



User Manual



Foreword

Thank you for purchasing Arkeia.

This software should be used in accordance with the terms and conditions of the following License Agreement.

This License is applied for all the network machines running Arkeia.

0.1. KNOX SOFTWARE LICENSE AGREEMENT

In order to preserve and protect its rights within the framework of currently applicable legislation, Knox Software Corp. and Knox Software SA, hereinafter referred to as KNOX, does not sell rights to this SOFTWARE, but grants the right to use this SOFTWARE, within the terms of this license agreement, hereinafter referred to as LICENSE AGREEMENT, and expressly retains ownership rights to all KNOX SOFTWARE. If you do not agree with all the terms and conditions of this LICENSE AGREEMENT you can obtain a refund by returning the SOFTWARE, all its manuals, its documentation and the original sealed license envelope, in salable condition, to the place you obtained them.

1. GRANT OF LICENSE. In return for payment of LICENSE fees included in the cost of the SOFTWARE and your commitment to comply with the terms and conditions of this LICENSE AGREEMENT as well as the limited warranty attached to, KNOX, the licensor, grants to you, the LICENSEE, the non-exclusive and non-transferable right to use the SOFTWARE on a single computer known as the backup server, here in after referred to as the SERVER, and its associated client computers, here in after referred to as CLIENTS, provided that the terms and conditions of the license are complied with.

If the SERVER or CLIENT on which the SOFTWARE is operated constitutes a system with several users, the LICENSE AGREEMENT shall apply to all such users without incurring additional costs.

KNOX reserves all rights that are not expressly granted to the LICENSEE.

2. **COPYRIGHT.** The beneficiary of the LICENSE is the owner of the magnetic media, or any other type of media on which the SOFTWARE is initially, or subsequently, recorded or stored. However, this License is granted on the express condition that KNOX retains copyrights to the SOFTWARE recorded on the original media as well as copyrights to all copies made, irrespective of the format and the media of said original media and said copies.

This LICENSE does not constitute a sale of the original SOFTWARE or of any copy thereof.

3. **REPRODUCTION RESTRICTIONS.** This SOFTWARE and the accompanying written materials are protected by copyright. Unauthorized reproduction of the SOFTWARE, including its modification, integration or inclusion in another software, or of the accompanying written materials is strictly forbidden. The LICENSEE is liable to legal sanctions for any copyright infringement caused or prompted by any breach, on the part of the LICENSEE, of the terms and conditions of this LICENSE AGREEMENT.

Subject to the above-mentioned restrictions, the LICENSEE is authorized to make one (1) backup copy of the SOFTWARE if said SOFTWARE is not copy-protected.

– Notice of copyright must appear on the backup copy.

4. **RESTRICTIONS OF USE.** The LICENSEE is authorized to physically transfer the SOFTWARE from one SERVER to another SERVER on condition that said SOFTWARE is completely and totally removed from the original SERVER. Electronic transfers of the SOFTWARE from one SERVER to another within a distribution network for the purpose of copying the SOFTWARE or the accompanying written materials are strictly forbidden. The LICENSEE shall not modify, adapt, translate, reverse engineer, decompile, disassemble or create written materials based on the SOFTWARE, and shall not modify, adapt, translate or write literature based on the written materials without the prior express written consent from KNOX.

5. **TRANSFER RESTRICTIONS.** No person whomsoever shall be authorized to operate this SOFTWARE without the prior express written consent from KNOX. Any beneficiary of a transfer thus authorized shall be bound by the terms and conditions of this LICENSE AGREEMENT and the limited warranty attached there. Under no circumstances shall the LICENSEE be entitled to transfer, convey, lease or sell the SOFTWARE, nor shall it be entitled to dispose thereof, in any manner whether temporary or permanent, except where otherwise expressly provided for herein.

6. **CANCELLATION.** This LICENSE AGREEMENT shall remain valid until its cancellation and shall be canceled, as a right without prior notice by KNOX should the LICENSEE fail to comply with the terms and conditions of this LICENSE AGREEMENT. In the event of cancellation, the LICENSEE shall immediately destroy all written materials and all copies of the SOFTWARE, including modified copies, where appropriate.

7. **MISCELLANEOUS.** This LICENSE AGREEMENT is governed by the laws of the State of California (USA) if the LICENSEE acquired the SOFTWARE in the USA with respect to KNOX, their successors and assigns. This LICENSE AGREEMENT is governed by the laws of the country of France if the LICENSEE acquired the SOFTWARE in any country except the USA, with respect to KNOX, their successors and assigns.

If you would like further information on this LICENSE AGREEMENT, please write to: KNOX SOFTWARE Corp.
1901 Camino Vida Roble – Suite 200 –Carlsbad CA 92008 – USA –

0.2. WARRANTY *

KNOX warrants its SOFTWARE for a period of ninety (90) days as of the date of delivery thereof. This warranty also includes reconditioning or replacing SOFTWARE media.

KNOX does not warrant and does not enter into any commitments regarding the content of the documentation and the software. KNOX further disclaims any implicit warranties tied to the sale of the right to use license of this SOFTWARE with respect to its quality, its results, its merchantability or its suitability for a particular purpose. Consequently, the license to use this SOFTWARE is granted “as is”, without any promise being made.

In the event of a defect in the software or in the documentation, the LICENSEE, and not KNOX, its dealers, distributors, agents, or employees shall bear all costs needed for servicing, repair or correction.

Under no circumstances shall KNOX, or anyone else participating in the design, production and delivery of this SOFTWARE, be liable for any damages, whether direct, indirect, secondary or incidental, including, but not limited to, damages caused by loss of profit, business interruption, loss of information or any other loss, resulting from the use of this SOFTWARE, even if KNOX has been informed of the possibility of such damages.

Information or advice given verbally or in writing by KNOX, its dealers, distributors, agents or employees shall not constitute a warranty, nor affect in any way this warranty, and as such, the recipient shall not in any way depend on any such information or advice.

Arkeia and KNOX SOFTWARE are registered trademarks (TM) of Knox Software, All Rights Reserved. All other trademarks mentioned in this documentation are the property of their respective owners.

(*) Warranty may vary according to local regulations.

CHAPTER 1

Introduction

I. About this manual

I.1. Who should read this manual?

This manual was written for system administrators, engineers and technicians who need to install, implement and operate Arkeia.

It is assumed that the user of the software, as well as the reader of this manual, has already acquired:

- Basic knowledge of the UNIX operating system
- Basic expertise in ftp/rcp and telnet/rlogin
- Some knowledge of cabling and how to operate a TCP/IP network
- Knowledge of UNIX tape devices.

The installation process itself does not require any prior programming knowledge, in shell or in the C language.

Once Arkeia is installed, no specific expertise is needed in UNIX to execute the software.

This manual can be used by non–specialist users who will find clear explanations whenever special techniques are to be implemented.

I.2. How to use this manual?

I.2.a. Introduction

This manual has been designed and written to allow any administrator without prior knowledge of Arkeia to have a good overview of its capability as well as configure and adapt it to fulfill specific needs.

It does not cover the installation. For more information on the installation, please refer to the “*Quick start guide*”.

All the Arkeia manuals are available on your CD or on our web site <http://www.arkeia.com/manuals>

1.2.b. What is Arkeia?

Arkeia is a very powerful and highly flexible tool, integrating state-of-the-art features. The engineers and designers of Arkeia have chosen the best default configuration, offering the best backup/restore functions to all users, whether or not they are familiar with the software.

Arkeia has been designed for very large networks and also provides many options that can be modified to offer both power and flexibility. This introduction describes the possible applications of this product, provides an overview of its operation and describes the contents of the package you have received.

- The section “*Before you begin*” is very important and should be fully understood, as this will ease the installation. The vast majority of problems arise from the configuration of existing networks, as well as the interfacing with backup devices.
- The “*Quick Start Manual*” will also be useful for checking that the software is properly installed.
- This manual provides a trouble-shooting guide which will help you find the potential problems and finding the right solution. This information has been supplied by the Technical Support Department of KNOX SOFTWARE.

II. General Concepts and Features

II.1. Concepts

Arkeia has been designed around a number of simple concepts. It is important to have a good understanding of these concepts before starting any backup configuration.

II.1.a. Architecture

Arkeia uses a modular client/server architecture. Each module has specific tasks and may be installed independently of the others, though the installation order is important, as stated in the “*Quick Start Manual*”.

Most components are described in the next section: “*Features*”.

II.1.b. Structure

Arkeia is not just a backup utility, it is a complete Backup Manager.

It is designed to help you manage your drives, your libraries and your tapes. It tracks which tape is needed for a backup and which one is required for a restore. It informs you on the tapes available for backup and won't erase a backup tape if it determines that the data on it is still valid.

To achieve this, Arkeia requests you to create a complete backup procedure: you have to define your tape drives, organize your tapes in tape pools and identify each tape by writing a specific label on it. You also have to define your libraries and inform Arkeia which tape is inserted in a given slot. You must plan your backup policy and have enough tapes, real and logical, to ensure its proper functioning. You have to evaluate the amount of data saved at each backup.

Therefore, it is quite important to plan ahead your final backup configuration. This manual is designed to help you carry out this planning. Reading it carefully will help you solve most of the issues you may run into.

II.2. Features

Arkeia is a high performance network backup product that supports a wide variety of operating systems, hardware platforms, tape drives and tape libraries.

It offers a lot of advanced and optimized features , in order to maximize network backup performance and stability:

- Parallel backup of multiple network architecture (up to 200 clients simultaneously)
- Parallel usage of multiple tape drives (up to 32 tape drives simultaneously)
- Tape index with file history and version information (0.5 % of saved data in size)
- Remote administration and operation
- Sustained throughput of 6 to 60GB/hour (depending on network)
- Precise restoration (complete or selective)
- File search engine
- Fast file restoration
- Tape library management (up to 10,000 tapes)
- Security features for both access and operation
- Automatic barcode recognition
- Client / Server architecture
- Supports for advanced encryption with no loss of throughput
- Dynamic connection to active tasks

III. Arkeia Overview

III.1. Introduction

Arkeia is a centralized backup software, designed to handle data on heterogeneous platforms. Its exclusive use of on the fly, massively parallel backup technology combined with data compression on the client side provides an unrivaled backup throughput while ensuring maximum levels of reliability. Typically, the average throughput achieved is 1.4 times of the network bandwidth, or 70 megabytes (MB) per minute for a standard Ethernet network (sustained throughput over several hours).

Arkeia consists of the following main function modules:

- The backup server
- The client module
- An X11 Interface or a Java User Interface
- A command line interface

III.2. The backup server module

This module manages:

- The Arkeia configuration database and the centralized backup index
- The tape drive(s) and tape library(s)
- The network connectivity
- The multiplex data stream to/from client machines
- The multiple simultaneous backup/restore processes
- The interactive backups and scheduled backups
- The user authentication
- The smart module for periodical (scheduled) backups
- The Arkeia journal

III.3. The client module

This module manages:

- The data transmission from the clients to the backup server during backups
- The data reception from the backup server during restore operations
- The data communication to the “Navigator” interface
- The data compression and encryption during the backup/restoration
- Access to the native file system of the client

- The integration with Open File Manager (OFM), by St. Bernard Software. OFM allows the open files to be backed up correctly (for Windows–based machines only)

During the backup and restoration operations, the client sends/receives data flows to/from the backup server in an open format, which contain general file information (for instance: the filename, the file size, data stream) as well as specific client data, for instance the security descriptors and the registry on Windows NT/2000, Trustees and extended attributes on Novell, or Access Control List (ACL) on HP/UX and Solaris.

In contrast with backup systems based on “adaptations” of the “*tar*” or “*cpio*” formats, which are antiquated and quite inflexible, this features allows:

- Restorations that provide rigorously identical files on the same OS, while respecting special functions and capabilities
- Restoration that minimizes the data loss on environments which have less features than the original operating systems (e.g. loss of UNIX file group or owner for a file restored on Windows)
- Easy backup and restoration of system features (Windows registry)
- Extensibility, compatibility and interoperability with different operating systems

Therefore, the long–term investment in this backup solution is guaranteed, not by “*patches*” but by its architecture and design.

III.4. The X11 graphical user interface module

This module is the interface between the backup server and the user. It may be installed on any machine equipped of the X11 interface, but is usually located on the backup server.

The interface is a group of X11 client programs that can be displayed on any X11 server (R4, R5, R6) or graphic workstation. Its characteristics (colors, fonts, images) can be fully configured for maximum user comfort. Its technology has been designed to prevent the excessive use of resources frequently found in X11/Motif applications. It provides unique functions such as animated icons, toolbar, context–sensitive menus, “*vu–meters*”, tips, and multiple languages. It manages:

- The backup administrator setup and operation screens
- The Tape drive definitions
- The Tape library definitions
- The Tape pool definitions
- The Tape definitions
- The Savepack definitions (logical backup group)
- The Scheduled backup definitions
- A simplified interface for desktop users who wish to restore their own files
- The interface for help–desk operators, to restore user files

III.5. The JAVA graphical user interface module

This module is another interface between the backup server and the user.

The interface is a group of JAVA clients, compiled for machines running Windows NT/2000, Windows 95/98 and Windows ME.

It can be fully configured to provide maximum user comfort. Its technology has been designed to prevent the excessive use of resources frequently found in Java applications. Just as the X11 GUI, the Java user interface provides unique functions such as animated icons, toolbar, context-sensitive menus, “*vu-meters*”, tips, and other functions.

III.6. The Arkeia command line Interface module (arkc).

The “*arkc*” utility is the Arkeia command line interface. It manages the Arkeia backup server and supports a large number of operations that are also supported by the graphical interface. The “*arkc*” command can be included in a script shell, and allows you to integrate scheduled actions to/from an another system tool.

III.7. Installation guide

A complete installation of Arkeia contains:

- One or more backup servers
- One or more clients
- One or more user interfaces

This section describes machines that operate simultaneously as backup servers and management servers, using the X11 interface. Arkeia has been designed to manage several backup servers on the same network. The operation of these servers can be controlled through several interfaces. Usually, a client should also be installed on the backup or management server (these servers should also be backed up!).

Follow these steps to perform installation:

- Install the client on a machine
- Install the backup server on the same machine
- Install the user interface on the backup server
- Configure the backup server (system settings)
- Run Arkeia, configure the NULL drive, and run a local backup test
- Install the other clients
- Test the installation thoroughly (access to clients and devices)
- Configure the software for its standard operation and backup procedure

The installation and configuration procedures for a network composed of five machines will take about one hour. This presupposes a properly managed TCP/IP network (coherent IP names and addresses), standard drives and libraries.

The installer must also know the “*root*” user password on all the machines, to be able to send files via ftp or rcp and to be able to connect using remote facilities via *telnet* or *rlogin*.

A physical access is required to install non-UNIX clients.

🔴 **Please note:** 99% of all installation problems arise because one of the above conditions has not been met.

III.8. Package content

Although the exact packages may vary, according to the products and the versions ordered, they usually contain the following items:

- One CD media containing the server components for Unix/Linux and a large variety of client components.
- A user’s manual
- The *Quick Start Guide*
- The serial numbers representing the licenses granted, and which must be entered using the graphical interface.

🔴 The media must be extracted into a temporary directory, or installed from the CD. A component usually contains five files:

- **The “INSTALL” program (in capitals):** this is the interactive installation program that must be launched from the current directory. It asks the user a series of questions to install the components and displays messages to the standard output during installation. This output can be saved into a file with the ‘-jlog_name’ option. The ‘-llog_level’ option can be used to increase the level of information contained in the log. This last instruction can accept values from “0” (silent installation, except for errors and questions) to 90 (extremely detailed), with the default set at 20. For example: `./INSTALL -j/tmp/journal -l30`
- **The “aiinfo.lst” file:** this is an ASCII file, which contains the default installation settings.
- **The “alias” file:** this is another ASCII file, which contains the list of names of the components
- **The “XXXtar.Z” file:** XXX depends on the component name. This file contains all the component files in a compressed “tar” format.
- **The “uncompress” program:** this program is used by “INSTALL” to uncompress data.

CHAPTER 2

Before you begin***I. Platform availability*****I.1. Available clients and servers**

Arkeia supports a wide variety of hardware and operating system platforms. Additional platforms are added regularly. Please refer to the Arkeia web site (<http://www.arkeia.com>) for the most latest information and versions.

OS Version	Server	Client	GUI	arke	ORACLE
AIX 3.2		YES			
AIX 4.x	YES	YES	YES		S2
BSDi 3.0, 4.0		YES			
DGUX AviiON		YES			
DRSNX 7 (ICL Intel x86)		YES			
DRSNX 7 (ICL Sparc)		YES			
FreeBSD 2.2.6 (Intel x86)		YES			
FreeBSD 3.2 (Intel x86)		YES			
HP-UX 11	YES	YES	YES	YES	S2
HP-UX 10	YES	YES	YES	YES	S2
HP-UX 9		YES			
IRIX 4		YES			
IRIX 5.3		YES			
IRIX 6.x	YES	YES	YES	YES	
Linux 2.x (Intel x86)	YES	YES	YES	YES	S1+S2
Linux 2.x (MIPS,Cobalt)		YES			

OS Version	Server	Client	GUI	arke	ORACLE
Linux 2.x (Alpha)		YES			
Linux 2.x (ARM,Netwinder)		YES			
MaxOS 4.2 (MIPS)		YES			
MaxOS 4.2 (PowerPC)		YES			
Novell 4.11 (Intel x86)		YES			
DEC Alpha Unix 3.2		YES			
SCO v5 (Intel x86)		YES			
Solaris 2.5, 2.6 (Intel x86)		YES			
Solaris 2.5, 2.6 (Sparc)	YES	YES	YES	YES	S2
SunOS 4.1		YES			
Compaq True 64	YES	YES	YES	YES	S2
Unixware 2.x		YES			
Windows 95/98		YES	YES		
Windows NT Server 4.0 (Alpha)		YES			
Windows NT Workstation 4.0 (Intel x86)		YES	YES		
Windows NT Server 4.0 (Intel x86)		YES	YES		
Windows 2000		YES	YES		

I.2. Oracle clients

Arkeia provides two modules to backup Oracle databases “on–line”.

- The first one (S1 in the table above) is “*Arkeia for Oracle RMAN*”.

This solution is an interface between Arkeia and Oracle’s RMAN. RMAN is the Recovery MANager, provided by Oracle – see <http://www.arkeia.com/oracle>.

- The second one (S2 in the table above) is “*Arkeia’s Oracle Assistant*”.

This assistant helps you to backup Oracle database online. Please contact sales@arkeia.com for more information.

II. Hardware requirements and prerequisites

II.1. Hardware requirements

Arkeia has the following minimum hardware requirements

- A computer capable of running your operating system.
- 64 MB RAM.
- SCSI tape drive.
- A SCSI tape drive is required for high speed positioning of the tape during restore operations.
- It is strongly recommended to connect the SCSI tape drive on a dedicated SCSI host adapter.

II.2. Prerequisites

II.2.a. Memory

Arkeia uses large amounts of memory, especially if the number of computers to be backed up simultaneously is important. Knox Software recommends a minimum of 128 MB of RAM. 256 MB of RAM are recommended on large networks.

II.2.b. Network cards

As a network based backup solution, Arkeia need a top quality Network Interface Card (NIC). Therefore, it is strongly advised to use supported devices from well-known manufacturers.

II.2.c. SCSI host adapters

Many errors that can be encountered with Arkeia come from the SCSI configuration. Please make sure the following items are checked before installing the software:

- *It is strongly advised that you plug your tape drive or your library on a separate SCSI adapter.*
- *Make sure your tape drives are not plugged in the adapter along with hard disk drives or CD-ROM drives.*
- *In any case, do not mix SCSI peripherals that are not of the same SCSI generation.*

Your SCSI board BIOS should be configured with the following options:

- Reconnect/Disconnect: Disable
- Multiple LUNs support: Enabled

II.2.d. Tape drives

Currently, Arkeia only supports **SCSI tape drives**. IDE tape drives are **not** supported.

Please note: some tape drives (Travan and others) exhibit very poor performance during the restoration operations, due to a lack of standard fast positioning features. These features exist, but are not supported by generic drivers. Specific features, which are only supported by the drivers supplied by the manufacturer of these drives are not currently supported by Arkeia.

III. Software requirements and Prerequisites

III.1. Requirements

III.1.a. Reliability

The backup server is the main component of the network and needs to be installed on a reliable machine if it is to function properly. Machine reliability can be further enhanced by the use of uninterruptible power supplies (UPS), by implementing mirrored disks or RAID disks, by controlling the physical access and by using the machine specifically for backup applications.

III.1.b. Available disk space

Installation needs a temporary space of 20 MB, depending on the modules selected (server, graphics or client).

The final directory, with all three modules installed, will take up a maximum of 30 MB.

III.1.c. Backup catalog/Index database

The backup catalog, or index of backups, is modified with each backup. For every megabyte of data backed up, Arkeia allocates an average of 5 KB in the catalogue. This average may increase when different trees are backed up, or decrease when identical trees are backed up.

By default, a minimum of 30 MB is required, which should be able to contain a list of backups of up to 6 GB.

It is strongly recommended to have enough inodes available on the partition where Arkeia is installed.

III.1.d. Workload

Arkeia offers high performances and many advanced functions to its users. However, since these functions make a very intensive (and parallel) use of all the components (memory, CPU, peripherals) of a server, it is highly recommended to dedicate the machine used as a backup server to backup-only tasks. Using a non-dedicated machine as a backup server may increase its CPU load while performing a large backup operation.

III.1.e. IP bandwidth

The backup server is the central point of the backup. Therefore, it must be positioned at a strategic node in the network. Arkeia uses a Client/Server architecture, which means an Arkeia server communicates, through the network, with the clients that should be backed up.

Arkeia provides data multiplexing, parallel device management as well as simultaneous access to backup machines, with compression on the client side. This design provides speeds that are on average 1.4 times faster than the overall network bandwidth.

III.1.f. ROOT account for installation

Before the software can be installed, you need to know the ROOT password for all the machines where Arkeia should be installed.

III.2. Prerequisites

III.2.a. Drivers

In general, Arkeia uses one device for standalone tape drives and two or more for libraries and autoloaders:

- The standalone tape drives are controlled through the standard SCSI tape drive device of the Operating System (for example, “/dev/st0” under Linux). You have to make sure that your kernel is compiled to support those drivers.
- The libraries and autoloaders use one standard tape device for each tape drive, and one generic SCSI device for the library medium changer (for example, “/dev/sg0” under Linux).

Tape devices

The tape drive installation is generally straightforward **once the kernel is correctly set to support such devices**.

The correct device can be detected using the following “mt” command and trying to connect to available tape devices:

mt -f [tape device] status

Depending on the Operating System you use, the devices can be the following (where * ranges from 1 to 9):

<i>Operating System</i>	<i>Tape device</i>
AIX	/dev/rmt*
Digital Unix / True64 Unix	/dev/rmt*h
HP/UX v10	/dev/rmt/*mb (b=Berkeley Mode)
Linux v2.X	/dev/st*
SGI Irix	/dev/rmt/tps*d3
Sun Solaris	/dev/rmt/*hb (b=Berkeley Mode)

An example of correct “mt” output is given below:

```
[root@betelgeuse office52]# mt -f /dev/st0 status
SCSI 2 tape drive:
File number=0, block number=0, partition=0.
Tape block size 1024 bytes. Density code 0x24 (DDS-2).
Soft error count since last status=0
General status bits on (41010000):
  BOT ONLINE IM_REP_EN
```

Figure 1: “mt” correct output

Libraries and Autoloaders

Usually, the relevant differences for the OS between a library and an autoloader are the following:

- Library: Tape drive and medium changer have different SCSI IDs.
- Autoloader: Tape drive and medium changer have the same SCSI ID but have different LUNs.

The tape drives of libraries and autoloaders are defined and detected as standalone tape drives.

The medium changer is generally a generic SCSI device. **You need to make sure that your kernel supports generic SCSI devices.**

Furthermore, as multiple LUNs can be involved, you have to make sure that your kernel probes all LUNs on an SCSI bus.

Standard Generic SCSI devices for various Operating Systems are (where * is a value from 1 to 9):

<i>Operating System</i>	<i>Medium changer device (control device)</i>
AIX	/dev/pthru* (provided by Knox Software)
Digital Unix / True64 Unix	/dev/b*t*l* (b=bus, t=target ID, l=lun)
HP/UX	/dev/scsi* or /dev/scsi/c*t*d* (c=instance, t=target ID, d=lun)
Linux v2.X	/dev/sg*
SGI Irix	/dev/rmt/tps*d3
Sun Solaris	/dev/rsst* (provided by Knox Software)

The correct device can be detected using the “stks” command provided with Arkeia and trying with available control devices:

```
stks -v -d[control device] -i
```

You’ll find an example of a correct “stks” output on the next page.

It’s **essential** that the inquiry type is “8” in the first part of the output and that the “s_getstkcapp” shows the correct number of slots, drives and medium changers (See next page for correct output).

A correct output must also show one ITEM entry for each slot, drive and medium changer. There must not be any SENSE keys, otherwise it means either a firmware issue or a SCSI problem.

III.2.b. Network

As a network backup solution, it is obvious that Arkeia can only run properly on a well-configured TCP-IP network.

Check your network, verify its IP addresses and machine names, entries in the file /etc/hosts and in DNS. Ping back and forth all machines from the future backup server on both IP addresses and domain names. Make sure you obtain a correct IP and name resolution.

- ✿ If you run Arkeia (server and/or client) on a machine with multiple NIC, please refer to Chapter 4, “Arkeia Initial Configuration”, in the “*Specific Name Resolution and servers with multiple Network Interface cards*” section, for more information.
- ✿ If you plan to backup a multi–domains network, please refer to the Chapter 4, “Arkeia Initial Configuration”, in the “*How to configure Arkeia with multi–domains network architecture*” section.

III.2.c. SCSI support in kernel

You need to make sure that your kernel support the following options:

- SCSI Tape Drive support
- Generic SCSI support
- Probe all LUNs support

Figure 2: Correct output of stks:

```

=====
 stks on device /dev/sga
Dumping stacker status in file '-'
s_open('/dev/sga') OK
s_open done
s_inquiry:entering
s_inquiry:continuing
s_inquiry:result=0x0 addtl length=51 NO SENSE
s_inquiry:type 8:'EXABYTE EXB-210          3.113.11.013
inquiry:type 8:'EXABYTE EXB-210          3.113.11.013
sense: NO SENSE
Type=3

getelst: entering
s_test_unit_ready: result=0 NO SENSE
s_getstkcapp: result=0 [0]=0x17 [1]=0x00 NO SENSE
s_getstkcapp: 2 drive(s) (@@2), 11 slot (@@), 1 arm(s) (@@6)
fillelmt: entering with type: 1
s_test_unit_ready: result=0 NO SENSE
s_read_element_status: entering with
      from: 86,number: 1,type: 1
s_read_elm_stat(from:86,nb:1,type:1,voltag:1) res=0x0 NO SENSE
getelst: asked:1, read:1
getelst: Pvoltag:-1
fillelmt: entering with type: 4
s_test_unit_ready: result=0 NO SENSE
s_read_element_status: entering with
      from: 82,number: 2,type: 4
s_read_elm_stat(from:82,nb:2,type:4,voltag:1) res=0x0 NO SENSE
getelst: asked:2, read:2
getelst: Pvoltag:-1
fillelmt: entering with type: 2
s_test_unit_ready: result=0 NO SENSE
s_read_element_status: entering with
      from: 0,number: 3,type: 2
s_read_elm_stat(from:0,nb:3,type:2,voltag:1) res=0x0 NO SENSE
getelst: asked:11, read:3
getelst: Pvoltag:-1
s_read_element_status: entering with
      from: 3,number: 3,type: 2
s_read_elm_stat(from:3,nb:3,type:2,voltag:1) res=0x0 NO SENSE
getelst: asked:8, read:3
getelst: Pvoltag:-1
s_read_element_status: entering with
      from: 6,number: 3,type: 2
s_read_elm_stat(from:6,nb:3,type:2,voltag:1) res=0x0 NO SENSE
getelst: asked:5, read:3
getelst: Pvoltag:-1
s_read_element_status: entering with
      from: 9,number: 2,type: 2
s_read_elm_stat(from:9,nb:2,type:2,voltag:1) res=0x0 NO SENSE
getelst: asked:2, read:2
getelst: Pvoltag:-1
NB_SLOTS 11
NB_DRIVES 2

ITEM {
  NAME "P"
  STATUS "EMPTY"
}
ITEM {
  NAME "1"
  STATUS "FULL"
  LAST_IN "A"
}
ITEM {
  NAME "2"
  STATUS "FULL"
  LAST_IN "A"
}
ITEM {

```

IV. Platform Specifics

IV.1. General information

Arkeia uses shared memory and message queuing for inter process communication (IPC) in order to ease and optimize parallel flow management.

All the values given are average values, which enable Arkeia to function properly in a conventional environment. When the backup server is also used as the database server (ORACLE), the figures should be increased.

IV.2. Configuring IPC (shared memory and message queue)

IV.2.a. COMPAQ TRUE64 UNIX / Digital Unix DEC OSF

- Edit the “/etc/sysconfigtab” file

```
cd /etc
vi sysconfigtab
```

- go to the “ipc” section (create it if it does not exist) and enter:

```
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) and enter:

```
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
```

Changing the 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
```

Viewing IPC and process settings

- Enter these commands:

```
/sbin/sysconfig -q ipc      to view “ipc” configuration
/sbin/sysconfig -q proc    to view “process” configuration
```

IV.2.b. Hewlett–Packard HP/UX

Modifying IPC and process settings

The SAM management tool configures IPCs.

You may run SAM in text or graphics mode using the command: **sam**

Accessing IPC and process settings via the menu:

- Configuring messages queues:

<i>[Kernel configuration]</i>	<i>[Configurable Parameter]</i>
<i>[msgmap]</i> →102	
<i>[msgmax]</i>	8192
<i>[msgmnb]</i>	65535
<i>[msgmni]</i>	200
<i>[msgseg]</i>	2048
<i>[msgssz]</i>	32
<i>[msgtql]</i>	100

- Configuring the semaphores:

<i>[Kernel configuration]</i>	<i>[Configurable Parameter]</i>
<i>[sema]</i>	1
<i>[semaem]</i>	16384
<i>[semmap]</i>	66
<i>[semni]</i>	64
<i>[semmns]</i>	200
<i>[semmnu]</i>	30
<i>[semume]</i>	10
<i>[semvmx]</i>	32767

- Configuring the shared memory segments:

<i>[Kernel configuration]</i>	<i>[Configurable Parameter]</i>
<i>[shmem]</i>	1
<i>[shmmax]</i>	67108864
<i>[shmmni]</i>	200
<i>[shmseg]</i>	120

- Configuring the number of processes:

<i>[Kernel configuration]</i>	<i>[Configurable Parameter]</i>
<i>[nproc]</i>	316

To modify each of these parameters, use the menu

[Kernel configuration] / [Configurable Parameter] / [Action] / [Modify configurable parameter...]

Changing the settings

To have these settings take effect, use the menu

[Action] / [New Kernel]

The machine will reboot automatically to execute the new kernel.

Viewing IPC and process settings

IPC can be viewed using the SAM management tool.

You may run SAM in text or graphics mode using the command given here: *sam*

IV.2.c. IBM AIX

There are no specific settings for the AIX system (3.2 and 4.1)

IV.2.d. LINUX

There are no specific settings for Linux.

IV.2.e. SGI IRIX

Modifying IPC and process settings in kernel 6.4 and lower

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

<i>/var/sysgen/mtune/shm</i>	for IPCs
<i>/var/sysgen/mtune/sem</i>	for semaphores
<i>/var/sysgen/mtune/msg</i>	for message queues
<i>/var/sysgen/mtune/kernel</i>	for the process number

- Should you need to change these values, create the following file “*/var/sysgen/stune*” (if it does not exist):

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

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

```
shmmax=0x20000000  
shmmmin=1  
shmmni=100
```

```
shmseg=512
shmall=512
```

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

```
semmni=10
semmns=60
semmnu=30
semmsl=25
semopm=10
semume=10
semvmx=32767
semaem=16384
```

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

```
msgmax=8192
msgmnb=65535
msgmni=64
msgssz=8
msgtql=1000
msgseg=1536
```

- Configuration processing is automatic.

Changing the settings

Enter the command: */usr/sbin/autoconfig*

Then reboot the machine with the command: **reboot**

Viewing IPC and process settings

Enter one of the following:

```
/usr/sbin/systune | grep shm    to view ipc configuration
/usr/sbin/systune | grep proc   to view process configuration
/usr/sbin/systune | grep sem    to view semaphore configuration
/usr/sbin/systune | grep msg    to view message configuration
```

Irix Kernel 6.5.X

For local backups to run with good performance, follow this procedure:

Change the following parameters in the kernel:

```
tcp_recvspace = 184320
```

```
tcp_sendspace = 61440
```

in the file `/var/sysgen/master.d/bsd` or `/var/sysgen/mtune/kernel`. It is possible to list and modify dynamically these parameters with the command: `systune`.

IV.2.f. Sun SOLARIS

Modifying IPC and process settings

- If you need to change these values, create the file “**/etc/system**” (if it does not exist)

```
cd /etc
```

```
vi system
```

- Insert the following lines (if they do not 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 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
```

Changing the settings

Reboot the machine by entering the command: **reboot**

Viewing IPC and process settings

Enter the following command: *sysdef -i*

CHAPTER 3

Arkeia's conventions

I. Convention used for commands and keys

I.1. Graphical User Interface (GUI)

Arkeia Version 4.2 has two graphical user interfaces:

- XWindow [X11] (UNIX).
- JAVA (for Windows).

All the screens of Arkeia, regardless of the graphical interface used, have the same layout:

- A drop-down menu bar (top of the screen).
- A tool bar (icons at the bottom of the screen showing the most commonly used functions).
- Contextual menus (different menus appear according to the position of the cursor when it is clicked).
- Help tips appear when the mouse pointer is on an icon.
- Keyboard shortcuts
- Contextual help button.



I.2. Function keys

Arkeia is fully operational with a mouse or with the keyboard.

Hit:	For:
<F1>	Help
<F2>	Drop-down menus
<F3>	Ends the task being processed. (Replaces the “Quit” command).

I.3. Keyboard shortcuts

You can access shortcuts by pressing ALT and a letter. For instance: ALT+B for backup.

From the navigator, you can enter the first letters of the name of a file or directory to jump to that file or directory.

I.4. Arrow key

Use the arrow keys to move forward, backward, up and down, and confirm by pressing the “Enter” key.

I.5. Context-sensitive menus

You can access a context-sensitive menu by clicking with the right mouse button.

This menu will change, based on the position of the mouse, when you click on the right mouse button



I.6. Tool bar buttons

The toolbar buttons allows you to access the main functions of the current window.

Move your mouse cursor over a button to see a contextual help (tip) appear.



I.7. Context-sensitive help

Click on the “Help” (?) button (lower right) or hit the F1 key to display the help screen.



Context-sensitive help means that you can first click on a zone or a field then click on the HELP button, or hit the F1 key, to get further information.

I.8. Copy/paste with the mouse

You may use the mouse to copy/paste in the editing fields. With the left mouse button pressed, select the text you want to copy. To paste the selected text, move the mouse pointer to the desired insertion point and press the middle mouse button.



Arkeia Initial Configuration

I. Drives and Devices

I.1. Introduction

Defining a “*device*” allows Arkeia to obtain the characteristics of a peripheral (single tape drive or library) and to assign it a logical name. The backup peripheral devices must be connected to the backup server. They must be declared by specifying the name of the (device) driver for each one.

Every tape drive (whether a single drive or a drive configured within a library) and every loading robot needs a logical input. Therefore, drives that are physically linked to a robot should be connected logically via Arkeia.

🔴 Please make sure you read and understand this section before proceeding. A summary of the possible devices, listed by operating systems, can be found in chapter 2 : “*Before you begin*”.

I.2. Drive management screen

From the main screen click on the [Devices] menu then select the [Drives management] option.



You can also use the “Drives management” button in the Toolbar.



Name:

Name of the drive definition

Type:

Type of device

Usage:

Device operating time

Bef. clean:

Time remaining before cleaning

of loads:

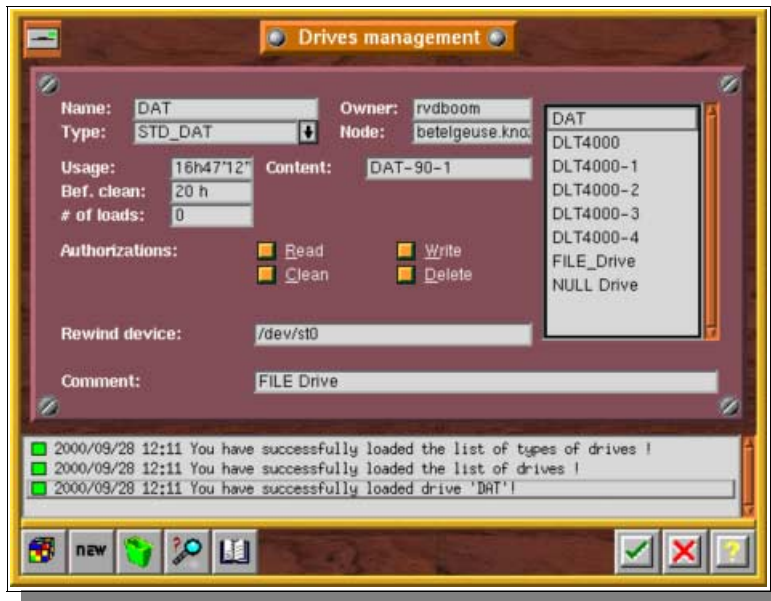
Number of times the tapes were loaded

Authorizations:

Actions allowed on the drive: read/write/clean/delete drive

Rewind device:

Name of the rewindable device



Owner:

Name of device creator

Node:

Name of backup server to which the drive is connected

Content:

States whether drive is empty or not (label displayed)

Library:

Name of robot to which drive is “connected” (displayed only if a library is defined)

Drive num:

Drive number in the robot (displayed only if a library is defined)

I.3. NULL drive creation

I.3.a. What is a NULL drive and what are its uses?

The “NULL” device type simulates a tape device but data cannot be written to it. As a result, the restore function is impossible.

This device allows you to quickly test your backup configuration without using a tape and without updating the Arkeia database. Nevertheless, connections can still be established between clients and server like in a real backup. Thus, it is a good test of network configuration.

On Unix systems, the control device is “/dev/null”.


Please note: When using a NULL type, we recommend that you configure it for “Write” and “Delete” authorizations only.

I.3.b. Drive creation (NULL drive)


Click the “New” button in the Toolbar of the “Drives Management” window.



Give the drive a name of your choice.



Select the drive type (“NULL”) in the drop–down list:



Set the authorizations on this drive.

Please note: this is usually done automatically.



Enter the full path to the NULL driver.



Confirm your choices by clicking on the “checkmark” (OK) button.



I.4. Tape drive creation

I.4.a. What is a tape drive?

As previously stated, you need to configure a separate tape drive in Arkeia for each tape drive that you plan to use for your backups.


🔴 Please refer to the “*Before you begin*” section to learn how to detect the correct devices for your drives.

I.4.b. Drive creation


Click the “New” icon in the Toolbar of the “Drives Management” window



Give the drive a name of your choice



Select the drive type from the list:



Set the authorizations on this drive.

Please note: this is usually done automatically.



Enter the full path to the driver

The detection of the correct tape device is described in the “*Before you begin*” chapter.



Confirm your choices by clicking on the “checkmark” (OK) button




1.4.c. Reading the tape label

This functionality is used to manually read the label of a tape present in a single tape drive. (Don’t use this function if you have a library). You can also use this function to validate your rewindable device.

Select the drive that contains the tape you want to read.

To read the tape, click on the right mouse button and select the drive.



1.4.d. Possible messages

The “Read Tape” option can give the following messages:

Message	Explanation
Drive is empty	There is no tape in the drive
Can’t read device	There is no label on the tape, your device is malfunctioning or your device is not of the correct type
Success + Unknown TPID in the content zone	The tape contains an Arkeia label but this label cannot be found in the database
Success + tape label in the content	The tape contains a label, which can be found in the database

1.5. File drive creation

1.5.a. What is a File drive and what are its uses?

If you plan to make backups to block devices (Hard drives, Floppy disks, Zip disks, Jaz and MO disks), then you have to create a File drive.

File drives and “tapes” are the only way Arkeia can backup on mass-storage, non-tape devices. A drive of this type can simulate a real tape drive on block devices.

- 🔴 The File-type “tapes” need to be defined first in order to use this type of device. Please refer to the “Where to Backup” chapter. When defining these “tapes”, state the disk location of the data you need to have backed up.
- 🔴 File drives must be attached to a File library to be used.
- 🔴 This feature is only available on the Professional Version of Arkeia.

I.5.b. Drive creation (File drive)

Click on the “New” button in the Toolbar of the “Drives Management” window



Give the drive a name of your choice



Select the “FILE” drive type from the list:



Set the authorizations on this drive



Confirm your choices by clicking on the “Checkmark” (OK) button



I.6. Drive deletion

Select the drive you want to delete in the “Drives management” window.



Click on the “Trashcan” button to delete the drive.



Confirm the drive deletion.



Confirm the deletion by clicking on the “checkmark” (OK) button



I.7. The Library Management Screen

From the main screen click on the [Devices] menu then on the [Libraries management] option.



Or click on the “Libraries management” button in the Toolbar:



Name:

Logical name of the library

Owner:

Name of the library creator

Node:

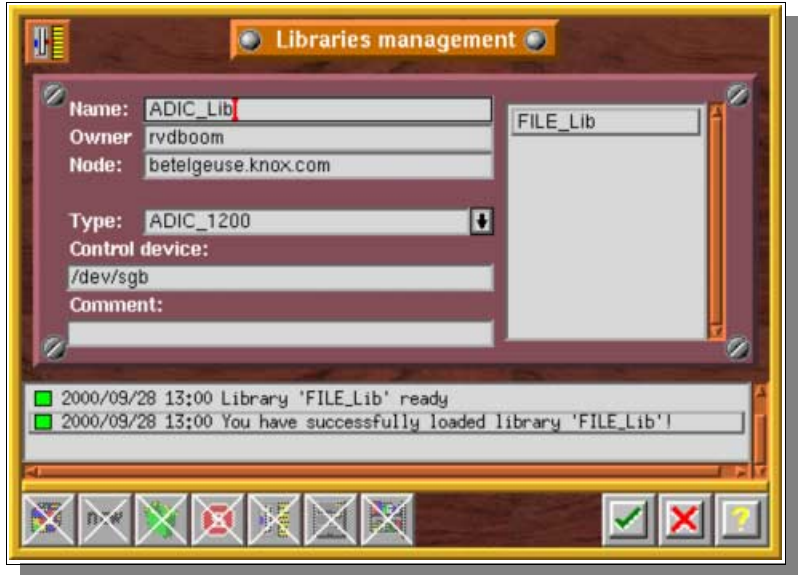
Name of the backup server to which the library is connected

Type:

Type of device

Control device:

Name of the control device for the operating system on which Arkeia is installed



I.8. Tape Library creation

I.8.a. What is a Tape Library?

A tape library is a robot that can automatically load and unload tapes to/from tape drives. It’s sold as a complete device that include one or more drives, a media changer and several slots where tapes are inserted.

For Arkeia, libraries are defined by their logical devices, one for each tape drives and one for the media changer.

The drives in a library are created in the way described above. The library itself is created through the “*Libraries Management*” screen.


🔴 Check the “*Before you begin*” chapter to detect the correct devices for tape drives and medium changer.

I.8.b. Library creation

Click on the “New” button in the Toolbar of the “Libraries Management” window



Give the library a name of your choice.



Select the library type from the drop-down list.



Enter the full path to the library device

The detection of the correct library device is described in the *“Before you begin”* chapter.



Confirm your choices by clicking on the “Checkmark” (OK) button



To attach the drives, click on the “Drive options” button in the Toolbar of the “Libraries Management” window



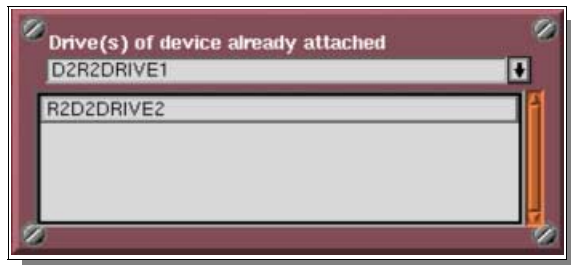
Select one of the drive slot named “No drive attached”



Click on the “Attach drive” button in the Toolbar of the “Drive options” window



Double-click on the drive of the Library you created in “Drives Management”




Confirm the drive attachment by clicking on the “Checkmark” (OK) button.



Confirm the library creation by clicking again on the “Checkmark” (OK) button



 You will also need to configure the tapes available in the library slots to be able to use your library. This is described in the following *“Where to backup”* chapter, in the *“Configure a tape library”* section.

I.9. File library creation

I.9.a. What is a File library?

A File drive can only be used when attached to a pseudo-library. This is actually convenient, since it is possible in this way to simply add a “File” drive to a File library if there is not enough space to complete a backup operation.

The creation of such a library proceeds exactly like a tape library creation, using File drives and “tapes”.

✿ Make sure you already have created the File tapes (see the “Where to backup” chapter) and the File drive before creating your File library.

I.9.b. Library creation (File library)


Click on the “New” button in the Toolbar of the “Libraries Management” window



Give the library a name of your choice



Select, for the library, the “File” type from the drop-down menu



Confirm your choices by clicking on the “checkmark” (OK) button



To attach a File drive to the File library, click the “Drive options” icon in the Toolbar of the “Libraries Management” window



Select one of the drive slots labeled “No drive attached”



Click on the “Attach drive” button in the Toolbar of the “Drive options” window



Double-click on the File drive you created in “Drives Management”



Confirm the drive attachment by clicking on the “checkmark” (OK) button.




Confirm the library creation by clicking on the "checkmark" (OK) button a second time.



🔴 You will also need to configure the tapes available in the library to be able to use it. This is described in the following chapter “Where to Backup”, in the “Tape creation” section.

I.10. Library deletion


Select the drive you want to delete in “Library management”.



Click on the “Trashcan” icon to delete it



Confirm the library deletion.



Confirm your choices by clicking on the “Checkmark” (OK) button



🔴 You will also need to remove the tapes available in the slots to be able to delete the library. This is described in the following chapter “Where to backup”, in the “Tape deletion” section.

II. Specific Configuration

II.1. Specific Name Resolution and servers with multiple Network Interface Cards (NIC)

II.1.a. Introduction

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.

II.1.b. Client machine configuration

On the client, the backup server name must be entered in the following file: */usr/knox/nlp/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*

II.1.c. Arkeia backup server machine configuration

Each client is declared in the following file on the Arkeia backup server: */usr/knox/nlp/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.

II.1.d. Hosts file used by Arkeia

By default, name resolution is based on the local system configuration. Generally, the resolution is done with:

1. The local host file (*/etc/hosts*)
2. The DNS (Domain name server)

Check */etc/nsswitch.conf* (and, on Linux, */etc/hosts.cfg*) 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:

1. 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.
2. 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.

II.1.e. Setting the 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 **/usr/knox/nlp/hosts.cfg**, in the following file: **/usr/knox/nlp/nlp.cfg**

Example	<code>NLP_HOSTFILE</code>	<code>"/usr/knox/nlp/hosts.cfg"</code>
----------------	---------------------------	--

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 “yes”.

Example	<code>NLP_ONLYHOSTFILE</code>	<code>"YES"</code>
----------------	-------------------------------	--------------------

Please note:

If the network does not use the “DHCP” (dynamic address allocation), it is not necessary to set the NLP_ONLYHOSTFILE variable, because if the FDDI network is not available, it will use the Ethernet network with /etc/hosts file. However, if DHCP is used, you must set the NLP_ONLYHOSTFILE.

II.1.f. 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.

Example	<code>"193.90.5.234"</code>	<code>"hp_test"</code>
----------------	-----------------------------	------------------------

II.1.g. NLP_HOSTNAME usage

If the NLP_HOSTNAME variable is uncommented, 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: **/usr/knox/nlp/nlp.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.

A short configuration example

The hostname command returns the host portion of the fully qualified name exactly as specified in the /etc/hosts file.

The table below shows the possible scenarios based on the comparison of the gethostname and the NLP_HOSTNAME.

hostname	gethostname	NLP_HOSTNAME	Comparison	Possible actions
hp67	hp67	hp67	Identical	Different IP addresses are possible
hp67	hp67	HP67	Different	Same IP addresses are mandatory
HP67	hp67	hp67	Identical	Different IP addresses are possible
HP67	hp67	HP67	Different	Same IP addresses are mandatory

II.1.h. Complete configuration example

The following example shows how to select a specific NIC when a machine has 2 NICs. This example uses a LAN, which contains two machines, each machine being equipped with 2 cards.

One card is a standard 10 BaseT and the other is a FDDI.

Arkeia backup server (rs6000)		Client machine (haley)	
hostname: rs6000		hostname: haley	
Address: 192.90.5.33 (10BaseT)		Address: 192.90.5.43 (10BaseT)	
/etc/hosts file:		/etc/hosts file:	
192.90.5.33	rs6000	192.90.5.33	rs6000
192.90.6.34	rs6000_fddi	192.90.6.34	rs6000_fddi
192.90.5.43	haley	192.90.5.43	haley
192.90.6.44	haley_fddi	192.90.6.44	haley_fddi

Here are the client and server configurations to use for FDDI network:

Arkeia backup server (rs6000)	Client machine (haley)
admin.cfg: rs6000	admin.cfg: rs6000
"NLP_HOSTNAME" "rs6000"	"NLP_HOSTNAME" "haley"
"NLP_HOSTFILE" "/usr/knox/nlp/hosts.cfg"	"NLP_HOSTFILE" "/usr/knox/nlp/hosts.cfg"
Its /usr/knox/nlp/hosts.cfg contains:	Its /usr/knox/nlp/hosts.cfg contains:
"192.90.6.34" "rs6000"	"192.90.6.34" "rs6000"
"192.90.6.44" "haley"	"192.90.6.44" "haley"

II.2. How to configure Arkeia with a multiple domains network architecture

By default, Arkeia truncates the domain name from each machine's name.

This can create problems if you plan to backup a network with multiple domain names on a single backup server.

To solve this issue, use the following procedure:

- Add/Uncomment the STRIP_DOMAIN parameter in the /usr/knox/nlp/nlp.cfg file, on the backup server as well as on the client machines:

STRIP_DOMAIN "0" (Don't strip domain name)

- Restart NLSERVD on all backup server and client modules:

NLSERVD restart

- Use the GUI to open the navigator, and check that all the clients appear with a complete domain name.
- Try a small backup, including a directory of each client machine.

II.3. How to use different TCP ports

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: */usr/knox/nlp/nlp.cfg*

For the JUI (Java User Interface), the file to edit is: *knox/arkeiaUI/login.prf*

Uncomment the line NLPPORTNUM and set it to the right port. Add the line “ARKJ_USE_PORTNUM” as shown below:

```
ARKJ_USE_PORTNUM          "0"  
NLPPORTNUM                "[New port number]"
```

Restart the **nlsvrd** process after modifying nlp.cfg, using the following command: *NLSERVD restart*

- ❖ **Please note:** if you plan to change the port value, you will also have to change it on all your clients. All the machines must use the same Arkeia port number, even those on which you only run the GUI or JUI.
- ❖ 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, ...).
- ❖ If the GUI/JUI is installed on another machine (administration server) than the backup server, it will talk to nlsvrd on the backup server directly on the first port available from 1024 and downward (1023, 1022, ...).

II.4. How to configure Arkeia to work from behind a firewall

II.4.a. Introduction

If you are using a firewall, you may have some connection problems between the backup server, the clients and the GUI/JUI if it is on another machine, because only specific ports are actually opened.

As stated above, Arkeia (more exactly the **nlsvrd** process) uses the TCP port **617** for backup and restore and for all operations or processes communications.

If you want to change the port value, use the information given in the “*How to use different TCP ports*” section.

II.4.b. Standard procedure

To get a backup working across the firewall, the port **617** (or any other port reserved for Arkeia) has to be open on the firewall, in both directions. For this, do the following:

- Allow connections on port 617 from the client to the backup server
- Allow connections on port 617 from the backup server to the client
- Remove "[1]" from `/usr/knox/nlp/auth_OPBS.cfg` on clients and backup server so it reads:

OPBS.* ALLOW * *

instead of:

OPBS.* ALLOW *[1] *

This allows the access to OPBS from a non reserved port.

II.4.c. 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 [your.firewall.address]`

which 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".

Short example

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 Server: 111.222.111.222 myclient.extdomain.com myclient
- Duplicate a block entry for "myclient" on the Server in `/usr/knox/nlp/rhost.lst` using its expected IP of 111.222.111.222
- Restart NLSERVD 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 NLSERVD
- Remove "[1]" from /usr/knox/nlp/auth_OPBS.cfg on client:

OPBS.* 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 aka 192.168.1.50). Any reverse traffic via port 617 are done through the generic internal-to-external MASQuerading.

CHAPTER 5

Where to backup?

I. The Tape Pools and Drivepacks concepts

I.1. Understanding the issues

Any backup administrator may want to reserve, in his backup policy, some of his tapes and drives for specific purposes.

He may, for example, want to use half of his tapes and drives to backup all the workstations on his network, while keeping the rest to backup the servers data, as he has a very large amount of data to be backed up on these servers.

A professional backup solution should allow the administrator to assign his resources as he sees fit. It has to be very flexible and allow multiple configurations.

I.2. Arkeia's approach

Arkeia has a very simple and straightforward approach to this problem. It asks you to organize your tapes and drives in "Tape Pools" and "Drivepacks". Please Note: this is mandatory. For instance, a tape must be created in a Tape Pool, while a drive should be configured in a Drivepack before they can be used for a backup or a restoration operation.

Basically, a pool is a group of tapes that are used together in the same backup strategy. In most cases, they are of the same type. Tapes are always created within a pool.

In the same way, a Drivepack is a group of drives that are used together in the same backup. One specific drive can be included in multiple drivepacks, according to your backup strategy.

Tape pools and Drivepacks are a complete part of Arkeia's structure and allow the Administrator to configure precisely how Arkeia is supposed to perform the backups.

II. Tape Pools

II.1. Definition and uses

A Tape Pool is simply a group of tapes that will be used for backups. The administrator can create a single pool or can dispatch his tapes into various pools, used respectively for a specific backup.

A Tape Pool will contain a certain number of tapes, created according to the next section, *Tapes*.

Though the Administrator has complete liberty over the way he creates his pools, we highly recommend to create a Tape Pool for each backup you plan to make, to avoid mixing problems that can occur after several backups have been performed. This is particularly true when the “Validity” of the backups are different. For more information, please refer to the “*Backup*” chapter of this manual.

Once a Tape Pool has been defined, Arkeia will automatically manage the tapes created in it.

II.2. The “Pools management” screen

From the main screen click on the [Tapes] menu then on the [Pools management] option.



Or use the “Pools management” icon in the Toolbar:



The following dialog box should now be displayed on the screen:



II.3. Pool creation

Click the “New” icon in the toolbar of the “Pools Management” window



Enter the name of the pool. Enter a comment.
The owner is the current Arkeia user.




Confirm your choice by clicking on the “checkmark” (OK) button.



II.4. Pool deletion

It is not possible to delete a Pool that still contains tapes. To do this, move the tapes to another Pool (in the menu “Tapes”, then the option “Pools management”) or delete them before deleting the tape pool.

Select the Tape Pool you want to delete in the “Pools management” screen.



Click on the “Trashcan” button to delete



Confirm the tape pool deletion



Close the dialog box by clicking on the “Checkmark” (OK) button



II.5. The Pool management window

The pool management screen allows you to manage the tapes in a Pool. Through this screen, you can check to see which tapes are used, which tapes will be used in the next backup, etc...

To see the tapes contained in a specific Tape Pool, double click on its name in the “Pool Management” screen or select it and click on the “Magnifying glass” button.



First column:

Tape names

Second column:

Free label or thread number

Third column:

Number of the tape in its thread



II.6. Thread and tape order

In a specific pool, for each backup that uses the “Always use a new tape” strategy, at least one thread is created (For more information, please refer to the “Backup” chapter). If you use more than one drive, several threads are created.

The next tape to be used with the strategy “Always use a new tape” is the first “free” tape, in the list of tapes.

Example If, in a backup, you have the following pool:

NAME	Thread	Usage order
tape4	free	00001
tape5	free	00002
tape1	001	00001
tape3	001	00002
tape2	002	00001

that means 3 tapes have been used (*tape1*, *tape2* and *tape3*) for two drives or two backups (thread number 1 and 2) and *tape1* is full because the *tape3* is on the same thread with a “used” value of 2.

The next tape to be used with the “new tape” strategy will be *tape4*.

II.7. Pool statistics

This screen allows you to view the statistics on a given Pool, as well as information on the tapes contained in that Pool.

To see the properties of a specific Tape Pool, double click on its name in the “Pools Management” screen or select the pool and click on the “Graph” button.



Backup Tapes

Total:

Number of Tapes currently in the Pool

Free:

Number of free Tapes in the Pool

In use:

Number of Tapes in use in the Pool

Full:

Number of filled–up Tapes in the Pool

Worn–Out:

Number of Tapes that should be replaced

Pool Space

Total:

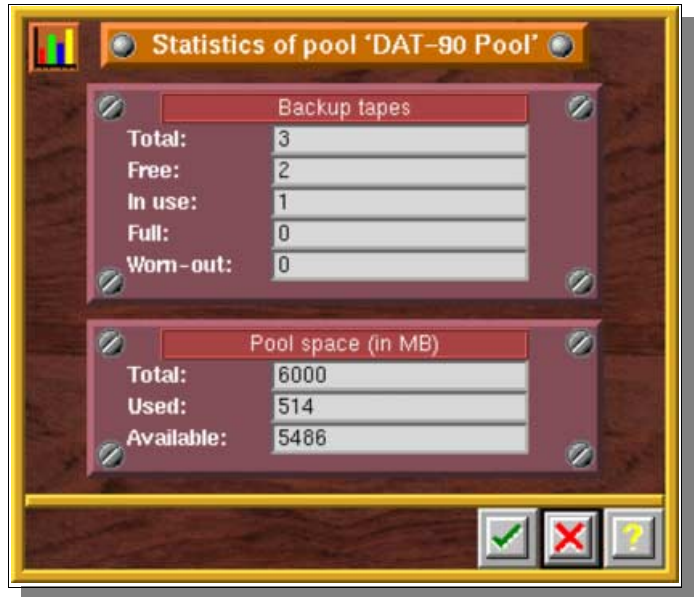
Total tape space available in the Pool (in megabytes)

Used:

Current space occupied by data

Available:

Remaining free space for backups



II.8. The Scratch Pool

There is a pre–defined “Tape Pool” in Arkeia named the *Scratch Pool*.

This pool can be used to contain tapes for other pools. Once a pool has run out of free tapes, it can get a new tape from the *Scratch Pool*, which will be temporarily assigned to it.

This is particularly useful when you have a large number of tapes and you can’t really keep up with tape management. It can be a policy to create tapes only in the *Scratch Pool* and to have other tape pools use these available tapes as needed.

See the “*Tape Recycling*” section and the “*Periodic Backup Policy*” chapter for more information on the Scratch Pool.

III. Tapes

III.1. Introduction

Arkeia will usually make a backup on a specific Tape Pool, using all the tapes in it sequentially, according to a tape strategy. This means you’ll have to create enough tapes in your tape pool to fulfill your backup policy.

The tapes created in Arkeia will match your real tape reserve, each tape will be labeled and clearly identified.

III.2. The “Tapes management” screen

This screen enables the Administrator to create, delete, modify and consult tapes used for the backup and restore.

From the main screen click on the [Tapes] menu then select the [Tapes management] option.

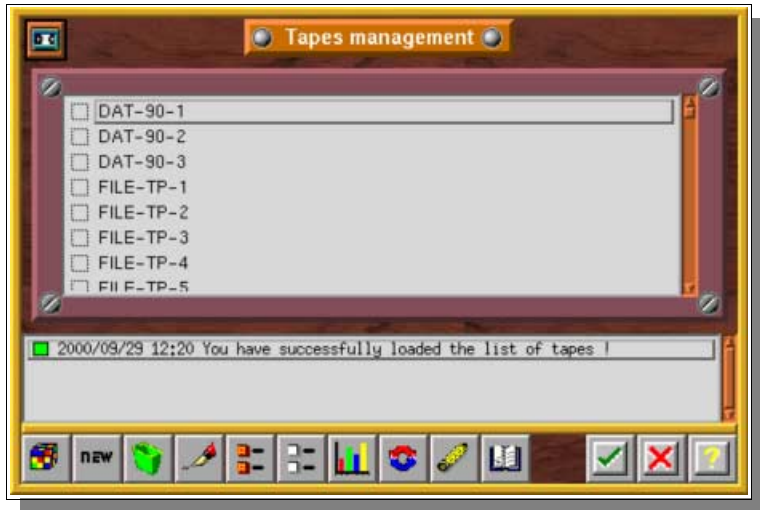


Or use the “Tapes management” icon in the Toolbar:



This window is essentially a list of the available tapes, all Tape Pools included.

Various operations can be run on tapes by selecting at least one tape then clicking on the icons of the toolbar.



III.3. Tape creation

III.3.a. Introduction

Tapes are created in Arkeia to link the logical names used by tape pools and the physical tapes in the devices.

A label is inserted at the beginning of the tape, to allow the identification and management of the tape by Arkeia.

Arkeia manages bar codes to locate tape in a library. This feature is essential if your library handles a large number of tapes to reduce the search time.

Tapes created by Arkeia will always be assigned to a specific tape pool. When a backup is launched, only the backup administrator decide which tape pool will be used by the backup. Arkeia then manages the tape(s) used in the tape pool, according to the tape strategy chosen by the administrator.

Creating your tapes is a step you really must prepare for: the amount of tapes you need will depend on the backup policy you will configure. Try to evaluate your needs as precisely as possible.

III.3.b. Standard tape creation

Click on the “New” button in the toolbar of the “Tapes Management” window



The “Create Tape(s)” window is then displayed on the screen.




Give the tape a name of your choice




Enter the first and last numbers (optional fields for a single tape). For more information, please see the section “How does Arkeia create tape names”.



Select the type of your tape in the drop-down menu.



Set the tape authorizations



Select the recycling destination (See the “Recycling” section)



Select the tape-recycling mode (See the “Recycling” section)



Select the tape pool for the tape or tapes created



Confirm your choices by clicking on the “Checkmark” (OK) button



III.3.c. “NULL” tape creation

“Null” tapes may be used for tests (see the “Quick Start Manual” and the “Arkeia initial configuration” chapter of this manual). Backups will be directed to an empty device. **Please note that no restorations are possible on this type of device.**

A “Null” tape is created just as a standard tape. The specific configuration items that are different between the two types of tapes are the following ones:

Select the “Null” type for your tape in the scroll down menu.



Set the tape authorizations



Confirm your choices by clicking on the "Checkmark" (OK) button




III.3.d. “FILE” tape creation

“File” tapes are volumes used for backups to a block device, such as a hard disk drive. For more information on these devices, please refer to the “Initial Configuration of Arkeia” chapter.

A “File” tape is created just as a standard tape. The specific configuration items that are different between the two are the following ones:

Select the approximate size of the “File” device you want.



In the barcode field, enter the path of the directory where the “File” tapes are going to be located. This can be on any block device (hard disk or magneto–optical disk: MOD).



Confirm your choices by clicking on the "Checkmark" (OK) button



Please note: Arkeia does not check (or fill) the available disk space when “File” devices/tapes are created. You must create a file library to manage “File” tapes.

III.4. Tape Recycling

One of the most important issue about tapes is recycling.

Once written, a backup Tape is supposed to be reused at some point. With Arkeia, a tape can be reused once it has been **recycled**. Therefore, the recycling function must be an integral part of a tape strategy.

As described in the “*Periodic backup*” chapter, a tape is recycled once it has reached its “*Retention Date*”. This date is determined by the backup run on that tape.

The way the tape is recycled is determined by the “*Recycling mode*” and “*Recycling destination*” fields in the T ape definition. The following options are possible:

Recycling mode	Meaning
FIFO	The first tape recycled will be reused first
LIFO	The last tape recycled will be reused first (risk of using the same tape)

Recycling destination	Meaning
Current pool	The tape can be recycled in its pool
Scratch pool	The tape will be recycled in the scratch pool and will be accessible to any backup.

Your tapes should be recycled very carefully, as you can run into quite a lot of problems once all of them are filled. Please note also that, if you force a tape to be recycled, all data backed-up on it is erased.

III.5. How does Arkeia create tape names?

Arkeia creates tapes with a generic name (the “Tape name” field) on which it add a number chosen in the range given by the “*first number*” and “*last number*” fields.

It is possible to specify formatting parameters in the name, as shown in the following examples:

Example 1 If the first number = 1 and the last number = 3

Tape name:

Tape names will be:

dat1 dat2 dat3

Example 2 If the first number = 1, and the last number = 15

Tape name:

Tape names will be:

dat01 dat02 dat15

Example 3 If the first number = 1, and the last number = 120

Tape name:

Tape names will be:

- dat 1* (2 spaces between “t” and “1”)
- dat 20* (1 space between “t” and “2”)
- dat120* (no space character)

III.6. Tapes deletion

Select in the “Tapes Management” window the tape(s) you want to delete.

(The button outline, in front of the tape name, should be filled–up for the tape to be selected).



Click on the “Trashcan” button to delete



Confirm the tape deletion.



Confirm your choice by clicking on the “checkmark” (OK) button



Please note: this action deletes tapes on a logical level, which means that tapes do not have to be physically present in the drive.

III.7. Modifying tapes: the Tape(s) modification window

It is possible to modify certain tape characteristics from the *Tape Management* screen.

Select in the “Tapes management” window the tape(s) you want to modify.

(The button outline, in front of the tape name, should be filled–up for the tape to be selected).



Click on the “Pen” button to modify the tape characteristics



The Tape(s) Modification window will be displayed on the screen.

You can now modify the options in this screen as you see fit. To modify an option, the button outline in front of the desired option should be selected (and filled–up).



Confirm your choices by clicking on the “Checkmark” (OK) button



Authorizations

Authorizations are used to reduce tape access in order to eliminate the risk of erasing data (Write), erasing the tape (Delete), recycle data written on the tape (Recycle), or prevent restoration from the tape (Read).

The “Clean” option enables to declare a cleaning tape.



Read is enabled (authorized).



Read is disabled (no authorization).

Recycling pool

By default, a tape is recycled (expiration date) in its pool but may also be recycled in the “Scratch Pool”.




- ✿ The “Scratch Pool” is the reserved Pool containing tapes that can be accessed by a backup when all the tapes in its current pool have been used and filled with backup data.

Recycling mode


By default, Arkeia uses all the tapes in the same pool before beginning to write over the first tape used (FIFO: First In / First Out).

It is also possible to reuse the same tapes as soon as they are recycled (LIFO: Last In / First Out)



Tape access mode

This option defines the way the tape is accessed (manual or automatic).



- *Automatically mountable*: The tape is in a tape library or changer, automatically accessible during backup or restore. The tape switches automatically to this mode when Arkeia detects it in a library or changer.
- *Manually mountable*: Arkeia will ask the user to insert the tape for all backup and restore operations. By default, a tape is created in this mode.
- *On line*: The tape is already in the drive, accessible for backup and restore operations. The tape switches automatically to this mode after a positive tape access.
- *Physically damaged*: The tape has to be switched manually to this mode when the user wishes to declare it physically damaged.

- *Secured in vault:* The tape has to be switched manually to this mode when the user wishes to declare it secured in vault.


Tape assignment pool

To change the Pool that is assigned to tapes, select another Pool from the list.



Comment zone

In this zone, you can write a personal note on your tape and backup.




III.8. Detailed tape information: the tape screen

This screen is used to display the detailed attributes of a tape.

Click on the right mouse button and choose [Examine current tape], or double click on the tape.

The detailed tape information dialog box is then displayed on your screen.



Bar code:

Bar code number or path for "File" tapes

Type:

Type of tape

Media life:

Number of tape usage

Recycling destination:

Recycling Pool

Recycling mode:

FIFO or LIFO

Access:

Arkeia tape access type

Pool:

Name of the pool the tape is assigned to

Space used/remaining:

(self-explanatory)



Creation date:

Tape creation date

Retention date:

Retention date (date when a tape is recycled)

Owner:

Name of the tape creator

Status:

Tape status (new, in use, full)

Life remaining:

Remaining tape usage number

Authorizations:

Actions authorized on the tape

Last write date:

(self-explanatory)

III.9. Writing the label on a tape

Use the “Tapes management” screen to label the tape manually.

Select in the “Tapes Management” window the tapes you want to label.

(The button outline, in front of the tape name, should be filled–up for the tape to be selected).



Click on the “Yellow Pen” (Modify) button to modify the tape characteristics.



Arkeia displays all compatible devices to label your tape. Double–click the device of your choice.



Make sure the correct tape is in the tape drive then confirm your choices by clicking on the “Checkmark” (OK) button.



Please note: manual tape labeling is an optional task. Arkeia labels tapes automatically during the first backup.

III.10. Tape recycling

A tape is usually recycled when the retention date has been reached.

However, you can recycle a tape manually. This will erase all references to this tape in the database index. All data written on this tape will be lost, once a backup has used the tape. After recycling, a tape is reused from its beginning and any previous data it may contain is lost.

Select in the “Tapes Management” window the tapes you want to recycle.

(The squares in front of the tape names have to be filled–up for the tape to be selected)



Click on the “Recycle” button to recycle the tapes



Confirm the tape recycling



Confirm your choice by clicking on the “Checkmark” (OK) button



IV. Drivepacks

IV.1. Description and use

“Drivepacks” are a group of drives used for a specific task.

While most Administrators will only have one drivepack that contains all their drives, many will want to use part of their drives for specific backups, or will want to keep a single drive for restoration, and all the others for backup, etc. This is the use of drivepacks.

- ✿ **Please note:** It is important to know that, while tapes can only belong to a specific pool, drives can belong to different drivepacks at the same time.
- ✿ Backups are always done on drivepacks, not on individual drives.

IV.2. Description of the “Drivepacks” screen

From the main screen click on the [Devices] menu then click on the [Drivepacks] option



Or you can click on the “Drivepacks” button in the toolbar:



Name:

Name of the drivepack

Owner:

Name of the creator of the drivepack

Number of drives:

Maximum number of drive to be used for a backup

List of drives:

List of drives included in the savepacks (those marked with a filled-up square next to their names)

Drive priority:

Priority of use of the selected drive.



IV.3. Creating a Drivepack:

Click the “New” button in the toolbar of the “Drivepack” window



Enter the name of the Drivepack. Enter a comment.
The owner is the current Arkeia user.



Select the drives you want to include in this new Drivepack.



Confirm your choice by clicking on the “Checkmark” (OK) button



IV.4. Drive Priority

When a backup is launched on several drives, Arkeia automatically manages the drives that will be used. The user may wish to select a particular drive first: this is the concept of drive priority.


In the “Drivepacks” window, select the drivepack



Select the appropriate device



Modify the priority of the drive



Check that the desired setting appears in the window.



Confirm your choice by clicking on the “checkmark” (OK) button



IV.5. Number of drives


When a backup is launched on several drives, Arkeia automatically manages the drives that will be used.

The user may wish to select a specific number of drives for a backup. Once this is done, multiple backups can share the same drivepack.

In the “Drivepacks” window, select the Drivepack you want to update.



Select the number of drives in the Savepack used by a single backup.



Confirm your choice by clicking on the “Checkmark” (OK) button



IV.6. Deleting a Drivepack


Select the Drivepack you want to delete in the “Drivepack” dialog box.



Click on the “Trashcan” button to delete it.



Confirm the Drivepack deletion.



Confirm your choice by clicking on the “Checkmark” (OK) button



CHAPTER 6

What to backup?

I. The “Savepack” concept

I.1. Understanding the issue

A backup operation is not as straightforward as it seems at first sight: an administrator may want to backup specific tree structures, complete file systems, or create images of complete partitions or disks or backup specific data such as the Windows Registry.

A backup system must be able to take into account all those specific cases, and define the data to backup, in a manner as general and as flexible as it can. It must also provide simple ways to apply encryption and compression on selected trees or files.

I.2. Arkeia’s approach

Arkeia defines the data to be backed up as Savepack(s). This highly flexible definition system allows the backup of various machines, various trees, complete raw devices or specific data.

With a Savepack, you are able to specify the machine(s), directory(-ies) and file(s) to backup. Several Savepacks can be defined for different tasks. You can even define Savepacks that include other Savepacks, making custom configuration of backups a very simple task.

Savepacks can be of various types, whether you plan to backup trees, objects or raw devices.

Each Savepack has options that include encryption, compression, pre-processing and post-processing, among others. These options are available either for the entire Savepack or for each tree structure included in the Savepack.


You may set priorities and chain options in a Savepack in order to create dependencies between the Savepack or the trees it contains.

Thus, defining a savepack is not only defining what to backup, but also how to perform the backup.

II. Savepack management

II.1. Description of the “Savepacks management” screen

From the main screen click on the [Backup] menu then on the [Savepacks] option.



Or use the “Savepacks” button in the Toolbar:




List of tree:

Trees included in the selected Savepack to be backed up.

Savepacks:

List of created Savepacks

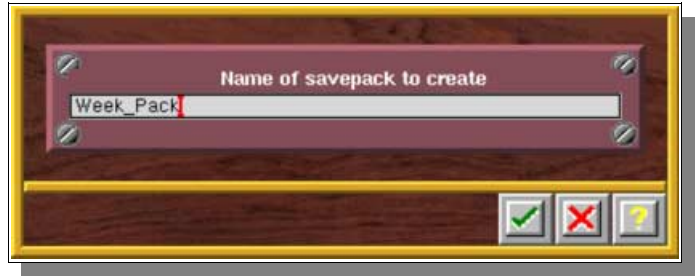


II.2. Savepack Creation

Click on the “New” button in the Toolbar of the “Savepacks management” window



Enter the name of the Savepack.



To confirm your choices, click on the “Checkmark” (OK) button.



II.3. Savepack Deletion

Select the Savepack you want to delete in the “Savepacks” window.


(The square next to a savepack’s name should be filled-up for the savepack to be selected)



Click on the “Trashcan” button called “Delete Savepack” to delete the savepack.



Confirm the Savepack deletion.



To confirm your choices, click on the “checkmark” (OK) button.



II.4. Adding a tree in the Savepack: the Network Navigator

Arkeia offers a powerful utility to select the trees and files you want to include in your Savepacks: the Network Navigator. It allows you to display all the available machines and to navigate within their tree structures to select the desired items.

Once you have selected a Savepack, click the “Navigator” button in the Toolbar of the “Savepacks management” window



The Network Navigator is then displayed on your screen.



Double-click on the appropriate icon to navigate.

Select the trees you want to backup by clicking on the empty square in front of the icon. Selected trees are indicated by filled-up squares next to their names.

You can also navigate back to select other trees or machines.



Click on the “Checkmark” button to confirm your choices.




The selected trees are now added to your Savepack.

Please note: you can add as many trees as you want in your Savepack, provided there are no more than 200 clients.

II.5. Deleting a tree in a Savepack

Select the correct Savepack, then the tree you want to delete:



Click on the “Trashcan” icon called “Delete Tree” to delete the desired tree.



Confirm the tree deletion.



Confirm your choices by clicking on the “Checkmark” button.



The selected trees are then deleted from your Savepack.

II.6. Inserting a Savepack in a Savepack

Select the Savepack you want to modify in the “Savepacks” window



In the “List of trees to backup” window, click with the right button of your mouse to obtain the contextual menu, then choose the “Add Savepack” option.



Select the Savepack you'd like to add in the menu



To confirm your choices, click on the “Checkmark” (OK) button.



Please note: make sure you never create co-dependencies between Savepacks. For instance: do not create a Savepack “A”, which includes Savepack “Z”, which, in turn, includes Savepack “A”, creating a codependency “loop” between the two Savepacks.

II.7. Advanced Savepack options

The following screen enables you to view and modify the backup parameters associated with the entire Savepack. All the actions (filter, compression, encryption) will apply to all the trees by default.

To see the properties of a specific Savepack, double click on its name in the “Savepacks” screen or select it and click on the “Menu” button.



Command before:

Shell command, or DOS batch, to be run on the Client before the backup.

Command after:

Shell command, or DOS batch, to be run on the Client after the backup.

Number of retries:

Number of times the backup is restarted in case errors are encountered.

Compression type:

Compression requested for this savepack.

Encryption type:

Encryption requested for this savepack



Follow symbolic links:

Requests the copy of symbolic links as standard directories or files.

Follow file systems:

Type of File Systems allowed.

Reset access times:

Reset or not the last access date of each file after a backup.

Find filter:

Restrict backup to files matching a UNIX find criteria

Inclusion filter:

Backup only files represented by UNIX regular expressions

Exclusion filter:

Backup all files except those represented by UNIX regular expression

These options are explained in the following sections.

II.7.a. Command before backup

A command executed before a savepack backup is a shell script on Unix systems, or a batch file on DOS-based systems, which should be run, on the client machine, before the backup itself.

This command can be used to determine if the backup must be started. To do this, check the box “Backup savepack if command fails” (see table below). The command must send a return code.

The syntax is as follows:

[Machine name]:/[path]/[command]




Table of backup execution conditions

Result of Command	Savepack Option	Backup Execution
OK	<input type="checkbox"/> Backup savepack if command fails	YES
OK	<input checked="" type="checkbox"/> Backup savepack if command fails	YES
Failed	<input type="checkbox"/> Backup savepack if command fails	NO
Failed	<input checked="" type="checkbox"/> Backup savepack if command fails	YES

Please note: the command can only be executed on a machine where the Arkeia client has been installed. Even if you don't backup this client machine, a license is needed if a command is entered.

II.7.b. Command after backup

A command after backup is a shell script on Unix systems or a batch file on DOS-based system which should be run on the client machine after the backup. The result of the backup may determine the execution of this command if the “execute if savepack backup fails” box has been checked (see table below)

The syntax is as follows:

[Machine name]:/[path]/[command]





Table of backup execution conditions

Backup Results	Savepack Option	Command execution
OK	<input type="checkbox"/> Execute if savepack backup fails	YES
OK	<input checked="" type="checkbox"/> Execute if savepack backup fails	YES
Failed	<input type="checkbox"/> Execute if savepack backup fails	NO
Failed	<input checked="" type="checkbox"/> Execute if savepack backup fails	YES

Please note: the command can only be executed on a machine where the Arkeia client has been installed. Even if you don't backup this client machine, a license is needed if a command is entered.

II.7.c. Number of retries

If Arkeia encounters a problem (For instance: a lost connection) during a backup, it automatically retries the operation. This option specifies the number of retries to be respected. The value ranges from “0” to “10”. A value of “0” prevents Arkeia from retrying. The default value is 3 retries.




II.7.d. Compression

Enables/Disables compression on the client. A drop-down menu offers different compression types:

LZ1 (Lempel Ziv): the compression speed is faster but the compression level is low. Best suited to computers with slow CPUs.

LZ3_3: the compression speed is slower but the compression obtained is better and results in smaller files.




- LZ1 or LZ3_3: Arkeia determines the compression used according to a preference file (by default LZ1) that must be present on the client itself.
- Only LZ1: Arkeia uses only in LZ1.
- Only LZ3_3: Arkeia uses only in LZ3_3.
- No compression: Arkeia does not compress data.

II.7.e. Encryption

Enables/Disables data encryption. Encryption is done on the client. A drop-down menu offers different types of encryption.

DES: for a very good level of security (DES size 56 bits)

Blowfish: for fast, simple encryption



- DES or Blowfish: Arkeia determines the encryption to be used according to a preference file that must be present on the client itself
- Only DES: Arkeia uses only DES
- Only Blowfish: Arkeia uses only Blowfish
- No encryption: Arkeia will not use any encryption.

🔴 Please note: you can find more information on the encryption functions of Arkeia in Chapter 12 of this manual: “Security in Arkeia”.

II.7.f. Follow symbolic links

Enables the backup to “follow” Unix symbolic links.
By default, Arkeia follows the symbolic links.



Example

If we have 4 directories (Dir1, Dir1/Dir11, Dir2, Dir2/Dir21) and a symbolic link “Link1” to “Dir1” in “Dir21”:

- *If the backup of “Dir2” is launched and the symbolic links are NOT followed: Arkeia will save “Dir2”, its sub-directory “Dir21” and the fact that there is a link, “Link1”, in the directory “Dir21”.*
- *If the backup of “Dir2” is launched and the symbolic links ARE followed. Arkeia will save the directory “Dir2”, its sub-directory “Dir21”, as well as the directory “Dir1” and its entire sub-directory “Dir11”.*

Please note: Arkeia does not detect recursive links.


II.7.g. Follow file systems

Enables the backup to “follow” local partitions.
By default Arkeia follows file systems.



II.7.h. File systems allowed

Allows the configuration of the types of file systems to backup.



On Unix


- Normal File systems: All normal file systems with the exception of CD drives, special partitions and NFS.
- All File systems: Normal file systems, NFS, CD-ROM.
- All except NFS: All the normal file systems, including CD-ROM but excluding NFS.

On Novell, Windows 9X, ME, NT and 2000

- Normal File systems: All “file systems” with the exception of shared CD-ROM drives, and network drives.
- All File systems: Normal file systems, network drives, CD-ROM
- All except NFS: All the normal file systems, including CD-ROM but excluding network drives.

II.7.i. Reset access times

When Arkeia reads files on the client, the system automatically modifies the last access time of the file. If you select “Yes”, Arkeia will reset this date to its original value.



II.7.j. Find filter

Backs up files meeting the UNIX “find” criteria.

❗ Please note: There is no need to type the “find” command itself, only its arguments.

Example *If you want to back up all files beginning with b or e, enter:*

`-type d -o -name [be]*`

Find filter: `-type d -o -name [be]*`

II.7.k. Inclusion filter

This option restrict the back up to specific files, matching a UNIX regular expression.

Example *If you want to include only the *.a and *.o files, enter:*

`-type d -o -name [be]*`

Inclusion filter: `.*\.[ao]$`

II.7.l. Exclusion filter

This option excludes from the back up the specific files, matching a regular expression.

Example *If you want to exclude the *.a and *.o files, enter:*

`-type d -o -name [be]*`

Exclusion filter: `.*\.[ao]$`

II.7.m. Regular expression

Regular expressions can be written as follows:

Symbol	Meaning or use
Normal characters	Used in a regular expression “as-is” and self-explanatory, except for the following characters: “.” (dot), “*” (asterisk), “^” (caret), “\$” (dollar sign), “+” (plus), “[...]” (brackets). For these characters, see the explanations below.
. (<i>dot</i>)	Represents any character, except when placed between brackets.

Symbol	Meaning or use
* (<i>asterisk</i>)	Indicates a repetition of 0 or <i>n</i> times of the regular expression preceding it.
+ (<i>plus sign</i>)	Indicates a repetition of 1 or <i>n</i> times.
^ (<i>caret</i>)	Represents the start of a line except when between brackets (see below).
\$ (<i>dollar sign</i>)	Represents the end of a line.
[] (<i>brackets</i>)	Represent alternatives: [ab] means “a or b”, [a–z] means any ASCII character from “a” to “z”. If the open bracket is followed immediately by the caret: “^”, then the alternatives are exclusions, if the dash “–” is placed right after the open bracket, it represents itself.
[^]	Means “except”

Examples

Here are some examples of regular expressions:

- If you wish to exclude all *.o and *.a files, the exclusion filter should be stated as follows:*

`^\.*\.[oa]$`

- If you wish to exclude all *.gif and *.jpg files, the exclusion filter should be stated as follows:*

`^\.*\.[gjl[ip]][fg]$`

- If you wish to qualify this exclusion by including all *.o and *.a files in the directory /export/home/dev/goodobjects, the inclusion filter should be stated as follows:*

`^/export/home/dev/goodobjects/.*\.[ao]$`

II.8. “Advanced Tree options” screen

The “Tree options” screen is used to configure tree settings.

To see the properties of a specific Tree, double click on its name in the “Savepacks” screen or select it and click on the “Tree” button.



Type:

Type of tree (see below)

Multiflow:

Flow number used for that tree

Priority:

Priority of the tree in the savepack

Chain:

Create dependencies between trees of different machines. Allow to control backup order with Priority setting.

Command before:

Shell command, or DOS batch, executed on the Client before the backup.

Command after:

Shell command, or DOS batch, executed on the Client after the backup.

Number of retries:

Number of times the backup is restarted in case errors are encountered



Compression type:

Compression type applied to this savepack.

Encryption type:

Encryption type applied to this savepack

Follow symbolic links:

Configures the inclusion of symbolic links as standard directories or files

File systems allowed:

Type of file systems saved.

Reset access times:

Resets or not the last access date of each file after a backup.

Find filter:

Restricts backup to files matching a UNIX “find” criteria

Inclusion filter:

Backup only files represented by a regular expression

Exclusion filter:


Backs up all files except those represented by a regular expression

Most of those options are set by default to the same settings of the Savepack in which the tree is included.

You can change them individually if needed. The specific options are detailed below:

II.9. Type of Trees

This option enables you to choose the type of tree you wish to back up. The main practical applications is the back-up of databases (TABLESPACE in RAW mode) and special files. The different types are: Tree, Object, Raw, Savepack



TREE

A “Tree” can be:

- a hostname: in this case, Arkeia will backup the entire machine.
- a directory: in this case the entire directory and its subdirectories will be backed up.
- a file: in this case, Arkeia will only backup the selected file.



OBJECT

An “Object” file will back up and restore the result of the standard output of a command, entered in the object backup/restore command field.

The object file does not exist in the machine tree itself but is created via Arkeia by specifying the machine name, then the object name.

The screenshot shows the configuration window for an OBJECT file type. The window title is "betelgeuse.knox.com:/home/rvdboom". The configuration fields are as follows:

- Type: OBJECT (dropdown menu)
- Multiflow: 0
- Priority: 50
- Chain: 0
- Comment: (empty text field)
- Command before tree: Backup tree if command fails
- Command after tree: Execute if tree backup fails
- Number of retries: As savepack option (dropdown menu)
- Compression type allowed: As savepack option (dropdown menu)
- Encryption type allowed: As savepack option (dropdown menu)
- Object backup command: /usr/bin/PROGRAM_BACKUP
- Object restore command: /usr/bin/PROGRAM_RESTORE

At the bottom right, there are three buttons: a green checkmark, a red X, and a yellow question mark.

❗ **Please note:** in contrast with the before/after backup commands, do not enter a machine name in the “Object backup/restore command” fields.

RAW

The “Raw” file is an entire hard disk (image copy) or a partition.

Use the navigator to select the name of the partition or the disk appearing in the “/dev” directory of your UNIX machine.

The screenshot shows the configuration window for a RAW file type. The window title is "betelgeuse.knox.com:/home/rvdboom". The configuration fields are as follows:

- Type: RAW (dropdown menu)
- Multiflow: 0
- Priority: 50
- Chain: 0
- Comment: (empty text field)
- Command before tree: Backup tree if command fails
- Command after tree: Execute if tree backup fails
- Number of retries: As savepack option (dropdown menu)
- Compression type allowed: As savepack option (dropdown menu)
- Encryption type allowed: As savepack option (dropdown menu)

At the bottom right, there are three buttons: a green checkmark, a red X, and a yellow question mark.

SAVEPACK

The “SAVEPACK” type screen can only be displayed if you have selected a tree which is a Savepack



II.9.a. Multiflow (parallel processing on one machine)

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 parts by entering a number made in the "MULTIFLOW" field.



Please note: Trees with the same Multiflow number are backed up sequentially.

II.9.b. Priority

The priority level ranges from 1 to 100 and is used to specify the order in which the trees are backed up. Trees with a priority of “1” are the first to be saved and those with a priority of “100” will be backed-up last.

The default setting for all trees is 50.



Please note: the priority level is only taken into account when trees have the same Multiflow number.

II.9.c. Chain

The “Chain” field is used to make dependencies between trees from different machines.

A chain is created between two trees when you put the same value in the chain field.

You also have to set the “priority” level for each chained tree to control the backup order

The default setting for all trees is “0” (no chaining).



Example



In this case:

- *libra.knox.com:/boot and neon.knox.com:/boot are chained*
- *libra.knox.com:/boot has a lower priority*

That means that the machine named “libra” will be backed up once the backup of “neon” has been completed.

II.9.d. Command before tree backup

A command before tree backup is a shell script on UNIX–based system that should be run before backup.

This command may be used to determine if the backup must be started. For this check the box “Backup tree if command failed” (see table below). The command must provide a return code (0=OK, other=Failure).

The syntax is as follows:

[Machine name]:/[path]/[command]




Table of backup execution conditions

Result of command execution	Option selection	Backup execution
OK	<input type="checkbox"/> Backup savepack if command fails	YES
OK	<input checked="" type="checkbox"/> Backup savepack if command fails	YES
Failed	<input type="checkbox"/> Backup savepack if command fails	NO
Failed	<input checked="" type="checkbox"/> Backup savepack if command fails	YES

Please note: the command can only be executed on a machine where the Arkeia client has been installed.

Example

On a UNIX system:

Command before tree: Backup tree if command fails
 neon.knox.com:/home/script/command.sh

On a Windows system:

Command before tree: Backup tree if command fails
 dune.knox.com:c:/script/command.exe

II.9.e. Command after the backup of a tree

A command after backup is a shell script on UNIX-based system or a batch file on DOS-based system that should be run after the backup.

The result of the backup may determine the execution of this command if the “Execute if savepack backup failed” box has been checked (see table below)

The syntax is as follows:

[Machine name]:/[path]/[command]

Command after savepack: Execute if savepack backup fails
 mars:/opt/oracle/bin/startdb

Table of backup execution conditions

Backup execution	Selected option	Command execution
OK	<input type="checkbox"/> Execute if savepack backup fails	YES
OK	<input checked="" type="checkbox"/> Execute if savepack backup fails	YES
Failed	<input type="checkbox"/> Execute if savepack backup fails	NO
Failed	<input checked="" type="checkbox"/> Execute if savepack backup fails	YES

Please note: the command can only be executed on a machine where the Arkeia client has been installed. Even if you don't backup this client machine, a license is needed if a command is entered.

Example

On a UNIX system:

Command after tree:	<input type="checkbox"/> Execute if tree backup fails
neon.knox.com:/home/script/command.sh	

On a Windows system:

Command after tree:	<input type="checkbox"/> Execute if tree backup fails
dune.knox.com:c:/script/command.exe	

II.9.f. Compression

By default, the options are the same as for the Savepack (check the option applied to the Savepack).

II.9.g. Encryption

By default, the options are the same as for the Savepack.

II.9.h. Follow symbolic links

By default, the options are the same as for the Savepack.

II.9.i. Follow file systems

By default, the options are the same as for the Savepack.

II.9.j. File systems allowed

By default, the options are the same as for the Savepack.

II.9.k. Reset access times

By default, the options are the same as for the Savepack.

II.9.l. Find filter

Back up files matching a UNIX “find” criteria (do not type the find command, only the arguments).

Example *If you want to back up all files beginning with b or e, enter:*
*–type d –o –name [be]**

Find filter:

II.9.m. Inclusion filter

Back up only the files matching a regular expression.

Example *If you want to include the *.a and *.o files, enter:*
*–type d –o –name [be]**

Inclusion filter:

II.9.n. Exclusion filter

Back up everything, except the files matching a regular expression.

Example *If you want to exclude the *.a and *.o files, enter:*
*–type d –o –name [be]**

Exclusion filter:

II.9.o. Regular expression

Regular expressions can be written as follows:

Symbol	Meaning or use
Normal characters	Used in a regular expression “as-is” and self-explanatory, except for the following characters: “.” (dot), “*” (asterisk), “^” (caret), “\$” (dollar sign), “+” (plus), “[...]” (brackets). For these characters, see the explanations below.
. (<i>dot</i>)	Represents any character, except when placed between brackets.
* (<i>asterisk</i>)	Indicates a repetition of 0 or <i>n</i> times of the regular expression preceding it.
+ (<i>plus sign</i>)	Indicates a repetition of 1 or <i>n</i> times.
^ (<i>caret</i>)	Represents the start of a line except when between brackets (see below).
\$ (<i>dollar sign</i>)	Represents the end of a line.

Symbol	Meaning or use
[] (<i>brackets</i>)	Represent alternatives: [ab] means “a or b”, [a–z] means any ASCII character from “a” to “z”. If the open bracket is followed immediately by the caret: “^”, then the alternatives are exclusions, if the dash “–” is placed right after the open bracket, it represents itself.
[^]	Means “except”

Example	<p><i>Here are a few examples of filters:</i></p> <ul style="list-style-type: none"><i>If you wish to exclude all the *.o and *.a files, the exclusion filter should be stated as follows:</i> <code>^.*\.[oa]\$</code><i>If you wish to exclude all *.gif and *.jpg files, the exclusion filter should be stated as follows:</i> <code>^.*\.[gjlip][fg]\$</code><i>If you wish to qualify this exclusion by including all *.o and *.a files in the directory /export/home/dev/goodobjects, the inclusion filter should be stated as follow:</i> <code>^/export/home/dev/goodobjects/.*\.[ao]\$</code>
----------------	--

III. Specific examples and cases

III.1. How to prevent the backup of a specific directory: the .OPB_NOBACKUP file

You may want to make sure that some of your directories will never be backed up, even if they are included in trees selected in savepacks. This can be useful for temporary directories, for example.

There is a very simple way to set this up: just create an empty file called, on Unix: *OPB_NOBACKUP*

On Windows systems, the name of the file will be: *NOBACKUP.OPB*

All the directories that include such a file will be ignored by any backup.

CHAPTER 7

Backups

I. Interactive backups

I.1. Introduction

The interactive backup is essentially a backup you start only one time, when you feel the need to save specific data or to make sure that sensitive data is saved if a periodic backup failed.

This function allows you to start a backup operation very fast, by reusing elements created previously:

- SAVEPACK: the data to be backed up.
- DRIVEPACK: the drive(s) the data will be saved to.
- and POOL: the tape(s) that will be used by the drive to save the data to.

Interactive backups are “one time” backups. If you want to create a backup that executes at a specific time, you should create a *Periodic Backup* (described later in this chapter).

Please note:

- ✚ After launching the backup via the interface, you may close the backup screen or even exit the program completely without interrupting the backup operation in progress.
- ✚ Once closed, you can reopen the backup window by double-clicking on the correct backup, in the backup list of the main window.

I.2. The “Interactive backup” screen

This is the main screen on which you define your interactive backup.

As soon as you validate it, the backup will proceed.

From the main screen click on the [Backup] menu, then on the [Interactive Backup] option.



Or click on the “Interactive Backup” button in the Toolbar:



Savepack:

Savepack to be backed up.

Drivepack:

Drivepack used for the backup

Pool:

Tape Pool used for the backup operation

Type:

Type of Backup: “Total”, “Incremental” or “Archive”. You can also select “Standard” or “Continuous”.

Tape strategy:

Policy used for the tapes: use a new tape for each backup or complete an existing tape.

Valid for:

Validity of the backup. Determines the date on which the tape used for the backup will be recycled.

Parallelism:

The number of flows used for this backup

Use emails:

Defines if an email is sent to the system administrator to indicate the final status of the backup.

Tag:

Reference to a specific backup (incremental backup).



I.3. Starting an Interactive Backup

Select the “Savepack” you want to backup

Select the “Drivepack” you want to backup to.

Select the “Tape Pool” you want to backup on.

Select the Backup Type, according to the policy you decided (“Total”, “Incremental” –based on a previous backup –or “Archive”)

Select the tape strategy

Select the validity of your backup (How long you want to keep the data on the tapes used).

After this period, no reference of the data is kept in the database and the tape is recycled automatically (except for Archive backup)

Set other parameters as needed (Parallelism and comment)

Confirm your choices by clicking on the “Checkmark” (OK) button.



I.4. Specific Options

I.4.a. Types of Backup

Total backup

All selected files and directories will be backed up.

Incremental (or differential)

Incremental backups are based on a previous backup (which can be itself a total or incremental backup). Arkeia reads the modification date to determine if it should backup a given file. Therefore, a specific file is backed up if it has been modified since the reference date.

Archive

Total backup with no retention date (the tape(s) will never be recycled automatically)

“Standard” or “Continuous”

A continuous backup is a backup that does not stop even once all the machines and trees are backed up. A continuous backup keeps on running, the drives and tapes are still reserved and can't be used by another backup operation. In effect, Arkeia is waiting until you add a new Savepack or you cancel the job.

Click on the “Standard” icon or on the “Continuous” icon to change the mode.

I.4.b. Tape strategy

Complete existing tapes: Backup will place data at the end of the current tape.

Use new tapes: Backup will place data at the start of a new tape.

I.4.c. Parallelism

The number of trees backed up at the same time may be changed. By default, Arkeia will use the maximum number of flows you have configured.

I.4.d. Use email

Arkeia will send you an email report at the end of the backup. Arkeia will also send an email if a new tape is needed.

I.4.e. Tag (Optional)

Backup may be “tagged”, to allow incremental backups to use the tag when using “arkc”, the Arkeia command line interface.

1.5. Monitor the Backup: the “Backup” screen

This screen allows you to monitor the backup operation in real time.

Counter:

Backup time

Server:

Name of the backup server

Savepack:

Savepack name

Drivepack:

Drivepack name

Pool:

Tape Pool name

Cruise:

Maximum throughput. Arkeia

Upper gauge:

Backup or Network Speed.

Lower gauge:

Speed of a drive, displaying the name of the tape (there will be as many gauges as drives used)



Speed in MB/min:

Click on this button to change the display unit

Instant:

Instant backup speed

Average:

Average backup speed

files:

Number of files saved

drive:

Number of drives used.

flows:

Number of flows used

Lines under “# flows”:

Details of each flow: current tree backed up, instant speed, drive used, name of the machine backed up and flow number.

The yellow “Magnifying Glass” button opens a window that shows details on the flow.

1.6. Specific options of the “Backup” Screen

1.6.a. Adding a Savepack

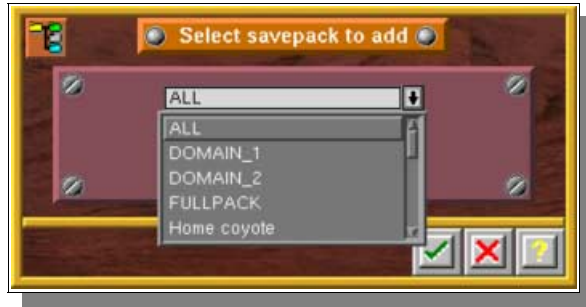
You can add a Savepack while a backup is running. This allows you to interactively modify the list of machines and trees of your current backup.

Click on the “Tree” button in the Toolbar of the “Backup” window



The “Select savepack to add” window is then displayed on the screen.

Select the Savepack you want to add to your backup.



Confirm your choice by clicking on the “checkmark” (OK) button.



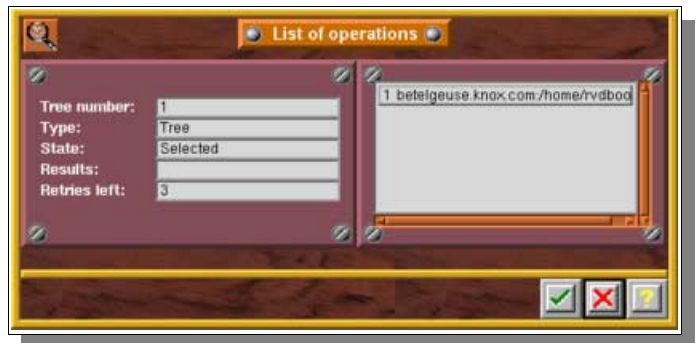
1.6.b. Tree status

This screen displays the status of each tree managed by the current backup.

Click on the “Magnifying Glass” button in the Toolbar of the “Backup” window



The “Tree Status” window is then displayed on the screen.



Exit this window by clicking on the “checkmark” (OK) button.



Here are the different values of the status, as well as a short explanation for each of them:

Status	Explanation
“Running”	Arkeia is backing up the tree.
“Waiting”	There are more trees than flows, meaning some operations are waiting for a flow to be freed
“End of tree”	The tree is backed up or the machine can’t be reached
“Aborting”	The backup operation for this tree is being aborted
“Selected”	The tree has been selected for a backup operation


“Retry” is the number of retries left for an operation. See “Advanced savepack option” and “Advanced tree option” to modify the default value. The default number of retry is “3”. For instance, if the value is left “1”, that means Arkeia has already restarted 2 times this operation due to connection problems.

I.7. Connection to backup or restore

After closing the backup screen, you may reopen it by double-clicking on the current backup or restore, in the “List of jobs” field of the main Arkeia window.



✚ Please note: do not pay attention to the message below. It is only displayed because the backup process is finished:



II. Periodic backups

II.1. Introduction

A periodic backup is programmed to start at a specific time and at regular intervals.

As with the “Interactive” backup feature, a periodic backup is a combination of the three basic elements of a backup operation:

- **SAVEPACK**: the data (files, directories, etc.) you want to backup.
- **DRIVEPACK**: the list of drives that will be used for the backup.
- **POOL**: the list of tapes available for the backup.

A fourth element also comes into play in periodic backups:

- **PERIODICITY**: the time interval between two backup executions.

The unit of each interval may be a day, a week, a month or a year.

A periodic backup is configured with a maximum of three levels. The first backup level has priority over the second level that has priority over the third.

Periodic backup manages:

- Exceptions (ex: backup on three non–consecutive days over seven days for example).
- “Before” and “after” backup commands.
- The setting change for a specific day of the week, month or year. It is also possible to run a backup every third day of the month or every Tuesday.

II.2. The “Periodic Backup” window

This screen allows you to define your periodic backups.

From the main screen click on the [Backup] menu then select the [Periodic Backup] option.



Or click on the “Clock” button in the Toolbar.



Status:

Enables or disables the backup.

Owner:

Owner of the Backup

Levels:

Number of levels of the backup.

E-mails:

Use emails to warn the backup owner.

Savepack:

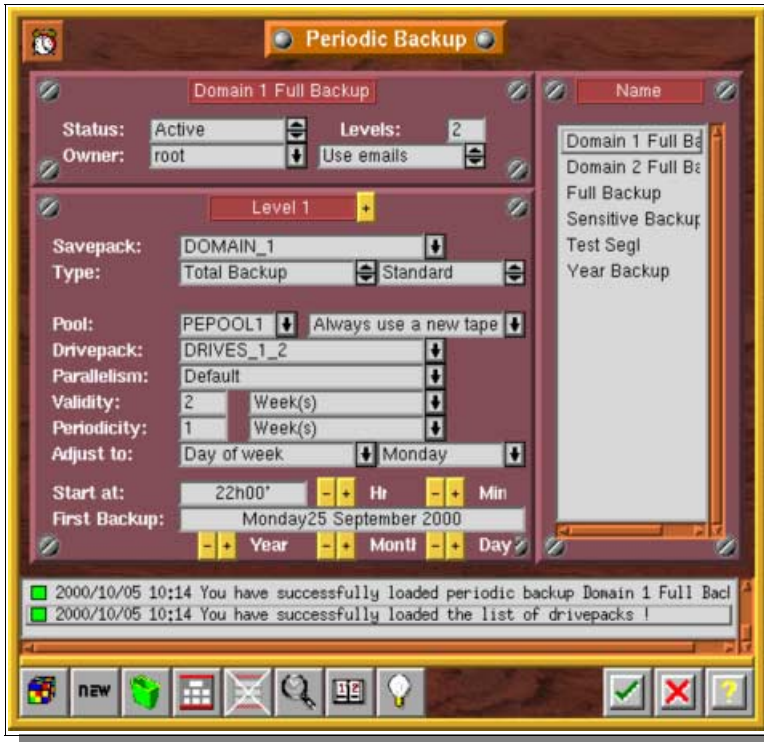
Savepack to be backed up.

Type:

Type of Backup: Total, Incremental or Archive. You can also select: “Standard” backup or “Continuous”.

Pool:

Tape Pool used to backup. Also set the Tape Strategy



Drivepack:

Drivepack used for backup

Parallelism:

The number of flows used for the current backup

Validity:

Validity of the backup. Determines the date at which the backup tape will be recycled.

Periodicity:

Periodicity of the backup.

Adjust to:

Adjust the backup date to a certain day of the week, month or year.

Start at:

Backup start hour.

First Backup:


Date of the first backup

II.3. Create a Periodic Backup

Click the “New” button in the Toolbar of the “Periodic Backup” window



Enter the name of the Periodic Backup.



Confirm your choices by clicking on the “checkmark” (OK) button.




II.4. Setting the standard Periodic Backup options


Select the Periodic Backup you want to modify




If you want to temporarily disable a Periodic Backup, without losing its definition, set it to “Inactive”




If you want to change the owner of the backup, use the “Owner” menu



If you don't want to get emailed reports as owner of the backup, choose the “Don't use emails” option




Select the “Savepack” that will be saved by this Periodic Backup



Choose backup Type (“Total” or “Incremental”, “Standard” or “Continuous”)




In the case of an Incremental Backup, set the backup used as reference.



Choose the “Tape Pool” used for backup and the Tape Strategy used.



Select the “Drivepack” used by the backup.




Set the number of simultaneous flows to be used.



Set the Validity of the Backup: how long the data on the tapes used in the backup is kept before the tapes are recycled.



Set the Periodicity of the Backup.



If the backup should be run on a certain day, select the reference (day of the week, of the month or of the year) then set the correct day in the field that is then displayed on the right.




Select the backup starting time with the yellow “+” and “-” buttons.



Select the date of the first backup with the yellow “+” and “-” buttons.

If an “Adjustment” is selected, the actual date will be the adjusted day that follows the date of first backup.

Confirm your choices by clicking on the “checkmark” (OK) button.



II.5. Periodic Backup deletion

Select the Backup you want to delete in the “Periodic Backup” window



Click on the “Trash can” icon to delete the backup.



Confirm the backup deletion.



Confirm your choices by clicking on the “checkmark” (OK) button.



II.6. Managing the Periodic Backup levels

An administrator may want to be able to run various backups that are more or less dependent on each other. For instance, he may want to make sure two backups with different periodicity or type cannot run the same day. He may need to base one backup on another, when doing backups of Incremental type.

The way Arkeia handle this is to offer various levels in a backup definition, each being a specific backup that can or cannot be dependent on the backups of lower level.

Arkeia manages three levels for each backup. In any case, the “Level 1” backup has priority over the “Level 2”, which has priority on “Level 3”.

The first level has the longer time interval in comparison with the second level and so on, as illustrated below:

Example


- Level 1: one backup per month
- Level 2 one backup per week
- Level 3 one backup every day

A single level is activated each day. The first level is executed before the second, which in turn has priority over the third level. The higher level sets the number of occurrences (or instances) of the lower level, for example:

Assuming a level 1 with a monthly interval, a level 2 with a weekly interval, a level 3 with a daily interval. Level 3 will run 6 times per week before level 2 is launched. Level 2 itself will run 3 times per month before level 1 is launched.

The result is a diagram for each level represented as follows:


Programmed execution diagram (without any exceptions):



Arkeia manages the number of operations and allows the user to invalidate some of these.

The programmed backups can be illustrated as follows:

Programmed execution diagram (with exceptions):



In this example, backup did not take place:

- on every fourth day (level 3)
- on the third week (level 2)

II.6.a. Adding a level

Select the Periodic Backup to which a level should be added.

You may add up to two levels to an existing Periodic Backup.



Click on the “Add level” button.



A level is added and Arkeia goes immediately to that level



Confirm your choices by clicking on the “checkmark” (OK) button.



Please note: by default, the fields in the added level retain the values of the lower level.

II.6.b. Changing a level

Select the Periodic Backup on which you want to update.



Use the “+” and “-” keys to go to the desired level



II.6.c. Deleting a level

Select the Periodic Backup level(s) which should be deleted.



Go to the last level of the backup (with the + / - buttons).



Click on the “Remove level” button.



Confirm your choices by clicking on the “checkmark” (OK) button.



II.7. The “Advanced options” screen

The "Advanced options" feature allows you to manage the executions of periodic backups and to insert commands before and after backup. Each backup level allows the use of advanced options.

Select the Periodic Backup on which you want to update.



Go to the desired level of the backup (with the + / - buttons).



Click on the “Magnifying Glass” button to open the “Advanced Option” screen



Enabling:

Management of instances

System command:

A shell script which is executed on the backup server before starting the backup, loading tapes or making connections.

Command before:

Shell command to be run before the backup on the Backup Server.


Command after:

Shell command to be run after the backup on the Backup Server.



II.7.a. Exception (occurrences) management

Select the type of validation.



By default, all occurrences are enabled (“*Always enabled*”). By default, scheduled backup operations are always executed by Arkeia, when their due date/time are reached. This is equivalent to the “*Always enabled*” option that is visible in the screen shot above.

However, it is also possible to configure the execution of these scheduled instances:

- By choosing which scheduled execution should be carried out (which is the equivalent of the option “*Enabled for instances*” in the screen shot above).
- By choosing which scheduled execution should NOT be carried out (option: “*Enabled, except for instances*”, above).

Configure the occurrences, which should or should not take place, separated by spaces.

In this example, only the second and the fourth occurrence of the level 2 backup will not take place.



❖ **Please note:** instances can only be modified on Level 2 or Level 3. You have to create a backup with multiple levels if you'd like to disable some instances.

Example

You want to create a backup that will only run on Mondays and Fridays:

Create a Weekly Level 1 backup that run on Monday. Then, add a new level.

The Level 2 backup, that start on Tuesday, is a daily backup which should be configured to execute on Fridays.

As Friday is the 4th occurrence of the Level 2 backup (since Tuesday), we set the “Enabling” parameter to “Enabled for instances”, the instance being “4” (=Friday).

II.7.b. System command

The user has the possibility of running a local server command for each level of the periodic backup function.

The command is executed first (tapes not yet loaded and connections not yet established)




❖ **Please note:** contrary to the “before” and “after” commands used for the Savepack level, the name of the machine is not included in the SYSTEM command path.

II.7.c. Command before


For each level of a periodic backup, it is possible to define a local server command to be run before the backup operation.

This command may determine the backup execution.



II.7.d. Command after

For each level of a periodic backup, it is also possible to define a local server command to be run after the backup operation.



II.8. The Schedule Viewer

To check the scheduling of your periodic backups, you can use the schedule viewer.

Click the “Diary” button in the Toolbar of the “Periodic Backup” window.



The “*Schedule viewer*” window is then displayed on the screen.



Set a time interval (in days) from the current day to check the schedule.



Click on the “Magnifying Glass” button to display more information on a selected periodic backup.



The “Schedule Information” window is then displayed on the screen.



Exit this window by clicking on the “checkmark” (OK) button.



🔴 Please note: to display the schedule of one periodic backup, set the other backups to “Inactive” mode.

II.9. The “Periodic Backup Assistant”

One easy way to create a standard Periodic Backup is to use the “*Periodic Backup Assistant*”.

It is an easy-to-use wizard that will help you set up simple periodic backups, asking simple questions about the policy you want.

You still have to define tape pools, drivepacks and savepacks before starting the “*Assistant*”

Click the “Light Bulb” button in the Toolbar of the “Periodic Backup” window



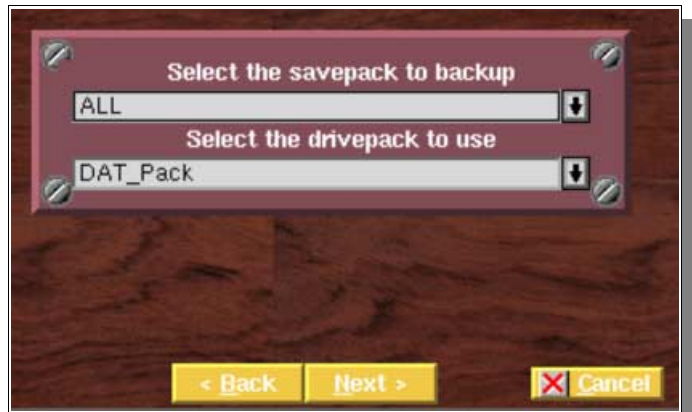
The first screen of the “Periodic Backup Assistant” is then displayed on the screen.

Enter a name for your Periodic Backup then click on the “Next >” button.



Select the Savepack to be backed up and the Drivepack to be used.

Then click on the “Next >” button



Select the type of Periodic Backup you want to create among the examples provided.

Then click on the “Next >” button.



Select the dates and hours when the backup is to be run.
Then click on the “Next >” button.



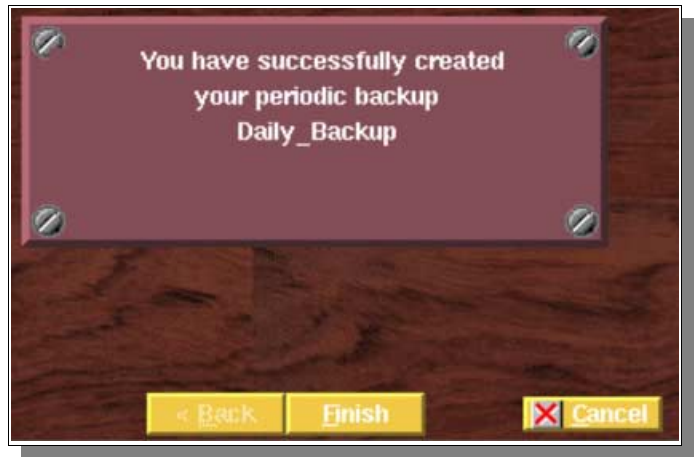
Select the time you want to keep each backup tape.
Then click on the “Next >” button.



Select the Tape Pool that will be used by this Periodic Backup.
Then click on the “Next >” button.



The Periodic Backup has been created.
Click on the “Finish” button to validate your choices.



CHAPTER 8

Periodic Backup policy

I. Periodic Backup Proceedings

I.1. Introduction

In this section, you'll find some advice on how to define a standard periodic backup.

This section describes a standard way to define and configure a periodic backup policy. It also contains several examples to help you achieve an optimum configuration.

I.2. Methodology

There are a few important points in the creation of a Periodic Backup, that should be clearly thought out, before anything else. Once the Policy is put into place, it can be difficult to change it: you will probably lose the backups already made by any previous policy.

If you want a complex backup strategy, try to divide it into several periodic backups that are more simple and configure these independently.

Most, if not all, backup policies simply group daily, weekly or monthly backups. We will see these first.

1.2.a. Simple Periodic Backups

Tapes

1. Evaluate how much data you want to backup, and particularly the number of tapes needed for any single type of backup. You also have to consider the amount of data that could have changed between two backups in the case of “Incremental” type backups.
2. Then, you have to define the schedule of your backup: daily, weekly, monthly, yearly.
3. Choose a tape policy: “*Always use a new tape*” or “*Complete existing tapes*”.
4. Define the “*Validity*” of the backup. This, in turn, determines the “*Retention Date*” of any single tape. After this date, the tape is automatically recycled.
5. Then, define the tape pool needed according to the above elements. Read the following examples for more information. Most of the time, try to put more tapes than the minimum needed. This will avoid mixing tapes in case a problem appears (SCSI errors, for instance).
6. Remember to create as many tape pools as there are backups to be made, taking into account the different validities.

Drives

If you have more than one drive, you may want to set apart some drives for specific backups:

1. Create a Drivepack and put in all the drives you want to use for the backup you are creating.

Savepack

1. Create the Savepack you need for your Periodic Backup

Create a Periodic Backup

1. Create the new Periodic Backup
2. Choose the Tape Pool, Drivepack and Savepack, the type of the backup, the Tape policy, the Validity and Periodicity.
3. Choose the date of the first backup and adjust the day if necessary.
4. Verify the settings in the Scheduler.
5. Congratulations: your Periodic Backup is now created.

1.2.b. Semi-periodic backups


Semi-periodic backups have a more complex periodicity: the first five days of the week every week, on Monday and Tuesday every week, etc.

Mostly, those backups fall into two categories:

1. Those which have irregular instances, like the examples above.
2. Those where the type of backups can change: that can happen if you want the first backup of the week to be “Total” backup and the six remaining ones to be “Incremental”.

For those backups, you may have to define different levels of backup: most of the time, you have to add one level for each particularity in the scheduling.

You can have up to three (3) levels on a given backup.

 **Please note:** two levels of backup should not be started on the same day. This could prevent the correct backup of data over a complete week.

1.2.c. Periodic Backup Assistant

Arkeia provides a Periodic Backup assistant that can help you to set your backup policy up.

If you don't feel confident in setting your policy manually, try configuring a periodic backup with the assistant first.

1.2.d. Tips and techniques

Here are some tips to avoid problems in setting up your backup policy:

- If you plan to remove some instances in the scheduling screen, you need to create a multi-level periodic backup.
- If you create a multi-level periodic backup, make sure each level starts on a different day than the others. If you start all your levels on the same day, your scheduling will “break down” very quickly.
- If you plan to run, for example, a backup on the first Monday of the month (by adjusting to “Day of the week” to create a monthly backup), start the backup the 1st of the month. As Arkeia adjusts to the next instance that follows the regular date, you can be sure that your backup will be started the first Monday of each month. Do not start at the precise date of the first Monday: as a new month may bring a very different date for the 1st Monday of the month.

II. Examples

II.1. A complete backup each day

II.1.a. Policy definition

A “Total” backup should be made everyday of the week. Each backup should be kept for two weeks.

For the sake of convenience, let's admit that each backup fills precisely one tape. A new tape should be used for each backup and one tape drive is available.

II.1.b. Solution and Analysis

This is the most simple backup possible: no exception and a simple tape retention policy.

Step 1: Tape pool evaluation

We first create a pool, simply named FULLPOOL.

Pool name: FULLPOOL
 Owner: root
 Comment:

We want to make one backup a day, which uses one tape per day.

As we want to keep each one 2 weeks before recycling it, we need enough tapes for 2 weeks of backup: thus, we need at least 14 tapes at least in our tape pool.

15 tapes have been added to FULLPOOL. This way, one spare tape is available for convenience.

Tape name: FULLTAPE-
 Bar code:
 First number: 1 Last number: 15
 Type: DLT 4000
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: FULLPOOL
 Comment:

Step 2: Drivepack selection

As only one drive is available, this selection is easy: our Drivepack is called "ONEDRIVE", and contains the unique tape drive available. In this example, the drive is a DLT 4000.

Name : ONEDRIVE
 Owner : root
 Comment
 Number of drives : All
 List of drives

<input type="checkbox"/>	DAT	
<input checked="" type="checkbox"/>	DLT4000	1
<input type="checkbox"/>	FILE_Drive	
<input type="checkbox"/>	NULL Drive	

 Current drive priority: 1

Step 3: Savepack selection

A Savepack is created, including all our servers. It is named "FULLPACK".

Name of savepack to create
 FULLPACK

Step 4: Periodic Backup creation

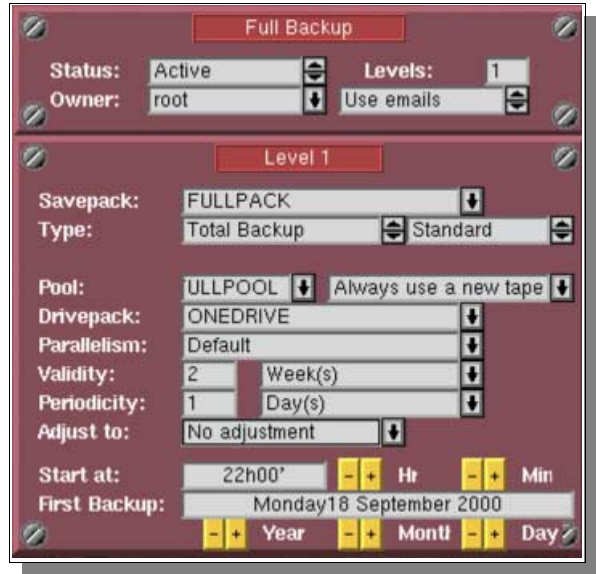
A Periodic Backup is created, named "Full Backup".

Name of periodic backup to create
 Full Backup

To this backup, we add “FULLPACK”, “ONEDRIVE” and “FULLPOOL”.

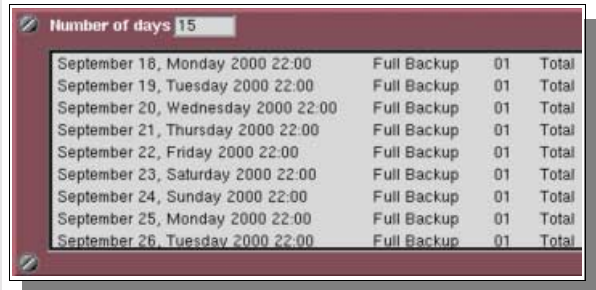
The Periodicity is set to “1 day”, and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start next Monday.



Step 5: Check the settings

The settings should be checked in the Scheduler to make sure they respect the policy defined above.



II.1.c. Conclusion

As the validity of the backup is two (2) weeks, the tape used during the first backup will be automatically recycled on Monday the 15th.

II.2. A “Total” Backup from Monday to Friday

II.2.a. Policy definition

A “Total” backup should be done everyday, from Monday to Friday, every week and each backup should be kept for two weeks.

For the sake of convenience, let’s admit that each backup requires one tape and that a new tape should be used for each backup. There is one tape drive available.

II.2.b. Solution and Analysis

This example is a little bit more complicated than the first one we have seen. This is a typical “semi-periodic” backup. Two levels of backup are needed here, to disable some executions of the daily backup (enabling or disabling periodic execution requires two levels of backup).

In the following steps, a first level weekly backup is going to be created for an execution on Monday, and a second level backup, based on the first one, will be created to be run on Tuesday, Wednesday, Thursday and Friday.

Step 1: Tape pool evaluation

A tape pool is first created, named “FULLPOOL”.

With the defined policy, five tapes are used per week.

As we want to keep each one 2 weeks before recycling it, we still need enough tapes for two weeks of backup: therefore, we need at least 10 tapes in our tape pool.

We will create eleven tapes in “FULLPOOL”. In this way, we have one spare tape for convenience.

Step 2: Drivepack selection

As one drive only is available, the selection is easy: our Drivepack is called “ONEDRIVE” and contains the unique drive available. In this example, the drive is a DLT 4000.

Drive Name	Count
<input type="checkbox"/> DAT	
<input checked="" type="checkbox"/> DLT4000	1
<input type="checkbox"/> FILE_Drive	
<input type="checkbox"/> NULL Drive	

Step 3: Savepack selection

One Savepack is created, named “FULLPACK”, which contains all the servers available.

Step 4: Periodic Backup – first level creation

A Periodic Backup is created, which is named “Full Backup”.

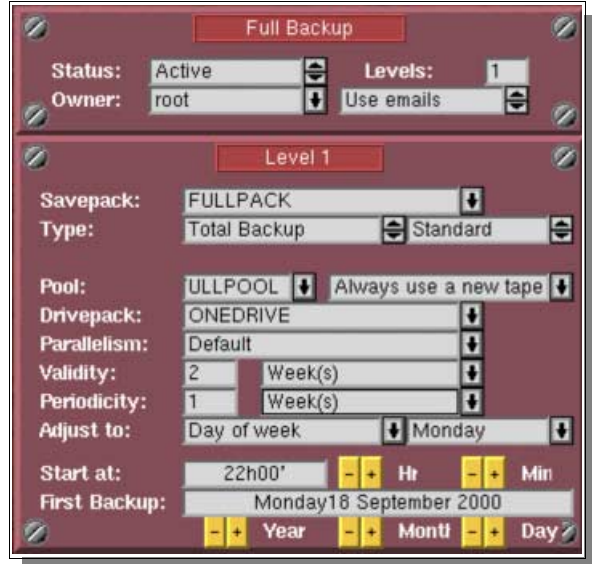


“FULLPACK”, “ONEDRIVE” and “FULLPOOL” are added to this Periodic Backup.

The type is set to “Total Backup”, the Periodicity to “1 week” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start on the next Monday.

The backup is then adjusted to “Day of week”: Monday.



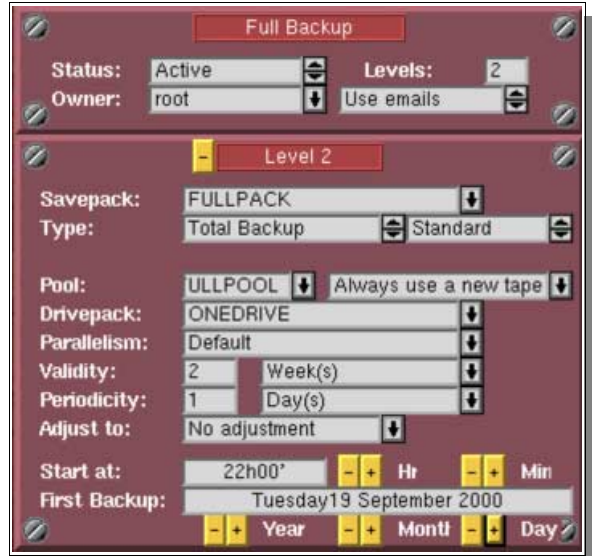
Step 5: Periodic Backup – second level creation

We add a level to “Full Backup”.


“FULLPACK”, “ONEDRIVE” and “FULLPOOL” are added to this new level.

The type is set to “Total Backup”, the Periodicity to “1 day” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start the next Tuesday.



Using the “Advanced Options”, “Enabled for instances” is selected and “1 2 3 4” are added to the field that appears below the drop-down box.



Step 6: Settings verification

The settings should be checked in the Scheduler to make sure they respect the policy defined above.

Number of days 15			
September 18, Monday 2000 22:00	Full Backup	01	Total
September 19, Tuesday 2000 22:00	Full Backup	02	Total
September 20, Wednesday 2000 22:00	Full Backup	02	Total
September 21, Thursday 2000 22:00	Full Backup	02	Total
September 22, Friday 2000 22:00	Full Backup	02	Total
September 25, Monday 2000 22:00	Full Backup	01	Total
September 26, Tuesday 2000 22:00	Full Backup	02	Total
September 27, Wednesday 2000 22:00	Full Backup	02	Total
September 28, Thursday 2000 22:00	Full Backup	02	Total
September 29, Friday 2000 22:00	Full Backup	02	Total
October 02, Monday 2000 22:00	Full Backup	01	Total

II.2.c. Conclusion

As the validity of the backup is two weeks, the tape used during the first backup will be automatically recycled two weeks later, (on Monday the 2nd of October in our example).

II.3. “Total” Backup on Monday, and “Incremental” Backup from Tuesday to Friday

II.3.a. Policy definition

A “Total” backup should be made on Monday, an “Incremental” backup should be made from Tuesday to Friday, every week. Each “Total” backup should be kept for two (2) weeks and each “Incremental” backup for five (5) days.

For the sake of convenience, let us assume that a “Total” backup fills one tape. A new tape should be requested for each “Total” backup and “Incremental” backups should complete existing tapes. There is one tape drive available.

II.3.b. Solution and Analysis

This is almost the same backup as in example two above. The two important changes are: the level two backup is “Incremental” instead of “Total”, and we want to “Complete existing tapes”.

The actual size of the Incremental backups can vary greatly over time. It is more secure to create a specific Tape Pool for the “Incremental” backups, to avoid mixing tapes with the “Total” Backups. The “Complete existing tapes” policy makes this more difficult because it changes the retention date of the tape at each backup.

The validity of the “Incremental” backup is not trivial. It is also important to have an idea of the number of “Incremental” backup needed to fill a tape. As long as it is used by an “Incremental” Backup, a tape has its Validity pushed back in time. The Validity must be set precisely to ensure that new “Incremental” backups will be completed on a tape when the previous tape is recycled.

In the example below, we will assume that three Incremental Backups fill one tape and set the validity to 5 days.

We are going to create a first level weekly backup on Monday and a second level of Incremental backups, based on the complete backup, that will run on Tuesday, Wednesday, Thursday and Friday.

Step 1: Tape pools evaluation

The tape pool for the “Total” backups is first created, simply named TOTALPOOL.

Pool name: TOTALPOOL
 Owner: root
 Comment:

For the “Total” Backup, one tape is used per week.

As each tape should be kept two weeks before it is recycled, the number of tapes should be enough for two weeks of backup. Therefore, at least two tapes are needed in the “Total” tape pool.

Tape name: TOTAL-
 Bar code:
 First number: 1 Last number: 3
 Type: DLT 4000
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: TOTALPOOL
 Comment:

Three tapes will be created in “TOTALPOOL”. This way, one spare tape is always available for convenience.

The tape pool for “Incremental” backups is then created, named “INCRPOOL”.

Pool name: INCRPOOL
 Owner: root
 Comment:

Incremental backup take much less space than Full backup. Furthermore, the “Complete existing tapes” policy ensure that the tapes will be completely used.

Tape name: INCR-
 Bar code:
 First number: 1 Last number: 3
 Type: DLT 4000
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: INCRPOOL
 Comment:

As three backups fill up a tape and as each tape is kept for only five days, only three tapes are necessary.

Three tapes are created in “INCRPOOL”.

Step 2: Drivepack selection

As only one drive is available, our Drivepack contains only this drive and is named (surprisingly) “ONEDRIVE”.


Name : ONEDRIVE
 Owner : root
 Comment
 Number of drives : All
 List of drives

<input type="checkbox"/>	DAT	
<input checked="" type="checkbox"/>	DLT4000	1
<input type="checkbox"/>	FILE_Drive	
<input type="checkbox"/>	NULL Drive	

 Current drive priority: 1


Step 3: Savepack selection

A Savepack is created, which includes all the servers. It is named “FULLPACK”.



Step 4: Periodic Backup – first level creation

A Periodic Backup is created, named “Full Backup”.

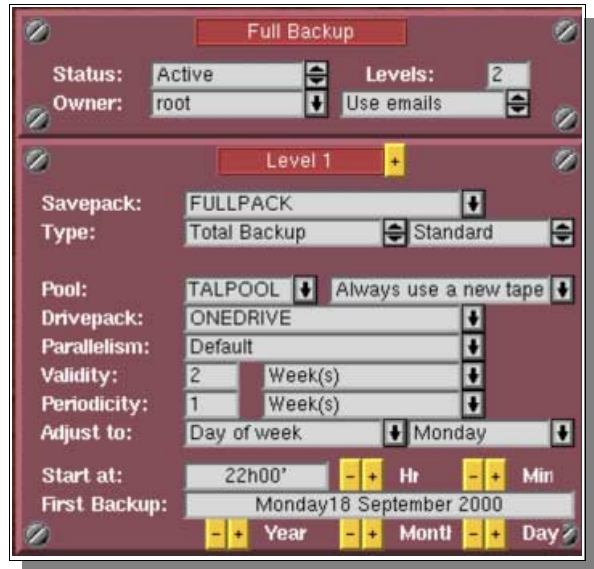


“FULLPACK”, “ONEDRIVE” and “TOTALPOOL” are added to this Periodic Backup.

The type of backup is set to “Total Backup”, the Periodicity to “1 week” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start on the next Monday.

The backup is then adjusted to a “Day of week”: Monday.



Step 5: Periodic Backup – second level creation

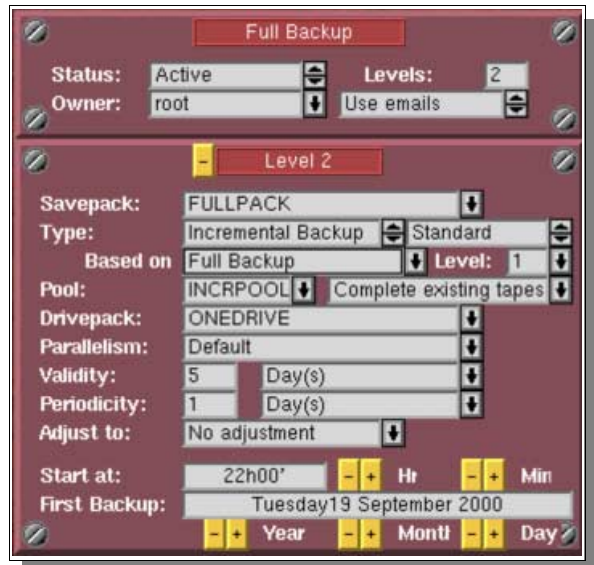
A level is added to the “Full Backup”.

“FULLPACK”, “ONEDRIVE” and “INCRPOOL” are added to this new level.


The level type is set to “Incremental Backup”, based on the level 1 “Full Backup”.

The Periodicity is set to “1 day”, and the Validity to “5 days”.

The tape policy is set to “Complete existing tape” and the backup is set to start on the next Tuesday.



Using the “Advanced Options”, the “Enabled for instances” option is selected and “1 2 3 4” are added in the field that is displayed under the drop–down menu.



Step 6: Verify the settings

The settings should be checked in the Scheduler to make sure they respect the policy defined above.



II.3.c. Conclusion

As the validity of the backup is two weeks, the tape used during the first “Total” backup will be automatically recycled two weeks later (on Monday October the 2nd, in our example).

The first “Incremental” tape is finished on Thursday 21st and recycled on Tuesday 26th. The second starts on Friday 22nd, is finished on Wednesday 27th and is recycled on Monday 2nd. The third is started on Thursday 28th, finished on Tuesday 3rd and recycled on Sunday 8th.

II.4. A more complex and complete backup Policy

II.4.a. Policy definition

A “Complete” backup should be made of the complete network once per month, with a new tape for each backup, and the backup should be kept for a whole year.

Once per week, a “Total” backup should be made of the most important machines and directories. This backup should take place on a Monday, with a new tape for each backup, and the data of this backup should be kept for one month.

An “Incremental” backup should be made every day, Tuesday to Friday, which completes the existing tape, and the data saved by this backup should be kept for one week.

We will suppose that the “Complete” Backup will take three tapes, while the backup of the most sensitive machines requires only one tape. One tape drive is available for all these operations.

II.4.b. Solution and Analysis

The above backup policy is a perfect example of a policy that actually needs two distinct periodic backups:

1. A monthly backup with a Validity of a “Year”.
2. A first level weekly “Total” backup on Monday, with a Validity of a Month, and a 2nd level of “Incremental” type based on it that will run on Tuesday, Wednesday, Thursday and Friday, with a “Week” Validity, as in Example 3.

II.4.c. Part one: the “Yearly” backup

Step 1: Tape pool evaluation

First the tape pool for the yearly “Total” backups is created, with a name of “YEARPOOL”.

Pool name: YEARPOOL
 Owner: root
 Comment:

With this policy, twelve backups are made per year.

Since each backup tape should be kept for one year before recycling, a large number of tapes is needed for one year of backup: 36 tapes at least in the tape pool.

Forty (40) tapes are created in our YEARPOOL. This way four spare tapes are available for convenience.

Tape name: YEARTAPE-
 Bar code:
 First number: 1 Last number: 40
 Type: 3590 CART
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: YEARPOOL
 Comment:

Step 2: Drivepack selection

Since only one drive is available, a Drivepack is created which contains it, and is named “ONEDRIVE”.


Name : ONEDRIVE
 Owner : root
 Comment
 Number of drives : All
 List of drives

<input type="checkbox"/>	DAT	
<input checked="" type="checkbox"/>	DLT4000	1
<input type="checkbox"/>	FILE_Drive	
<input type="checkbox"/>	NULL Drive	

 Current drive priority: 1


Step 3: Savepack selection

One Savepack is created, including all the computers on the network: it is named “SAVEYEAR”.



Step 4: Periodic backup creation

A Periodic backup is created and named “Year Backup”.



“SAVEYEAR”, “ONEDRIVE” and “YEARPOOL” are added to this Periodic backup.

The type of the backup is set to “Total Backup”, the Periodicity to “1 month” and the Validity to “1 year”.

The tape policy is set to “Always use a new tape”.

The backup is set to start on the 1st day of the current month (whatever this day is, it will automatically adjusted to the first Saturday by the following option).

It is then adjusted to a “Day of the Week”: Saturday.



Step 5: Verify the settings

The settings should be checked in the Scheduler to make sure they respect the policy defined above.



II.4.d. Part two: Sensitive machines backup

Step 1: Tape pool evaluation

A tape pool is then created for the monthly “Total” backups, named TOTAL_SENSITIVE.

Pool name: TOTAL_SENSITIVE
 Owner: root
 Comment:

For the “Total” Backup, one tape is used per week.

Since each backup tape should be kept for one month before recycling it, a sufficient number of tapes should be added for one month of backup. Therefore, at least four or five tapes should be created in the tape pool, depending on the month.

Tape name: TOT_SENS-
 Bar code:
 First number: 1 Last number: 6
 Type: DLT 4000
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: TOTAL_SENSITIVE
 Comment:

In the example on the right, six tapes are created in the “TOTAL_SENSITIVE” tape pool. This way, we have one spare tape for convenience.

Finally, the tape pool for the daily “Incremental” backups is created and named “INCR_SENSITIVE”.

Pool name: INCR_SENSITIVE
 Owner: root
 Comment:

“Incremental” backups take less space than a “Total” backup. The “Complete existing tapes” policy ensure that the tapes will be completely used. As in Example 3, we will suppose that only three tapes are necessary.

Tape name: INCR_SENS-
 Bar code:
 First number: 1 Last number: 3
 Type: DLT 4000
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: INCR_SENSITIVE
 Comment:

Therefore, three (3) tapes are created in “INCR_SENSITIVE”.


Step 2: Drivepack selection

Since only one drive is available, it is added to the Drivepack, which is named “ONEDRIVE”.




Step 3: Savepack selection

A Savepack is created, including all the sensitive machines. It is named “SENSITIVE”



Step 4: “Sensitive Backup” – first level creation

A Periodic Backup is created and named “Sensitive Backup”.



“SENSITIVE”, “ONEDRIVE” and “TOTAL_SENSITIVE” are added to this Periodic backup.

The type of the backup is set to “Total”, the Periodicity to “1 week” and the Validity to “1 month”.

The tape policy is set to “Always use a new tape” and the backup is set to start next Monday.

The backup is adjusted to “Day of week”: Monday.

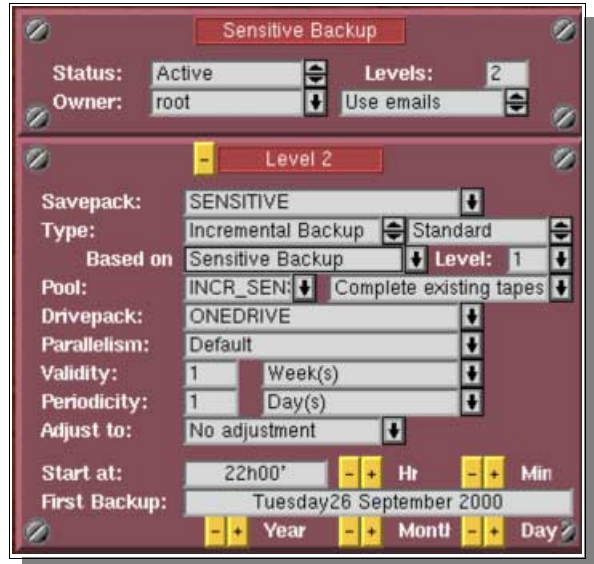


Step 5: “Sensitive Backup” – second level creation

A level is added to “Sensitive Backup”.

“SENSITIVE”, “ONEDRIVE” and “INCR_SENSITIVE” are then added to this new level. We set the type of the backup to “Incremental”, based on the level 1 of “Sensitive backup”, the Periodicity to “1 day” and the Validity to “1 week”.

The tape policy is set to “Complete existing tape” and the backup is set to start next Tuesday.

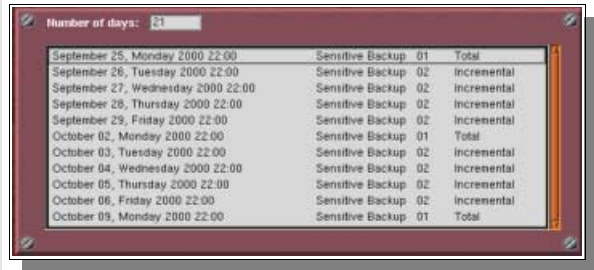


Using the “Advanced Options”, “Enabled for instances” is selected and “1 2 3 4” are added in the field that is displayed below the drop-down box.



Step 6: Verify the settings

The settings should be checked in the Scheduler to make sure they respect the policy defined above.



II.4.e. Conclusion

Most policies can be split into simpler backups. Separate tape pools should simply be created to handle for different backups.

II.5. Two simultaneous “Total” Backups on two different domains, from Monday to Friday

II.5.a. Policy definition

A network exists, with two different domains. A backup should be made of the machines of both domains without mixing the tapes of these two different backups. Two specific sets of tapes should be created, one for each backup. For both domains, a “Total” backup should be made every day, from Monday to Friday, every week and each backup should be kept for two weeks.

For the sake of convenience, let us assume that each backup fills about one tape and a new tape should be used for each “Total” backup. A library is available, which contains four tape drives.

II.5.b. Solution and Analysis

We have here two backups like the one detailed in example 2.

Since more than one drive is available, though, both backups can be run at the same time.

Two identical backups are going to be created: a first level weekly backup on Monday and a second level based on it that will run on Tuesday, Wednesday, Thursday and Friday.

A specific tape pool, Drivepack and Savepack will be created for each backup.

Step 1: Tape pool evaluation

Two tape pools, “TOTALPOOL1” and “TOTALPOOL2” are created, one for each of the weekly backups.




With the policy defined above, five tapes are used per week by each backup.

As each tape should be kept two weeks before recycling it, enough tapes should be defined for two weeks of backup: ten (10) tapes should be defined in each tape pool.

Eleven (11) tapes will be created in each tape pool. In this way, one spare tape will be available per backup for convenience.

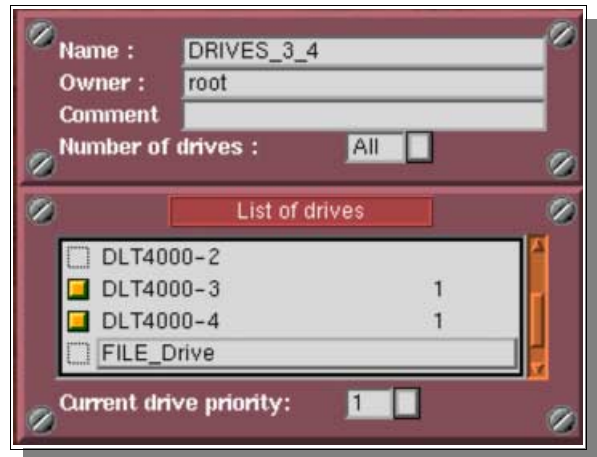


Step 2: Drivepack selection

Two Drivepacks should be defined, to allow a dual, simultaneous, backup operation.




Two drives are added to each Drivepacks. These Drivepacks are then named "DRIVES_1_2" and "DRIVES_3_4".




Step 3: Savepack selection

One Savepack is created for each domain, named: "DOMAIN_1" and "DOMAIN_2".



Step 4: "Domain 1 Full Backup" – first level creation

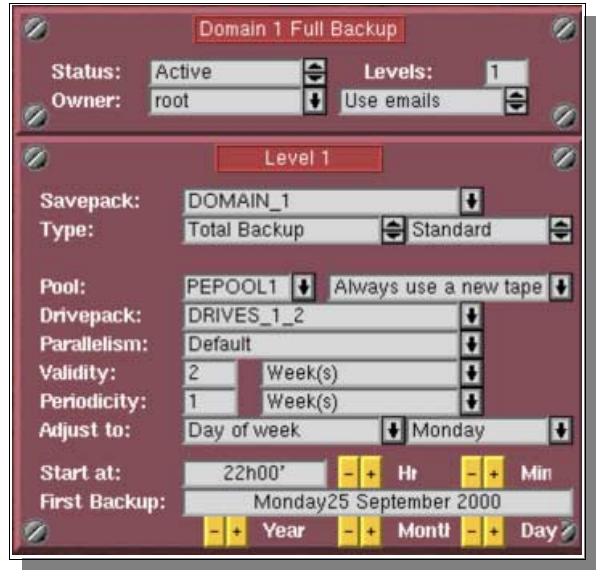
A Periodic Backup is created for the first domain, appropriately named: "Domain 1 Full Backup"



“DOMAIN_1”, “DRIVES_1_2” and “TOTALPOOL1” are then added to this first level.

We set the type of the backup to “Total Backup”, the Periodicity to “1 week” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start on the next Monday.



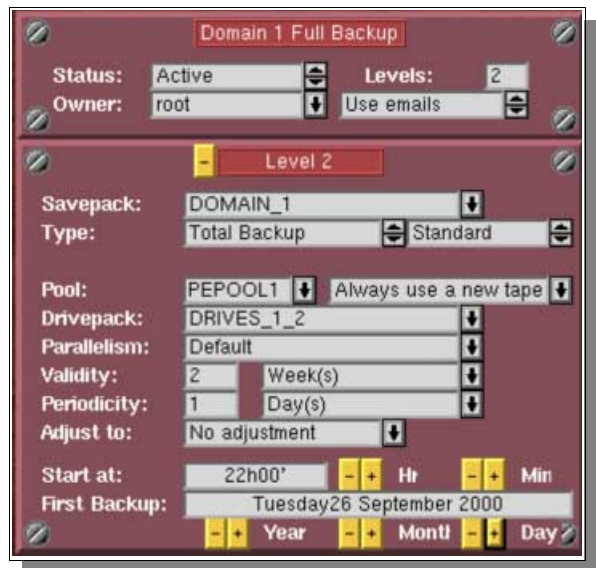
Step 5: “Domain 1 Full Backup” – second level creation

A new level is added to “Domain 1 Full Backup”.


“DOMAIN_1”, “DRIVES_1_2” and “TOTALPOOL1” are added to this level.

The backup type is set to “Total Backup”, the Periodicity to “1 day” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start next Tuesday.




“Enabled for instances” is selected in the “Advanced Options”, and “1 2 3 4” is added in the field that appears under the drop-down menu.



Step 6: “Domain 2 Full Backup” – first level creation

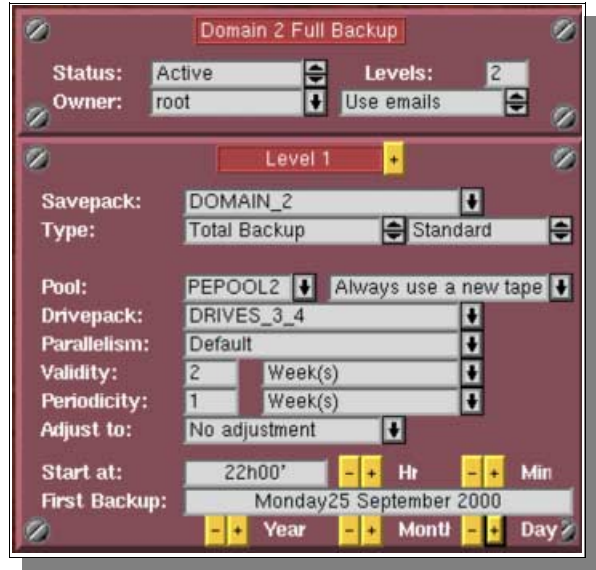
We create our Periodic Backup called “Domain 2 Full Backup”



“DOMAIN_2”, “DRIVES_3_4” and “TOTALPOOL2” are added to this backup.

The backup type is set to “Total Backup”, the Periodicity to “1 week” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start on the next Monday.



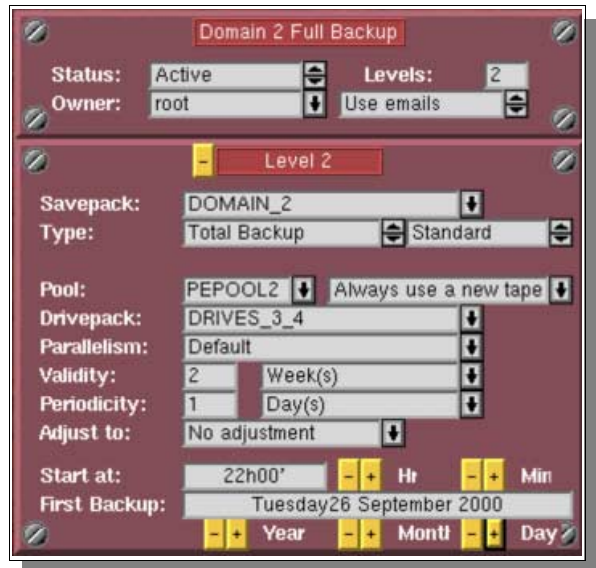
Step 7: “Domain 2 Full Backup” – second level creation

A level is added to “Domain 2 Full Backup”.

“DOMAIN_2”, “DRIVES_3_4” and “TOTALPOOL2” are added to this new level.

The backup type is set to “Total Backup”, the Periodicity to “1 day” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start on the next Tuesday.



“Enabled for instances” is selected, using the “Advanced Options”, “1 2 3 4” are then added in the field that appears under the drop-down box.



Step 8: Verify the settings

The settings should be checked in the Scheduler to make sure they respect the policy defined above.



Number of days: 30

September 25, Monday 2000 22:00	Domain 2 Full Ba 01	Total
September 25, Monday 2000 22:00	Domain 1 Full Ba 01	Total
September 26, Tuesday 2000 22:00	Domain 2 Full Ba 02	Total
September 26, Tuesday 2000 22:00	Domain 1 Full Ba 02	Total
September 27, Wednesday 2000 22:00	Domain 2 Full Ba 02	Total
September 27, Wednesday 2000 22:00	Domain 1 Full Ba 02	Total
September 28, Thursday 2000 22:00	Domain 2 Full Ba 02	Total
September 28, Thursday 2000 22:00	Domain 1 Full Ba 02	Total
September 29, Friday 2000 22:00	Domain 2 Full Ba 02	Total
September 29, Friday 2000 22:00	Domain 1 Full Ba 02	Total
October 02, Monday 2000 22:00	Domain 2 Full Ba 01	Total
October 02, Monday 2000 22:00	Domain 1 Full Ba 01	Total
October 03, Tuesday 2000 22:00	Domain 2 Full Ba 02	Total
October 03, Tuesday 2000 22:00	Domain 1 Full Ba 02	Total

II.5.c. Conclusion

In this example, both backups are actually identical.

They could be completely different, while still managed simultaneously. For instance, one of these backups could have a second level of an “Incremental” type.

II.6. A “Total” backup once a week plus a daily backup of the modified files

II.6.a. Policy definition

Every week, a “Total” backup should be made on Monday and another backup, that saves only the modified files since the day before, should be executed from Tuesday to Friday. Each “Total” backup should be kept for two (2) weeks and each daily backup should be kept for five (5) days.

For the sake of convenience, let us assume that a “Total” backup and a daily backup fill almost one full tape each.

A new tape should be used for each “Total” backup and the daily backup should complete existing tapes. One tape drive is available.

II.6.b. Solution and Analysis

This problem is very similar to the example 3 above. The only real difference is in the way the “Incremental” backup is defined.

In example 3 we have seen above, the incremental backup was based on the “Total” backup. Here, it will be based on itself and will run independently of the “Total” backup (except that it won’t run on the “Total” backup’s day). This is called a “differential” backup.

Step 1: Tape pool evaluation

A tape pool is created for the “Total” backups, it is named “TOTALPOOL”.

Pool name: TOTALPOOL
 Owner: root
 Comment:

For the “Total” Backup, one tape is used per week.

Since each tape should be kept for two weeks before recycling it, enough tapes should be created for two weeks of backup: Therefore, the “Total” tape pool requires at least two tapes.

Tape name: TOTAL-
 Bar code:
 First number: 1 Last number: 3
 Type: DLT 4000
 Owner: root
 Authorizations: Read Write Recycle
 Delete Clean
 Recycling dest.: current pool
 Recycling mode: FIFO
 Current pool: TOTALPOOL
 Comment:

Three tapes are created in “TOTALPOOL”. In this way, one spare tape is always available for convenience.

The Tape pool for incremental backups is then created and named “INCRPOOL”.

Pool name: INCRPOOL
 Owner: root
 Comment:

“Incremental” backups take less space than “Total” backups, particularly when they are based on themselves. The “Complete existing tapes” policy ensures complete use of tapes.

As a week of backups fills up a tape and as a tape is kept for five days, only one tape is really necessary.

For the sake of convenience, two tapes are created in our “INCRPOOL”.




Step 2: Drivepack selection

As only one tape drive is available, one Drivepack is created, including this one drive, and is called “ONEDRIVE”.



Step 3: Savepack selection

A Savepack is created, including all the servers to be backed up. It is named “FULLPACK”.



Step 4: “Full Backup” – first level creation

A Periodic Backup is created, called “Full Backup”.



“FULLPACK”, “ONEDRIVE” and “TOTALPOOL” are added to this first level.

We set the type of the backup to “Total Backup”, the Periodicity to “1 week” and the Validity to “2 weeks”.

The tape policy is set to “Always use a new tape” and the backup is set to start on Monday.

The backup is adjusted to a “Day of week”: Monday.



Step 5: “Full Backup” – second level creation

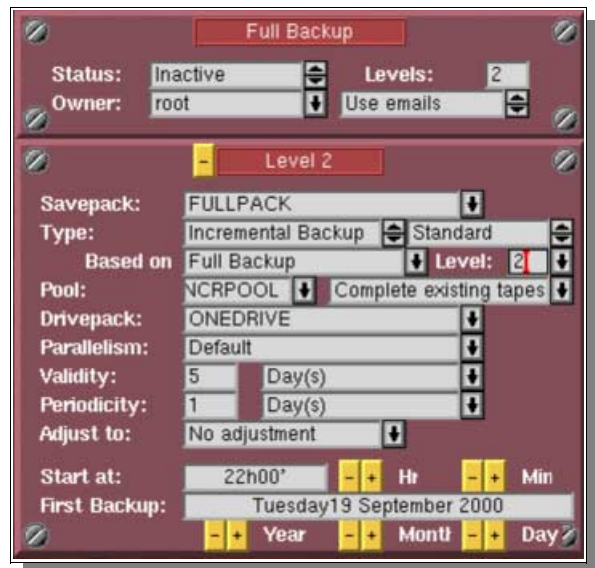
A level is added to “Full Backup”.

“FULLPACK”, “ONEDRIVE” and “INCRPOOL” are added to this new level.


The backup type is set to “Incremental Backup”, based on the “Full Backup”, **level 2**.

The Periodicity is set to “1 day” and the Validity to “5 days”.

The tape policy is set to “Complete existing tape” and the backup is set to start on the next Tuesday.



“Enabled for instances” is selected in the “Advanced Options”, and “1 2 3 4” are added in the field that appears below the drop-down menu.



Step 6: Verify the settings

The settings should be checked in the Scheduler to make sure they respect the policy defined above.



Number of days 15			
September 18, Monday 2000 22:00	Full Backup	01	Total
September 19, Tuesday 2000 22:00	Full Backup	02	Incremental
September 20, Wednesday 2000 22:00	Full Backup	02	Incremental
September 21, Thursday 2000 22:00	Full Backup	02	Incremental
September 22, Friday 2000 22:00	Full Backup	02	Incremental
September 25, Monday 2000 22:00	Full Backup	01	Total
September 26, Tuesday 2000 22:00	Full Backup	02	Incremental
September 27, Wednesday 2000 22:00	Full Backup	02	Incremental
September 28, Thursday 2000 22:00	Full Backup	02	Incremental
September 29, Friday 2000 22:00	Full Backup	02	Incremental
October 02, Monday 2000 22:00	Full Backup	01	Total

II.6.c. Conclusion

The first time the Level 2 backup will be run, a Total backup will actually be made, as no reference backup will be available. Then it will proceed by only saving files modified since the day before.

CHAPTER 9

Restoration

I. Principles of Restoration

I.1. Definition

Restoration is an interactive function that can be used to recover a complete machine or a single file. It may be selective depending on the “role” of the user.

The large number of files makes this potentially difficult. It can become quite hard to correctly identify the tape that contains the latest version of a file.

I.2. Arkeia’s approach

To overcome this problem, Arkeia integrates a database index. The index allows you to use the navigator to select the files/directories to restore and Arkeia then determines the backup media to be used for restoration.

The database is updated with each backup. It is an interactive procedure that ensures that data already archived can be retrieved even if the backup operation is suddenly interrupted. The database can only be modified by Arkeia.

Files to be restored are selected with the help of a time navigator that allows the user to view files backed up at different dates. Search operations are available to display the objects backed up during the period.

If the data to be restored is divided into different backups (total and incremental), Arkeia will perform tape scan selection without assistance. Restoration is automatic, if the required tapes are accessible. If the last version of the files to be restored cannot be accessed, Arkeia will indicate which older files that can be accessed.

A log is displayed during the restoration. It shows the operation under way, the tree–structures restored and the errors encountered. As with backup, the user may stop the Arkeia graphical interface without interrupting the restoration and can reconnect at any time.

II. Restoration Management

II.1. The “Restoration” screen

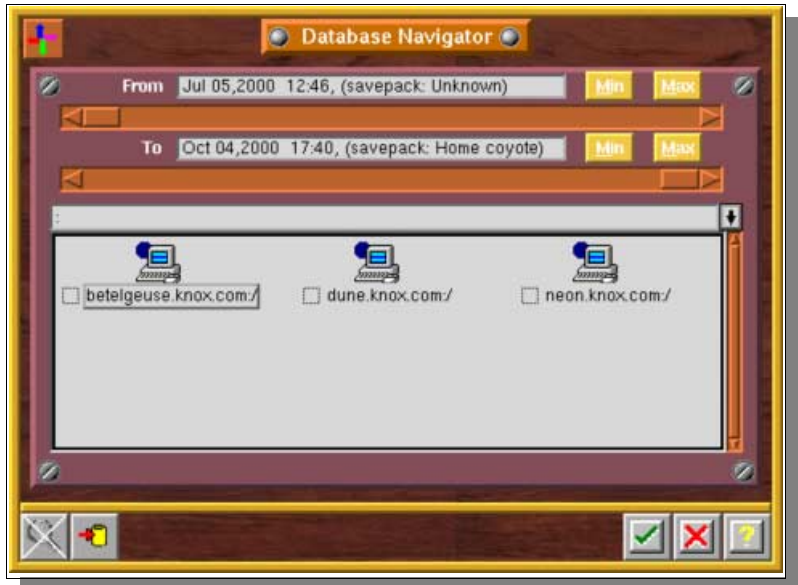
From the main screen click on the [Restoration] menu then select the [Restoration] option.



Or use the “Restoration” button in the Toolbar:



If no tree was previously selected in the GUI session, Arkeia automatically launches the Database Navigator.



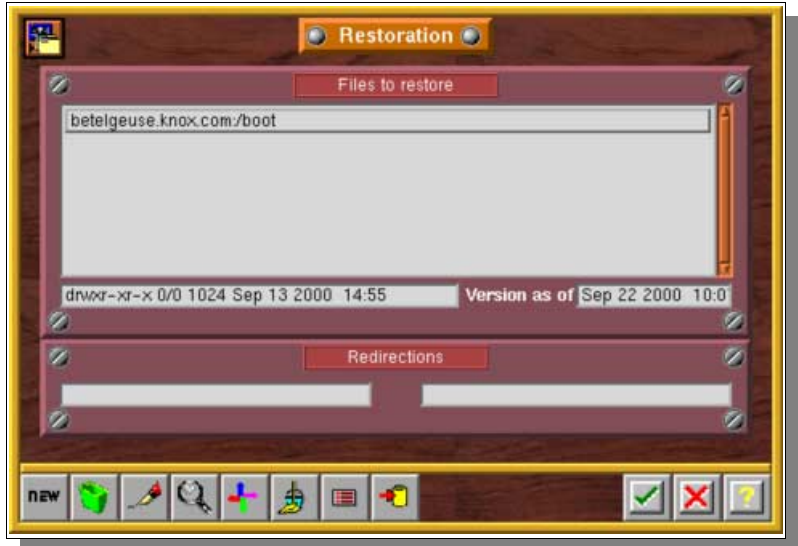
Once a tree is selected, the restoration screen displays the tree(s) selected.

Files to restore:

Selected trees or files to be restored

Redirection:


Path for the restoration



II.2. Select what to restore: the “Time Navigator”

The “Time Navigator” (also called the “Database Navigator”) is used in the same way as the network navigator. It has two additional scroll bars for time navigation.


To go to this screen from the main screen "Server administration", click on the [Restoration] menu then on the [Restoration] option.



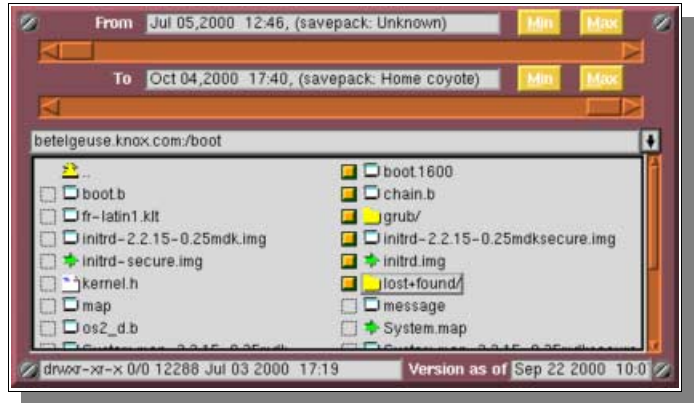
You can also open the “Time Navigator” by clicking on the “Navigator” button in the “Restoration” screen.



Use the “Time sliders” to choose an amount of time between two backup dates or a specific date (see below).



Select the files and directories you want to restore. It is possible to select an entire tree if needed.



Confirm your choices by clicking on the “checkmark” (OK) button.



Please note: when the restoration function has been started and no tree-structure has been selected, Arkeia automatically runs the restoration navigator, otherwise the restoration screen shows the tree-structure(s) already selected.

II.3. Select a single tree or file

An alternative way to path selection is to press the “New” button of the “Restoration” screen



Enter the tree or the complete path of the file to be restored.



Confirm your choices by clicking on the “checkmark” (OK) button.



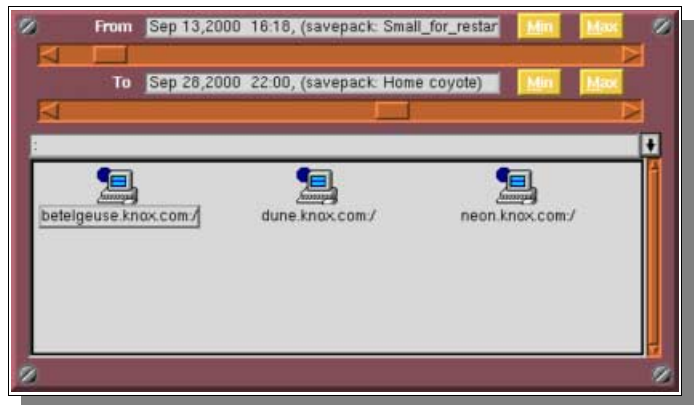
II.4. Using the “time sliders” in the Time Navigator

The “time sliders” are a convenient way to select a specific version of a file or tree. Each slider’s position corresponds to a specific backup that was previously run. The “time sliders” can be used in two ways:

You can either select a period of time between two dates, including all the backups made in that period.

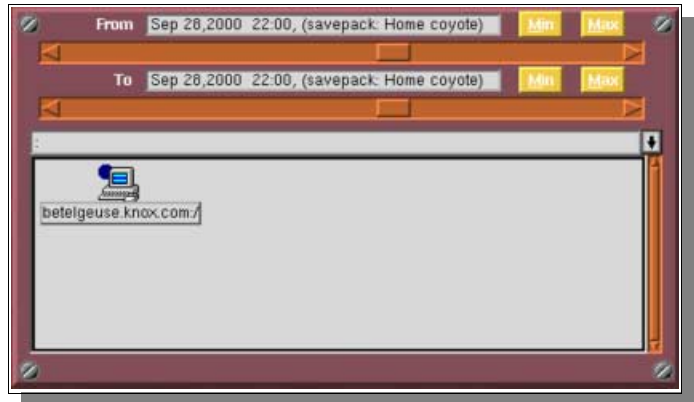
In that case, all the files saved during that period of time will be restored.

You can automatically select the complete backup period using the “Min.” and “Max.” buttons.



Or you can select a specific backup, by setting both sliders on the same position (= on the same backup).

In this case, the files which were backed up at that precise date will be restored. If the backup was incremental, Arkeia will automatically search through older incremental backups to restore the files that were not backed up on the date selected.

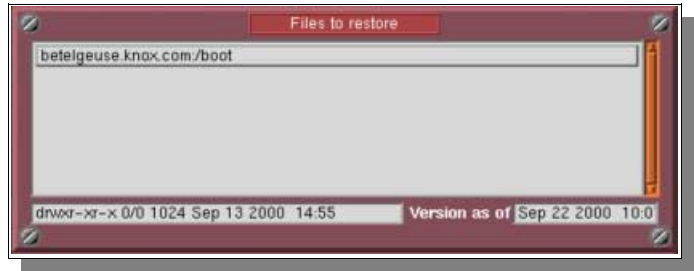


- Please note:** in the navigator window itself, the available trees are displayed according to the backup selected with the “time sliders”. Once the sliders are set, you can navigate among the trees and select the files you want to restore.
- Please note:** The name of the Savepack that corresponds to the backup is displayed in the date field. If this Savepack has been deleted from the list of Savepacks, the message displayed is “UNKNOWN”.

II.5. Modify a path or a name

You can modify a selected path or name for restoration:

Select the path you want to change.



Click on the “Pen” button to modify the path



Modify the path.



Confirm your choices by clicking on the “checkmark” (OK) button.



II.6. Paths syntax

The standard syntax of trees in Arkeia is the following: *[Machine name]: [Path]*

It uses the UNIX “/” (slash) separators.

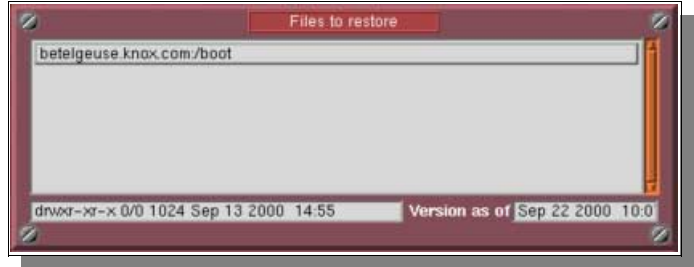
Windows Paths have to include the letter of the disk (“c:”, “d:”, etc.)

Example	<i>mars:/usr/home</i>	(unix)
	<i>silos:c:/windows/tmp</i>	(windows)

II.7. Backup information of a file or tree

It is possible, for each file you want to restore, to know which backup operation(s) saved that file, on what date, to what tape, etc. To do this, use the following procedure:

Select the file or tree you want to display information on.



Click on the “Magnifying Glass” button to open the information screen



Name:

Name of the file or tree

Type:

Type of data (file, tree)

Size:

Size of the file or tree

Date:

Date of last modification on the host machine



Label:

Label of the tape on which the selected version of the file or tree is recorded

From... To...:


Segments of the tape on which the file or tree are saved

List of Backups:

List of the backups that include a version of the file or tree.

This information can also be displayed from the “Time Navigator”:

Open the “Time Navigator” and select the file or tree you want to have information on.



Click on the “Magnifying Glass” button to open the information screen



The information dialog is then displayed on your screen.



II.8. Applying Redirection

Arkeia allows you to redirect a tree or a file restored:


- to another directory
- under another name
- to another machine

Everything stated after the path in the first field will be redirected to the directory given in the second field. The substitution takes place at character level.

The syntax is the following for the source and the destination (the name of the volume is needed for W95, NT, Novell): *machine:[volume:]/...*

Example Redirection to another directory

The file rs6000:/usr/bin/ls should be restored:



The file will be restored to: rs6000:/tmp/bin/ls

Example Redirection to another name

The file rs6000:/usr/bin/ls should be restored:



The file will be restored as: rs6000:/tmp/bin/ls.back

Example Redirection to another machine

The file rs6000:/usr/bin/ls should be restored:



The file will be restored to: mars:/tmp/bin/ls

II.9. Searching a file to restore

This feature helps you find files in the database.

Click on the search files button to open the “Search files” screen



Start of search:

Tree to start search

Criteria:

Search criteria (name, strings in name, etc.)



II.9.a. Search criteria

Filename matching exactly

The search keyword must be a complete file name or directory name (respect lowercase and uppercase)

Filename containing:

The keyword must exist in the file name or directory name (*Keyword*).

Filename ending by

The keyword must exist at the end of the file name or directory name (*Keyword).

Filename starting with

The keyword must exist at the beginning of the file name or directory name. (Keyword*).

When a file or directory is found, Arkeia adds it in the restoration screen list or selects the file if the navigator is open.

II.10. The “Restoration” monitor

This screen allows you to monitor the restore function in real time:

A dynamic events log provides information on the files being restored and the problems encountered.

As with backup, you can stop the Arkeia interface without interrupting the restoration in progress and restart the interface at any time without having any influence on the operation(s) in progress.

Several restorations and backups can take place at the same time.

A restoration must be launched to go to this screen from the main screen “Server administration”



II.11. Index browser

The index browser is the consultation mode for restoration. You can directly start it from the GUI without starting a restore.


From the main screen click on the [Restoration] menu, then click on the [Index Browser] option.



Navigation in the index browser is identical to the time navigator

Index browser can be used to query the Arkeia base, by viewing:


- All backups performed.
- The different file instances.
- File attributes
- The tape containing file instances.
- The segment number with the file instance.



II.12. Restoration options

The restoration option allows you to configure the file attributes to restore.

From the main screen click on the [Restoration] menu then click on the [Restoration options] option.



Alternatively, you can click on the “Restoration options” button of the “Restoration” screen



Access rights:

Restore original file access rights or use current access rights mask

Modification date:

Restore original date or local date

Files modified since backup date:

Erase file existing by restored file

By user name:

Arkeia will try to find the same user name on the destination

By user ID:

Arkeia will try to find the same user ID on the destination



II.13. List of tapes used for the restoration

Before each restoration, Arkeia displays the list of tapes used for the operation. In the event you unselect one of the tapes, Arkeia will automatically search for other tapes containing the file(s) to restore.



You can list the files saved on that tape by clicking on the “Magnifying glass” button.



The “Files in tapes” window is then displayed on the screen.



II.14. Who has access to the restore function?

II.14.a. The administrator

An Arkeia administrator can restore any file to any machine by default. For more information on this subject, please refer to chapter 12, “Security”, of this manual.

II.14.b. The operator

An operator can restore his or her own files on the machine if he/she has been given the rights to do so.

II.14.c. The user

A user can restore files if he/she has access rights to them (for example the user’s own files).

I. How to increase performance

I.1. Introduction

As a general rule, the only 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.

Particularly, 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.

I.2. Performance expectations

On a 100Mb network, you can expect backup speed anywhere from 300 – 1000 MB/minute. Backup speed will depend on the client, server, network and tape drive(s)/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/s, with a 100Mb/s Ethernet network, it gives 12,5 MBytes/s theoretical traffic, meaning 7 to 8 MB/s of real throughput.

You can expect to reach this value on the network with 10 to 15 flows of 0,83 MB/s. If you have 25 clients machines, you can divide this 25 clients in two group of 12 and 13 clients and with a license of 15 flows you will be able to optimize your backup performance.

I.3. Use of multiflows

I.3.a. Basics

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 multiflow, 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.

I.3.b. Parallelism of multiple machines

By default, Arkeia backs up multiple machines using one flow per machine. The parallelism is automatically configured and performance is optimum.

I.3.c. Parallelism of a specific machine

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 a flow number.

Then trees with the same “Multiflow” number are backed up sequentially.

To change the “Multiflow” field, open [Backups], [Savepacks], [Tree options].

Then edit the Multiflow field as necessary.



Please note: this feature is useful mainly on machines with several disk drives. Usually, the idea is to use multiflow on each file system. Computers with a single disk drive will generally not benefit too much from multiflow.

Please note: the Priority field, in *Tree options*, only comes into play between trees of the same Multiflow number.


1.3.d. Last settings

Once the *Tree Options* are set, there is something else to take into account: you may have to indicate Arkeia that the backup is to use a certain number of flows.

By default, Arkeia will use all your flows for the backup and if you have another backup launched at the same time, or just after, this second backup operation will have to wait the end of the first one.

You can also decide to limit the number of flows that will be used:

In the “Backup definition” window, change the “Parallelism” backup field.



✚ Please note: the number of flows you can use simultaneously is determined by your software license. If you need more flows, please send a request to “sales@arkeia.com”.

Example

Let’s assume a machine with four disk drives, three of them being partitioned. That means there are seven file systems /u1, /u2, /u3, /u4, /u5, /u6 and /u7. Sixteen (16) flows are available, but the limit should be 4 flows for this machine.

In the Savepack, using “Tree Options”, each partition can be assigned to a specific flow:

Flow 1: /u1 and /u2

Flow 2: /u3 and /u4

Flow 3: /u5 and /u6

Flow 4: /u7

When the backup is defined, the Parallelism is set to four, limiting the use of flows to four.

II. How to limit backup speed

II.1. Why should backup speed be limited?

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.

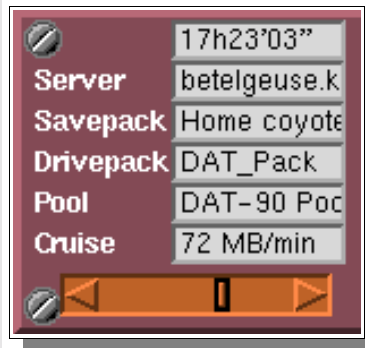
In this case, it can be useful to control the backup speed, thus making sure that your network won't be completely saturated.

II.2. How to limit backup speed

II.2.a. The graphical “cruise control”

With the graphical interface, you can limit the backup speed by using the cruise control slider.

In the “Backup Monitor” window, use the slider to change the backup speed.



II.2.b. 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 `/usr/knox/arkeia/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”.

All the backups will then use this preference.

You can still use the cruise control from the GUI to change the current speed of a specific backup.

Example	<i>SOCK_SPEED</i> “1024” means a throughput limit of 1024kB/s, i.e. 60MB/minute.
----------------	---

III. Priority

III.1. Introduction

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 (see the next section for more information on chaining).

Priority is set between trees saved by the same flow. Trees backed up by different flows are saved simultaneously.

III.2. How to use Priority

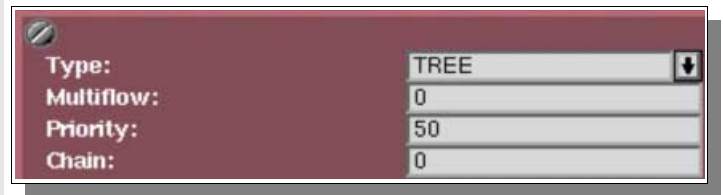
The use of Priority is extremely simple.

In the “Tree” Options, for each tree of the Savepack, you can add a Priority field to your 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”)



The screenshot shows a dialog box with a dark red header and a light gray body. It contains four rows of settings:

Type:	TREE
Multiflow:	0
Priority:	50
Chain:	0

Trees will then be processed from the highest priority down to the lowest.

🔴 Please note: this 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.

IV. Chaining

IV.1. Introduction

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.

IV.2. How to use Chaining

The use of Chaining is very simple.

In the “Tree Options”, for each tree of the Savepack, you can 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.



Type:	TREE
Multiflow:	0
Priority:	50
Chain:	0

❗ Please note: Chaining has to be used with the “Priority” function. These two functions are complementary and are actually almost useless when set alone.

V. Configuring compression in Arkeia

V.1. Introduction

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.


You may want to increase compression, or adjust it according to the type of file being backed up.

V.2. Compression settings

V.2.a. 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.



V.2.b. 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 /usr/knox/nlp/nlp.cfg:

```
“PREFERRED_COMPRESSION” “[value]”
```

Where [value] can be: 0, 1 or 2. This code correspond to the following compression:

Code	Compression algorithm
0	No compression
1	LZ1, fast compression
2	LZ3, maximum compression ratio (slower)

V.2.c. Setting a compression algorithm to a specific file type

It is possible to configure a client to use a specific algorithm according to the type of file backed up. This avoids losing CPU time to process files which are already heavily compressed (like JPEG or video files). This also enable to select a compression type appropriate to some type of files.

Example

Exclude compression for .mov files but compress .txt files using LZW3_3.

- *Add “.mov” to /usr/knox/obs/nocompr.ext*
- *Add “.txt” to /usr/knox/obs/lzrw3_3.ext*

Arkeia Management and Administration

1. How to move the index database to another directory

1.1. Problem definition

There are some situations when a System Administrator may have to move Arkeia's index database to another directory on the same machine. For instance, these situations include a full file system, reorganization of the tree structure of the backup server, etc.

To move the database is a delicate task: as some information is contained and optimized within the software itself to increase its speed, it has to be updated with the new Arkeia location to work properly.

The goal of this section is to provide you with the correct procedure.

1.2. Procedure

Follow the steps of the procedure below to move the database to another location:

1. Stop all Arkeia processes and close the graphical interface if it is currently opened. Stop MDL if necessary.
2. Use full path names for each of the commands below.
3. Create a tarball of `/usr/knox/arkeia/dbase/o3dbtree` and all of its subdirectories.
4. Make two copies of this tarball for backup.
5. On a different disk drive, create a new directory, for instance `/u5/arkeia/o3dbtree`
6. In this directory, create an empty `.OPB_NOBACKUP` file.
7. Rename `/usr/knox/arkeia/dbase/o3dbtree` in `/usr/knox/arkeia/dbase/o3dbtree.old`

8. Create a symbolic link to `/u5/arkeia/o3dbtree/` (new directory) from `/usr/knox/arkeia/dbase/o3dbtree` (old directory).
9. Expand the tarball created in step 3 into `/u5/arkeia/o3dbtree`
10. List the contents of `/usr/knox/arkeia/dbase/o3dbtree` and `/u5/arkeia/o3dbtree` directories, the results should be the same. If not, please check the steps listed above.
11. Create a TEST file in your `/home` directory
12. Start Arkeia and perform a small interactive backup of the TEST file in your `/home` directory
13. Delete the TEST file you just backed up
14. Restore the TEST file to verify that restore works
15. You can now perform regular backups and delete the `/usr/knox/arkeia/dbase/o3dbtree.old` directory. Alternatively, you may want to wait until the next “Total” backup to perform this last step.

II. Move the Arkeia backup server to a new machine

II.1. Problem definition


The next step for Administrators can be to upgrade their system to a new backup server. In order to get back the complete database, follow the procedure below.

II.2. Procedure

Follow this procedure to move the database to another machine:

1. On the old backup server, stop all Arkeia processes and close the GUI.
2. Use full path names for each of these commands
3. On the the old backup server, create a tarball from `/usr/knox/` and all of its subdirectories.
4. Make two copies of this tarball for backup
5. On the new backup server expand the tarball created in step 3 into the directory `/usr/knox`
6. On the new backup server, launch a normal installation:
 - 6.a CLIENT installation
 - 6.b SERVER installation
 - 6.c GUI installation
7. Configure the Arkeia clients to work properly with the new backup server:
 - 7.a Remove the `/usr/knox/nlp/rhost.lst` file
 - 7.b Restart NLSERVD
8. **On each client:**
 - 8.a Modify the administration server name in the file `/usr/knox/nlp/admin.cfg`, and set the new backup server name
 - 8.b Restart NLSERVD.
 - 8.c Check that the clients have correctly declared themselves to the backup server
9. Perform a backup on a NULL device

10. Set the drive and the library on the new backup server
11. Perform a backup using the tape drive or library
12. Perform the restoration of an old backup. Check that you can browse the database index using the graphical interface in the restoration screen
13. You can now perform regular backups on the new backup server

 **Please note:** you should keep the old backup server in activity until you have made sure that your new backup server is working correctly

III. Removing an installed client

III.1. Problem definition

If you replace an old computer by a new one with a different name, you may need to uninstall the client to avoid running into license issues.

To uninstall a client is, more exactly, to remove references to this client on the backup servers, in the Savepacks and in the database. It will also allow you to get back some disk space, as the information in the database related to this machine will also be deleted.

Removing a client is a simple operation, but it requires some operations nonetheless.

III.2. Procedure

Follow this procedure to remove an installed client:

1. Stop the *nlservd* daemon on the backup server using the “*NLSERVD stop*” command.
2. Edit the */usr/knox/nlp/rhost.lst* file and remove the *ITEM* that corresponds to the client.
3. Go to */usr/knox/arkeia/dbase/o3dbtree*
4. Remove the directory of the host you want to delete with “*rm -fr*”
5. Remove the file */usr/knox/arkeia/dbase/o3dbtree/o3_cpnt*
6. Restart the *nlservd* daemon with “*NLSERVD*”
7. Install the new clients if necessary
8. Start an interactive backup of a small file (Ex: */etc/hosts*) on each machines installed onto the network. The interactive backup must be done to a real tape. DO NOT use a NULL drivepack.

IV. Creating Arkeia users

IV.1. Problem definition

By default, Arkeia creates a **ROOT** user who has ADMINISTRATOR privileges, meaning he can backup and restore any machine on the Arkeia network.

However, the Administrator may want to create other users, with various rights.

This section shows how to create new accounts.

Please note: the “root” login does not have a password after the installation has been completed. You should always assign a password to this account **immediately**.

IV.2. Creating a user

IV.2.a. The Users Management window

To create a new user, go into the “Users Management” window.

1. From the main screen click on the [Utilities] menu then on the [Users Management] option.



Name:

Name of the user

Role:

Role of the user: Administrator, Operator or User.

Email:

Email address of the user for reports

List of machines:

Machines the user has access to for backup and restore.

List of users:

Self-explanatory.



IV.2.b. User creation

In the “Users Management” window, click on the “New” button.



Enter the name of the user, choose his/her “Role”, enter his/her email address and select what machines are accessible to him/her among the ones who have been backed up.



❖ **Please note:** you must have ADMINISTRATOR rights to create a user.

❖ As soon as a user is created, you should assign him/her a password.

IV.2.c. User modification

In the “Users Management” window, select the user you want to modify, then click on the “Pen” (modify) button.



IV.2.d. User deletion

In the “Users Management” window, select the user you want to delete, then click on the “Trashcan” icon.



V. Setting up the email feature

V.1. Problem definition

A System Administrator may want to get reports of backups anywhere, sent by email to a specific address. Arkeia provides such a feature.

V.2. Procedure

To get emailed reports, there are two parameters to be set: the “Owner email” and the “User email”.

V.2.a. The “Owner email”

The “*Owner email*” is the email of the user who created the backup. If not configured, when the user was created, it can be added in the “*User Management*” window, using the “*Modify a user*” procedure (see above).

V.2.b. Configure the email feature in Backups

The next step is to set the email feature (“*User emails*”) in the Backup definition. For more information on this feature of the backup, please refer to the *Backups* chapter.

V.2.c. Notifications provided

Arkeia will send emails to the defined user when:

- Arkeia needs a tape during a backup
- Backup is finished

🔴 **Please note:** in the case of an interactive backup, only the current user can receive emails.


I. Encryption Configuration

I.1. Introduction

Some System Administrators may be worried by the regular transfer of critical or confidential data on the network. For instance, some intruders may try to intercept network packets and use the information uncovered.

To provide a high level of security, even during local transfers, and to allow protection for the tape itself, Arkeia offers a data encryption feature. Files are encrypted on the client during the backup, before being sent to the backup server through the network.

Therefore, encryption should be configured on the client side. You'll find below the procedure to do this.

 **Please note:** a special license key is necessary to use the encryption feature.

I.2. Encryption configuration

I.2.a. Encryption Type

Arkeia features two encryption algorithms. The choice between these algorithms should be made according to the level of security you need and the CPU time you can afford to encrypt the data on the clients.

The algorithms are:

- DES
- BLOWFISH

1.2.b. Configuration

There are two steps in the encryption configuration. Though most of the configuration is on the client, there is a parameter that must also be set on the backup server.


Arkeia provides multiple option for the encryption of data. For instance, you can apply encryption to a whole tree or to a single directory. For both these cases, a specific key for encryption can be used.

The two steps are:

1. Select a type of encryption in the Savepack
2. Define the encryption key

1.2.c. Encryption selection

To select the encryption, go in the savepack screen or in the tree option screen and set the desired encryption type for that client.



1.2.d. Encryption key

As stated above, Arkeia can encrypt either whole trees or specific directories. In both cases, a specific key will be used. The encryption key has a minimum size of 8 Bytes.

Encryption of a whole tree

A. Create a key file in the directory `/usr/knox/obs`:

This file contains an encryption key for each type of encryption, DES or BLOWFISH, which will be used.

- The name of the file doesn't matter
- Its format is standardized as follow:

DES "[DES Encryption Key]"

FISH "[Blowfish Encryption Key]"

B. Create a file ***cryptkey.tree*** in the directory `/usr/knox/obs`

This file contains the names of the trees which should be encrypted and the name of the key file to be used for this encryption:

`"/home"` `"/usr/knox/obs/encryption_file"`

Example For example, for a machine named Orion, we create the file *orion_key_file* in */usr/knox/obs*:

```
#Key file for machine Orion
#Orion_key_file file
FISH      "12e25rth92178i"
DES       "125dhjeo16954h2"
```

Then, we add a */usr/knox/obs/cryptkey.tree* file with the following content:

```
"/etc"      "/usr/knox/obs/.cryptkey.opb"
(see "encryption of a directory" below)
"/usr/home" "/usr/knox/obs/orion_key_file"
```

Encryption of a directory

A. Create a key file in the directory, with the name *.cryptkey.opb*

This file contains an encryption key for each type of encryption, DES or BLOWFISH, which will be used. Obviously, it is possible to create a key file in each directory, with different keys.

- It uses a standard format:

```
DES      "[DES Encryption Key]"
FISH     "[Blowfish Encryption Key]"
```

The local key file, *.cryptkey.opb*, is never backed up. This is done to ensure that a backup operator cannot decrypt files inappropriately.

B. If the local *.cryptkey.opb* file is empty, Arkeia uses the key defined in */usr/knox/obs/.cryptkey.opb* to encrypt files of the local directory.


Example On a machine named Proxima, instead of a general encryption, we prefer to set a local encryption. we want to encrypt the directories: */home/rvdboom* and */usr/src*

We create in */home/rvdboom* directory a file named *.cryptkey.opb* with the following parameters:

```
FISH      "51654dfe21fd53d"
DES       "fde45823d6484d"
```

while the one in */usr/src* is empty.

During backup, Arkeia will use the *.cryptkey.opb* in */home/rvdboom* for the encryption of this directory, while it will use the *.cryptkey.opb* in */usr/knox/obs* for encryption of */usr/src*.

 **Please note:** the encryption of a given directory does not apply to its subdirectories. You have to encrypt whole tree structures to apply a uniform encryption to a directory and all its subdirectories.

II. Users authentication and ROLES

II.1. 3.1 Introduction

A backup software should apply security rules to avoid the manipulation of sensitive files by unauthorized users. It must be possible to tell which user can backup or restore which file. However, this must be kept simple enough so that an Administrator can backup a whole network, while he/she may not have access as a user to all machines.

This chapter provides an overview of the security features built into Arkeia, which control and prevent unauthorized use of the software and access to different machines in an Arkeia configuration.

II.2. Authentication and authorization

II.2.a. Definition

The Arkeia login process needs to be carefully controlled in order to maintain overall system security. Arkeia has a number of built-in features to maintain this security.

Arkeia is able to monitor and control the origin of any incoming requests before providing any services. The policy for Arkeia to provide or deny access is based on access control lists.

Security in Arkeia is based on various types of files which define authorized users, file access rights, accessible services and, eventually, user impersonation:

- *auth.cfg files:* service execution authorization
- *proxy.cfg files:* file access rights
- *Role management:* secure product access

II.2.b. Authorization files

Authorization files are used on every client machine to restrict user access to the various Arkeia services installed on that client.

In an environment with multiple backup servers, the authorization files can be used to configure client machines in such a way that only specific backup servers can access them. The authorization files can also be used to prevent access to the Arkeia services on the client machines and on machines not running Arkeia.

Each process has an authorization file associated with it. Additionally, there is an overall authorization file, which applies if the individual file is not present. **The default is to deny all access.**

The files are simple text files with one line per entry. Each line specifies the service, the access/deny value, and a list of authorized machines, reserved/non-reserved port requirements and a list of users to which the rule applies. Multiple entries are allowed; Arkeia scans the file and stops at the first matching entry.

These files are located in the */usr/knox/nlp* directory. They are named *auth_XXX.cfg*.


They are all formatted as follows:

```
#####  
# File: auth_MDR.cfg  
# Authorization file for Knox Arkeia Restoration Agent. (C) KNOX 97  
#  
# The values in this file are used to allow or deny access to the  
# browser server.  
# For every connexion the file is scanned sequentially. The first line whose  
# left hand side part matches the connection is used for granting or denying  
# access to the server. If no line matches the connection, access is denied.  
#  
# The syntax of this file is:  
# SERVER_LIST.SERVICE_LIST ALLOW/DENY HOST_LIST[RESERVED PORT]  
USER_LIST  
#  
# SERVER_LIST is a list of one or more servers separated by '|'  
# SERVICE_LIST is a list of one or more services separated by '|'  
# ALLOW_DENY is "ALLOW" or "DENY"  
# HOST_LIST is '*' or a list of one or more hosts separated by '|'  
# [RESVPORT] optional: [1] if connected from a reserved port, else [0].  
# USER_LIST is '*' or a list of one ore more remote user names  
# separated by '|'  
#####  
  
MDR.* ALLOW * *  
  
#####  
# Keep this last line, otherwise CR+LF environnements will fail
```

In this example, all users and all hosts have access to MDR, whatever the port used to access.

You normally shouldn't change anything to those files, as they are configured for optimal security and use. If you restrict some user or machines, you may not be able to backup them.

Multiple entries are allowed. Administrators may create their personal configuration, totally denying access to some machine if necessary.

 **Please note:** use of a reserved port for the connection means that the connection session must be running as root.

II.3. Restricting access of a server/client to a specific client

II.3.a. Basics

Additional security can be implemented by modifying configuration files on the client machine. The `/usr/knox/nlp/auth_XXX.cfg` can be modified in such way that the client machine will only allow connections from the configured Arkeia backup server. This will prevent a random Arkeia backup server from connecting to one of the client machines.

II.3.b. Restrict navigation

Change the following on the client machine, in the *auth_RNV.cfg* file:

```
RNV.* DENY aaaa.bbb.ccc * *  
RNV.* ALLOW * *
```

where “*aaaa.bbb.ccc*” is the fully qualified domain name of the Arkeia backup server machine.

After you have modified the client machine in this manner, you should not be able to use the navigator to see the directory structure on the client machine.

II.3.c. Forbid access by a specific backup server

Modify *auth_OPBS.cfg* on the client machine as follows.

```
OPBS.* DENY aaaa.bbb.ccc * *  
OPBS.* ALLOW * *
```

Any backup from the “*aaaa.bbb.ccc*” server will fail. All other backup servers will have access to the client.

II.3.d. Restrict the access rights to a specific server

Conversely, the following changes will only grant access to the “*aaaa.bbb.ccc*” server to the client machine.

```
OPBS.* ALLOW aaaa.bbb.ccc * *
```

II.4. Proxy or relation files

Proxy files are used on each machine running the Arkeia client, the server or the graphical interface in order to restrict access to certain users.

Again, like authorization, each Arkeia’s process has a proxy file associated with it. The proxy file effectively controls what user ID (and what access rights) a remote user will obtain on the client machine.

Multiple entries are allowed: Arkeia will scan the file and stops at the first matching entry.

Proxy files actually contain the rules used by Arkeia users to impersonate standard users, in order to be allowed to execute backups and restore operations. Thus, any user with “USER ROLE” will be able to restore his own machine, as if he was “root” on his machine.

These files are located under */usr/knox/nlp* and they are named *proxy_XXX.cfg*

They are all formatted as follows:

```
#####  
# File: proxy_MDS.cfg  
# Proxy file for Knox Arkeia Backup Agent. (C) KNOX 97  
#  
# Values in this file are used to determine the id on a local system  
# of a user login on from a remote machine.  
# For every connection the file is scanned sequentially. The first line whose  
# left hand side part matches the connection is used for specifying the local  
# authorizations (uid and gid).  
#  
# The syntax of this file is:  
# [SERVER_LIST,<RESERVED PORT>]HOST_LIST,USER_LIST  
LOCAL_USER,LOCAL_GROUP  
#  
# [SERVER_LIST] is an optional list of servers separated by '|'  
# ,<RESERVED PORT>' is optional. Reserved port equals to 1 if the  
# comes from a reserved port, else by default 0.  
# HOST_LIST is '*' or a list of hosts separated by '|'  
# USER_LIST is '*' or an optional list of remote user names separated by '|'  
# LOCAL_USER is '*' or the name or uid of a local user  
# LOCAL_GROUP (optional) is '*' or the name or gid of a local group  
#  
# If the user and/or group are not specified on the left hand side, any  
# user logging in on the machine will take the uid of the user specified  
# on the right hand side.  
# If the user specified on the right hand side is '*', a user logging from  
# the machine specified on the left hand side retains his user name on the  
# local machine. All users corresponding to the specification on the  
# left hand side must then have equivalent user names on the local machine.  
#  
# Example:  
#  
# *,* nobody  
# Means: Any user (root or standard user) from any machine asking for  
# any service on the local machine will execute it as local user nobody.  
#  
# [ORAS,1]*,* root  
# Means: Any user connecting from a reserved port will execute the server ORAS  
# as root on the local machine. (NOTA: In order to call from a reserved port  
# either the remote user is super-user (uid=0) or the calling program belongs  
# to root and has the setuid bit set).  
#  
# [RNV|ODTS]orion|vega|cassio,* guest  
# Means: Any user (Privileged or not) connecting from one of the three hosts  
# (orion or vega or cassio) and asking for server RNV or ODTS will execute  
# it as the local user guest.  
#####  
  
[MDS,1]*,* root  
[MDS,0]*,* root  
  
#####  
# Keep this last line, otherwise CR+LF environments will fail
```

In the previous example, all users connecting through the Reserved Port (“,1]”) will give MDS root privileges. All users connecting through the other ports (“,0]”) will also give MDS root privileges

❖ **Please note:** after the installation process, access to Arkeia and all the backed up files is left “open”: the root user is authorized to access all clients, regardless of the machine used as a server. Once Arkeia is installed, the security should be tightened down as required.

❖ **Please note:** Arkeia maintains its own list of users, passwords and roles, independent of the operating system used.

II.5. ROLES management

II.5.a. Basics

The backup administrator can create different users within Arkeia. This can be used to delegate responsibility and prevent inappropriate use of the system.

There are three predetermined roles, “ADMINISTRATOR”, “OPERATOR” and “USER”.

By default a single user (root), exists after the installation has been completed. There is no initial password. This user has full read/write rights across the Arkeia backup network.

II.5.b. Arkeia backup and restore configuration management, using the ROLES

Administrator

An Arkeia “Administrator” is able to perform all the operations offered by Arkeia: backup/restore, creation and modification of Savepacks/Drivepacks/Tape pools/Libraries, other users, etc.

Operator

An Operator can backup and restore all machines as an Administrator but he cannot modify Arkeia’s configuration. An Operator has to be able to backup and restore as needed but he cannot modify the configuration that the Administrator has established.

User

A user can only restore all the files on the machine he/she is logged on (even root files), with the GUI of Arkeia.

APPENDIX A

Troubleshooting***I. Arkeia configuration and usage***

<i>Issue</i>	<i>Possible causes and solutions</i>
Arkeia server freeze under Linux kernels 2.2.14 / 2.2.15	<ul style="list-style-type: none"> • Update to the most recent kernel.
Can't connect to a client	<ul style="list-style-type: none"> • Check the TCP/IP connectivity • Check if NLSERVD is running on both client and server • Check that the 617 port is accessible to Arkeia on both client and server on both directions
Can't find library control device under AIX	<ul style="list-style-type: none"> • Download Knox the "/dev/pthru0" driver from our website (http://www.arkeia.com) or install it from the Arkeia's CD–Rom.
Can't log into GUI	<ul style="list-style-type: none"> • Check the DNS • "Ping" the Arkeia server from another machine. Make sure ping resolves the correct domain name. • Ping the Arkeia server from the Domain name server. Make sure ping resolves the correct IP address. • Make sure the server is correctly defined in the /usr/knox/nlp/rhost.lst • Check the /usr/knox/nlp/services file. It must include all Arkeia's services.
Can't start Knox Network Service on Windows NT	<ul style="list-style-type: none"> • Make sure you are running at least Windows NT4 SP3. • Make sure the installed client is the correct version for Windows NT, Server or Workstation • Make sure you are logged as Administrator to start the KNS service
After a domain change, the backups won't start	<ul style="list-style-type: none"> • Edit the domain name in /usr/knox/arkeia/dbase/f3lib/libXXXXXXXXX/*.lst and /usr/knox/arkeia/dbase/f3drv/drvXXXXXXXXX/*.lst. Restart NLSERVD.

Issue	Possible causes and solutions
Navigator connection is very slow	<ul style="list-style-type: none"> • Check, if you have a Multi–NIC configuration, your NLP_HOSTNAME and your NLP_HOSTFILE on both server and clients • Check, if you have a Multi–NIC configuration, your hostfile entries (often /usr/knox/nlp/hosts.cfg) on both server and clients
Drive is already reserved	<ul style="list-style-type: none"> • Delete the PID line in the /usr/knox/arkeia/dbase/f3drv/drvXXXXXXXX of the corresponding drive • Close the Arkeia GUI on all machines
E–mail notification doesn’t work	<ul style="list-style-type: none"> • Check if sendmail is correctly configured and running • Check your e–mail address in Utilities – Users Management
JUI doesn’t start	<ul style="list-style-type: none"> • Reinstall the Microsoft Virtual Machine and reboot.
Navigator doesn’t display the backup server	<ul style="list-style-type: none"> • Check that the /usr/knox/nlp/admin.cfg correctly points to server. Restart NLSERVD. • Be sure you can make an accurate ‘nslookup’ on the server.
Navigator doesn’t display the clients	<ul style="list-style-type: none"> • Check the connection between the server and the clients • Check that the file /usr/knox/nlp/admin.cfg correctly points to the server • Restart NLSERVD on both client and server • Check that the client is defined in /usr/knox/nlp/rhost.lst • Use /usr/knox/bin/chknlp on the client to verify connection
Navigator shows duplicate entries	<ul style="list-style-type: none"> • Delete any extra entry in /usr/knox/nlp/rhost.lst. Restart NLSERVD.
Not enough resources to run Arkeia on Unix OS	<ul style="list-style-type: none"> • Check the Advanced Manual, section 2: ‘Configuring the backup server ’
Periodic backups don’t start	<ul style="list-style-type: none"> • CRON is not started or not running • ARKPER line missing in Crontab • Local time is not set
GUI or JUI displays an incorrect time	<ul style="list-style-type: none"> • Create a symbolic link to /etc/localtime in /usr/lib/zoneinfo. Create the directory if needed. Restart NLSERVD and GUI.
Periodic backup kill the network functionalities of the machine	<ul style="list-style-type: none"> • Change the Network Adapter
STKS does not detect correctly my library under Linux	<ul style="list-style-type: none"> • Check ‘dmesg’ for correct hardware detection • Check your ‘/proc/scsi/scsi’ and verify you have a sgX device • Check that your kernel supports ‘Generic SCSI’ and ‘Probe All LUNs’ • Check that your SCSI BIOS has the “Multiple LUNs support” enabled and “Disconnect/Reconnect” disabled
Tapes are not processed in sequence	<ul style="list-style-type: none"> • Check if all your tapes are full • Check the retention dates of your tapes and make sure one is always free (recycled) in time for a new backup • Recreate your tapes and pre–label them. Use FIFO policy.
Tape doesn’t function under Irix 6.3	<ul style="list-style-type: none"> • You have to disable the ‘mediad’ daemon, using the “mediad off” command.
Tape is labelled ‘Full’ but there is remaining space	<ul style="list-style-type: none"> • Arkeia labels tape ‘Full’ when SCSI errors happen while writing. If this happens too often, check your SCSI configuration and hardware.

II. Error codes and messages

II.1. Introduction

Arkeia displays error codes and standard messages when an error occur. Those codes and messages can help you debug your system.

II.2. Error Codes

Arkeia uses error codes to indicate where the issue is. Those codes are explained below:

Error 0:	OK (end of tree)
Error 1:	OK (end of tree + database updated)
Error 2:	Connection error
Error 3:	Network error during backup
Error 4:	Remote command failed
Error 5:	End of flow (abort, kill, ...)
Error 6:	Backup with errors

In case “Error 3” is displayed, check all network–related logs. In case “Error 5” is displayed, check the logs of the O3flow process.

II.3. Error messages

Here are common error messages in Arkeia and their possible causes and solutions.

<i>Error message</i>	<i>Possible causes and solutions</i>
“Backup criteria not met”	<ul style="list-style-type: none"> • Check the Savepack for non–valid file system, particularly NFS. Change “File System Allowed” to “All file systems”.
“Can’t find available tape for drive”	<ul style="list-style-type: none"> • If you are using a library, there are probably no tapes defined for the slots • All the tapes in the selected pool are full • The backup is defined with a tape pool of the wrong tape type • The tapes were created with the wrong tape type.
“Can’t find drive”	<ul style="list-style-type: none"> • Check if the drive is associated with a DrivePack
“Can’t find LIBID”	<ul style="list-style-type: none"> • Please refer to our web site for the latest information and troubleshooting options. • As a last resort: clean all libXXXXXXXX files and all files with LIBID=XXXXXXXX in /usr/knox/arkeia/dbase/ then rebuild a new library
“Can’t find tape ID 0xXXXXXXXX”	<ul style="list-style-type: none"> • Search in /usr/knox/arkeia/dbase/f3tape/tpmaster.lst the name of the corrupted tape, note the TPID value and erase the TPID line. Delete all files in sub–directories with the same TPID. Restore the database with the ‘arkrstdb’ command.

<i>Error message</i>	<i>Possible causes and solutions</i>
“Can’t find public symbol NWDSLogin” message under Novell	<ul style="list-style-type: none"> • You have to load the netnlm32.nlm module.
“Can’t load backup master list “	<ul style="list-style-type: none"> • Make sure your Arkeia directory has enough free inodes. • Check for corrupted periodic backup information (in /usr/knox/arkeia/dbase/f3per)
“Licence server badly declared”	<ul style="list-style-type: none"> • Update /usr/knox/nlp/admin.cfg with the correct hostname of your backup server. Restart NLSERVD.
“Please insert Tape ‘XXXXX’ in drive ‘YYY” message while the tape is in the drive	<ul style="list-style-type: none"> • Check the control device of the library
“Slot X is not full”	<ul style="list-style-type: none"> • You have not configured any tape for the slots in “Library Management – Slot Usage” • Your library control device is not correct. Check kernel logs for the correct control device.
“TO15 protocol error”	<ul style="list-style-type: none"> • Make sure your Arkeia directory has enough free inodes. • Please refer to our web site for the latest information and troubleshooting options.
“X! You are limited to Y standard machine(s)”	<ul style="list-style-type: none"> • Search the offending computer (set MDSLOGLEVEL to 60 in /usr/knox/arkeia/arkeia.prf; run a backup; grep “get_hosts_nonpresents” /usr/knox/log/mds.lg*). Beginning from Arkeia 4.2.7, all machines are listed with their types: finding the offending one is therefore easy. • Check the number of Level 1 licences (Commercial UNIX, Windows NT Server, Novell, etc.) • Check for missing entries in the file /usr/knox/nlp/rhost.lst of the computers you’d like to backup. • Watch for misspelled client names in “command before” or “command after”. • Check for client name changes and disappearances • Check for client IP changes and update rhost.lst

APPENDIX B

Log Management

I. The Arkeia Journal

I.1. Introduction

Arkeia maintains an ASCII file in which a record is kept of all operations carried out:

- \$KNOX/arkeia/arkeia.jnl (Arkeia version 4.0)
- \$KNOX/arkeia/arkeia.jl2 (Arkeia version 4.2)

This file can be filtered by using different menus of the graphical interface so that only certain sections are shown (backup, restoration, and tape labeling...)

By default, the user has access to the log file for the current month. At the end of each month, the log file is copied and kept for one year (unless overwritten).

I.2. The “Log Consult” screen

From the main screen click on the [Utilities] menu then on the [Log Management] option, and finally, on the [Log Consult] menu item.



Arkeia then displays the requested log.



You'll find below a short description of the graphical conventions used in the log:


A GREEN square means an INFORMATION (tree backup...)



A YELLOW square is a WARNING (start, end of backup...)



A RED square indicates an ERROR (no tape, no network connection...)



I.3. Display filter

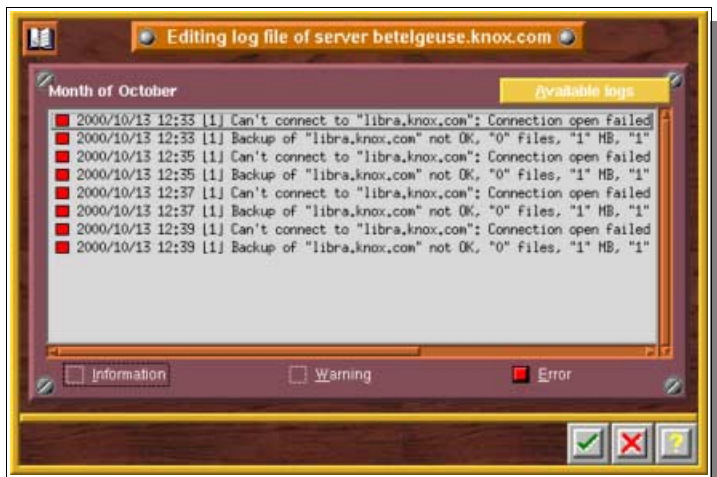
It is possible to display only "Error", "Warning" or "Information" messages by selecting or not the relevant buttons and then clicking on the "Checkmark" (OK) button.



If all buttons are selected, the screen displaying the log looks like this:



With only the "Error" button selected, the log display is similar to this screen shot:



I.4. Change the displayed month

It is possible to display other months than the current one. To do this, click on the "Available logs" button.



The “Available logs” window is then displayed on the screen.



Select the desired month and confirm your choice by clicking on the “Checkmark” (OK) button.



II. The “Backup done” screen

II.1. Introduction

This is a display configuration of the Journal to include only the information on the successful backups.

To go to the logs from the main screen, click on the menu [Utilities], then on the [Log management] option and, finally, on the [Backup done] menu item.



You can also use the “Magnifying glass” button in the main Arkeia window



III. Other logs

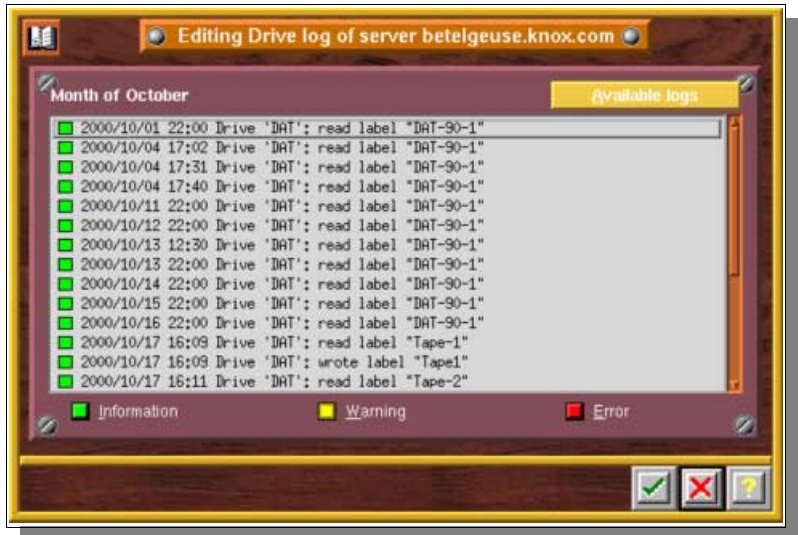
III.1. The “Drive” log

This log contains the tasks completed by drives over a period of one month.

To go to this log from the main screen, click on the [Utilities] menu, then on the [Log management] option and, finally, on the [Drive Log] menu item.



The “Drive Log” screen is then displayed on the screen.



You can change the month to be displayed by clicking on the “Available Logs” button.



The “Available Logs” screen is then displayed on the screen.

Double-click on the desired month to display its log.




III.2. The “Tape” log

This screen lists the tasks completed by the tapes over a period of one month.

To go to the logs from the main screen, click on the [Utilities] menu, then on the [Log management] option, and, finally, on the [Tape Log] item.



The “Tape Log” window is then displayed on the screen.



You can change the month to be displayed by clicking on the “Available Logs” button.



The “Available Log” screen opens. Double-click on a month to display its log.



III.3. The “Restore” log

This screen lists all the restorations completed over a period of one month.

To go to the logs from the main screen, click on the [Utilities] menu, then on the [Log management] option, and, finally, on the [Restore Log] menu item.



The “Restore Log” window is then displayed on the screen.



You can change the month to be displayed by clicking on the “Available Logs” button.



The “Available Logs” screen is then displayed on the screen.



Double-click on a given month to display its log.

Backup of Open Source Databases

I. Important notice to Oracle users

🔴 **Please note:** this appendix is only focused on Open Source database solutions. If you'd like to use Arkeia to backup your Oracle databases, Knox Software can supply a backup Assistant and an interface with Oracle's RMAN. For more information on these products, please refer to chapter 2 of this manual, and contact: sales@arkeia.com.

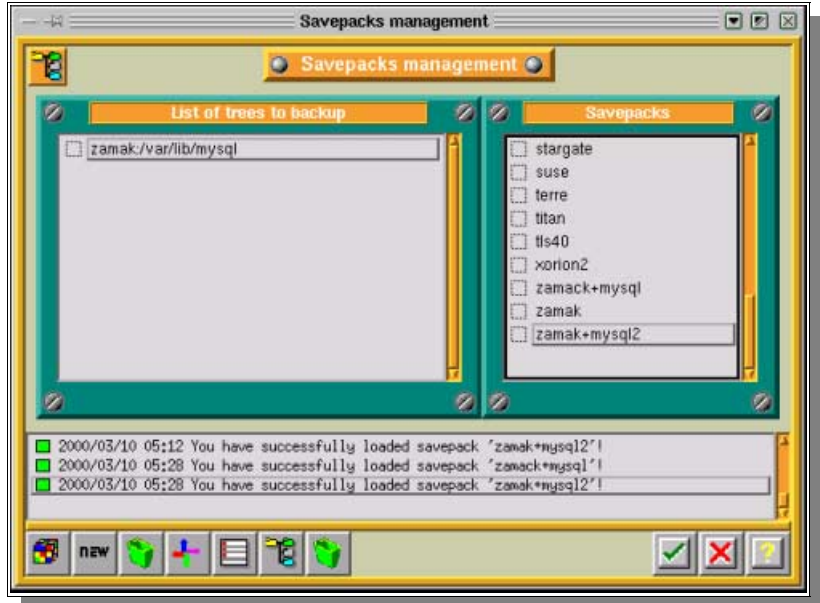
II. MySQL

II.1. How to backup a MySQL database offline

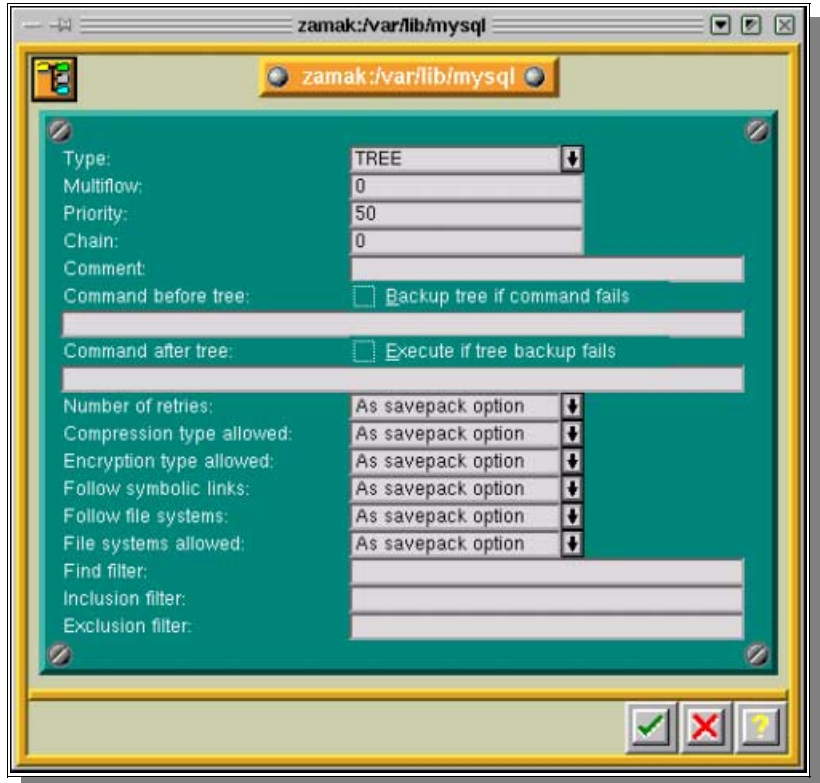
The first way to easily backup a database is to perform this operation while MySQL is off-line. This simple solution is most of the time inadequate, as a database may have to stay online 24 hours a day.

Here is the procedure to do a backup while the MySQL database is off-line:

Create a Savepack that will save only the directory where the MySQL database is located.



Open the tree options (by double-clicking on it).

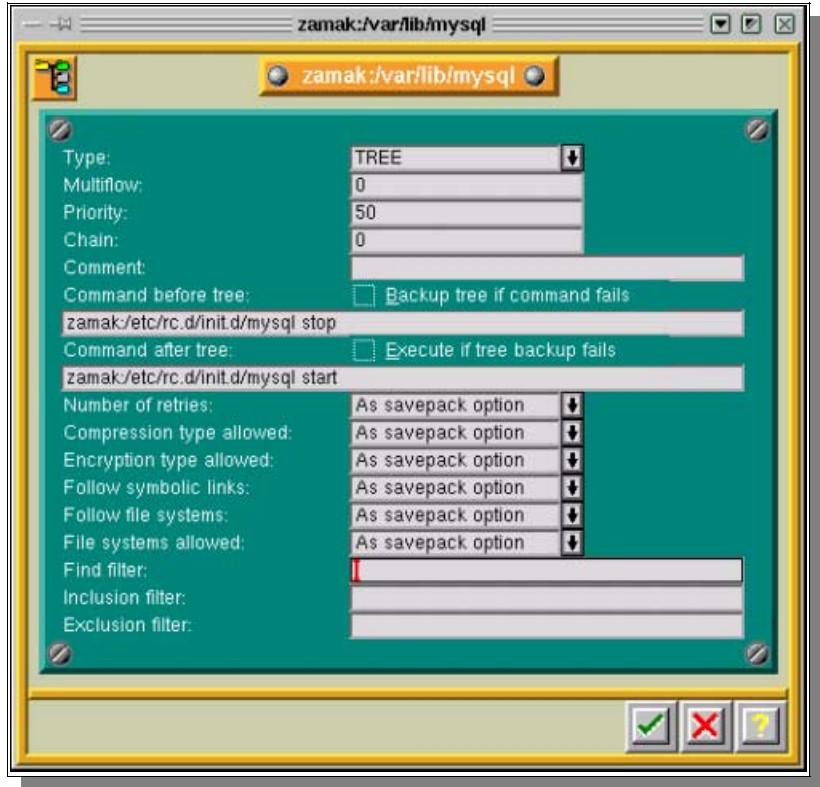


Here, “zamak” is the machine where MySQL is running, complete the fields "command before" with:

`“zamack:/etc/rc.d/init.d/mysql stop”`

and "command after" with:

`“zamack:/etc/rc.d/init.d/mysql start”`



When the backup of this Savepack is launched, your database will be stopped before the backup of the Savepack and will be restarted just after the Savepack has been backed-up.

➤ **Please note:** the “command before” and “command after” may vary, depending on your OS. Refer to your MySQL documentation to check.

II.2. How to backup a MySQL database online

II.2.a. First method

Create a script like the one below and place it in the “command before” field, in the options of the “tree” included in the Savepack you backup. Have a look at the offline mode

```
SHELL>mysqldump -l -q -T/xxxxxxx/database_1/ database_1
```

-l will lock all the tables in the database

-q will backup without using buffers

/xxxxxxx/database_1/ must be an existing directory

This way, your are sure to have a consistent export of your data before backing it up.

II.2.b. Second method

First of all, your database must be in the “mode `--log-update`”.

Please note: to do this, stop mysqld if it is running, then restart it with the “`--log-update`” option.

You will get log files named “hostname.n”, where “n” is a number that is incremented every time you execute “`mysqladmin refresh`” or “`mysqladmin flush-logs`”, the FLUSH LOGS statement, or restart the server.

These log files provide the information you need to replicate the changes made to the database after the execution of “`mysqldump`”.

A dump of the database will be done to a file rather than in a directory.

The logs will also be flushed at the time of the backup, in order to have logs only since the last backup.

To backup manually, enter the following command:

```
SHELL>mysqldump -l -q --add-drop-table --flush-logs database_1 >/xxxxxxxx/database_1.date_and_time.backup
```

-l will lock all the tables in the database

-q will backup without using buffers

--add-drop-table will recreate the tables without obtaining “table already exists” errors

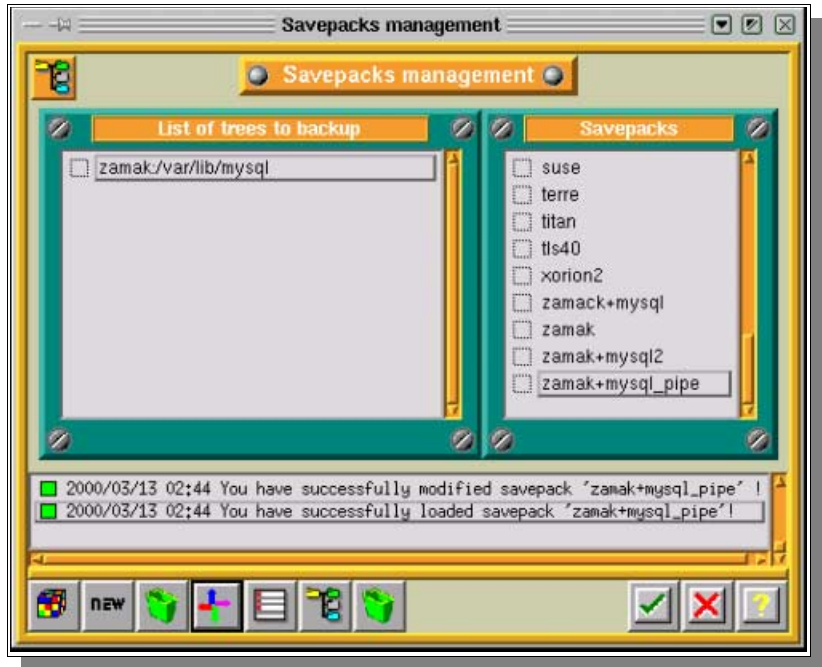
--flush-logs or -F will flush logs file in server before starting dump

To restore manually:

```
SHELL>mysql database_1 </xxxxxxxx/database_1.date_and_time.backup
```

II.2.c. Now with ARKEIA!

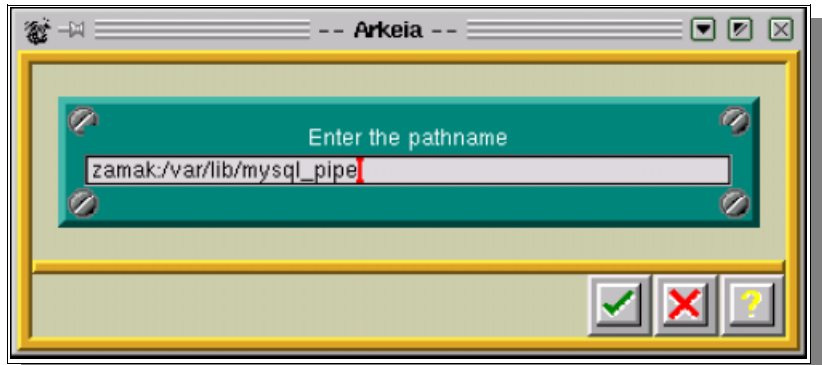
Create a Savepack that will back up only the directory where MySQL database is located.



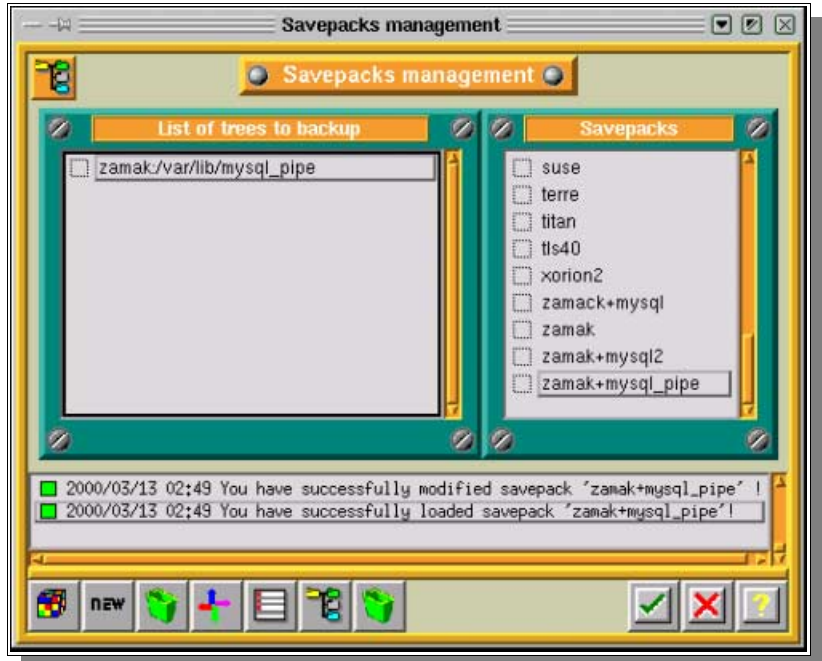
Modify the tree by right-clicking on it.



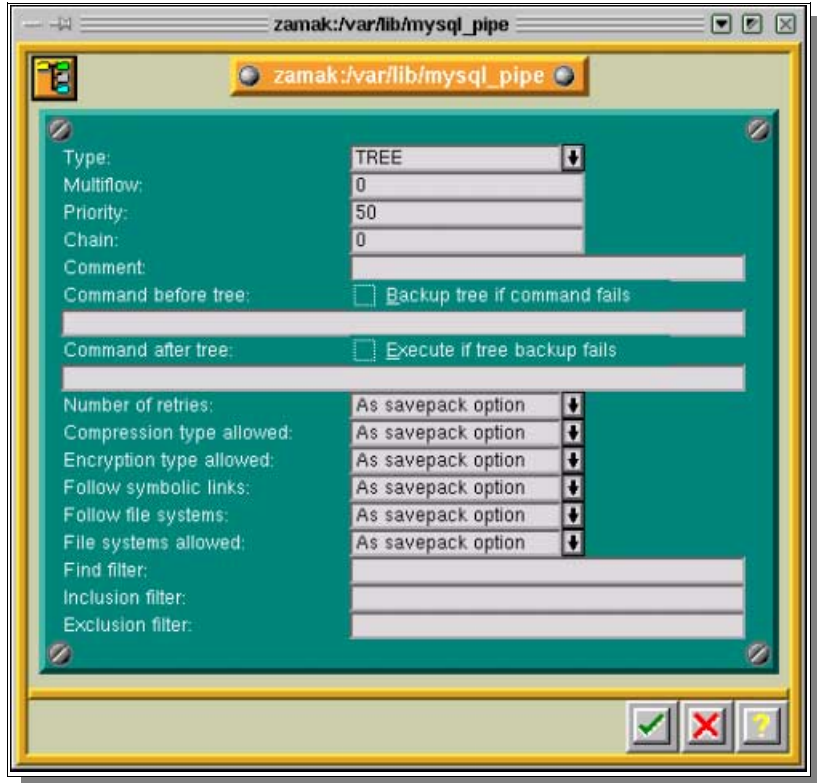
Modify the tree.




The tree is now modified



Open the options of the tree (by double-clicking on it)



Change the type to OBJECT, and in the field “Object backup command” add your mysqldump command, and in the object restore command add a “cat > mysql databasename”



Please note: If you do a restoration, do not forget to apply the last logs since the backup.

I. Extracting data from tape: *readarkeia*

I.1. Introduction

readarkeia 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.

I.2. Getting *readarkeia*

A binary version of *readarkeia* is installed in the `/usr/knox/bin` 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.

To extract the C code, first rewind the tape using the following `mt` command on the rewindable device defined in the *Initial Configuration* chapter. 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`

Then edit the *readarkeia.c* file and delete the first line (it is actually the tape label). Then replace the very first space character by a `/"` character.

Then compile *readarkeia.c* using the following command: `cc readarkeia.c -o readarkeia`

I.3. Using readarkeia

readarkeia is more or less 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 ...]

Option	Use
l	Please note: this character is the lower case letter “L”. Change the verbose level. Use it to obtain detailed information on what is really on the tape. Value can be from 10 to 80.
x	Extract the named files from the tape. If no files are specified then all the files on the tape are extracted.
t	List the table of contents of the tape.
i	Get information about the tape.
m	m tells <i>readarkeia</i> 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 ‘-’, <i>readarkeia</i> reads tape data from standard input
v	Verbose.

I.3.a. List tape’s content

To list the files present on a tape, type: *./readarkeia tvf /dev/st0*

I.3.b. Extract all the files on a tape

To extract all the files present on a tape, type: *./readarkeia xvf /dev/st0*

🔴 **Please note:** *readarkeia* extracts data locally and creates a subdirectory for each host backed up.

🔴 **Please note:** *readarkeia* scans the tape completely. This operation can take some time.

I.3.c. 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 the tape completely, even if the requested file is at the beginning. This operation can take some time.

Example	<i>If you want to extract a /usr/src/kernel/make file from a machine named ariane, you have to type (under Linux): ./readarkeia xfv /dev/st0 ariane/usr/src/kernel/make</i>
----------------	--

II. Restoring the database: arkrstdb

II.1. Introduction

In rare occasions, your backup server will crash and the Arkeia database may be corrupted. In order to keep the usability of your backup history, you may have to rebuild the index, in whole or in part, from what's been saved on tape.

This is the purpose of the *arkrstdb* utility.

You need to be aware that this restoration is a long process, since this utility needs to read the entire tape pools to completely rebuild Arkeia database.

II.2. Usage

Type *arkrstdb -usage* for more details

The general command line of *arkrstdb* is:

```
arkrstdb -d{rewindable device} -t{tape type} -r{drive type} [-D {database}] [-p(pool)] [-v(voltag)] [-s(silent)] [-l <Loglevel>] [-j(journal)]
```

Option	Use
-d	Specify the rewindable device of the drive as you entered it in drive management.
-t	Specify the tape type
-r	Specify the drive type
-D	The pool name (<i>optional</i>)
-p	The path to your database (<i>optional</i>)
-v	The voltag of your tape (<i>optional</i>)
-s	To run <i>arkrstdb</i> in silent mode (<i>optional</i>)
-l	Log level to get more information. Possible value are 10, 20 ,30 ,40 ,50 ,60 ,70 ,80 ,90. (<i>optional</i>)
-j	(journal)

Example	<p><i>You plan to restore a tape's content to the database from a DLT drive. Use the following command:</i></p> <pre><i>/usr/knox/bin/arkrstdb -d/dev/st0 -t"DLT 4000" -rSTD_DLT</i></pre>
----------------	--

III. Testing the connectivity: chknlp

III.1. Introduction

Networking problems are the most frequent 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 very simple to solve.

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. It will help you find the weak spots of your network.

III.2. Usage

The *chknlp* program is a binary provided with client and server to check network connectivity between client machine and backup server. It is currently provided only on UNIX systems and you must be root to use this command.

To start *chknlp*, just type the following command: */usr/knox/bin/chknlp -pause*

“*chknlp*” then runs four or five tests, the last test being run only on the backup server.

III.3. Nature of the tests

III.3.a. Test 0

chknlp tries to locate the Knox directory.

In case of error

If this test fails, an installation problem occurred: retry installing the software.

Correct output

```
>>> chknlp 4.2.3-1 <<<

>>> TEST FOR CERTIFICATION LEVEL ONE <<<
*** TEST 0 IN PROGRESS...
>>> Getting LANG: >>> Getting KNOX variables:
    KNOX = /usr/knox
    NLPDIR = /usr/knox/nlp
    DIRAPPLIC = /usr/knox/nlpc
    PATHAPPLIC = /root/.nlpc:/usr/knox/nlpc:/usr/knox/nlp
*** TEST 0 OK
```

III.3.b. Test 1

The “*chknlp*” software reads network preferences like

- The port number used for communications (PORT_NUMBER)
- 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)

In case of error

If *chknlp* cannot reach the ADMINSERVER, check that the file */usr/knox/nlp/admin.cfg* contains the name of the backup server.

If *chknlp* 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* does not reach your local hostname or local address verify your network configuration.

Correct output

```
*** TEST 1 IN PROGRESS...
>>> Testing variables in configuration file
    NLP_AUTH = /usr/knox/nlp/auth.cfg
    NLP_PROXY = /usr/knox/nlp/proxy.cfg
    PORTNUMBER = 617
    ADMINSERVER = jupiter.knox.com
    NLP_TIMEOUT = 60
    STRIP_DOMAIN = -1
    Internet address for Licence server 'jupiter.knox.com' = 192.168.8.22
!!! Warning – Domainname knox.com in host name jupiter.knox.com discarded   HOSTNAME =
jupiter (default value)
    Internet address for local HOST 'jupiter' = 192.168.8.22
*** TEST 1 OK
```

III.3.c. Test 2

This test checks that the local daemon or service is running correctly

In case of error

If this test fails, the local daemon or service is not running correctly, or does not accept connection to port 617. Restart the daemon with the NLSERVD command or restart the Knox Network Service

Correct output

```
>>>> Testing local (jupiter) NLSERVD daemon
    Message from NLSERVD daemon:   Test OK
*** TEST 2 OK
```

III.3.d. Test 3

This test checks the network addresses of the client and the backup server as seen by each other and compare both.

In case of error

It is quite common to configure the ADMINSERVER as the backup server – although this can be configured differently.

The host name of the ADMINSERVER on the local machine (*admin.cfg* file) must be the name of your Arkeia ADMINSERVER. The ADMINSERVER is typically the backup server.

If the ADMINSERVER's 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 verify the IP addresses of the local host on both machine.

If the local host address is different on ADMINSERVER, verify 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.

Correct output

```
*** 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
```

III.3.e. Test 4

This test only runs on the server side. It verifies the network addresses of all the clients listed in the */usr/knox/nlp/rhosts.lst* file, and of the backup server as seen by each other.

In case of error

The parameters you have to check are basically the same as those in Test 3, though you should have a complete set of information about your network in this way.

As *chknlp* stops at each error, you need to rerun it after resolving an issue, to make sure there are no other problems with another machine.

Correct output

```
>>> 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
```

IV. Duplicate a tape: *tpdup*

IV.1. Introduction

Some System Administrators would like to be able to duplicate tapes, the tape copy being stored outside of the main facility, in a safe, for instance. As this cannot be done with standard tools, due to Arkeia specific format, a dedicated tool is needed. This tool is *tpdup*.

🔴 **Please note:** this tool was originally supplied with Arkeia v4.0. It is not supported anymore in version 4.2

IV.2. Requirements

tpdup requires its own license and the availability of a library. It won't work if you only have access to a tape drive.

IV.3. Installation and use of *tpdup*

IV.3.a. Installation

The *tpdup* software is installed like most Arkeia packages: when the package is uncompressed, a *tpdup/* directory is created. Just enter this directory and run *.install*.

The installation script then creates a *tpdup/* directory in */usr/knox/arkeia*. This directory contains two files, *tpused.lst* and *tpused.lck*, that will be used in the configuration. The “*tpdup*” command is then created in */usr/knox/bin*.

IV.3.b. Configuration and usage

When Arkeia uses a tape and *tpdup* is installed, the name of the used tape is written in the */usr/knox/arkeia/tpdup/tpused.lst* file. Only tapes referenced in the *tpused.lst* file can be duplicated.

Then you need to create, in the */usr/knox/arkeia/tpdup/* directory, a duplication directory named *config_tpdup*. In this directory, two files must be created:

drv.lst

tpclone.lst

drv.lst

This file contains a list of the drive pair (source/destination) that will be used for duplication.

Its general layout is the following:

```
ITEM {
SRC_DRVNAME      "UNIT_1"
DST_DRV_NAME     "UNIT_2"
}
```

In most cases, the `drv.lst` file will only contain one ITEM.

tpclone.lst

This file contains a list of the tape pairs (source/destination slot number) that will be used for duplication.

```
ITEM {
SRC_TPNAME       "TAPE_DLT1"
DST_SLOT        "5"
}
ITEM {
SRC_TPNAME       "TAPE_DLT2"
DST_SLOT        "6"
}
ITEM {
SRC_TPNAME       "TAPE_DLT3"
DST_SLOT        "7"
}
```

- 🔴 **Please note:** each source must be a tape listed in the `tpused.lst` file and each slot must be filled with an empty tape.
- 🔴 **Please note:** the slot must be labeled as “*Reserved*” in Arkeia.

IV.3.c. Force end of tapes

It is mandatory to force the end of tape for all the tapes used in duplication. This is done in the following way:

Edit the `/arkeia/dbase/f3tape/tpypes.lst` file. Add the following entry to the used tape type ITEM:

“USE_CAPACITY” “YES”

IV.3.d. Duplication start

`tpdup` is started by entering the following command: `/usr/knox/bin/tpdup -v -lconfig_tpdup`

IV.3.e. Reuse of duplicata

To use the duplicated tapes instead of the original ones (in case they are corrupted, for example), just “remove” the original ones in the graphical interface, get them out of the library, put the duplicated ones in place and “add” them, through the graphical interface.

APPENDIX E

Glossary

ARKC

Command line interface to Arkeia, supports most operations normally done by the GUI interface.

ARKPER (IN /USR/BIN)

ARKPER checks the database to see if a periodic backup should be executed. This utility is executed every five minutes by *cron*.

ARKRSTDB

Utility program used to recreate the index database corresponding to a given tape.

ATAPI

An IDE extension that provides support for tape drives and CD-ROM drives.

AUTOLOADER

Low-end combination tape drive / library unit, typically holds less than 10 tapes.

CONTROL DEVICE

Device file that interfaces the operating system to the robot arm of the library unit.

DAEMON

Server processes in UNIX that run in the background until needed by the system. These processes work like the Services under Windows.

DRIVE

Actual tape drive, or tape drive component of a library unit.

DRIVEPACK

The tape drive or the set of tape drives to be used by a specific Arkeia backup.

FILE LIBRARY

Configuration used to backup files to a disk drive instead of tape/cartridge.

FLOW

An individual stream of data that is sent from the client machine to the backup server.

INODE

A special data structure on a disk that maintains the attributes and location of files on the disk (on UNIX systems).

INTERACTIVE BACKUP

A non-scheduled backup that is launched manually.

LIBRARY

A multi-tape storage device that loads and unloads its own tape drive. The capacity of a library varies by model and some units can contain hundreds of tape slots and one or more tape drives.

LUN

Logical Unit Number, allows multiple devices to share one SCSI ID.

MULTIFLOW

Arkeia's ability to correctly manage several data flows from the same machine at the same time.

NLSERVD (IN /USR/BIN)

Shell script to start, stop, and restart the main Arkeia network daemon (nlservd).

PERIODIC BACKUP

A scheduled unattended backup that can be configured three levels deep (For example: Monthly, Weekly and Daily levels).

PERIODICITY

How often a particular backup is scheduled (Daily, weekly, etc.).

RANDOM MODE

Setting on a tape library unit allowing the software to select any tape in the library (also called "SCSI mode").

READARKEIA

A utility program used to extract data directly from an Arkeia data tape.

RECYCLE

Removing all database references to a tape prior to reuse/recycle (this last operation overwrites all data on a tape).

REWIND DEVICE

Device file that interfaces the operating system to a tape drive.

ROBOT

The “arm” mechanism that moves tapes between the drive and the library slots. Also called “*Medium Changer*”.

RPM

“Red Hat Package Manager”, a Linux utility used to install/uninstall software.

SAVEPACK

The designated files, trees or machine(s) to be archived by a specific Arkeia backup.

SCSI

High speed interface standard for disks, tape drives and other peripherals (*Small Computer Systems Interface*).

SLOT

Location of an individual tape within a library.

STKS

Utility used to inventory a library

TAPE POOL

The designated tapes to be used by a specific Arkeia backup.

TAR.GZ

A Tape ARchive file (tar) that has been compressed (using gzip).

TYPE 1 MACHINE

Type 1 refers to the operating system configuration Type 1 systems include UNIX machines running: AIX, Drsrx7, IRIS, IRIX, HP–UX, SCO, SunOS, Solaris, Maxion, OSF, Unixware, Novell, Windows NT Server and Windows NT Workstation on an Alpha platform.

TYPE 2 MACHINE

Type 2 refers to the operating system configuration running BSD, Linux, Windows 95/98/ME, Win NT 4.0 Workstation and Windows 2000.

VALIDITY

The period of time tape indexes are maintained in the Arkeia database (Also referred to as retention period).

VOLTAG

Bar code (tape label) information assigned to tapes, either manually or automatically. The library unit must have a barcode reader to be able to use voltags (*Voltag* = “*Volume Tag*”).

TABLE OF CONTENTS

FOREWORD.....	2
0.1. KNOX SOFTWARE LICENSE AGREEMENT.....	2
0.2. WARRANTY *	4
INTRODUCTION.....	6
I. About this manual.....	6
I.1. Who should read this manual?.....	6
I.2. How to use this manual?.....	6
I.2.a. Introduction.....	6
I.2.b. What is Arkeia?.....	7
II. General Concepts and Features.....	7
II.1. Concepts.....	7
II.1.a. Architecture.....	7
II.1.b. Structure.....	7
II.2. Features.....	8
III. Arkeia Overview.....	9
III.1. Introduction.....	9
III.2. The backup server module.....	9
III.3. The client module.....	9
III.4. The X11 graphical user interface module.....	10
III.5. The JAVA graphical user interface module.....	11
III.6. The Arkeia command line Interface module (arkc).....	11
III.7. Installation guide.....	11
III.8. Package content.....	12
BEFORE YOU BEGIN.....	14
I. Platform availability.....	14
I.1. Available clients and servers.....	14
I.2. Oracle clients.....	15
II. Hardware requirements and prerequisites.....	15
II.1. Hardware requirements.....	15
II.2. Prerequisites.....	16
II.2.a. Memory.....	16
II.2.b. Network cards.....	16
II.2.c. SCSI host adapters.....	16
II.2.d. Tape drives.....	16
III. Software requirements and Prerequisites.....	17
III.1. Requirements.....	17
III.1.a. Reliability.....	17
III.1.b. Available disk space.....	17
III.1.c. Backup catalog/Index database.....	17
III.1.d. Workload.....	17
III.1.e. IP bandwidth.....	17
III.1.f. ROOT account for installation.....	18
III.2. Prerequisites.....	18
III.2.a. Drivers.....	18

Tape devices.....	18
Libraries and Autoloaders.....	19
III.2.b. Network.....	19
III.2.c. SCSI support in kernel.....	20
IV. Platform Specifics.....	22
IV.1. General information.....	22
IV.2. Configuring IPC (shared memory and message queue).....	22
IV.2.a. COMPAQ TRUE64 UNIX / Digital Unix DEC OSF.....	22
Changing the settings.....	23
Viewing IPC and process settings.....	23
IV.2.b. Hewlett–Packard HP/UX.....	23
Modifying IPC and process settings.....	23
Accessing IPC and process settings via the menu:.....	24
Changing the settings.....	25
Viewing IPC and process settings.....	25
IV.2.c. IBM AIX.....	25
IV.2.d. LINUX.....	25
IV.2.e. SGI IRIX.....	25
Modifying IPC and process settings in kernel 6.4 and lower.....	25
Changing the settings.....	26
Viewing IPC and process settings.....	26
Irix Kernel 6.5.X.....	27
IV.2.f. Sun SOLARIS.....	27
Modifying IPC and process settings.....	27
Changing the settings.....	27
Viewing IPC and process settings.....	28
ARKEIA'S CONVENTIONS.....	29
I. Convention used for commands and keys.....	29
I.1. Graphical User Interface (GUI).....	29
I.2. Function keys.....	30
I.3. Keyboard shortcuts.....	30
I.4. Arrow key.....	30
I.5. Context–sensitive menus.....	30
I.6. Tool bar buttons.....	30
I.7. Context–sensitive help.....	31
I.8. Copy/paste with the mouse.....	31
ARKEIA INITIAL CONFIGURATION.....	33
I. Drives and Devices.....	33
I.1. Introduction.....	33
I.2. Drive management screen.....	33
I.3. NULL drive creation.....	34
I.3.a. What is a NULL drive and what are its uses?.....	34
I.3.b. Drive creation (NULL drive).....	34
I.4. Tape drive creation.....	35
I.4.a. What is a tape drive?.....	35
I.4.b. Drive creation.....	35
I.4.c. Reading the tape label.....	36
I.4.d. Possible messages.....	36
I.5. File drive creation.....	36

I.5.a. What is a File drive and what are its uses?.....	36
I.5.b. Drive creation (File drive).....	37
I.6. Drive deletion.....	37
I.7. The Library Management Screen.....	38
I.8. Tape Library creation.....	38
I.8.a. What is a Tape Library?.....	38
I.8.b. Library creation.....	38
I.9. File library creation.....	40
I.9.a. What is a File library?.....	40
I.9.b. Library creation (File library).....	40
I.10. Library deletion.....	41
II. Specific Configuration.....	42
II.1. Specific Name Resolution and servers with multiple Network Interface Cards (NIC).....	42
II.1.a. Introduction.....	42
II.1.b. Client machine configuration.....	42
II.1.c. Arkeia backup server machine configuration.....	42
II.1.d. Hosts file used by Arkeia.....	42
Arkeia hosts file.....	43
II.1.e. Setting the NLP_HOSTFILE and NLP_ONLYHOSTFILE.....	43
Please note:.....	43
II.1.f. Syntax of the Arkeia specific hosts file.....	43
II.1.g. NLP_HOSTNAME usage.....	43
A short configuration example.....	44
II.1.h. Complete configuration example.....	44
II.2. How to configure Arkeia with a multiple domains network architecture.....	45
II.3. How to use different TCP ports.....	46
II.4. How to configure Arkeia to work from behind a firewall.....	46
II.4.a. Introduction.....	46
II.4.b. Standard procedure.....	46
II.4.c. SSH configuration.....	47
Short example.....	47
WHERE TO BACKUP?.....	49
I. The Tape Pools and Drivepacks concepts.....	49
I.1. Understanding the issues.....	49
I.2. Arkeia’s approach.....	49
II. Tape Pools.....	50
II.1. Definition and uses.....	50
II.2. The “Pools management” screen.....	50
II.3. Pool creation.....	50
II.4. Pool deletion.....	51
II.5. The Pool management window.....	51
II.6. Thread and tape order.....	52
II.7. Pool statistics.....	52
II.8. The Scratch Pool.....	53
III. Tapes.....	53
III.1. Introduction.....	53
III.2. The “Tapes management” screen.....	54
III.3. Tape creation.....	54
III.3.a. Introduction.....	54

III.3.b. Standard tape creation.....	55
III.3.c. “NULL” tape creation.....	56
III.3.d. “FILE” tape creation.....	56
III.4. Tape Recycling.....	57
III.5. How does Arkeia create tape names?.....	57
III.6. Tapes deletion.....	58
III.7. Modifying tapes: the Tape(s) modification window.....	59
Authorizations.....	60
Recycling pool.....	60
Recycling mode.....	60
Tape access mode.....	60
Tape assignment pool.....	61
Comment zone.....	61
III.8. Detailed tape information: the tape screen.....	62
III.9. Writing the label on a tape.....	63
III.10. Tape recycling	64
IV. Drivepacks.....	65
IV.1. Description and use.....	65
IV.2. Description of the “Drivepacks” screen.....	65
IV.3. Creating a Drivepack:.....	66
IV.4. Drive Priority.....	66
IV.5. Number of drives.....	67
IV.6. Deleting a Drivepack.....	68
WHAT TO BACKUP?.....	69
I. The “Savepack” concept.....	69
I.1. Understanding the issue.....	69
I.2. Arkeia’s approach.....	69
II. Savepack management.....	70
II.1. Description of the “Savepacks management” screen.....	70
II.2. Savepack Creation.....	70
II.3. Savepack Deletion.....	71
II.4. Adding a tree in the Savepack: the Network Navigator.....	71
II.5. Deleting a tree in a Savepack.....	73
II.6. Inserting a Savepack in a Savepack.....	73
II.7. Advanced Savepack options.....	74
II.7.a. Command before backup.....	75
Table of backup execution conditions.....	76
II.7.b. Command after backup.....	76
Table of backup execution conditions.....	76
II.7.c. Number of retries.....	77
II.7.d. Compression.....	77
II.7.e. Encryption.....	77
II.7.f. Follow symbolic links.....	78
II.7.g. Follow file systems.....	78
II.7.h. File systems allowed.....	78
On Unix.....	78
On Novell, Windows 9X, ME, NT and 2000.....	78
II.7.i. Reset access times.....	79
II.7.j. Find filter.....	79

II.7.k. Inclusion filter.....	79
II.7.l. Exclusion filter.....	79
II.7.m. Regular expression.....	79
II.8. “Advanced Tree options” screen.....	81
II.9. Type of Trees.....	82
TREE.....	82
OBJECT.....	83
RAW.....	83
SAVEPACK.....	84
II.9.a. Multiflow (parallel processing on one machine).....	84
II.9.b. Priority.....	84
II.9.c. Chain.....	85
II.9.d. Command before tree backup.....	85
Table of backup execution conditions.....	85
II.9.e. Command after the backup of a tree.....	86
Table of backup execution conditions.....	86
II.9.f. Compression.....	87
II.9.g. Encryption.....	87
II.9.h. Follow symbolic links.....	87
II.9.i. Follow file systems.....	87
II.9.j. File systems allowed.....	87
II.9.k. Reset access times	87
II.9.l. Find filter.....	88
II.9.m. Inclusion filter.....	88
II.9.n. Exclusion filter.....	88
II.9.o. Regular expression.....	88
III. Specific examples and cases.....	90
III.1. How to prevent the backup of a specific directory: the .OPB_NOBACKUP file.....	90
BACKUPS.....	91
I. Interactive backups.....	91
I.1. Introduction.....	91
I.2. The “Interactive backup” screen.....	92
I.3. Starting an Interactive Backup.....	93
I.4. Specific Options.....	94
I.4.a. Types of Backup.....	94
Total backup.....	94
Incremental (or differential).....	94
Archive.....	94
“Standard” or “Continuous”.....	94
I.4.b. Tape strategy.....	94
I.4.c. Parallelism.....	94
I.4.d. Use email.....	94
I.4.e. Tag (Optional).....	94
I.5. Monitor the Backup: the “Backup” screen.....	95
I.6. Specific options of the “Backup” Screen.....	95
I.6.a. Adding a Savepack.....	95
I.6.b. Tree status.....	96
I.7. Connection to backup or restore.....	97
II. Periodic backups.....	98

II.1. Introduction.....	98
II.2. The “Periodic Backup” window.....	99
II.3. Create a Periodic Backup.....	99
II.4. Setting the standard Periodic Backup options.....	100
II.5. Periodic Backup deletion.....	102
II.6. Managing the Periodic Backup levels	102
Programmed execution diagram (without any exceptions):.....	103
Programmed execution diagram (with exceptions):.....	103
II.6.a. Adding a level.....	104
II.6.b. Changing a level.....	104
II.6.c. Deleting a level.....	105
II.7. The “Advanced options” screen.....	105
II.7.a. Exception (occurrences) management.....	106
II.7.b. System command.....	107
II.7.c. Command before.....	107
II.7.d. Command after.....	107
II.8. The Schedule Viewer.....	108
II.9. The “Periodic Backup Assistant”	108
PERIODIC BACKUP POLICY.....	113
I. Periodic Backup Proceedings.....	113
I.1. Introduction.....	113
I.2. Methodology.....	113
I.2.a. Simple Periodic Backups.....	114
Tapes.....	114
Drives.....	114
Savepack.....	114
Create a Periodic Backup.....	114
I.2.b. Semi–periodic backups.....	114
I.2.c. Periodic Backup Assistant.....	115
I.2.d. Tips and techniques.....	115
II. Examples.....	115
II.1. A complete backup each day.....	115
II.1.a. Policy definition.....	115
II.1.b. Solution and Analysis.....	115
Step 1: Tape pool evaluation.....	116
Step 2: Drivepack selection.....	116
Step 3: Savepack selection.....	116
Step 4: Periodic Backup creation.....	116
Step 5: Check the settings.....	117
II.1.c. Conclusion.....	117
II.2. A “Total” Backup from Monday to Friday.....	117
II.2.a. Policy definition.....	117
II.2.b. Solution and Analysis.....	118
Step 1: Tape pool evaluation.....	118
Step 2: Drivepack selection.....	118
Step 3: Savepack selection.....	118
Step 4: Periodic Backup – first level creation.....	119
Step 5: Periodic Backup – second level creation.....	119
Step 6: Settings verification.....	120

II.2.c. Conclusion.....	120
II.3. “Total” Backup on Monday, and “Incremental” Backup from Tuesday to Friday.....	120
II.3.a. Policy definition.....	120
II.3.b. Solution and Analysis.....	120
Step 1: Tape pools evaluation.....	121
Step 2: Drivepack selection.....	121
Step 3: Savepack selection.....	122
Step 4: Periodic Backup – first level creation.....	122
Step 5: Periodic Backup – second level creation.....	122
Step 6: Verify the settings.....	123
II.3.c. Conclusion.....	123
II.4. A more complex and complete backup Policy.....	123
II.4.a. Policy definition.....	123
II.4.b. Solution and Analysis.....	124
II.4.c. Part one: the “Yearly” backup.....	124
Step 1: Tape pool evaluation.....	124
Step 2: Drivepack selection.....	124
Step 3: Savepack selection.....	125
Step 4: Periodic backup creation.....	125
Step 5: Verify the settings.....	125
II.4.d. Part two: Sensitive machines backup.....	126
Step 1: Tape pool evaluation.....	126
Step 2: Drivepack selection.....	127
Step 3: Savepack selection.....	127
Step 4: “Sensitive Backup” – first level creation.....	127
Step 5: “Sensitive Backup” – second level creation.....	128
Step 6: Verify the settings.....	128
II.4.e. Conclusion.....	128
II.5. Two simultaneous “Total” Backups on two different domains, from Monday to Friday.....	129
II.5.a. Policy definition	129
II.5.b. Solution and Analysis.....	129
Step 1: Tape pool evaluation.....	129
Step 2: Drivepack selection.....	130
Step 3: Savepack selection.....	130
Step 4: “Domain 1 Full Backup” – first level creation.....	130
Step 5: “Domain 1 Full Backup” – second level creation.....	131
Step 6: “Domain 2 Full Backup” – first level creation.....	131
Step 7: “Domain 2 Full Backup” – second level creation.....	132
Step 8: Verify the settings.....	133
II.5.c. Conclusion.....	133
II.6. A “Total” backup once a week plus a daily backup of the modified files.....	134
II.6.a. Policy definition.....	134
II.6.b. Solution and Analysis.....	134
Step 1: Tape pool evaluation.....	134
Step 2: Drivepack selection.....	135
Step 3: Savepack selection.....	135
Step 4: “Full Backup” – first level creation.....	135
Step 5: “Full Backup” – second level creation.....	136
Step 6: Verify the settings.....	137

II.6.c. Conclusion.....	137
RESTORATION.....	138
I. Principles of Restoration.....	138
I.1. Definition.....	138
I.2. Arkeia’s approach.....	138
II. Restoration Management.....	139
II.1. The “Restoration” screen.....	139
II.2. Select what to restore: the “Time Navigator”	140
II.3. Select a single tree or file.....	141
II.4. Using the “time sliders” in the Time Navigator.....	141
II.5. Modify a path or a name.....	142
II.6. Paths syntax.....	142
II.7. Backup information of a file or tree.....	143
II.8. Applying Redirection.....	145
Redirection to another directory.....	145
Redirection to another name.....	145
Redirection to another machine.....	145
II.9. Searching a file to restore.....	146
II.9.a. Search criteria.....	146
Filename matching exactly.....	146
Filename containing:	146
Filename ending by.....	146
Filename starting with.....	146
II.10. The “Restoration” monitor.....	147
II.11. Index browser.....	147
II.12. Restoration options.....	148
II.13. List of tapes used for the restoration.....	148
II.14. Who has access to the restore function?.....	149
II.14.a. The administrator.....	149
II.14.b. The operator.....	149
II.14.c. The user.....	149
TUNING ARKEIA.....	151
I. How to increase performance.....	151
I.1. Introduction.....	151
I.2. Performance expectations.....	151
I.3. Use of multiflows.....	152
I.3.a. Basics.....	152
I.3.b. Parallelism of multiple machines.....	152
I.3.c. Parallelism of a specific machine.....	152
I.3.d. Last settings.....	153
II. How to limit backup speed.....	154
II.1. Why should backup speed be limited?.....	154
II.2. How to limit backup speed.....	154
II.2.a. The graphical “cruise control”.....	154
II.2.b. Default Backup speed limitation.....	154
III. Priority.....	155
III.1. Introduction.....	155
III.2. How to use Priority.....	155
IV. Chaining.....	156

IV.1. Introduction.....	156
IV.2. How to use Chaining.....	156
V. Configuring compression in Arkeia.....	157
V.1. Introduction.....	157
V.2. Compression settings.....	157
V.2.a. Savepacks.....	157
V.2.b. Setting a default compression method for a particular client.....	157
V.2.c. Setting a compression algorithm to a specific file type.....	158
ARKEIA MANAGEMENT AND ADMINISTRATION.....	161
I. How to move the index database to another directory.....	161
I.1. Problem definition.....	161
I.2. Procedure.....	161
II. Move the Arkeia backup server to a new machine.....	162
II.1. Problem definition.....	162
II.2. Procedure.....	162
III. Removing an installed client.....	163
III.1. Problem definition.....	163
III.2. Procedure.....	163
IV. Creating Arkeia users.....	164
IV.1. Problem definition.....	164
IV.2. Creating a user.....	164
IV.2.a. The Users Management window.....	164
IV.2.b. User creation.....	165
IV.2.c. User modification.....	165
IV.2.d. User deletion.....	165
V. Setting up the email feature.....	166
V.1. Problem definition.....	166
V.2. Procedure.....	166
V.2.a. The “Owner email”.....	166
V.2.b. Configure the email feature in Backups.....	166
V.2.c. Notifications provided.....	166
SECURITY IN ARKEIA.....	167
I. Encryption Configuration.....	167
I.1. Introduction.....	167
I.2. Encryption configuration.....	167
I.2.a. Encryption Type.....	167
I.2.b. Configuration.....	168
I.2.c. Encryption selection.....	168
I.2.d. Encryption key.....	168
Encryption of a whole tree.....	168
Encryption of a directory.....	169
II. Users authentication and ROLES.....	170
II.1. 3.1 Introduction.....	170
II.2. Authentication and authorization.....	170
II.2.a. Definition.....	170
II.2.b. Authorization files.....	170
II.3. Restricting access of a server/client to a specific client.....	171
II.3.a. Basics.....	171
II.3.b. Restrict navigation.....	172

II.3.c. Forbid access by a specific backup server.....	172
II.3.d. Restrict the access rights to a specific server.....	172
II.4. Proxy or relation files.....	172
II.5. ROLES management.....	174
II.5.a. Basics.....	174
II.5.b. Arkeia backup and restore configuration management, using the ROLES.....	174
Administrator.....	174
Operator	174
User.....	174
TROUBLESHOOTING.....	175
I. Arkeia configuration and usage.....	175
II. Error codes and messages.....	177
II.1. Introduction.....	177
II.2. Error Codes.....	177
II.3. Error messages.....	177
LOG MANAGEMENT.....	179
I. The Arkeia Journal.....	179
I.1. Introduction.....	179
I.2. The “Log Consult” screen.....	179
I.3. Display filter.....	180
I.4. Change the displayed month.....	181
II. The “Backup done” screen.....	182
II.1. Introduction.....	182
II.2. Checking a specific backup: the “backup done log” screen.....	183
III. Other logs.....	184
III.1. The “Drive” log.....	184
III.2. The “Tape” log.....	185
III.3. The “Restore” log.....	186
BACKUP OF OPEN SOURCE DATABASES.....	187
I. Important notice to Oracle users.....	187
II. MySQL.....	187
II.1. How to backup a MySQL database offline.....	187
II.2. How to backup a MySQL database online.....	189
II.2.a. First method.....	189
II.2.b. Second method.....	190
II.2.c. Now with ARKEIA!.....	191
ARKEIA TOOLS.....	194
I. Extracting data from tape: readarkeia.....	194
I.1. Introduction.....	194
I.2. Getting readarkeia.....	194
I.3. Using readarkeia.....	195
I.3.a. List tape’s content.....	195
I.3.b. Extract all the files on a tape.....	195
I.3.c. Extract a specific file.....	195
II. Restoring the database: arkrstdb.....	196
II.1. Introduction.....	196
II.2. Usage.....	196
III. Testing the connectivity: chknlp.....	197
III.1. Introduction.....	197

III.2. Usage.....	197
III.3. Nature of the tests.....	198
III.3.a. Test 0.....	198
In case of error.....	198
Correct output.....	198
III.3.b. Test 1.....	198
In case of error.....	198
Correct output.....	199
III.3.c. Test 2.....	199
In case of error.....	199
Correct output.....	199
III.3.d. Test 3.....	199
In case of error.....	199
Correct output.....	200
III.3.e. Test 4.....	200
In case of error.....	200
Correct output.....	201
IV. Duplicate a tape: tpdup.....	202
IV.1. Introduction.....	202
IV.2. Requirements.....	202
IV.3. Installation and use of tpdup.....	202
IV.3.a. Installation.....	202
IV.3.b. Configuration and usage.....	202
drv.lst.....	202
tpclone.lst.....	203
IV.3.c. Force end of tapes.....	203
IV.3.d. Duplication start.....	203
IV.3.e. Reuse of duplicata.....	204
GLOSSARY.....	205
arkc.....	205
ARKPER (in /usr/bin).....	205
arkrstdb.....	205
ATAPI.....	205
Autoloader.....	205
Control Device.....	206
daemon.....	206
Drive.....	206
Drivepack.....	206
File Library.....	206
Flow.....	206
inode.....	206
Interactive Backup.....	206
Library.....	206
LUN.....	207
Multiflow.....	207
NLSERVD (in /usr/bin).....	207
Periodic Backup.....	207
Periodicity.....	207
Random mode.....	207

readarkeia.....	207
Recycle.....	207
Rewind Device.....	207
Robot.....	208
rpm.....	208
Savepack.....	208
SCSI.....	208
Slot.....	208
stks.....	208
Tape Pool.....	208
Tar.gz.....	208
Type 1 machine.....	208
Type 2 machine.....	209
Validity.....	209
Voltag.....	209