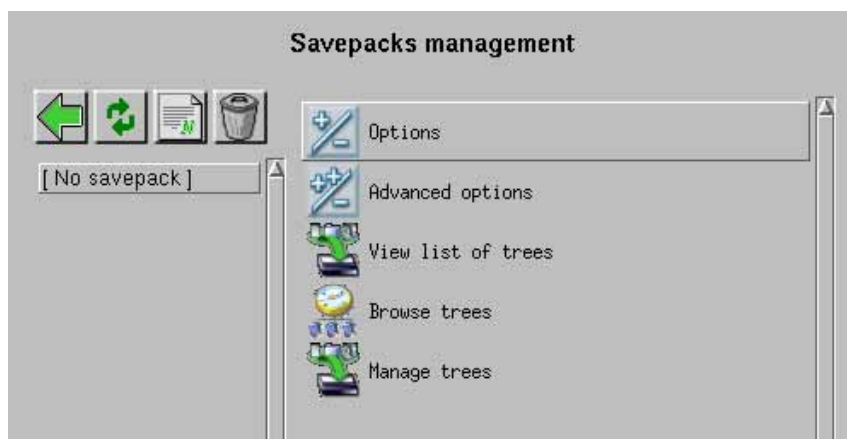
This section will deal with the first two classes, and the next will deal with the database and index backup.

- *How to backup the system (client/server) information*

For a Disaster Recovery procedure, it is essential that the first parameter to be restored (regardless of the client/server status of the machine) is the network configuration. This allows the Disaster Recovery procedure to connect to an Arkeia server in order to restore the data.

To backup the network configuration of a given machine, Arkeia version 5 supplies a plug-in, named “sysinfo”, which is designed especially for this task. If you are using “sysinfo” to backup the system configuration and parameters of an Arkeia server, Arkeia Corporation recommends that you use dedicated tapes and Savepacks to optimize this operation. If you are using “sysinfo” to backup the system configuration and parameters of an Arkeia client, using dedicated tapes or Savepacks is not necessary.

Within the Arkeia version 5 GUI, the “sysinfo” plug-in is available by using the function “Browse trees” within the definition of the Savepack:



Clicking on the icon will display the list of machines declared as client to the current Arkeia server. Clicking on the name of machine will then display the plug-ins available on the selected client:



The normal plug-ins that are available within a standard Arkeia version 5 are “file” (to backup the file systems of an Arkeia client) or “sysinfo”, which, as noted above, is the plug-in that is used to backup the system information of the client. The “/” after their respective names indicates that these elements contain other elements.

Clicking on “sysinfo” will then display the different classes of information that can be backed-up:



In the example above, which was taken from an Arkeia server, three standard classes are available: “config” (network parameters), “system” (hard disk parameters) and “dbase/” (Arkeia database and index).

To backup the “config” and “system” information, click on the check box next to their names. Once the check boxes have been “ticked”, the system information can be backed-up to tape by the Arkeia server.

Once selected, the system information items are displayed within the Savepack as a path that respects the following syntax:

```
machine_name!sysinfo:/config
machine_name!sysinfo:/system
```

Therefore, for a machine named “snafu.usa.acme.com”, the “sysinfo” information would be displayed as:

```
snafu.usa.acme.com!sysinfo:/config
snafu.usa.acme.com!sysinfo:/system
```

The same syntax will appear in the Arkeia logs.

As noted above, once the network parameters (the “sysinfo:/config” shown above) have been backed-up, they can be restored to an optional Disaster Recovery diskette. This diskette is quite useful if a system must be restored as quickly as possible. See above for more details.

- *How to backup the server database configuration*

Since an Arkeia version 5 server stores detailed information in both its database and its index each time an operation is performed (backup or restoration), including the machine name and target file, directory or path, restoring this information can be crucial to properly complete a Disaster Recovery procedure.

To save this information, all that is needed is to use the “`sysinfo`” plug-in, this time choosing the “`dbase/`” path instead of the two “`config`” and “`system`” files.



Clicking on the “`dbase/`” displays the two parts of the Arkeia database:



In the example shown above, the “`f3`” file is the configuration database itself, while the “`o3dbtree`” folder is the Arkeia index.

While the Arkeia database (or “`f3`”) is of a reasonable size, the index (which is contained in the directory “`o3dbtree/`”) is much larger, as it is divided in as many parts as there are clients declared to the current server. For instance, if a given Arkeia server has ten clients, named “`puter01`” to “`puter10`”, the “`o3dbtree/`” index will contain ten sub-directories, with the same names as the machines: “`o3dbtree/puter01`”, “`o3dbtree/puter02`”, etc... to “`o3dbtree/puter10`”.

Using the index, it is therefore easy to retrieve the server data, and only the server data, by restoring only the sub-directory that corresponds to the current server.



Please note: There are several important points to keep in mind when saving the database/index information of an Arkeia server:

- The database and the index should be saved at the same time. Restoration, though, can be partial (client by client).
- Since both the index and the database are updated every time Arkeia version 5 performs an operation on a client, it is suggested to backup both on a regular basis. However, even if a Disaster Recovery procedure has to be launched without the benefit of a complete database/index backup, the “`arkrstdb`” utility can help complete the missing entries within the database/index. For more information, refer to the section named “Recovering the Arkeia database with arkrstdb” on page 182 of this manual.

Please note that, whenever the Disaster Recovery applies to an Arkeia server, which performs backup operations on itself, it is recommended to backup the index of the server separately from the index of its clients. For instance, if the server is named “`plex`”, the index information that should be backed up is contained in “`sysinfo:/dbase/o3dbtree/f3`” and “`sysinfo:/dbase/o3dbtree/plex`”.

- If necessary, the database and index can be backed up and restored, even if no Disaster Recovery procedure is launched. This is possible only for the client index entries, not for the server-related information contained in the index. This may be useful if the database or the index become damaged.

2.3 Setting up a comprehensive backup procedure

Even though the Disaster Recovery procedure of Arkeia version 5 is sufficient to quickly repair a damaged machine, this procedure cannot do anything unless backups are performed on a regular basis on all machines.

Even if a machine crucial information has been saved to backup media, this will only restore the machine basic information: the data contained on its hard disk(s) should be saved, and not just the basic configuration information.

If the Disaster Recovery procedure is to be used on a site, Arkeia Corp. recommends the following backup factors to be reviewed and carefully applied on all the site's machines as well:

- Saving each machine configuration information should be done immediately. If necessary, Disaster Recovery diskettes should be created and duplicated immediately, with one kept on-site and another kept off-site.
- Saving the complete original state of the hard disk drive of all machines (client and server) should also be done as soon as possible. If necessary, duplicate of the backup tapes (or backup media) should be done, with one copy kept on-site and another kept off-site.
- A comprehensive backup policy should be established, using the differential/incremental backup functions of Arkeia version 5. These functions allow the system administrators to save only the files that have been modified since the last complete backup and therefore dramatically reduce the time necessary to perform backups. For more information on these functions, refer to chapters "7: Interactive backup" and "8: Periodic backup" of this manual.
- A second backup policy should be established, in order to save regularly the database and the index of the Arkeia server itself. Based on the importance of each client, as well as the importance of the data it contains, these backups can be done more or less regularly.



Please note: As noted above, it is possible to restore the data from backup tapes and update the database and index with this information. However, it is highly recommended to take into account the large sizes of the database and the index before setting the backup policy of these elements.

3 • Disaster Recovery procedure for a client



Please note: Reading and understanding the following section, on the hardware limitations of the Disaster Recovery procedure, is absolutely essential to the proper restoration of a system.

Attempting to restore a system without reading this section first is at the reader's own risk!

3.1 Hardware limitations

Whenever the Disaster Recovery procedure is launched, it is highly recommended to use a machine which is virtually identical to the original computer.

This includes:

- **Hard disk:** the machine used for the Disaster Recovery procedure should have a hard disk with a size at least as large as the original machine. If a larger disk is used, the unused space will be assigned as an unused (or: "free") partition. It is also recommended to obtain a disk with the same geometry as the original.
- **Network card:** The machine used for the Disaster Recovery procedure should have a network card equivalent to the original machine, if a Disaster Recovery diskette is used. However, the network card is automatically detected at the beginning of the procedure, and a different card can therefore be installed in a crashed machine.
- **SCSI card:** If the original machine used an SCSI card to connect to its hard disks, the system administrator should check that this card is in the list of supported SCSI cards (see below).

Supported SCSI cards:

- "aha152x" (Adaptec AHA152X/2825)
- "aha1542" (Adaptec AHA1542)
- "aha1740" (Adaptec AHA1740)
- "aic7xxx" (Adaptec AIC7xxx)
- "advansys" (AdvanSys)
- "AM53c974" (AM53/79C974 PCI)
- "BusLogic" (BusLogic)
- "dmx3191d" (DMX 3191 D)
- "dttc3x80" (DTC 3180/3280)
- "eata" (EATA ISA/EISA/PCI – DTP and generic EATA/DMA-compliant boards)
- "fdomain" (Future Domain 16xx SCSI/AHA-2920A)
- "g_NCR5380" (Generic NCR5380/53c400)
- "NCR53c406" (NCR 53c406a)
- "53c7,8xx" (NCR 53c7,8xx)
- "psi240i" (PSI 240i)
- "qllogicfas" (Qlogic FAS)

- “qlogicisp” (Qlogic ISP)
- “qlogicfc” (Qlogic ISP FC)
- “qla1280” (Qlogic QLA 1280)
- “seagate” (Seagate ST-02 and Future Domain TMC-8xx)
- “sim710” (Simple 53c710)
- “sym53c416” (Symbios 53c416)
- “tmcsim” (Tekram DC390(T) and Am53/79C974)
- “u14-34f” (UltraStor 14F/34F)
- “ultrastor” (UltraStor SCSI)

3.2 Disaster Recovery – step one: booting and network parameters

A Disaster Recovery CD-ROM is supplied by Arkeia Corp. with every copy of Arkeia version 5 that has a valid Disaster Recovery license. To start the Disaster Recovery procedure, the system administrator should power on the system, insert the CD-ROM in the drive of the crashed computer and make sure the BIOS of the system is set to boot on the CD-ROM drive.



Please note: At the time of this writing, the Disaster Recovery procedure does not support machines that either (a) do not have a CD-ROM drive or (b) that cannot boot on the CD-ROM drive.

Any attempt to recover these machines is probably doomed to failure, and Arkeia Corp. highly recommends that these machines should be upgraded to more modern PCs.

Once the machine has booted, the first screen of the Disaster Recovery procedure should be displayed.



Once the initialization stage of the Linux system has been completed, the Disaster Recovery procedure will ask if a recovery diskette is available.

```
-----  
Arkeia Crash Recovery Tool  
-----  
Do you have a Configuration Disk for this computer ? [y/N]  
(Say Yes if your backup administrator prepared a configuration  
floppy disk for this computer)  
-
```

If a Disaster Recovery diskette is available, enter “y” for “Yes”. If several Recovery files are present on the floppy, the Disaster Recovery procedure will then ask which file should be used for the computer being restored.

```
-----  
Arkeia Crash Recovery Tool  
-----  
Do you have a Configuration Disk for this computer ? [y/N]  
(Say Yes if your backup administrator prepared a configuration  
floppy disk for this computer)  
What is your network interface ? [eth0]  
-
```

The Disaster Recovery procedure will, at this stage of the process, ask a series of questions to initialize and restore the network connection of the computer. These questions are always followed by a default value, put between square brackets. For instance, as can be seen in the screen shown above, the default value for the first question is “[eth0]” (equivalent to the “Ethernet card 0” device) which is the default Linux network interface.



Please note: Even if a Disaster Recovery diskette has been prepared for the system being recovered, the Disaster Recovery procedure will ask the questions detailed below. The only difference is that the default values that will be offered at the prompt will be the correct ones for the system being restored. With a Disaster Recovery diskette, it becomes possible to just hit the “Enter” key to quickly select the correct values.

The questions asked by the Disaster Recovery procedure are the followings:

- Question 1: “What is your network interface?”
Enter the correct interface, or press the “Enter” key to select the default value, which is “eth0”.
- Question 2: “What is your network card?”
Enter the correct model, or the closest model you can find in the supported network card list, or press “Enter” to select the default value (see remark below).



Please note: The Disaster Recovery procedure will attempt to detect the correct model of the network card. The result of this detection will then appear as the default value.

If a Recovery diskette has been prepared, the default value will be always be the one that has been saved on the diskette.

If the network card default value does not seem correct, entering the letter “h” for “Help”, instead of a value, will display all the models that are supported by the Disaster Recovery procedure. The value of the correct model should then be entered after the question.

Once the network interface has been entered, the Disaster Recovery procedure will “probe” it in order to make sure it functions properly, as can be seen below:

```
Arkeia Crash Recovery Tool
-----
Do you have a Configuration Disk for this computer ? [y/N]
(Say Yes if your backup administrator prepared a configuration
floppy disk for this computer)

What is your network interface ? [eth0]

What is your network card ? [pcnet32/h for help]

pcnet32_probe_pci: found device 0x001022.0x002000
PCI: Enabling device 00:11.0 (0001 -> 0003)
PCI: Found IRQ 6 for device 00:11.0
IRQ routing conflict for 00:11.0, have irq 10, want irq 6
  ioaddr=0x0010a0  resource_flags=0x000101
eth0: PCnet/PCI II 79C970A at 0x10a0, 00 50 56 41 e3 78
pcnet32: pcnet32_private lp=c1cae000 lp_dma_addr=0x1cae000 assigned IRQ 10.
pcnet32.c:v1.25kf 26.9.1999 tsbogend@alpha.franken.de

Network card setup is ended.
Press ENTER to continue...
```

Once it has completed the network card initialization, the Disaster Recovery procedure continues by a series of question, designed to configure the network parameters themselves, this is the second step of the Disaster Recovery procedure:

```
-----
                        Arkeia Crash Recovery Tool
-----

What is the local computer's name
(in FQDN form like 'host.domain.tld') ? [cassioj.knox.com]
_
```

3.3 Disaster Recovery – step two: network configuration

Once the network card of the crashed machine has been properly initialized and configured, the Disaster Recovery procedure will ask a series of question in order to properly configure the network access of the machine.

- Question 3: “Local computer name”.
Enter the name of the computer to be restored in the “FQDN” (Fully-Qualified Domain Name) format. For instance: “`sigma.eu.acme.com`” is a valid name under this standard.
- Question 4: “Local IP address”.
Enter the IP address of the computer to restored. For instance, “`192.168.7.52`” is a valid IP address.
- Question 5: “Netmask”.
Enter the IP netmask of the local network.
- Question 6: “Gateway/Gateway IP address”.
The Disaster Recovery procedure asks if your network uses a gateway. If this is the case, it will then ask for the IP address of the gateway.
- Question 7: “DNS/IP address of DNS”.
The Disaster Recovery procedure also asks if the local network uses a dedicated DNS machine. If this is the case, the procedure asks the operator to enter the TCP/IP address of the DNS server.



IMPORTANT. Please note: At the time of this writing, the Disaster Recovery procedure does not support DHCP.

If the network the crashed machine is connected to uses DHCP, Arkeia Corporation recommends to configure the DHCP server in such a way that both the crashed machine and the Arkeia server used for its restoration always use the same “fixed” TCP/IP address.

For more information on this configuration, refer to the documentation supplied with your DHCP implementation.

Once the questions above have been answered, the Disaster Recovery procedure attempts to configure its connection to the Arkeia server by asking several questions to its operator:

```
What is the local computer's name
(in FQDN form like 'host.domain.tld') ? [cassiope.knox.com]
ulyse.knox.com
What is the local computer's IP address ? [192.168.7.52]
192.168.7.216
What is your netmask ? [255.255.255.0]

Do you have a gateway ? [Y/n]

What is your gateway's IP ? [192.168.7.1]

Do you have a domain name server ? [Y/n]

What is your DNS server's IP ? [192.168.7.1]

Is the local computer your Arkeia backup server ? [y/N]

What is your Arkeia backup server's name
(in FQDN form, like 'host.domain.tld') ? [parsec.knox.com]
gabriel.knox.com
What is your Arkeia backup server's IP ? [192.168.7.141]
192.168.7.3
What is the listening port of your Arkeia backup server ? [617]
```

- Question 8: “Is the local computer the Arkeia server?”
If you are using another machine as the Arkeia server to restore the crashed machine, answer “n” (for: “No”) in order to start the server configuration process. If you are trying to restore an Arkeia server, answer “y” (for: “Yes”). The default value is “No”. If the answer to this question is “Yes”, the Disaster Recovery continues with the process detailed in the next section: “Disaster Recovery procedure for a server” on page 228.
- Question 9: “Name of the Arkeia server?”
Enter the name of the Arkeia server in FQDN format. For instance “arkeia.nyc.stratus.com” is a valid name.
- Question 10: “IP address of the Arkeia server?”
Enter a valid IP address for the Arkeia server. For instance “192.168.7.141” is a valid IP address.
- Question 11: “TCP port used by the Arkeia server?”
Enter the TCP port used by the Arkeia server to communicate with its client. By default, the port used is “617”. This port may be changed by the system administrator, for instance if it conflicts with another network application or if a firewall has been setup to filter net-

work traffic.

- **Question 12: “TCP port used by local machine?”**
Enter the TCP/IP port used by the local machine to connect to the Arkeia server. By default, this value is “617”, but can be changed by the system administrator (see question 11 above).
- **Question 13: “Arkeia login/password?”**
Enter the login used to connect to the Arkeia server to retrieve data. By default, the login used is “root” and the password is empty. The login/password combination is defined when creating users on the Arkeia server. For more information on this subject, please refer to chapter: “2: Installation procedure” on page 21 of this manual, and to chapter: “6: Objects and users creation” on page 71 of this manual.

Once the parameters requested have been entered, the Disaster Recovery procedure will display the complete list of parameters and their values. The procedure then asks its operator to make sure that all parameters have been entered correctly.

```
-----  
Arkeia Crash Recovery Tool  
-----  
  
Configuration data  
Local Hostname: 'ulyse.knox.com'  
Local IP address: '192.168.7.216'  
Netmask: '255.255.255.0'  
Gateway: '192.168.7.1'  
Domain name: 'knox.com'  
DNS server: '192.168.7.1'  
Backup Server name: 'gabriel.knox.com'  
Backup Server IP address: '192.168.7.3'  
Backup Server listening port: '617'  
Local Arkeia port: '617'  
Login: 'root'  
Password: [no value]  
  
Is it correct ? [Y/n]  
-
```

If one of the parameters is incorrect, it is possible to answer “n” for “No”. The procedure then restarts from the beginning and asks again for the network card configuration, machine name, TCP/IP address, etc...

If the parameters are correct, the Disaster Recovery then attempts to initialize the network card that was configured on the first step of the procedure. If needed, the complete network configuration can then be displayed, as can be seen below:

```

Is it correct ? [Y/n]
Start IP config
Set the eth0 interface
Set the route
Show interfaces configuration ? [y/N]
y
eth0      Link encap:Ethernet  HWaddr 00:50:56:41:E3:78
          inet addr:192.168.7.216  Bcast:192.168.7.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:4 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          Interrupt:10 Base address:0x10a0

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:8 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0

Show route table ? [y/N]

```

If necessary the route table of the system can also be displayed, as can be seen below.

```

y
eth0      Link encap:Ethernet  HWaddr 00:50:56:41:E3:78
          inet addr:192.168.7.216  Bcast:192.168.7.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:4 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          Interrupt:10 Base address:0x10a0

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:8 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0

Show route table ? [y/N]
y
Kernel IP routing table
Destination Gateway         Genmask         Flags Metric Ref    Use Iface
localhost  *                255.255.255.255 UH    0      0      0 lo
192.168.7.0 *                255.255.255.0  U    0      0      0 eth0
default    chom.knox.com    0.0.0.0         UG    0      0      0 eth0
Press ENTER to continue...

```

Once these steps are completed, the restoration of the configuration by the Disaster Recovery procedure is almost completed. To finish this configuration restoration and start the system restoration itself, the “Enter” key should be pressed:

If, at any time during the restoration of the system by the Disaster Recovery procedure, an additional tape is necessary, the whole process will stop, display the name of the tape that it requires and the list of drives available. The operator should then insert the proper tape in one of the drives and indicate which drive contains the tape.

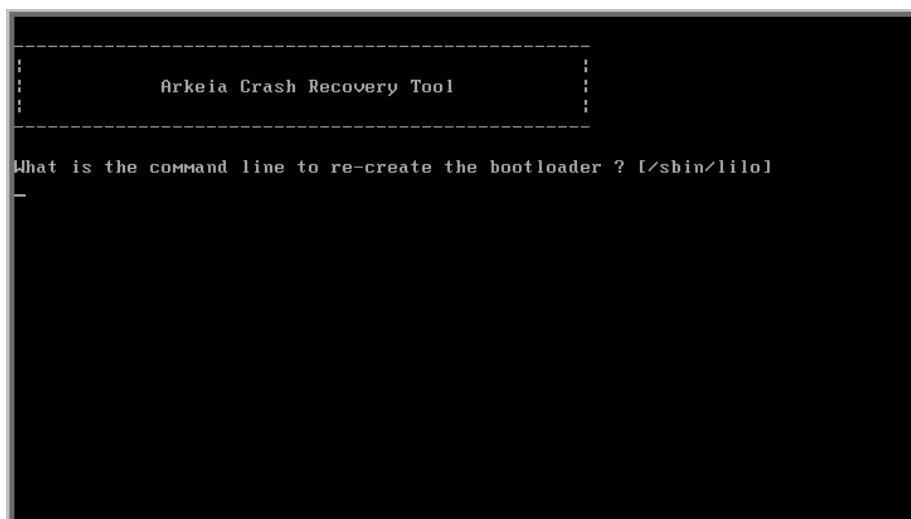
As can be seen in the screen shown above, the restoration itself is in two parts. The Arkeia Disaster Recovery will first rebuild the file system of the crashed machine, including the partitions and the formatting of the hard disk. Once the hard disk has been rebuilt, the partitions are mounted and the restoration of the data is launched.

Once the data restoration has been completed, the last step left in the Disaster Recovery process is the installation of the bootloader.

3.5 Disaster Recovery – step four: bootloader installation

Once the restoration of the client data has been completed, the fourth and final step of the Disaster Recovery procedure is the installation of the bootloader. Usually, the bootloader used by Linux systems is the program “`lilo`” (the “`LINUX LOADER`”). It may be possible that some Linux systems use “`grub`” or another bootloader.

If the user has created a Disaster Recovery diskette, the bootloader configuration may have been detected automatically. The two boot loaders that can be detected automatically by the “`sysinfo`” plug-in are “`grub`” and “`lilo`”. If the bootloader has been detected automatically, the default answer value will be the executable detected.



The Disaster Recovery procedure will ask which command should be used to install the bootloader. As can be seen above, the default value is “`/sbin/lilo`”. If the crashed machine used another bootloader, its complete path and name should be entered when requested. For “`grub`”, this path is usually “`/usr/sbin/grub`”.



Please note: If “grub” is used as the bootloader of the machine, the system administrator may have to enter the following command at the “grub” prompt:

```
root (hd0,0)
setup (hd0)
```

The command described above assumes that the machine boots from the first partition of the first disk: “(hd0,0)”. The first hard disk can be a SCSI disk (`/dev/sda1` under Linux) or an IDE (`/dev/hda1` under Linux).

After this last stage, the Disaster Recovery procedure should be completed and the computer should be completely restored, if a correct backup procedure has been put into place.

4 • Disaster Recovery procedure for a server

4.1 Overview

The Disaster Recovery procedure for a server is somewhat different from the “normal” Disaster Recovery procedure for a client. To be able to complete its restoration, the server must:

- Retrieve the parameters of its own disk(s) from its own backup tapes.
- Apply the saved parameters to the disk, to restore it to its original state.
- Once the disk parameters have been restored, the Arkeia database and index should be restored, but in a restoration limited to the Server backup information.
- Use this backup information contained in the database and index to completely restore its hard disk.
- Once the server has been completely restored, the client backup data, contained in the server database and index, can be restored client by client. This, in turn, allows the Arkeia server to restore the clients, using the normal functions of Arkeia version 5.

This section will deal with the different issues of this process and detail the procedure step by step.

4.2 Disaster Recovery – step one: hard disk parameters

As was explained earlier in this chapter, the system information of a given computer can be saved by Arkeia by using the “`sysinfo`” plug-in. This information is divided into two parts the “`sysinfo:/system`”, which contains the hard disk parameters for the Arkeia server, and the “`sysinfo:/config`”, which stores the network parameters for the server.

Since the “`sysinfo:/config`” can be saved to a Recovery diskette, in order to speed up the Disaster Recovery procedure, this section will not deal with it. For more information on the Disaster Recovery diskette, refer to the previous section of this chapter.

When the Disaster Recovery procedure is launched, it asks if the “local” machine is the Arkeia server. If this question is answered by “y” (for “Yes”), then the procedure prompts the user to insert the backup tape that contains the system information, stored in “`sysinfo:/system`”.

As noted above, this requires (a) that a “`sysinfo:/system`” backup has been performed before the crash/event that damaged the Arkeia server, and (b) that the tapes are correctly labeled in order to find the correct one quickly. For more information, refer to the section named “Creating a specific Disaster Recovery configuration” on page 211. Not labeling the Disaster Recovery tapes of an Arkeia server can be very costly, as it

increases the time of the restoration process, since Arkeia may have to read every tape in order to find the correct information!

Once the correct hard disk parameters have been retrieved from the Recovery tape, the Disaster Recovery process will apply them to the server hard disk and proceed with the next step.

4.3 Disaster Recovery – step two: server database and index

The second step of an Arkeia server Disaster Recovery is the restoration of the information contained in the server database and index. This is a necessary step, because, at this stage of the Disaster Recovery process, the only things that were recovered on the (damaged) Arkeia server are the network parameters and the disk parameters.

At this point, the Arkeia server itself does not know where the data it contains is stored, and not even what the data to be restored is. To be able to locate the data and restore it, it is therefore necessary to retrieve this information from the backup tapes.

Once the hard disk of the server has been restored, the Disaster Recovery will therefore prompt the operator to insert the tape(s) that contain the data backed up by the “`sysinfo:/dbase/f3`” file, which contain the database information.

Once the Arkeia server database has been restored, the Disaster Recovery process will prompt the operator to insert the tape(s) that contain the index of the Arkeia server. For instance, if the name of the server is “`apex`”, its index will be saved by the “`sysinfo`” plug-in as: “`apex!sysinfo:/dbase/o3dbtree/apex`”. Again, correctly labeling the backup tapes that contain this crucial data will save a lot of time for the Disaster Recovery process!

4.4 Disaster Recovery – step three: complete server restoration

Once the Arkeia database, and the index that contain the relevant information for the server, have been restored, the Disaster Recovery process can be completed and the Arkeia server restored to its pre-crash state.

During this restoration, the Disaster Recovery process will prompt the operator to insert the tapes needed to retrieve the data of the server. Once the restoration has been completed, the Disaster Recovery is terminated and the Arkeia server has been successfully brought back online.

Once this step has been completed, the bootloader should be installed on the server. For more information on this (important) step, refer to “Disaster Recovery – step four: bootloader installation” on page 226 of this manual.

4.5 Disaster Recovery – step four: client index restoration



Please note: Once the Disaster Recovery process has been completed, only the server database and index have been restored. To be able to restore data that was backed-up from a client machine, the client database and index should be restored.

To do this, assuming the client name is “`t e r r a`” and the server name is “`a p e x`” a backup of the following plug-in information should have been done on separate, clearly labeled tapes:

```
apex!sysinfo:/dbase/o3dbtree/terra
```

If this information has been correctly backed-up, restoring it is possible through the Arkeia GUI or through the “`arkrstdb`” utility. For more information on this utility, refer to chapter “12: Arkeia Utilities”.

CHAPTER 14:

ADVANCED NETWORK CONFIGURATION

- *Summary:*

This chapter details some of the advanced network configuration that may be required to allow Arkeia version 5 to function properly in specific network environments. For instance, this chapter will cover the subject of Arkeia servers with multiple network interface cards (NICs), how to backup Arkeia clients from behind a firewall and more.

1 • Name resolution and servers with multiple NICs

1.1 Overview

This paragraph explains how to use the “NLP_HOSTFILE”, “NLP_ONLYHOSTFILE”, and “NLP_HOSTNAME” parameters to configure the Arkeia network architecture.

These parameters can be used to specify which network interface card should be used, if several NICs are installed in a machine.

By default, the Arkeia name resolution is based on the local name resolution configuration. Machines are usually configured to use the local “hosts” file, and Arkeia always uses either DNS or the local name resolution. If your machine can resolve a name request and return the corresponding IP address (and vice-versa), Arkeia will be able to connect to the remote machine. If not, the connection problem is not a name resolution issue.

However, if Arkeia needs a specific network configuration, it is possible to separate the local network configuration and Arkeia’s own name resolution process.

1.2 Configuration for an Arkeia client

On the client, the backup server name must be entered in the following file:

```
/opt/arkeia/arkeiad/admin.cfg
```

This is, usually, the host name of the backup server. The IP address of the Arkeia backup server is usually located in the file:

```
/etc/hosts
```

1.3 Configuration for an Arkeia server

Each client is declared in the following file on the Arkeia backup server:

```
/opt/arkeia/arkeiad/rhost.lst
```

This file contains the machine name that will be shown in the navigator. The machine name in this file must be known by the backup server through its own name resolution configuration, as described in the introduction section.

Arkeia updates this information automatically. You should not have to update this file, except to possibly delete duplicate entries.

1.4 Hosts file used by Arkeia

By default, name resolution is based on the local system configuration. Generally, the resolution is done with:

- The local host file (“`/etc/hosts`”)
- The DNS (Domain name server)

Check “`/etc/nsswitch.conf`” (and, on Linux, “`/etc/hosts.conf`”) to check the order used by your backup server.

- *Arkeia hosts file*

You can set a specific Arkeia name resolution configuration. Two different methods are possible:

- Arkeia will use a specific hosts file first. (In this case, you need to set the “`NLP_HOSTFILE`” preference, see below). This file, like most hosts file, contains links between machine names and IP addresses. If a machine can't be found, Arkeia looks for the local system resolution.
- Arkeia will use only its own host file. If the search fails, then network connection is aborted. To set this you have to add “`NLP_ONLYHOSTFILE`” parameter.

1.5 Setting up “`NLP_HOSTFILE`” and “`NLP_ONLYHOSTFILE`”

On the Arkeia backup server and/or Arkeia client(s), you must configure the “`NLP_HOSTFILE`” variable with the name of the Arkeia hosts file, for example: “`/opt/arkeia/arkeiad/hosts.cfg`”, in the following file: “`/opt/arkeia/arkeiad/arkeiad.cfg`”. For instance:

```
NLP_HOSTFILE "/usr/knox/nlp/hosts.cfg"
```

Arkeia will search the addresses entered in the “`NLP_HOSTFILE`” first and then, if it cannot find the machine in question, it will request the local system name resolution configuration.

To exclusively use the “`NLP_HOSTFILE`” hosts, you must set the “`NLP_ONLYHOSTFILE`” variable to “1”. For instance:

```
NLP_ONLYHOSTFILE "1"
```


1.6 Syntax of the Arkeia-specific “hosts” file

The hosts file must contain all machine names in lower case. Do not use upper case names in this file. Additionally, the IP address must be surrounded by quotes. For instance:

```
"193.90.5.234" "hp_test.acme.com"
```

1.7 Usage of “NLP_HOSTNAME”

If the “NLP_HOSTNAME” variable is un-commented, it contains the machine name that Arkeia will substitute for the name returned by the “hostname” command.

This variable can be found in the following file:

```
/opt/arkeia/arkeiad/arkeiad.cfg
```

As part of the hostname substitution process, Arkeia performs a consistency check between the name specified by the “NLP_HOSTNAME” parameter and the name returned by the “gethostname” system call. It is possible to associate different IP addresses with a specific machine:

- If the name returned by the “gethostname” system call is different from the name defined on the “NLP_HOSTNAME” parameter, then the IP addresses must be identical.
- If the name returned by the “gethostname” system call and the name specified on the “NLP_HOSTNAME” parameter are identical then the IP addresses can be different.

The “hostname” command returns the host portion of the fully qualified name exactly as specified in the “/etc/hosts” file.

1.8 Configuring Arkeia not to use FQDN

By default, Arkeia uses Fully Qualified Domain Names (FQDN) for its client/server communications.

If for any reason, it is necessary not to use FQDN machine names, follow the procedure detailed below:

- Add/Un-comment the following “STRIP_DOMAIN” parameter in the “/opt/arkeia/arkeiad/arkeiad.cfg” file, on the backup server as well as on the client machines:

```
STRIP_DOMAIN "1" (Strip domain name)
```

- Restart “arkeiad” on all backup server and client modules with the following command, entered as “root”:

```
# arkeiad restart
```

- Use the GUI to open the navigator, and check that all the clients appear without a complete domain name.
- Make sure the Arkeia database is correctly updated by using the utility “`arkmvhost`”. For more information on this program, please refer to the section titled “Renaming servers/clients with `arkmvhost`” on page 198 of this manual.
- Try a small backup, including a directory of each client machine.

2 • How to use different TCP ports with Arkeia

By default, Arkeia uses the TCP port “617”. However, this can create some problems if another process tries to reserve the same port on your client.

The port used can be modified in the following file for client machines:

```
/opt/arkeia/arkeiad/arkeiad.cfg
```

The “`arkeiad.cfg`” contains many comments that can be used to obtain more information.

If the machine to be re-configured is the Arkeia server itself, the procedure is the same, but the file to be modified is:

```
/opt/arkeia/arkeiad/admin.cfg
```

In the files listed above, un-comment (or add) the line “`PORT_NUMBER`” and set it to the right port as shown below:

```
PORT_NUMBER "[New port number]"
```

Before modifying one of the files listed above, please enter the following command as “`root`” to stop the “`arkeiad`” Arkeia UNIX daemon:

```
# arkeiad remove
```

After one of the files detailed above has been modified, restart the “`arkeiad`” process, using the following command, entered as “`root`”:

```
# arkeiad start
```



Please note: if you plan to change the port value, you will also have to change it in the file named “`admin.cfg`”.

If a backup fails (due to a network problem, for example), Arkeia will try to recover and continue the backup operation by trying the other ports available from 1024 and upward (1025, 1026,...).

3 • Configuring Arkeia to work from behind a firewall

3.1 Introduction

If a firewall is present between an Arkeia version 5 server and a client machine, many problems can ensue, since most firewalls will, likely, block the ports used by Arkeia version 5 to establish a communication between two machines.

This section aims to solve these problems and explain how to establish correct communications between Arkeia servers and their clients, even through a firewall.

3.2 Standard configuration

- *Using the standard Arkeia port (617)*

The easiest way to configure Arkeia to work through a firewall, all that is needed is to allow connections on the “617” TCP port to/from the Arkeia client and the Arkeia server.

This means that your firewall should allow data to transit through its “617” port.

- *Using a different TCP port*

Of course, this port can be changed easily by following the information detailed in the previous section. This means that, on the Arkeia server and on the Arkeia client, the “PORT_NUMBER” variable should be changed in the file “/opt/arkeia/arkeiad/arkeiad.cfg”, to another value.

For instance, here is an example taken from a modified “arkeiad.cfg” file:

```
(previous lines deleted)

# The following field can be used to reconfigure the port
# number used by the Arkeia network programs as it known
# to the tcp/ip system and as, notably, it should be entered
# in /etc/services.
# Caution ! Do not modify this field unless you are absolutely
# sure you need to, and if so make sure that you attribute the
# same value to this port number on all hosts within the
# network which are to communicate using Arkeia services.

PORT_NUMBER    "500"

(following lines deleted)
```

In the case described above, all the machines that need to use Arkeia will have to have the same port modification (from port “617” to port “500”) applied to them, which can be a little bit long to set up.

Before modifying “`arkeiad.cfg`”, the system administrator should enter the following command, as the super user “`root`”:

```
./arkeiad remove
```

Once the file has been modified with a new port number, the system administrator should enter the following command as “`root`”, to restart “`arkeiad`” with the new port number:

```
./arkeiad start
```

Finally, another possibility is to use the port range management of Arkeia version 5, which is detailed below.

3.3 Using the port ranges

Arkeia version 5 manages all data exchange through three different ranges of TCP ports.

These ranges are usually limited to values higher than “1023”, since ports numbers smaller or equal to that limit are reserved by the system (see below).

The Arkeia TCP port ranges are the following:

- **Listening ports:** TCP Ports that are “listening” do just this – they allow an application to “listen” on them and wait for a connection to be established. In the Arkeia version 5 architecture, this usually means the processes that manage tape libraries. They can be also named “input” ports, because they receive, but do not send, data.
 - **Normal ports:** TCP ports that are “normal” are “output” ports – as opposed to the “listening” ports described above – meaning, they send data over the network to another machine. The difference between “normal” and “reserved” ports (see below) is that “normal” ports are assigned to normal users and applications, while “reserved” ports are assigned to “`root`” (the UNIX super user) and the applications “`root`” has started.
 - **Reserved ports:** TCP ports that are “reserved” can only be used by the “`root`” user and the programs started by this user. Like the “normal” ports, they are “output” ports, meaning they transmit – over the network – the data they receive from applications. Within the Arkeia version 5 architecture, the processes that use these ports are the most important processes of Arkeia, including “`arkeiad`”, the Arkeia UNIX daemon itself.
- *Control variables for port ranges*

Arkeia version 5 provides several port management variables, that can be used by both client and server to use ports other than port “617” to connect to each other and transfer data. This port management uses control variables to define which addresses can be used to establish communication between client and server.

The different ranges detailed above are controlled by the following variables:

- **Listening ports:** This port range is controlled by the following variables: “NLP_LISTEN_MIN_PORT” (for the minimum value) and “NLP_LISTEN_MAX_PORT” (for the maximum value). These two variables should be entered in the following file, for both the client and the server:

```
/opt/arkeia/arkeiad/arkeiad.cfg
```

Please note that this configuration is optional, since the ports that are really needed for client/server communications are the reserved and the normal ports.

- **Normal ports:** This port range is controlled by the following variables: “NLP_NORMAL_MIN_PORT” (for the minimum value, the default value being “1024” under Linux) and “NLP_NORMAL_MAX_PORT” (for the maximum value, the default value being “32000”). These two variables should be entered in the following file, for both the client and the server:

```
/opt/arkeia/arkeiad/arkeiad.cfg
```

- **Reserved ports:** This port range is controlled by the following variables: “NLP_RESV_MIN_PORT” (for the minimum value, the default value being “512”) and “NLP_RESV_MAX_PORT” (for the maximum value, the default value being “1023” under Linux). These two variables should be entered in the following file, for both the client and the server:

```
/opt/arkeia/arkeiad/arkeiad.cfg
```

Whenever a particular port, specified in these variables, is not available, Arkeia will try to use the next port in the range, until it finds a port which is both open and available.

You’ll find below some examples of configuration for these variables:

```
NLP_RESV_MIN_PORT "768"
NLP_RESV_MAX_PORT "800"
NLP_NORMAL_MIN_PORT "1024"
NLP_NORMAL_MAX_PORT "1048"
```

3.4 SSH configuration

Another option, if you want to increase security, would be to use SSH to tunnel the port securely through the firewall. Enter the following line on the client machine:

```
ssh -g -C -L 617:[backup_server]:617 [firewall_address]
```

This redirects local port 617 to the remote port specified above. Then all you need to do is start the Arkeia GUI and login to “localhost”.

- *Sample configuration*

Here’s how an external Arkeia server (192.168.1.50) is set to backup an internal Arkeia client (111.222.111.222) through a Linux firewall (eth0=111.222.111.2 and eth1=192.168.5.254):

- Allow connections on port 617 from the client to the backup server:

```
/sbin/ipchains -A input -d 111.222.111.2 -p tcp-s
111.222.111.222 -j ACCEPT ipmasqadm portfw -a -P tcp -L
111.222.111.222 617 -R 192.168.1.50 617
```

- Allow connections on port 617 from the backup server to the client:

```
/sbin/ipchains -A forward -s 192.168.1.0/24 -i eth0 -j MASQ
```

- Add an entry in “/etc/hosts” on the Arkeia Server:

```
111.222.111.222 myclient.extdomain.com myclient
```

- Duplicate a block entry for “myclient” on the Server in “/opt/arkeia/arkeiad/rhost.lst” using its expected IP of “111.222.111.222”

- Restart “arkeiad” on the backup server

- Remove the bogus entry for “myclient” in “rhost.lst”, as another should have been added that represents the real machine.

- Restart “arkeiad”

- Remove “[1]” from “/opt/arkeia/arkeiad/auth_ARKFS.cfg” on the client:

```
ARKFS.* ALLOW * *
```

Run Arkeia on the client, connect to the firewall machine and authenticate as if you were connected to the Arkeia server on “myserver.intdomain.com”.

It will appear as if your Arkeia server is now residing on the firewall, but the packets for port 617 are actually being forwarded to the internal Server port 617 (“myserver.intdomain.com” with the IP of “192.168.1.50”). Any reverse traffic via port 617 are done through the generic internal-to-external MASQuerading.

4 • Using Arkeia and DHCP

4.1 Overview

An Arkeia version 5 server assigns a name to each of its client, based on its FQDN host name and on its IP address. Since DHCP assigns a new IP address to each computer, this name assignement may fail with the result that the Arkeia server becomes hopelessly confused.

However, Arkeia version 5 provides a simple way to manage the changes of IP address that are the result of the use of DHCP on a network.



Please note: At the time of this writing, the Disaster Recovery function of Arkeia does not manage DHCP. Normal backup and restore operation, though, are able to manage DHCP.

4.2 Using the variable “VIRTUAL_HOSTNAME”

To properly manage DHCP and its attendant IP address changes, all that is needed is to enter the variable “VIRTUAL_HOSTNAME” in the “arkeia.cfg” file of the client machines.

This variable should contain the name of the client. This name will then be used by the Arkeia server to identify the client, store information in its logs and its backup database, display the client name in its network navigator, etc.

For instance, the following lines define two different clients that will be named: “fubar.acme.com” and “Jack's server”, respectively, for the Arkeia version 5 server it will connect to:

```
VIRTUAL_HOSTNAME "fubar.acme.com"  
VIRTUAL_HOSTNAME "Jack's server"
```



Please note: This type of configuration should never be used on an Arkeia server! The Arkeia clients needs to know the name of their server, and changing this name will prevent the clients from connecting to their server.

Of course, this procedure is useless and not recommended if the DHCP server assigns static IP addresses to the Arkeia clients.

APPENDIX 1: MYSQL BACKUP

- *Summary:*

This appendix contains important information on the “`arkmysql`” Arkeia plug-in. This plug-in is used to directly backup and restore MySQL databases with Arkeia version 5.

1 • Overview

The program “`arkmysql`” is a system plug-in for the Arkeia v5 backup program. Through this plug-in, system administrators can backup and restore MySQL databases without having to stop the MySQL server itself.

The plug-in itself should be installed after the installation of Arkeia version 5, and requires a specific license to work properly.

Once installed, a file named “`arkmysql.prf`” will be created in the directory named: “`/opt/arkeia/arkeiad`”.

This file contains the configuration of the plug-in. It should be modified before attempting any backup operation with the plug-in. To obtain more information on the contents of this configuration file, refer to the rest of this appendix.

2 • Configuration of the Arkeia plug-in

2.1 The “arkmysql.prf” preference file.



Please note: At the time of this writing, the only way to properly configure the MySQL plug-in for Arkeia version 5 is to edit manually the file “arkmysql.prf”, located in the directory: “/opt/arkeia/arkeiad”.

This file contains many comments that should help all users edit and modify the file, but you'll find a short description of each configuration field in the rest of this appendix.

Any UNIX text editor can be used to modify the file.

The format of the “arkmysql.prf” file is the following:

```
PARAMETER          "VALUE"
```

Where “PARAMETER” is one of the following:

- “ARKPMYSQLOGLEVEL”: Modifies the log verbosity.
- “ARKPMYSQL_USOCKET”: Modifies the socket used by the plug-in.
- “ARKPMYSQL_HOST”: Modifies the host name.
- “ARKPMYSQL_USER”: Modifies the MySQL user ID.
- “ARKPMYSQL_PASSWD”: Modifies the MySQL password.
- “ARKPMYSQL_PORT”: Modifies the TCP/IP port used by the plug-in.

2.2 Configuration of the log verbosity.

To configure the log verbosity, the parameter “ARKPMYSQLOGLEVEL” must be assigned a value.

The possible values for this parameter are “0”, “10”, “20” or “30”. A value of “0” means no messages, generated by the MySQL plug-in, will be stored in the Arkeia log.

A value of “10” will produce only exception errors in the log. This is the default value of the parameter “ARKPMYSQLOGLEVEL”.

A value of “20” will also include in the log the messages returned when an operation has been successfully completed.

A value of “30” (which is the maximum value for this parameter) will give detailed information on errors.

This parameter does not have to be set. If needed, it will mostly be used by the Arkeia Technical Support to troubleshoot problems that have been encountered at customers sites.

2.3 Configuration of the UNIX socket used by the plug-in

The parameter “ARKPMYSQL_USOCKET” is used to configure the UNIX socket (or the named pipe, under Microsoft Windows) that will be used to connect to the MySQL database.

The default value for this parameter is “/tmp/mysql.sock”.

2.4 Configuration of the host name or TCP/IP address.

The parameter “ARKPMYSQL_HOST” can contain either the name or the TCP/IP address of the computer running the MySQL database and the Arkeia backup client.

If this parameter is not set, the default value is “localhost” and assumes that the computer running the MySQL database is the local host.

2.5 Configuration of the MySQL user ID.

The parameter “ARKPMYSQL_USER” can be used to configure the name/user ID that will be used to connect to the MySQL database.

If this parameter is not set, the default value is “root”.

2.6 Configuration of the MySQL password.

The parameter “ARKPMYSQL_PASSWD” can be used to configure the password that will be used with the user ID to login on the MySQL database.

If this parameter is not set, the default value is a “blank” (or empty) value.

2.7 Configuration of the TCP/IP port.

The parameter “ARKPMYSQL_PORT” can be used to configure the TCP/IP port that will be used to connect to the machine running the MySQL database.

If this parameter is set with a value of “0” (or a blank value) then the port used is the default MySQL port.

3 • Using the plug-in to backup and restore MySQL

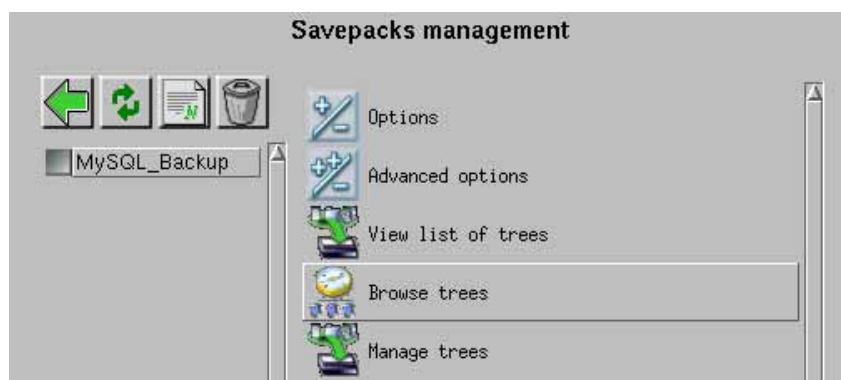
3.1 Backup with the MySQL plug-in.




To access a MySQL database, open the Arkeia version 5 GUI log in the server, and click on the following buttons: “Backup” then “Savepacks”.

These buttons are pictured on the left of this paragraph.

Then, create a new Savepack with a descriptive name, such as “MySQL_Backup”. Please refer to “Creating savepacks” on page 83 of this User Manual for more information on how to create a Savepack.




Once the Savepack has been created, click on: the option named “Browse trees” in the “Savepacks management” screen. This screen is shown above.

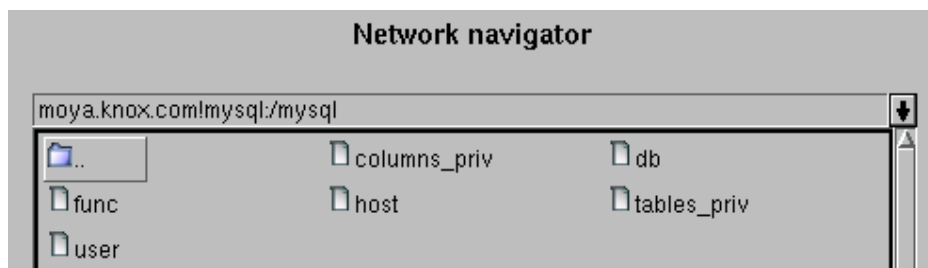


Then, select the client machine that contains the MySQL database. In the Navigator, you should find a list of all the plug-ins that are available to backup the client machine, usually “file/” and “sysinfo/”, with a plug-in added, named “mysql/” or “arkp-mysql/”. This is shown in the screenshot pictured above.

Select the MySQL plug-in. The GUI should now display the MySQL database as a standard tree structure. You can select the whole database, or only some specific tables for backup.



For instance, in the screenshot pictured above, there are two databases available within the “mysql/” plug-in, named: “mysql/” and “test/”. On a production machine, the “mysql/” plug-in can contain many databases, with less confusing names than in the example above, such as “accounting/”, “webstats/” and so on.



In the screenshot pictured above, the “mysql/” database has been selected by the user, and the Arkeia “Network Navigator” displays the tree structure of the database itself. This allows system administrators to select only some parts (or “tables”) of a database to be backed up by Arkeia version 5.

Once the database/tables to be backed up have been selected, right-click anywhere in the Arkeia v5 “Navigator” window and click on the “Validate” option. This operation adds the selected objects to the current Savepack.

3.2 Restoration with the MySQL plug-in.

If the backup of the MySQL database has been completed successfully with the Savepack created in the section above, the database or its internal tables should be available when clicking on “Restoration > Files”.

For more information on the restoration of files and objects saved by Arkeia version 5, please refer to chapter “9: Data restoration” on page 138 of this User Manual.



Please note: the restoration of the MySQL data should be done while MySQL itself is running. **Do not stop** the MySQL database software while attempting to restore its data.

APPENDIX 2: DB2 BACKUP

- *Summary*

This appendix describes how Arkeia version 5 can be used to backup an application's data (IBM DB2 database) by using its standard commands and functions of Arkeia version 5.



Please note: Arkeia Corporation is dedicated to bringing the easiest high-performance software to its clients. Therefore, this Appendix should not be understood as the “only” way to back-up an IBM DB2 database with Arkeia version 5. It is, however, a good example of the flexibility of the Arkeia functions.

A DB2 plug-in for Arkeia should be available in the near future from Arkeia Corporation. This plug-in will simplify the backup and restoration of the DB2 databases, by allowing the system administrators to “browse” and select the database from within the standard Arkeia interface. For more information, please check the web site, at the following address:

<http://www.arkeia.com>

1 • Introduction

1.1 How can Arkeia be used to backup an IBM DB2 Database?

Arkeia can use two methods to backup an IBM DB2 Database:

- Arkeia can backup the entire database structure “*Off-Line*” by using the standard Arkeia directory “tree” methods.
- Arkeia can backup a DB2 Database Online through several shell scripts that use the DB2 backup command-line facility.

1.2 Arkeia for IBM DB2 – requirements.

The following is required to perform a backup of the IBM DB2 Universal Database:

- The Arkeia Server, client and GUI are installed and properly configured.
- The IBM DB2 Database software is installed and functional, with at least one instance of the database created.



Please note: before attempting to backup the IBM DB2 Database, the Arkeia Server software installation should be validated by backing up and restoring files from/to local and remote machines.

2 • Offline Backup Configuration

2.1 “Offline” (Cold) Backup of the DB2 Database

The IBM Disaster Recovery Manual recommends to backup “*Offline*” the entire DB2 Database structure at least once a month to ensure that at least one month of logs are available in case the entire database requires a rebuild.

2.2 Prerequisite for an “Offline” backup

To accomplish the “*Offline*” backup of all database instances of the IBM DB2 Database, the DB2 Admin Server and all Database Instances must be Offline. To do this, follow the procedure outlined below:

- ***Stop the DB2 Database Admin Server:***

To do this, enter the following commands. These commands assume that you have a “root” access to the machine running the IBM DB2 database software:

```
su - db2as
db2admin stop
```

The IBM DB2 Database should then return the following message:

```
SQL4407W The DB2 Administration Server was stopped successfully.
```

- ***List all running DB2 database instances***

To do this, enter the following command (as “root”):

```
db2ilist
```

The instances of DB2 that are currently running should then be displayed. For instance:

```
db2inst1
db2inst2
db2inst3
```

- ***Stop all running database instances.***

To stop a running instance, enter the following commands:

```
su - db2inst1
db2 force application all
```

IBM DB2 should then return the following messages:

```
DB20000I The FORCE APPLICATION command completed successfully.
DB21024I This command is asynchronous and may not be effective immediately.
```



Please note: You have to repeat the commands detailed above for every instance of the IBM DB2 database that is currently running.

Therefore, if there are 4 instances of the database running, you'll have to enter:

```
su - db2inst1
su - db2inst2
su - db2inst3
su - db2inst4
```

- ***Stop the Database Manager.***

To stop the Database Manager, enter the following command:

```
db2stop
```

DB2 should then return the following message:

```
SQL1064N DB2STOP processing was successful.
```

- ***Check if all IBM DB2 processes have been terminated successfully***

To check if all the IBM DB2 processes have been terminated successfully, enter the following command:

```
db2 terminate
```

DB2 should then return the following message:

```
DB20000I The TERMINATE command completed successfully.
```

2.3 Arkeia Configuration for “Offline” backup

This section will provide an example of how to configure the Arkeia Savepack and Tree options to backup the database instances while in the Offline (“Cold”) mode.

- ***Savepack and Tree configuration***

Within Arkeia, create a Savepack designated for the Monthly “Offline” backup of all instances of the IBM DB2 Database and select the path to all database instance locations.



- ***Advanced Savepack Options Configuration.***

After the backup of all directory trees have completed, we will use the “Command After Backup” in the “Savepack Option” to start the IBM DB2 Database by placing the following command in the “Command After Savepack” field:

```
<hostname>:/etc/rc.db2 start
```



3 • “Hot” Backup Configuration

3.1 Online Backup of DB2 Database

When instances of the IBM Database are created, by default each instance that is created is given its own user name and password. The Super user (“root”) does not have permission to administer the database instance by default. Since Arkeia runs as “root”, Arkeia Corp provides the following method to backup a DB2 instance, while this instance is online.

The following example is based on a DB2 instance owner named “b2inst2” with a database named “sample”.

3.2 Setting the DB2 instance logging to “ON”

Make sure that the “LOGRETAIN” Parameter to is set to “RECOVERY” or “ON” for all instances of a database that will be backed up online.



Please note: by default this parameter is set to “OFF”

To do this, please enter the following commands:

```
db2 "update database configuration for sample using logretain on"
db2 "get db cfg for sample" | grep LOGRETAIN
Log retain for recovery enabled (LOGRETAIN) = RECOVERY
```

3.3 Creating the Backup Scripts

The sample scripts listed below are based on a single instance of an IBM DB2 database running.

As super user (“root”), create the following shell scripts in the “/tmp” directory.



Please note: For the example provided in the rest of this appendix, we will use the following parameters:

```
Instance Owner = db2inst1
Database Name = sample
```

The examples supplied also assume that you are using a Linux workstation.

3.4 The “create_lock.sh”

Create the “create_lock.sh” script, as the user “root”, in the directory “/tmp”.

```
#!/bin/sh
#
# file: create_lock.sh
# description: set the lock file and the named pipe
#
TMP=/tmp
BIN=/tmp
PIPE=$TMP/db2_pipe
LOCK=$TMP/db2_arkeia.lck
LOGFILE=$TMP/log
ECHO="/bin/echo $(date) "
$ECHO Entering create_lock.sh >> $LOGFILE

#
# create the named pipe first
#

if [ ! -e $PIPE ]
then
$ECHO creating the named pipe \($PIPE\) >> $LOGFILE
mkfifo $PIPE
fi

#
# test the pipe creation and set the lock
#

if [ ! -e $LOCK ]
then
$ECHO creating the lock file \($LOCK\) >> $LOGFILE
touch $LOCK
else
$ECHO current date is $(date) >> $LOGFILE
$ECHO a lock file is already set >> $LOGFILE
fi
```

3.5 The “delete_lock.sh” script

The “delete_lock.sh” script should be created as “root” in the directory “/tmp”.

```
#!/bin/sh
#
# file: delete_lock.sh
# description: unset the lock file and the named pipe
#
TMP=/tmp
BIN=/tmp
PIPE=$TMP/db2_pipe
LOCK=$TMP/db2_arkeia.lck
LOGFILE=$TMP/log
ECHO="/bin/echo $(date) "
$ECHO Entering delete_lock.sh >> $LOGFILE

#
# delete the lock file
#

if [ -e $LOCK ]
then
$ECHO removing the lock file \($LOCK\) >> $LOGFILE
rm $LOCK &> /dev/null
fi

#
# delete the named pipe
#

if [ -e $PIPE ]
then
$ECHO removing the named pipe \($PIPE\) >> $LOGFILE
rm $PIPE &> /dev/null
fi
```

3.6 The “db2_backup.sh” script

Create the “db2_backup.sh” script as user “root” in the directory “/tmp”.



Please note: modify the “DB2_USER” and “DB” parameter as required by your system.

```
#!/bin/sh
#
# file: db2_backup.sh
# description:
# Use the lock file for synchronization before
# starting the dump operation to the named pipe
#
TMP=/tmp
BIN=/tmp
PIPE=$TMP/db2_pipe
LOCK=$TMP/db2_arkeia.lck
LOGFILE=$TMP/log
ECHO="/bin/echo $(date) "
DB2_USER=db2inst1
DB=sample
$ECHO Entering db2_backup.sh >> $LOGFILE

#
# synchronization on the lock file
#

while [ ! -e $LOCK ]
do
sleep 5
done

#
# start the db2 database dump to the pipe
#

su - $DB2_USER -c "/usr/IBMdb2/V7.1/bin/db2 backup db $DB to $PIPE >>
$LOGFILE"

#
# exit
#

exit 0
```

3.7 The “db2_backup_cat.sh” script

Create the “db2_backup_cat.sh” script as user “root” in the directory “/tmp”.

```
#!/bin/sh
#
# file: db2_backup_cat.sh
# description: Send the data to the backup server through the
# standard output
#

TMP=/tmp
BIN=/tmp
PIPE=$TMP/db2_pipe
LOCK=$TMP/db2_arkeia.lck
LOGFILE=$TMP/log
ECHO="/bin/echo $(date) "
$ECHO Entering db2_backup_cat.sh >> $LOGFILE

#
# create the lock file and the named pipe
#

$BIN/create_lock.sh

#
# get the data from the named pipe
#

cat $PIPE

#
# remove the lock file and the pipe
#

$BIN/delete_lock.sh

#
# exit
#

exit 0
```


3.8 The “db2_restore.sh” script

Create the “db2_restore.sh” script, as user “root” in the directory “/tmp”.



Please note: modify the “DB2_USER” and “DB” parameter, as required by your system.

```
#!/bin/sh
#
# file: db2_restore.sh
# description:
# Use the lock file for synchronization before
# starting the dump operation to the named pipe
#
TMP=/tmp
BIN=/tmp
PIPE=$TMP/db2_pipe
LOCK=$TMP/db2_arkeia.lck
LOGFILE=$TMP/log
ECHO="/bin/echo $(date) "
DB2_USER=db2inst1
DB=sample
$ECHO Entering db2_restore.sh >> $LOGFILE

#
# synchronization on the lock file
#

while [ ! -e $LOCK ]
do
sleep 5
done

#
# start the db2 database dump from the pipe
#

su - $DB2_USER -c "/usr/IBMDB2/V7.1/bin/db2 restore db $DB from $PIPE
without prompting >> $LOGFILE"

#
# exit
#

exit 0
```

3.9 The “db2_restore_cat.sh” script

Create the “db2_restore_cat.sh” as user “root” script in the directory “/tmp”.

```
#!/bin/sh
#
# file: db2_restore_cat.sh
# description:
# Use the lock file for synchronization before
# starting the restoration to the named pipe
#
TMP=/tmp
BIN=/tmp
PIPE=$TMP/db2_pipe
LOCK=$TMP/db2_arkeia.lck
LOGFILE=$TMP/log
ECHO="/bin/echo $(date) "
$ECHO Entering db2_restore_cat.sh >> $LOGFILE

#
# start the db2 restoration operation
#

$BIN/db2_restore.sh &

#
# synchronization on the lock file
#

$BIN/create_lock.sh

#
# start the db2 database restoration to the pipe
#

cat - > $PIPE

#
# remove the lock file and the pipe
#

$BIN/delete_lock.sh

#
# exit
#
exit 0
```

3.10 The “db2_check.sh” script

Create the “db2_check.sh” script, as user “root” in the directory “/tmp”.

```
#!/bin/bash

#
#Set all VARIABLES
#

LOG=/tmp/backup_log
LAST=0
FILE=0
LASTID=0
ABORT=0
GOOD=0
LOG=/tmp/db2_abort.log
FILE=/opt/arkeia/server/server.jl2

#
#Get the latest Backup Job that was started.
#
LASTID=`tail -1 $FILE |awk '{print $4}'`

#
#Echo the start time to the log file
#

echo $(date) $LASTID db2 Backup has Started. >> $LOG

#
#Set the loop while the backup is in process

while [ "$LASTID" = "$LASTID" ]
do

#Grep for aborted job
ABORT=`cat $FILE|grep $LASTID|grep aborted |awk '{print $8}'`

#Grep for Backup Finished OK
GOOD=`cat $FILE |grep $LASTID |grep "End of backup for mount"|awk
'{print $5}'`

if [ "$GOOD" = End ]
then
echo $(date) $LASTID db2 Backup Finished OK >>$LOG
exit 0
else
if [ "$ABORT" = aborted ]
then
echo $(date) $LASTID db2 Backup was ABORTED!! >>$LOG
echo $(date) $LASTID Killing all processes for this backup!!>>$LOG

#If the backup is aborted, we need to kill the defunct processes

ARKSTATS=`ps -ef|grep arkstats|head -1|awk '{print $2}'`
ARKWRITE=`ps -ef|grep arkwrite|head -1|awk '{print $2}'`
ARKBFLOW=`ps -ef|grep arkbflow|head -1|awk '{print $2}'`
DBP1=`ps -ef|grep db2_backup_cat.sh|head -1|awk '{print $2}'`
DBP2=`ps -ef|grep db2_pipe|head -1|awk '{print $2}'`

#Give Arkeia 10 seconds to killoff it's own processes
sleep 10
#Arkeia is toooooo slow for the user, start the kills
```

```
/tmp/delete_lock.sh

kill -9 $ARKSTATS
kill -9 $ARKWRITE
kill -9 $ARKBFLOW
kill -9 $DBP1
kill -9 $DBP2

echo $(date) Executed /tmp/delete_lock.sh.
echo $(date) PID $ARKSTATS for arkstats was killed. >>$LOG
        echo $(date) PID $ARKWRITE for arkwrite was
killed. >>$LOG
        echo $(date) PID $ARKBFLOW for arkbflow was
killed. >>$LOG
echo $(date) PID $DBP1 for db2_backup_cat.sh was killed.>>$LOG
echo $(date) PID $DBP2 for db2_pipe was killed.>>$LOG
echo $(date) $LASTID All Processes have been killed!! >>$LOG
exit 0
        fi
fi

done

exit 0
```

3.11 Arkeia Configuration for “Online” Backup

Configure the Arkeia Savepack for an “object” Backup and Restore.

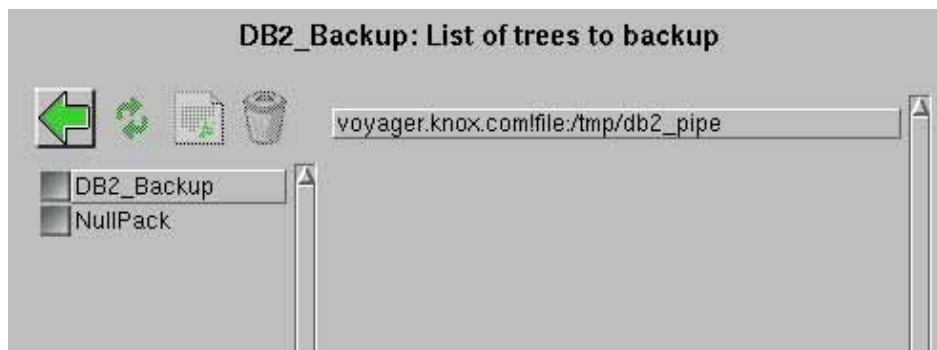
- *Savepack and Tree Configuration*

Create a Savepack that will be used to backup the Database Object that is created by the Backup Scripts. A good example would be to name this Savepack, “db2_backup”.

In the “List of Trees to Backup” window, right click and select “Create Tree”.

Enter the path of the created “db2_pipe” file per the “db2_backup.sh” script.

```
PIPE=$TMP/db2_pipe
```



- *Advanced Savepack Configuration*

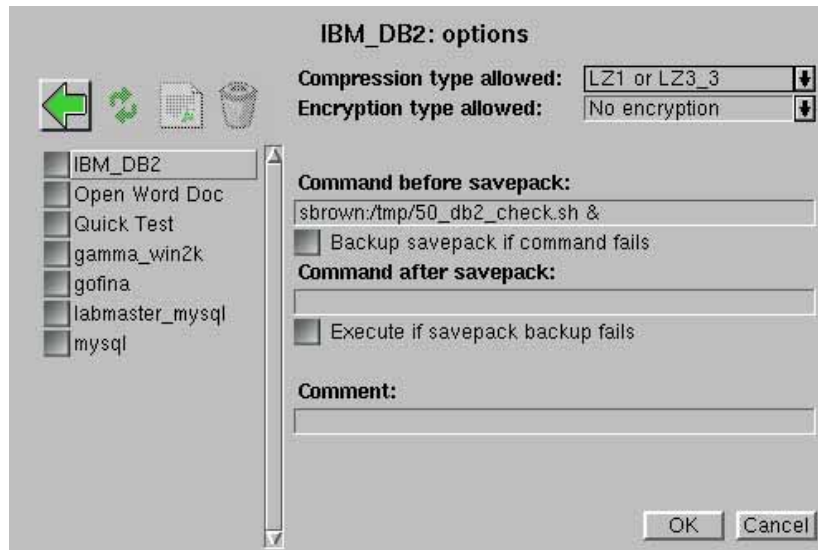
Double-click on the Savepack that you created.

Enter the “db2_check.sh” script into the “Command Before Backup” option. For instance:

```
Sbrown:/tmp/db2_check.sh &
```



Please note: this script must be run in the background.



3.12 Advanced Tree Configuration

Double-clicking on the “Tree List” in within the Savepack will provide access to the “Advanced Tree” options where the following changes must be made:

- Change the default Tree “Type” from “TREE” to “OBJECT”.
- Set the Multi-flow parameter to “1”.
- Enter the following path in the “Command Before Tree” option:

```
<machine name>:/tmp/db2_backup.sh &
```

- Enter the following path in the “Command After Tree” Option:

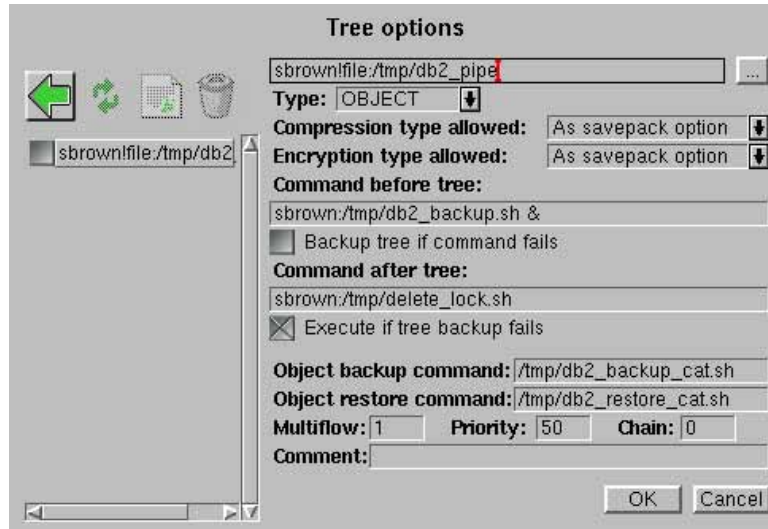
```
<machine name>:/tmp/delete_lock.sh
```

- Click on the “Execute if Tree Backup Fails” check box.
- Enter the following as the “Object Backup Command”

```
/tmp/db2_backup_cat.sh
```

- Enter the following as the “Object Restore Command”:

```
/tmp/db2_backup_cat.sh
```



APPENDIX 3: MIGRATING TO VERSION 5

- *Summary:*

The purpose of this appendix is to allow existing users of Arkeia version 4.2.x to migrate their installations and servers to Arkeia version 5, while retaining the backup information, index and database necessary to recover the information already saved to tape.



Please note: While Arkeia Corporation has tried to provide a complete and fairly “safe” migration tool for its products, the reader should be aware that this procedure can have potentially very destructive effects on existing Arkeia version 4.2.x servers.

ARKEIA CORPORATION RECOMMENDS THAT THIS APPENDIX IS READ AND FULLY UNDERSTOOD BEFORE ATTEMPTING TO RUN THE MIGRATION TOOL ITSELF.

1 • Overview of the migration tool

The Arkeia migration tool is designed to ease the transition from Arkeia version 4 to Arkeia version 5.

This utility is designed to transfer as many information and configuration parameters as possible, in order to allow all system administrators to transfer their existing 4.2.x configuration to the new architecture of Arkeia 5.

Applying the migration tool to an existing Arkeia version 4.2.x server allows the administrator to transfer the following information to the Arkeia version 5 server:

- Tape and Tape Pool configuration.
- Backup device (drive/library) configuration.
- User configuration.
- Backup configuration (except “*Periodic*” backups).
- Database, backup history and index. (These items are necessary for restorations).
- Partial network configuration.



Please note: As was briefly noted above, the migration tool detailed applies only to Arkeia servers. For more information on the best way to transfer a v4.2.x client machine to Arkeia version 5, please read the section titled “Migrating a client machine.” on page 272.

At the time of this writing, the migration tool has been tested on Linux, IBM AIX, Sun Solaris and SGI Irix. The procedure should be tested on other UNIX and operating systems shortly.

Please refer to the Arkeia Corp. web site (<http://www.arkeia.com>) for more information on the availability of this procedure.

2 • Obtaining the migration tool

It is possible to obtain the Arkeia migration tool from the following sources:

- The Arkeia Corporation web site, at the following URL:
`http://www.arkeia.com`
- The Arkeia version 5 CD-ROM distribution also contains the migration tool. It can be found in the following directory: “`/opt/arkeia/bin/migration`”.

This procedure is installed by default along with the Arkeia server. For more information, please refer to the rest of this document.



Please note: The Arkeia version 4.2.x license should be upgraded to an Arkeia version 5 license or the migration tool will not be able to work properly. Please contact the Arkeia Sales Team for more information.

3 • Running the migration tool on a server

3.1 Keep a copy of the Arkeia version 4.2 configuration

If you have made some changes to the configuration of your Arkeia version 4.2.x machine, make sure you keep a copy of these modifications!

Some modifications may not be taken into account by the migration tool from the version 4.2 to Arkeia version 5.

3.2 Required operations before the migration.

The following operations should ALWAYS performed BEFORE attempting the migration of an Arkeia version 4.2.x server to an Arkeia version 5:

- Stop all the Arkeia version 4.2.x processes:
 - a. Stop all backup/restore operations.
 - b. Stop the “mdl” library processes.
 - c. Quit the Arkeia version 4.2.x GUI (no “odts” processes should be left running).
 - d. Stop the “nlserverd” process by entering the command “NLSEVRD stop”.
 - e. Remove the “ARKPER” entry in the root crontab.
- Rename the Arkeia directory and make a backup:
 - a. The Arkeia version 4.2.x directory should be renamed from: “/usr/knox” to (for instance): “/usr/knox-4.2”
 - b. Make a complete backup of this directory to a tape in order to keep a copy of the Arkeia database and index, if necessary.



Please note: performing these operations before the migration tool itself is launched is absolutely crucial to the success of the overall migration. If you do not perform these steps, the migration tool may fail and all your previous backup data could be lost!

- Remove the references to Arkeia version 4.2.x in the “rpm” database (“.rpm” files only):

If you have used “.rpm” packages to install Arkeia version 4.2.x, please un-install Arkeia version 4.x with the following commands (entered as “root”):

```
rpm -e --justdb arkeia-server
```

```
rpm -e --justdb arkeia-client
rpm -e --justdb arkeia-gui
rpm -e --justdb arkeia-arkc
```



Please note: if you have installed the Arkeia version 4.2.x using the “tar.gz” package, you can safely disregard the above note. However, if you have used an “.rpm” package, using the commands indicated above is mandatory. Any other “rpm” command may well fail or delete important files that are required by the migration.

3.3 Install Arkeia version 5 (“.rpm” only):

Install Arkeia version 5 by entering the following command as “root”:

```
rpm -Uvh arkeia-5.0.2-1.i386.rpm
```



Please note: the exact name of the RPM file to be installed may be different from the example given above. If you are using the “tar.gz” package, just install Arkeia by using the procedure detailed in chapter “2: Installation procedure” on page 21 of this manual.

Once the installation has been completed on your system, don’t forget the step indicated below!

- Stop Arkeia version 5:
 - a. Arkeia version 5 must be stopped before the migration tool can be started. To do this, enter the command indicated below as “root”:

```
/opt/arkeia/bin/arkboot stop
```

- b. Check that all the Arkeia processes have been stopped by entering the following command as “root”:

```
ps -faxuw | grep ark
```

3.4 Running the migration tool.

To run the migration tool, simply enter the following command as “root”:

```
/opt/arkeia/bin/migration/arkmigrate [options]
```

The following options must be supplied on the command line:

- “--oldname [VALUE]”: Arkeia version 4.2.x server name.
- “--fqdn [VALUE]”: Arkeia version 5 “FQDN” server name.

- “`--path42 [VALUE]`”: Arkeia version 4.2.x root directory. (For instance: “`/usr/knox`”).
- “`--path50 [VALUE]`”: Arkeia version 5 root directory. (For instance: “`/opt/arkeia`”).

For instance, the following command, entered as root, would perform a valid migration from the Arkeia version 4.2.x “`foobar`” to the new Arkeia version 5 “`foobar.acme.com`”:

```
/opt/arkeia/bin/migration/arkmigrate --oldname foobar
--fqdn foobar.acme.com --path42 /usr/knox-4.2 --path50 /opt/arkeia
```

(For readability reasons, the options above appear on different lines – they should, of course, be entered on the same line as the command itself).

The migration tool performs automatically the tasks described in the section below. Follow the instructions printed on the screen to complete the migration.

3.5 Tasks performed by the migration tool

The “`arkmigrate`” procedure does the following:

- Make a backup copy of the version 5 database directory for reference.
- Copy and change the version 4.2 “`/usr/knox/nlp`” configuration files in the appropriate version 5 directory (“`/opt/arkeia/arkeiad`”).
- Copy and change the Arkeia version 4.2.x “`dbase/f3`” directories (except for the Periodic backup configuration).
- Copy and rename the Arkeia journal.
- Change the backup server name to its FQDN value. For instance: as we have seen above the server “`foobar`” becomes “`foobar.acme.com`”.
- Migrate the version 4.2.x index to reflect the version 5 structure.



Please note: The “`arkmigrate`” utility actually creates a symbolic link from the Arkeia version 5 location (`/opt/arkeia/server/dbase/o3dbtree`) to the version 4.2.x index, which is kept in its original directory (Please refer to the sections above for more information). A backup copy of the version 4.2.x index should be made before launching the “`arkmigrate`” procedure, since it will deeply modify the index to put it into the version 5 index format.

Keeping a backup of the original allows you to “roll back” the version 5 installation to the previous version 4.2.x installation if need be.



Arkeia Corporation recommends you to test the Arkeia version 5 server using the migrated information (see the section “Testing the migrated information.” on page 271, below). If the server seems to be working correctly, the converted index can be moved and the symbolic link: “`/opt/arkeia/server/dbase/o3dbtree`” should be updated to reflect this change.

A complete log of the operations of the `arkmigrate` utility is kept in the following file:

```
/opt/arkeia/log/arkmigrate.log
```

If, for any reason, the migration tool fails to perform correctly, it should be restarted from the very beginning, meaning a backup of the version 4.2.x index and a new installation of Arkeia version 5. If the migration tool fails, please contact the Arkeia Technical Support for more information.

3.6 Tasks to be performed after the migration:

Do not forget to install the version 5 license using the Arkeia GUI (Graphical User Interface).

As noted above, Arkeia version 5 should be restarted and tested thoroughly. The index, if correct, can then be moved from its original directory to another location if the symbolic link is updated accordingly.

The entire Periodic Backup configuration should be done again, but the Arkeia version 5 GUI makes this operation much easier.

3.7 Testing the migrated information.

Arkeia Corporation recommends to test the Arkeia version 5 server using the migrated information. To do this, follow the steps below:

- Library inventory:
Exchange two tapes within a library, then stop and restart the library processes and check that Arkeia version 5 reflects the changes you have just made.
- Restoration from version 4.2 tapes:
Check that you can restore files from tapes that were created with Arkeia version 4.2.
- New backup/restore using the Arkeia version 5 server:

Perform a small backup and restoration operation using the migrated Arkeia version 5 backup server.

4 • Migrating a client machine.

There is no need to migrate immediately the client to Arkeia version 5: an Arkeia version 5 server is backward-compatible and can backup and restore data to version 4.2.x clients.

However, it is recommended to migrate the clients to version 5 as soon as possible.

To do this, follow the steps indicated below:

- Stop Arkeia version 4.2.x on the client machines.
- Make a backup copy of the configuration files.
- Un-install Arkeia version 4.2.x.
- Install Arkeia version 5.
- Make sure that the “`admin.cfg`” file contains the FQDN name of the backup server.
- Restart Arkeia (“`arkeiad`” on UNIX machines, or “`Arkeia Service`” on MS Windows machines).
- On the Arkeia server, stop the library processes with the following commands (entered as “`root`”):

```
/opt/arkeia/bin/arkboot stop
```

- Once the library processes have been stopped, restart the `arkeiad` process by entering (as “`root`”):

```
/opt/arkeia/bin/arkeiad
```

- Update the backup server configuration with the FQDN names for the upgraded clients by entering the following commands as “`root`”:

```
/opt/arkeia/bin/migration/arkmvhost --allhosts
```

(to upgrade all modified clients)

```
/opt/arkeia/bin/migration/arkmvhost --host old_name  
--newhost new_name
```

(to upgrade the client “`old_name`” to “`new_name`”)

- Once the backup server configuration has been updated, restart the library processes with the command

```
/opt/arkeia/bin/arkboot start
```



Please note: You will find a complete description of the “`arkmvhost`” command in “Renaming servers/clients with `arkmvhost`” on page 198 of this manual. Typing this command with the option “`--usage`” displays the options that it can accept.

After performing the operations described above, do not forget to restart the library processes if the Arkeia server uses one to perform its backups!

5 • Known problems and issues.

The following problems and issues are known to exist with the current version of the migration tool:

- The “`arkmigrate`” procedure does not take into account the configuration of the Periodic backups. Arkeia Corporation recommends to re-configure the Periodic backups under Arkeia version 5, and to start a new cycle of Periodic backups with a new “`Total`” backup. Please refer to the rest of this User Manual for more information.
- In the case of Windows or Novell clients, “`arkmigrate`” only takes into account the “normal” files, and not special files such as Novell’s NDS or the Windows Registry. It is possible to backup these files through the Arkeia version 5 plug-ins. Please contact Arkeia Corporation for more information.
- Special network configurations for the version 4.2.x, such as:
 - a. Arkeia servers with several network cards.
 - b. Backups through firewalls.
 - c. Authentication and proxies configuration.

are not migrated by “`arkmigrate`”. These configurations should be completely re-done. Please contact the Arkeia Technical Support Team for more information.

- At the time of this writing, the Arkeia version 5.0.3 GUI does not support the “`roles`” configuration. This could be a problem if several different users and roles have been defined on a (migrated) Arkeia v4.2.x server.



Please note: The types should not be modified, UNLESS these modifications have been specifically recommended by the Arkeia Technical Support Team. You’ll find below a short list of all the tape changes recommended by the Technical Support for Arkeia version 4.2.x:

- *Recommended tape changes:*
 - In the file named “`f3drv/drvtypes.lst`”:
 - AIT-2 drive: variable tape block size.
 - LTO drive: 1024 tape block size (This value is the new default for Arkeia version 5).

- In the file named “`f3tape/tptypes.lst`”:

AIT-2 drive: enter the property “`USE_CAPACITY`”.

Create a new “`FILE`” tape type with a size of 1000 MB.

- In the file named “`f3lib/libtypes`”:

Enter the property “`LOCK_STK`” to prevent inventory problems with some Autoloaders.

Increase the number of slots in the “`FILE`” library.



Please note: Do not use the utility “`arkrstdb`” to transfer the information contained on a v4.2.x tape back into an Arkeia version 5 database.

Unfortunately, a bug prevents the “`arkrstdb`” utility from updating the Arkeia version 5 index when processing tapes created with Arkeia v4.2.x. This problem could potentially create licenses issues.

A solution to this problem is to use the “`readarkeia`” utility instead of the “`arkrstdb`” utility. Please contact the Arkeia technical Support Team if you have problems with this utility or for more information on the availability of a correction.

If you have ANY problem with the configurations detailed above, please contact the Arkeia Technical Support Team.

6 • Obtaining help on the migration tool.

For any and all problems and issues regarding the migration tool from Arkeia version 4.2.x to Arkeia version 5, please contact the Arkeia Technical Support Team. The web page for the Technical Support Team is:

`http://www.arkeia.com/support.html`

APPENDIX 4: MICROSOFT WINDOWS BACKUP

- *Summary*

This appendix details how to use Arkeia version 5 and its Windows plug-in to backup and restore system data from Microsoft Windows clients.

1 • Overview

The Windows System plug-in for Arkeia version 5 is designed to backup and restore the system information contained in a Windows NT, Windows 2000 and Windows XP servers, such as the “Primary Domain Controller”.

The system information that can be backed-up and restored by this plug-in can be divided into four distinct parts:

- The Windows Registry.
- The Windows Active Directory (Windows 2000/XP only).
- The Windows Certificates (Windows 2000/XP only).
- The Windows Event Logs.



Please note: Backing up and restoring the system information of a Windows server is not an easy task. Before attempting any operation with Arkeia version 5, its Windows client and the “sysinfo” plug-in, Arkeia Corporation recommends that the system administrators read and understand this README file.

Specifically, not following the procedures described in the rest of this file may result in severe data loss and/or damage to the Windows operating system!

The Windows Registry, Active Directory and Certificates backup and restore operations are supported only by Arkeia version 5.0.7 and higher. Users who have a previous version of Arkeia 5 (5.0.6 or earlier) should upgrade at least to Arkeia version 5.0.7 if they'd like to use these functions.

2 • Installation

At the time of this writing, the Arkeia plug-in is installed automatically with every copy of the Windows client. Therefore, no special manipulation is needed to install and use the specialized Windows plug-in.

The use of the plug-in requires a valid license for your Windows client(s). If you do not have a valid license for the Windows computers you'd like to backup, your Arkeia 5 server will not be able to connect and/or backup the data present on the Windows clients.

3 • Backup

3.1 Backup of the Registry through the “sysinfo” plug-in:

To backup the Registry of a Windows machine through the standard “sysinfo” plug-in of Arkeia version 5, follow the steps outlined below:




Click on “Backup” > “Savepack” > “New Savepack”. These buttons are displayed, in the order mentioned, on the left of this paragraph.

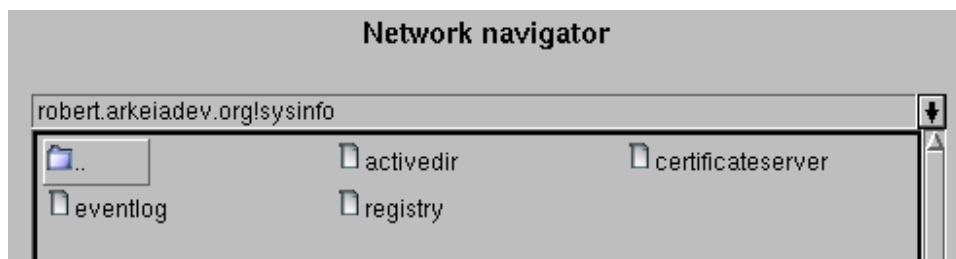
Enter a name for the Savepack of the system information, and press the “Enter” key (or click on the “OK” button) to create the Savepack and return to the main Savepack list.

Within the Savepack list, make sure the new Savepack is selected and click on the “Browse Trees” option. This should display a list of all clients registered on the Arkeia server.

Double-click on the Windows client you’d like to backup. This displays the list of plug-in available on the client, as can be seen in the screenshot below.



Then, click on the “sysinfo/” plug-in. The following screen should be displayed within the Arkeia version 5 window.



To backup the Windows Registry with Arkeia version 5, click on the “registry” file.

Once this is done, right-click anywhere in the list of plug-in and click on the “validate” option in the contextual menu that should be displayed.

The operation described above selects the Registry for backup within a Savepack. The next time the Savepack is chosen as the target of a backup operation (Interactive or Periodic backup), the content of the Registry of the Windows client will be saved to tape by the Arkeia server, along with all the other files/directories selected for backup.

When displayed, the tree path of the registry should be:

```
machine_name.domain.tld!sysinfo:/registry.
```

For instance, if three Windows machines are available, named: “foo”, “bar” and “rtfm” in a company named “acme.com”, the following trees are correct:

```
foo.acme.com!sysinfo:/registry
bar.acme.com!sysinfo:/registry
rtfm.acme.com!sysinfo:/registry
```

If a backup contains both the Windows Registry and the Active Directory, the restoration should always start by the Active Directory and finish with the Registry. Please see below, the section titled “Restoring the Active Directory with the “sysinfo” plug-in.” on page 284, for more information on the restoration of the Active Directory.

3.2 Backup of the Windows Active directory

The backup of the Windows Active Directory is managed in the same way than the backup of the Windows Registry. The only change is the name of the tree itself in the sysinfo plug-in, which is “activedir/”.

Therefore, to backup the Active Directory of a Windows machine, use the same procedure described above for the backup of the Windows Registry, but, within the “sysinfo” structure, select the “activedir/” directory instead.

For instance, the following tree path would be a correct one to backup the Active Directory of the Windows machine named “foobar”, within the domain named “acme.com”:

```
foobar.acme.com!sysinfo:/activedir
```


The procedure for backing up the Active Directory is, internally, slightly more complicated than the backup of the Registry, as Arkeia version 5 connects to the Active Directory service, saves the database and the transaction logs into the directory:

```
$ARKEIA_DIR/tmpreg/arkeia_activedir
```

Where “\$ARKEIA_DIR” is the directory in which the Arkeia client has been installed.

Once the database has been correctly saved to the temporary directory, the Arkeia Windows client will prepare the Active Directory data and send it to the Arkeia server, where it will be saved to tape.



Please note: Since the Arkeia Windows client will use a local directory to store the data to be backed-up before sending it to the Arkeia client, the local hard disk should have enough free space available to store a complete copy of the Active Directory database and logs. To check this, verify the size of the directory named “winnt/ntds”, which contains the Active Directory database (named “ntds.dit”) and its logs (named “edb*.log”).

Please read the Restoration sections of this appendix before attempting any restoration of the Windows Active Directory. Attempting to restore an Active Directory without reading the Restoration section may waste a lot of time if the restoration procedures are not followed.

3.3 Backup of the Certificates through the “sysinfo” plug-in.

The backup of the Windows Certificates is done by Arkeia version 5, through its “sysinfo” plug-in. The procedure to select the Windows Certificates is the same as the one detailed above for the Windows Registry, except that the directory to be selected within the “sysinfo/” directory is “certificateserver/”.



Please note: The directory “certificateserver/” will only appear within the tree structure of the “sysinfo/” plug-in if the service named “Certificate Server” (which is controlled by the binary named “certsrv.exe”) has been launched. If this entry does not appear while browsing the “sysinfo/” structure, make sure the corresponding service is launched. Arkeia version 5 can perform the backup of the Certificates only if the Certificates service is running.

At the time of this writing, Arkeia version 5 does not backup the private keys of the Certificates service. However, it is possible to save these keys, by using the Windows command “certutil” to create an “object” backup within Arkeia version 5. The Windows help files details the syntax of the “certutil” command. For more information on the “object” backup of Arkeia version 5, please refer to the Arkeia User Manual.

To check if the Certificates service is running on your Windows 2000/Windows XP server, use the Services MMC snap in, by clicking on “Start > Control Panel > Administrative Tools > Services”.

3.4 Backup of the Event Logs through the “sysinfo” plug-in.

The Windows Event Logs can be backed-up by an Arkeia v5 server by using the “sysinfo/” plug-in, in the manner described above for the Windows Registry.

To backup the Windows Event Logs, the system administrator who is creating the Save-pack should select the “eventlog/” entry in the Arkeia Navigator.

Like the other Windows information, the Windows Event Logs are saved to a temporary directory, on the hard disk of the Windows Arkeia client, before being sent to the Arkeia server. This temporary directory is named:

```
$ARKEIA_DIR/tmpreg/arkeia_eventlog_backup
```

Sufficient disk space should therefore be available to keep this temporary copy of the Windows Event Logs. For more information, please refer to the previous section of this appendix.

4 • Restoration of the Windows system information

4.1 Restoring the Windows Registry with the “sysinfo” plug-in.

Restoring the Windows Registry information can be done by selecting the correct, original path, within the Restoration screen. To do this, use either the “Navigator” function or the “Search file” function.

Once the correct backup has been selected, launch the Restoration from the Arkeia version 5 Server to the Arkeia version 5 Windows Client.

For more information on the Restoration functions of Arkeia version 5, please refer to the chapter “9: Data restoration” on page 138 of this manual.

4.2 Restoring the Active Directory with the “sysinfo” plug-in.

To restore the Active Directory backed up by Arkeia version 5:

- Restart the Windows machine in “Directory Services Restore mode”.

This reboot mode is available by pressing the “F8” key during the reboot of the machine to access the correct mode. Arkeia version 5 will NOT restore the Active Directory if this service is running.



Please note: While the Active Directory service is off-line, the authentications are still performed locally, through the SAM database of the Windows Registry. Never restore the Registry database while restoring the Active Directory, or before rebooting the machine hosting the Active Directory service. The restoration of the Active Directory should always be performed before the restoration of the Registry, and the Active Directory should be completed, and the server rebooted, BEFORE restoring the Registry.

Once the Restoration has been completed by Arkeia, the machine hosting the Active Directory service should be immediately rebooted in normal mode. Once the machine has completed its reboot, the Active Directory should detect the restoration and start a recovery process. During this recovery process, it will check the internal structure of the Active Directory database and adds the different logs. If the Windows NT/XP/2000 server restarts normally, the restoration has been successful.



Please note: The Restoration process, like the backup is a two-step procedure and will create a copy of the Active Directory files on the local hard disk drive. Please make sure there is enough free space on the drive to hold a complete copy of the Active Directory database and logs before attempting a restoration.

4.3 Restoring the Certificates with the “sysinfo/” plug-in.

The Windows Certificates should be restored while the Certificates service itself is off-line. To do this, click on “Start > Control Panel > Administrative Tools > Services” and use the MMC snap-in to list the services and stop the Windows Certificates service.

If the Certificates service is not stopped when the restoration is launched, Arkeia version 5 will refuse to restore the data to the Windows client.

Like noted above, the restoration of the Windows Certificates may require a large amount of disk space to be available on the Windows client machine. Please make sure there is enough free disk space on the client before attempting to launch the operation.

Once the restoration has been completed, the Certificates service can be restarted through the Microsoft Management Console by using the procedure detailed above.

Once restarted, the Certificates service will recover the data from its restored database. Once this recovery is completed, a new Certificates backup should be performed immediately, as the recovery process itself may modify a lot of information within the Certificates log.

4.4 Restoring the Event Logs with the “sysinfo/” plug-in.

The Windows Event Logs can be restored to the Windows client of Arkeia without any special manipulation, simply by selecting the correct files in the Restoration Navigator of Arkeia (see above for more information). The files are restored to the temporary directory used for the backup (see above for more details).

However, the Windows Event Logs, once they have been restored to the Windows client, should be specifically opened within the Windows Event Viewer, in order to be displayed. Otherwise, the Windows Event Viewer will only open the standard Windows log files and NOT the restored files.

To open the restored Windows Event Logs, open the files ending with the extension “.evt” in the directory named:

```
$ARKEIA_DIR/tmpreg/arkeia_eventlog_backup
```

Replace “\$ARKEIA_DIR” with the name of the directory that contains the Arkeia Windows client.

INDEX

- *Symbols*

.tar.gz 22, 24, 25, 27, 147

/dev/cdrom 23, 25

/dev/nst 41

/dev/pthru 13

/dev/rmt 12

/dev/rmt/tps 13

/dev/rsst 13

/dev/scsi 13

/dev/sg 13, 41, 162, 163, 188

/dev/sg0 11

/dev/st 12, 41, 179, 180, 181, 183, 187, 189

/dev/st0 11

/etc/hosts 9, 233, 234

/etc/hosts.cfg 233

/etc/nsswitch.conf 233

/etc/sysctl.conf 17

/mnt/cdrom 23, 25

/opt/arkeia/arkeiad/admin.cfg 232, 236

/opt/arkeia/arkeiad/arkeiad.cfg 234, 236, 239

/opt/arkeia/arkeiad/rhost.lst 232

/opt/arkeia/bin/ 184

/var/sysgen/stune 18

- **A**

AIT 37, 274, 275

AIX 7, 12, 13, 16

arkc 7, 23, 27, 63, 191, 192, 193

arkdetect 44, 184, 186, 187, 188, 189, 190

Arkeia backup server 10

arkeia.cfg 241

arkeiad 234, 236, 238

arkeiad remove 236, 238

arkeiad start 236, 238

arkmigrate 269, 270, 271, 274

arkmvhost 198, 200

arkpmysql 242, 243, 246

arkpmysql.prf 243, 244

arkvlib 167, 205

Authorizations 48

Auto-detection 46

auto-detection 44, 210

autoloader 11, 12

autoloaders 10, 11

- **B**

backup 272

Backup configuration 266

backup configuration 270

Backup device 266

backup device 11, 36, 42, 52, 155

backup hardware 37, 38, 39, 40, 42, 44, 46, 52, 56, 184, 185, 187, 189

backup history 266

backup peripheral 11

backup server 270, 271, 272

BIOS 10

bottleneck 10

BSDi 7

- *C*

CD-ROM drives 10

Change password 101, 102

Client 7

client 9

Client/Server 52

client/server architecture 9

Comment 48, 76, 81, 93, 113, 162

Confirm password 102

Control device 162

control device 13, 163

Create pool 81

Creating a drive 46

- *D*

DAT 37, 48, 191

DB2 248, 249, 250, 251

DB2 Admin Server 250

DB2 Administration Server 250

DB2 database 248, 249, 250, 251, 252, 253

DB2 instance 253

DB2 plug-in 248

DB2 processes 251

db2 terminate 251

db2_backup 262

db2_backup.sh 256, 262, 263

db2_backup_cat.sh 257, 263

db2_check.sh 260, 262

db2_pipe 262

db2_restore.sh 258

db2_restore_cat.sh 259

DB2STOP 251

db2stop 251

DDS 37

DEC OSF 14

default password 29

default pool 76

Detect 44, 185

Detected devices 44

DEVICE 189, 190

device 10, 11, 12, 13, 23, 25, 39, 159, 160, 162, 163, 168, 179, 180, 187, 188, 189, 190, 219

device detection 187

device drivers 42

device management 41

device number 41, 42

DGUX 7

DHCP 241

Digital UNIX 12, 13, 14, 15

Digital Unix 14

DLT 37, 183, 191

dmesg 42

DNS 9, 194, 195, 210, 221, 232

drivers 10, 42

Drives 47, 49, 159

drives 38, 47, 60

Drives management 47

DRSNX 7

DTF 37

- *E*

Exabyte 37

Exceptions 136

- *F*

f 189

f3lib 26, 275

f3tape 26, 275

File libraries 163, 168

file libraries 168

file library 168

firewall 222, 237, 274

Firewire 36, 41

flow 51, 57, 58, 98, 112, 117, 134, 170, 171, 174, 175

flow number 98, 171

FQDN 153, 221, 222, 234

FreeBSD 7

- ***G***

GNU 45

GNU-ZIP 24

grub 226, 227

GUI 7, 22, 23, 27, 28, 30, 46, 62, 63, 66, 72, 108, 161, 173, 184, 185, 189, 191, 213, 230

gui 7

gzip 24

- ***H***

hard disk 10

Hardware 9, 44, 46, 47, 69, 70, 74, 75, 77, 80, 161, 185, 217

Hardware checks 42

hosts 232

HP-UX 7, 12, 13, 15

- *I*

IBM 37, 42, 56, 209

IDE 10, 36, 209, 227

index 215

Installation 22

Interface 7, 36, 53, 62

interface 4, 7, 22, 23, 27, 41, 51, 53, 54, 63, 72, 77, 113, 172, 176, 191, 219, 220

interfaces 41

IP address 232

IP addresses 9

IPC 14, 15, 18, 19

IPC settings 14

IRIX 8, 18, 19, 56

Irix 12, 13

- *K*

kernel 11, 12

Kernel configuration 11

- *L*

LAN 52, 54

LANs 54

LIBRARIES 188, 190

Libraries 159, 185

libraries 10, 11, 39, 40, 46, 69, 71, 158, 159, 160, 162, 163, 168, 186, 187, 193

Libraries management 162, 163, 164, 165, 166, 185, 186

LIBRARY 188

Library 12, 39, 167

library 10, 11, 12, 13, 40, 41, 59, 159, 160, 161, 162, 163, 165, 167, 170, 184, 186, 189, 190, 191

Library Manager 205

Library manager 167

library manager 167

lilo 226

LINUX 17

Linux 8, 12, 13, 22

Linux device 23, 25

Listening ports 238

lspci 43

LTO 37, 42, 274

LUN 10, 11

- *M*

Mammoth 37

man 23, 24, 25, 44, 95, 212

man rpm 22

Manage drives 163, 164, 165

Mandrake 22, 25

MaxOS 8

media changers 11

Memory 9

Message queues 18

Microsoft Exchange 54

Microsoft Windows 54, 56, 93, 277

MODEL 189, 190

model 159, 190, 220

Mounting 12

mt 12, 44, 45, 179

Multiple LUN support 10

MySQL 242, 243, 244, 245, 247

MySQL database 246

MySQL password 245

MySQL plug-in 244, 246, 247

MySQL port 245

MySQL used ID 245

- *N*

Name 48, 162

network adapter 10

network adapters 9

Network bandwidth 9

New 47, 68, 74, 80, 82, 85, 87, 125, 127, 161

New password 102

NLP_HOSTFILE 232, 233

NLP_HOSTNAME 232, 234

NLP_LISTEN_MAX_PORT 239

NLP_LISTEN_MIN_PORT 239

NLP_NORMAL_MAX_PORT 239

NLP_NORMAL_MIN_PORT 239

NLP_ONLYHOSTFILE 232, 233

NLP_RESV_MAX_PORT 239

NLP_RESV_MIN_PORT 239

Node 48

Normal ports 238

Novell 8

- *O*

Old password 102

OpenUnix 8 8

Owner 48, 76, 81, 162

- *P*

password 17, 23, 25, 26, 29, 72, 73, 101, 102, 210, 223

periodic backup 266, 274

Periodicity 128, 129, 130, 131, 135, 136

Ping 9

plug-in 57, 83, 112, 122, 133, 153, 212, 213, 214, 215, 226, 228, 229, 230

Pool 60, 76, 110, 116, 134

pool 59, 76, 79, 80, 81, 82, 171, 182

Pool creation 80

pool list 80

Pool management 60

pool management 81

Pool name 81

pool name 182

Pools 82

Pools Management 82

Pools management 80, 82

PORT_NUMBER 236, 237

Primary Domain Controller 278

Priority 174, 175

probing 12

process 16, 18

proxies 274

- *Q*

QIC 37

- *R*

RAID 55

RAM 9

Red Hat 22

Reserved ports 238

Rewind device 48

roles 274

root password 23, 25, 29

root" password 99

RPM 22

- *S*

SAM 15, 16

SAM database 284

Savepack 280, 281, 283

Savepacks 51

SCO 8

Scratch Pool 59, 60, 76

Scratch pool 80, 82

SCSI 11, 36, 38, 41, 42, 43, 56, 168, 170, 209, 217, 227

SCSI adapter 10

SCSI adapters 10

SCSI bus 11

SCSI chain 42

SCSI Configuration 10

SCSI configuration 10

SCSI controller 10

SCSI device 11, 13

SCSI devices 11, 13

SCSI ID 11

SCSI interface 41

SCSI peripherals 10

SCSI tape 12

SCSI tape devices 13

SD-3 37

Semaphore configuration 16

Semaphores 18

Server 7

server 9

server capacity 9

Slackware 27, 89

SLR 37

Solaris 8, 12, 13, 19, 54, 56

SSH 239

ssh 239

Stackers 39

stackers 40, 46

Statistics 75, 81, 82, 163

statistics 75, 82

STRIP_DOMAIN 234

su root 23

SunOS 8

SuSE 22

sysdef 20

sysinfo 213, 214, 215, 226, 228, 229, 278, 280, 281, 282, 283, 284, 285

/config 212, 214

/dbase 215, 229, 230

/system 214, 228

- *T*

Tape 266

tape 58, 59, 60, 268, 271, 274

tape block size 274

tape changes 274

Tape device 12

tape device 11, 12

tape drive 10

tape drives 10

Tape Pool 59, 60, 110, 134, 135, 168

Tape pool 59, 60, 79, 107, 113, 123, 134, 266

tape pool 58, 59, 79, 80, 81, 82, 116, 182, 211, 212

tape type 275

tar 22, 24, 25, 93, 180

TCP 236

TCP/IP 54, 55, 210, 221, 222, 223

total backup 274

Travan 10

troubleshooting 42

Tru64 8, 12, 13, 14, 15

Type 48, 76, 120, 121, 162

- *U*

UNIX 3, 7, 22, 36, 41, 42, 44, 48, 54, 87, 93, 95, 97, 99, 152, 153, 161, 189, 205

Unixware 8

Unixware 7.1 8

USB 36

- *V*

VENDOR 189

Vendor 42, 43

VIRTUAL_HOSTNAME 241

VXA 37, 43, 189

- *W*

Windows 8, 277

Windows 2000 8, 278, 283

Windows Active Directory 278, 281, 282

Windows Certificates 278, 282, 285

Windows Client 284

Windows client 274, 277, 278, 279, 280, 281, 282, 285

Windows command 282

Windows computer 279

Windows Event Log 278, 283, 285

Windows Event Viewer 285

Windows help file 282

Windows information 283

Windows Machine 284

Windows machine 272, 280, 281

Windows NT 8, 278

Windows operating system 278

Windows plug-in 277, 278, 279

Windows Registry 274, 278, 281, 282, 283, 284

Windows registry 284

Windows server 278

Windows system information 284

Windows XP 8, 278, 283

- *X*

xarkeia 27