





Booting Procedure Guide For AIX

December 2022

Version 9.1

Cristie Software Ltd. New Mill Chestnut Lane Stroud GL5 3EW United Kingdom Tel:+44(0)1453 847000 Fax:+44(0)1453 847001 support@cristie.com Cristie Data Products GmbH Nordring 53-55

63843 Niedemberg Germany Tel: +49 (0) 60 28/97 95-0 Fax: +49 (0) 60 28/97 95 7-99 cbmr@cristie.de Cristie Nordic AB Knarramäsgatan 7 164 40 Kista Sweden Tel:+46(0)8 718 43 30 cbmr@cristie.se

Copyright © 2013-2022 Cristie Software Ltd. All rights reserved.

The software contains proprietary information of Cristie Software Ltd.; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

Due to continued product development this information may change without notice. The information and intellectual property contained herein is confidential between Cristie Software Ltd. and the client and remains the exclusive property of Cristie Software Ltd. If you find any problems in the documentation, please report them to us in writing. Cristie Software Ltd. does not warrant that this document is error-free.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior written permission of Cristie Software Ltd.

- IBM Tivoli Storage Manager (TSM), AIX and TIVOLI are trademarks of the IBM Corporation.
- IBM Spectrum Protect is a trademark of the IBM Corporation.
- IBM Virtual I/O Server (VIOS) is a trademark of the IBM Corporation.
- NetWorker and Avam ar are tradem arks of the Dell EMC Corporation.
- vSphere, vCenter and vCloud are tradem arks of VMware Inc.
- Hyper-V is a trademark of Microsoft Corporation.
- Azure is a trademark of Microsoft Corporation.
- Am azon Web Services (AWS) and Am azon Elastic Compute Cloud (EC2) are trademarks of Am azon. com, Inc.
- Cohesity DataProtect is a trademark of Cohesity Inc.
- CloneManager® is a registered tradem ark of Cristie Software Ltd.

PC-BaX, UBax, Cristie P4VM (Protect for VMs), Cristie Storage Manager (CSM), SDB, ABMR (Bare Machine Recovery for EMC Avamar), NBMR (Bare Machine Recovery for EMC NetWorker), TBMR (Bare Machine Recovery for Spectrum Protect/TSM), CBMR (Cristie Bare Machine Recovery), CoBMR (Bare Machine Recovery for Cohesity DataProtect) and CRISP (Cristie Recovery ISO Producer) are all trademarks of Cristie Software Ltd..

> Cristie Software Ltd New Mill Chestnut Lane Stroud GL5 3EW UK

Tel: +44 (0) 1453 847009 Email: support@cristie.com

Contents

1	Introduction	4
	1.1 Document Notes	
	1.2 Overview	
2	Network Booting	5
3	Creating the Environment ISO	6
4	Creating a Bootable CD	7
5	Booting using BOOTP	8
	5.1 Unpacking the ISO	
	5.1.1 AIX 6.1 and later	
	5.2 Setting up	
	5.2.1 Example	
6	Booting using DHCP	10
	6.1 Unpacking the ISO	10
	6.1.1 AIX 6.1 and later	
	6.2 Setting up	10
7	Configuring the DHCP Server	11
	7.1 Example	11
8	Troubleshooting	13
	8.1 If creating the ISO fails	
	8.2 NFS Server	13
	8.3 TFTP Server	13
	8.4 DHCP Server	
	8.5 BOOTP	
	8.7 Other booting failures	
	8.7.1 Phase 1	
	8.7.2 Mount net procedure	
	8.7.3 Mount CD procedure	
	8.7.5 Notes	
	Fallback	
9	Cristie Technical Support	18

1 Introduction

This document describes the steps required to boot the recovery environment for AIX and a troubleshooting guide for booting problems. The AIX recovery environment may boot over the network using either Dynamic Host Control Protocol (DHCP) or Bootstrap Protocol (BOOTP) or directly from a CD.

1.1 Document Notes

Commands to be typed at the console are preceded by '>', ie:

```
> echo "this is a command"
```

Output from commands will be preceded by '\$', ie:

```
> echo "this is a command"
$ this is a command
```

The syntax "< some text >" is used in a command to indicate a value that should be replaced with the appropriate value for your environment.

1.2 Overview

The boot procedure is split into two phases - **Phase I** and **Phase II**. The first phase is run from an executable file containing a kernel and ramdisk called '**bootfile.exe**'. As this file is quite small, approximately 12Mb, it contains only the necessary drivers to access the CD and network to begin Phase II.

Phase I will perform some initialisation, then attempt to mount a Network File System (NFS) over /SPOT. If this operation fails, it will attempt to mount the CD booted from the CD drive 0 and CD drive 1. The address of the NFS server is supplied either using DHCP or added to the environment during CD creation.

Once the /SPOT directory has been mounted, the files necessary to begin Phase II are copied or linked to the bootable environment. Once complete, Phase II begins.

Phase II loads more drivers from the /SPOT file-system and initialises hardware missed during Phase I. It then sets up networking, user administration and the daemons necessary for performing recovery.

2 Network Booting

A network boot environment consists of at least two systems - a client and a server. The server supplies the files that the client requests when it boots.

When building a recovery environment, it is important to note that this environment is created **for the client**. As such, the recovery environment must be built on the client, then transferred to the server where it can be supplied to booting systems.

6

3 Creating the Environment ISO

The AIX recovery environment is created using the **mkdrcd** command, which can be accessed from the GUI or run directly from the command line.

Whilst there are many options for this command, this document will only describe the options required for the purposes of creating and troubleshooting a bootable environment.

4 Creating a Bootable CD

This is the simplest method of booting. The mkdrcd program is used to create a bootable ISO file and this is burned to a blank CD. For example, to create the bootable ISO file / tmp/myisofile.iso, then burn this to a CD, the following commands would be used:

- > mkdrcd -o /tmp/myisofile.iso
- > burn cd -d cd0 /tmp/myisofile.iso

5 Booting using BOOTP

8

When booting using the **BOOTP** protocol, the address of the NFS server must be hardcoded into the recovery environment. This is done by supplying the mkdrcd program with the address using the '-static' option.

For example, to boot from the server 192.168.1.1 with the environment in the directory /export/recoverycd, the following commands would be run on the client:

> mkdrcd -o /tmp/myenvironment.iso --static 192.168.1.1:/export/recoverycd

This file must then be transferred to 192.168.1.1 (using scp for example) and then used to set up a boot environment on 192.168.1.1. Here we assume this is being performed on an AIX server - if another type of server is used, please refer to that operating system's documentation for setting up BOOTP and NFS. This operation consists of four steps:

- 1. Unpack the ISO file to the directory to be exported
- 2. Export the directory using NFS
- 3. Copy the bootfile.exe to the TFTP server
- 4. Add an entry to the BOOTP table

5.1 Unpacking the ISO

The ISO file contains the bootable recovery environment to be loaded during Phase II. This is unpacked to a directory which is shared over NFS.

The later versions of AIX have the ability to mount the ISO file directly. However, earlier versions require a more complicated procedure.

5.1.1 AIX 6.1 and later

```
> mkdir /mnt/iso
```

```
> loopmount -i <path to iso> -o "-o ro -V cdrfs" -m /mnt/iso
> cp -r /mnt/iso/* <shared directory>
```

5.2 Setting up

The directory used must now be exported using NFS. This is performed by adding an extra export directive to the /etc/export file and informing the NFS server:

```
> echo "<shared directory> -ro, anon=0" >> /etc/exports
> /usr/sbin/exportfs -a
```

Note: see 'man exportfs' for more details

The next step is to copy the bootfile.exe from the shared directory to the tftp server location (usually /tftpboot/):

> cp <shared directory>/ppc/chrp/bootfile.exe /tftpboot/

Finally, an entry must be added to the /etc/bootptab file:

```
> echo "<client name>: bf=/tftpboot/bootfile.exe: ip=<IP address of
client>: ht=ethernet: ha=<MAC adress>: sa=<IP address of server>:
gw=<Gateway>: sm=<subnet mask>" >> /etc/bootptab
```

Note: see 'man bootptab' for more details

5.2.1 Example

In the following example, the client 'testclient' with IP address '192.168.1.50' and MAC address '001234567890' is booted from the server '192.168.1.1' using /export/recoverycd as the shared directory. Both machines are on the same subnet.

```
> mkdrcd -o /tmp/dr.iso --static 192.168.1.1:/export/recoverycd
> mkdir /mnt/iso
> loopmount -i /tmp/dr.iso -o "-o ro -V cdrfs" -m /mnt/iso
> cp /mnt/iso/* /export/recoverycd
> cp /export/recoverycd/ppc/chrp/bootfile.exe /tftpboot/
> echo "/export/recoverycd -ro,anon=0" >> /etc/exports
> /usr/sbin/exportfs -a
> echo "testclient: bf=/tftpboot/bootfile.exe: ip=192.168.1.50:
ht=ethernet: ha=001234567890: sa=192.168.1.1: sm=255.255.255.0" >> /etc/
bootptab
```

6 Booting using DHCP

10

When booting using **DHCP**, the address of the NFS server need not be hard-coded in the recovery ISO. This is because the DHCP protocol is more flexible and allows a greater range of values to be passed during configuration.

The recovery environment is easily created on the client using the '**mkdrcd**' tool: > mkdrcd -o /tmp/myrecoveryenvironment.iso

This file must then be transferred to 192.168.1.1 (using scp for example) and then used to set up a boot environment on any server. Here we assume this operation is being performed on an AIX server - if another type of server is used then refer to that operating system's documentation for setting up DHCP and NFS.

This operation consists of four steps:

- 1. Unpack the ISO file to the directory to be exported
- 2. Export the directory using NFS
- 3. Copy the bootfile.exe to the TFTP server
- 4. Add an entry to the DHCP configuration

6.1 Unpacking the ISO

The ISO file contains the bootable recovery environment to be loaded during Phase II. This is unpacked to a directory which is shared over NFS.

The later versions of AIX have the ability to mount the ISO file directly. However, earlier versions require a more complicated procedure.

6.1.1 AIX 6.1 and later

```
> mkdir /mnt/iso
```

```
> loopmount -i <path to iso> -o "-o ro -V cdrfs" -m /mnt/iso
> cp -r /mnt/iso/* <shared directory>
```

6.2 Setting up

The directory used must now be exported using NFS. This is performed by adding an extra export directive to the /etc/export file and informing the NFS server:

```
> echo "<shared directory> -ro, anon=0" >> /etc/exports
> /usr/sbin/exportfs -a
```

Note: see 'man exportfs' for more details

The next step is to copy the bootfile.exe from the shared directory to the tftp server location (usually /tftpboot/):

```
> cp <shared directory>/ppc/chrp/bootfile.exe /tftpboot/
```

7 Configuring the DHCP Server

The DHCP server configuration is stored in the file /etc/dhcpsd.cnf. The following information is set for each client that must be booted:

- IP address
- hostname
- MAC address
- Boot file location
- Server address
- NFS Server
- NFS Directory

This is specified by adding an entry of the following form:

```
<client name> 1 <MAC Address> <IP address>
{
    option 12 <hostname>
    option bf /tftpboot/bootfile.exe
    option sa <server address>
    option 151 <NFS server address>
    option 152 <NFS server directory>
}
```

Unfortunately, the AIX DHCP system is not very user friendly; the values for options 151 and 152 must be specified in 'unsigned byte form'. The NFS Server address may be converted as follows:

```
> echo "<NFS server address>" | tr . " " | xargs printf "%x "
```

The NFS Server directory may be converted as follows:

```
> echo "<NFS server directory>" | od -t x1 -A n
```

For example:

```
> echo "192.168.1.1" | tr . " " | xargs printf "%.2x "
$ c0 a8 01 01
> echo "/export/recoverycd" | od -t x1 -A n
$ 2f 65 78 70 6f 72 74 2f 72 65 63 6f 76 65 72 79 63 64 0a
```

7.1 Example

In the following example, the client 'testclient' with IP address '192.168.1.50' is booted from the server '192.168.1.1' using /export/recoverycd as the shared directory. Both machines are on the same subnet.

> mkdrcd -o /tmp/dr.iso

12

```
> mkdir /mnt/iso
> loopmount -i /tmp/dr.iso -o "-o ro -V cdrfs" -m /mnt/iso
> cp /mnt/iso/* /export/recoverycd
> cp /export/recoverycd/ppc/chrp/bootfile.exe /tftpboot/
> echo "/export/recoverycd -ro,anon=0" >> /etc/exports
> /usr/sbin/exportfs -a
> echo "192.168.1.1" | tr . " " | xargs printf "%.2x "
$ c0 a8 01 01
> echo "/export/recoverycd" | od -t x1 -A n
$ 2f 65 78 70 6f 72 74 2f 72 65 63 6f 76 65 72 79 63 64 0a
Edit /etc/dhcpsd.cnf to contain the following information:
testclient 1 001234567890 192.168.1.50
{
option 12 testclient
option bf /tftpboot/bootfile.exe
option sa 192.168.1.1
```

option 151 c0 a8 01 01 option 152 2f 65 78 70 6f 72 74 2f 72 65 63 6f 76 65 72 79 63 64 0a }

8 Troubleshooting

Please view the following sections for troubleshooting guidance.

8.1 If creating the ISO fails

Check that the 'mkdrcd' program is running on the client machine. If the program stops with an error, then using the '-v' flag may provide some more information.

8.2 NFS Server

Before attempting to boot a client, it is useful to check that it would have access to the shared NFS directory.

On the client, type:

```
> mkdir /mnt/tmp
> mount <NFS server address>:<NFS server directory> /mnt/tmp
```

If this fails, then you need to check the NFS services are started on the server:

```
> lssrc -g nfs
```

If they are not started:

> startsrc -g nfs

You can also check that the directory is correctly exported:

```
> showmount -e <NFS server>
```

8.3 **TFTP Server**

The **TFTP** server is responsible for serving the first file required for booting to the client. It is therefore very important this is working.

From the client run:

```
> echo "get /tftpboot/bootfile.exe" | tftp <server IP address>
> lssrc -s tftpd
```

If it is not started, run:

> startsrc -s tftpd

TFTP is usually started via the 'inet' service, so make sure the inetd service is running:

> lssrc -s inetd

Finally, check the inet.conf file to ensure that the tftp service is enabled:

```
> grep tftp /etc/inetd.conf
$ tftp dgram udp6 SRC nobody /usr/sbin/tftpd tftpd -n
```

If this is commented out (using '#'), then uncomment the line and restart inetd:

> refresh -s inetd

8.4 DHCP Server

If DHCP is in use, then check the DHCP server by running:

```
> lssrc -s dhcpsd
```

If it is not started, run:

> startsrc -s dhcpsd

It is possible to check DHCP operation on the running client by adding the line 'interface <connected ethernet interface>' to /etc/dhcpcd.cnf:

```
> echo "interface en0" >> /etc/dhcpcd.cnf
```

Running the dhcpcd daemon should then output information to the /usr/tmp/dhcpcd. log file.

```
> /usr/sbin/dhcpcd -r
```

The DHCP server must also respond to BOOTP requests from booting clients:

```
> grep supportBOOTP /etc/dhcpsd.cnf
$ supportBOOTP yes
```

8.5 BOOTP

If using only the **BOOTP** protocol, then the AIX machine must respond to BOOTP requests from the client. This is configured using inetd.conf:

```
> grep bootps /etc/inetd.conf
$ bootps dgram udp wait root /usr/sbin/bootpd bootpd /etc/bootptab
```

If this is commented out (using '#'), then uncomment the line and restart the service: > refresh -s inetd

8.6 Network booting multiple systems

In many cases a single recovery image may be used to multiple systems - for example all AIX 6.1 machines. Unfortunately, if the network hardware differs between thee machines, then an environment created on one machine may not be used to boot another.

To avoid this situation, the mkdrcd command provides the '--driver' switch to add driver packages to the boot image.

To determine the driver package required for a given interface, you must examine the output of the configuration manager. For example, to find the driver used for ent0:

```
> cfgmgr -v | grep invoking | grep ent0
$ invoking /usr/lib/methods/cfgvioent -2 -1 ent0
```

It is then possible to search for the package that provides this file:

```
> lslpp -fc | grep /usr/lib/methods/cfgvioent
$ /usr/lib/objrepos:devices.vdevice.IBM.l-lan.rte
7.1.0.1:/usr/lib/methods/cfgvioent
```

It is then possible to make a recovery image that supports this system on another system using the command:

> mkdrcd -o <recovery image> --driver devices.vdevice.IBM.l-lan.rte

The image created will be able to network boot the machine it was created on and any machines whose network interfaces use the 'devices.vdevice.IBM.l-lan.rte' package.

8.7 Other booting failures

It is possible to track the progress of the boot sequence using the LED codes on the front panel of the machine.

The following is a table of LED codes that will be cycled through as the machine boots. The 'MOUNT NET' and 'MOUNT CD' codes represent the start of procedures that are covered in separate tables.

The LED codes will display 50 plus the current code for indicating the action failed. For example, if DHCP (103) fails then 153 will be displayed immediately afterwards.

000	STARTING	
001	DEV CFG 1 START	About to run cfgmgr to initialise hardware for boot
002	start ifaces	Scans available network interfaces
003	MOUNT NET	Mounts network boot (see mount net procedure)
004	MOUNT CD	Mounts CD (called if MOUNT NET fails)
005	LINKING BINARIES	Links binaries from CD/Net to root FS
006	COPYING BINARIES	Copies binaries from CD/Net to root FS (for those that cannot be linked)
007	COPYING SBINARIES	Copies secure binaries from CD/Net to root FS (for those that cannot be linked)

8.7.1 Phase 1

16

008	COPYING LIBRARIES	Copies libraries from CD/Net to root FS (for those that cannot be linked)
009	LINKING LIBRARIES	Links remaining libraries from CD/Net to root FS
010	REPAIRING LIBC	Replace boot libc with full libc
011	LINKING DRIVERS	Links device drives to root FS
049	end phase 1	

8.7.2 Mount net procedure

100	LOOPBACK	Starts up loopback interface
101	BOOTINFO	Calls bootinfo -c to find bootp parameters
102	DEFAULTS	Accesses the /etc/PXEVARS file - this file will only exist if the ISO was created using - static
103	DHCP	Runs dhcpd to get an IP address, only called if DEFAULTS or BOOTINFO failed
104	FINDING SERVER	Use either PXEVARS or DHCP to locate the NFS server
105	IFCONFIG <dev> <ip> <subnet></subnet></ip></dev>	Sets the interface to the parameters retrieved from BOOTINFO if DHCP is not in use or failed
106	ROUTING <server> <gateway></gateway></server>	Adds a default route to access the NFS server if bootp or dhcp sets a gateway
107	MOUNT <server> <dir></dir></server>	Mounts the NFS server
108	CHECK SPOT	Checks to determine the mounted folder is good
149	END MOUNT NET	

8.7.3 Mount CD procedure

May be run multiple times.

200	MOUNT <device></device>	Mounts the device at /SPOT
201	CHECK SPOT	Checks to determine the mounted folder is good
249	END MOUNT CD	

8.7.4 Phase 2

500	STARTING DR	
501	INITIALISING ENV	Sets up library paths
502	INIT FILES	Copies startup files from /SPOT
503	STRLOAD	Runs 'strload' to load drivers for cfgmgr
504	CFGMGR II	Runs 'cfgmgr -p 2' to initialise hardware
505	LIB PATH	Links library paths from /SPOT
506	LINKING	Links other files from /SPOT
507	USERS	Sets up root and system users
508	NETWORKING	Checks addresses
509	DHCP P2	Attempts to determine an IP address using DHCP if networking has failed
510	SSHD	Starts ssh daemon if present
511	CONSOLE	Starts the console selector
512	STARTUP SCRIPT	Called if a startup script was included on the CD
549	END PHASE II	

8.7.5 Notes

If **PXEVARS** fails, it indicates that 'mkdrcd' was not run with the '--static <address>: <directory>' flag, or the creation process failed. Running 'mkdrcd' using the '-v' flag may provide more information.

If any of the copying or linking stages fails, that indicates a possible problem with the NFS server or shared directory. In this case, it may be useful to check the /etc/exportfs file and ensure the contents of the shared directory matches that of the CD.

8.7.5.1 Fallback

If booting from either DHCP or BOOTP fails, the machine will attempt to boot from cd0 then cd1. It is then possible to burn the recovery environment to a CD to diagnose booting problems.

Once the machine has fallen back to the CD, select '**Open Shell**' from the 'Tools' menu. The network booting log can be found in /tmp/iflog and contains the output of the commands run during network boot.

9 Cristie Technical Support

If you have any queries or problems concerning your Cristie Bare Machine Recovery product, please contact Cristie Technical Support. To assist us in helping with your enquiry, make sure you have the following information available for the person dealing with your call:

- CBMR Version Number
- Installed OS type and version
- Any error message information (if appropriate)
- Description of when the error occurs
- All Cristie log files relating to the source or recovery machine. This is very important to help us provide a quick diagnosis of your problem

Contact Numbers - Cristie Software (UK) Limited

Technical Support	+44 (0) 1453 847 009
Toll-Free US Number	1-866-TEC-CBMR (1-866-832-2267)
Knowledgebase	<u>kb.cristie.com</u>
Forum	forum.cristie.com
Sales Enquiries	<u>sales@cristie.com</u>
Email	support@cristie.com
Web	www.cristie.com

Support Hours

05:00 to 17:00 Eastern Standard Time (EST) Monday to Friday

Out-of-Hours support available to customers with a valid Support Agreement - Severity 1

issues* only

UK Bank Holidays** classed as Out-of-Hours - Severity 1 issues only.

*Severity 1 issues are defined as: a production server failure, cannot perform recovery or actual loss of data occurring.

**For details on dates of UK Bank Holidays, please see <u>www.cristie.com/support/</u>

Cristie Software Ltd. are continually expanding their product range in line with the latest technologies. Please contact the Cristie Sales Office for the latest product range.