Backup and Recovery of RDM 4.20

A White Paper

January 25, 2005

Notes:

Visit www.ibm.com/pc/safecomputing periodically for the latest information on safe and effective computing.

Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services.

Before using this information and the product it supports, read the general information in “Notices,” on page 17.

Contents

1 Preface

This White Paper gives suggestions on how to backup and recover IBM® Remote Deployment Manager (RDM) 4.20 (or a later RDM version) servers.

It is important to understand that RDM has no built-in high-availability functionality. This document gives suggestions on how to use standard backup/recovery techniques to protect your RDM servers and how to recover from outages. It describes how we used these techniques in one specific environment.

The procedures described in this paper accomplish their desired functions in a variety of ways. There are alternate techniques available for doing all of these functions. The intent is to illustrate various methods, as well as to describe a way to implement these particular functions. To use these procedures in your own environment will probably require some extrapolation on your part.

You can use this White Paper to learn how to do the following:

- Replace an RDM server with a backup server on a different subnet.

In addition, this White Paper will help you learn some of the typical techniques that you can use to extend and customize RDM:

- Customizing RDM *Windows Clone Install* tasks.
- RDM cloning without using Microsoft SYSPREP.EXE.

1.1 Who should read this White Paper

This White Paper is intended to help skilled RDM administrators to create deployment procedures and to understand the concepts involved. To effectively use this White Paper, you should already have an extensive knowledge of your Network environment, your RDM environment, and DOS batch files.

1.2 Further reference

In addition to this paper, there are various other sources of information that you can consult for RDM and for RDM Custom tasks.

Guides

The following product documentation is available for RDM:

- *Remote Deployment Manager 4.20 Users Guide* – The main reference manual for RDM
- *Remote Deployment Manager 4.20 Installation Guide* – Describes the complete installation process of RDM
- *Remote Deployment Manager 4.20 Compatibility and Configuration Guide* – Lists RDM-supported hardware and software

Check the IBM Web site at http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-50575 to get the current versions of the above documents.

White papers

The various RDM white papers are available on the IBM Web site at http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-53487.

Online help

In general, most RDM windows have online help available (except for some message windows or other windows where no help is applicable), either using a **Help** menu or a **Help** button.

Links

The following links are available for further information:

- Support is available for supported systems (IBM and non-IBM) through e-mail or fee-based telephone support. Telephone support is not available in all countries. For more information about the fee-based telephone support, go to http://www.ibm.com/support or http://service.software.ibm.com/supportline.html. For more information about e-mail support, refer to the RDM home page.

Important: Before using RDM 4.20, check the compatibility test results and browse the rest of the RDM Web site for additional information and tips concerning the installation and use of RDM.

2 Test Environment

This section describes the environment used while creating this document. It is a lab environment that, although it is sufficient to illustrate the principles, does not necessarily reflect a typical production environment (due to a limited supply of hardware being available for testing). We point out contrivances due to the lab environment in the text, as appropriate.

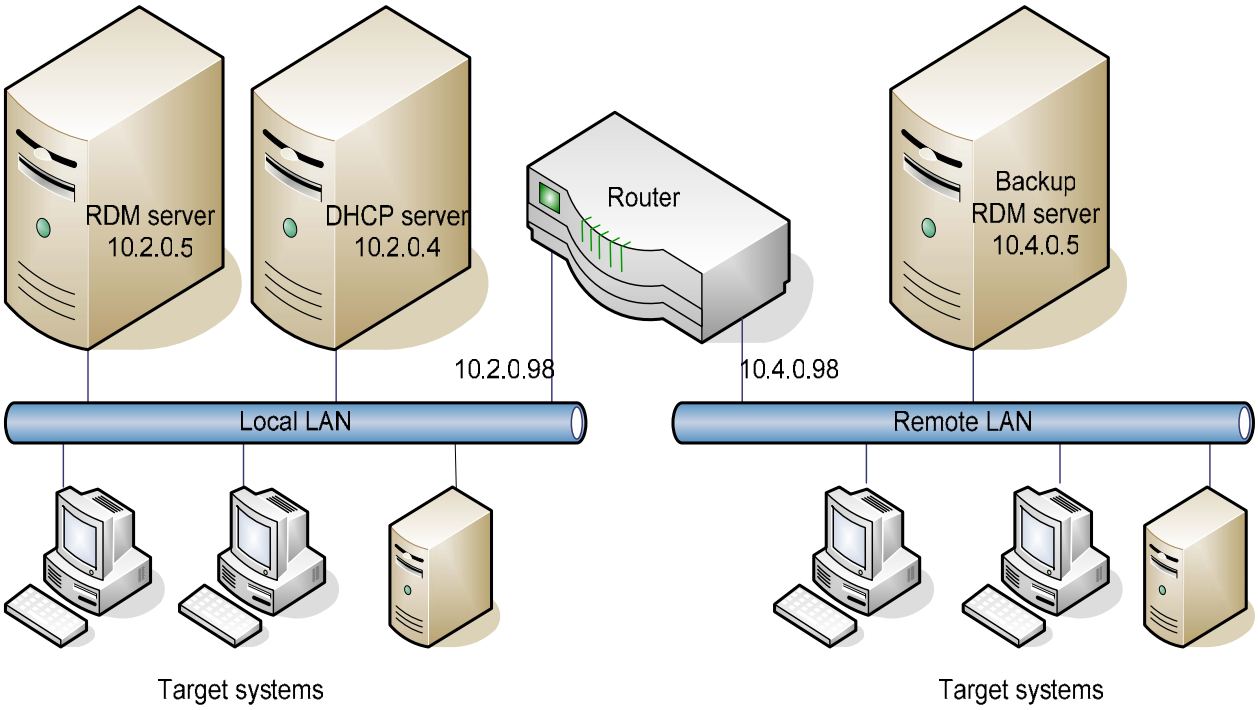

The test environment used for this document is an isolated Ethernet network. It contains a single Ethernet 10/100 switch, configured as follows:

- IP address 10.2.0.99 with 20-bit subnet mask 255.255.240.0
- VLAN 1: ports 1 through 12
- VLAN 2: ports 13 through 24

There are 2 subnets (which are actually the VLANs configured on the switch). The following describes the subnets and the systems that are connected to each VLAN.

- IP address 10.2.0.0 with 20-bit subnet mask 255.255.240.0
  - Primary RDM server (IP address 10.2.0.5)
    - 2 disk drives, configured as RAID 1
    - IBM Director 4.20.2 installed, using a Microsoft Jet database
    - RDM 4.20 installed
  - DHCP server (IP address 10.2.0.4)
    - Address pool 10.2.0.121 through 10.2.0.160 (option 3 configured as 10.2.0.98)

        - Address pool 10.4.0.121 through 10.4.0.160 (option 3 configured as 10.4.0.98)
        - IBM Director 4.20 installed
        - RDM 4.11 with Update 3 installed
    - Several target systems
- IP address 10.4.0.0 with 20-bit subnet mask 255.255.240.0
    - Standby RDM server
        - Identical hardware as primary RDM server
        - Powered off
        - Not deployed (i.e., initially contains no operating system)
        - IP address 10.4.0.5 assigned
    - Several target systems

The 2 subnets are connected by a router:

- IP addresses 10.4.0.98 and 10.2.0.98
- DHCP Relay Agent with pointers to all DHCP and Proxy DHCP servers:
    - 10.2.0.4 (DHCP server)
    - 10.2.0.5 (RDM Master Deployment Server)
    - 10.4.0.5 (standby RDM Master Deployment Server)
- Allow subnet-directed broadcast messages to be forwarded
- Allow multicast (IGMP) messages to be forwarded
- Allow Proxy ARP messages to be forwarded

3 Procedure

The following is a summary of the procedure described in this section:

a. Build a test environment that includes a primary RDM server.
b. Capture an image of the primary RDM server (without running Microsoft SYSPREP.EXE).
c. Deploy an image of the primary RDM server to the backup RDM server.
d. Change the test environment to a production-like state.
e. Configure the backup RDM server as appropriate for its subnet.
f. Test.

3.1 Build the test environment

How to accomplish the steps in this section depend on the particular hardware and software you use. We describe the steps generically, based on section 2 above, leaving the user to define and implement the actual details.

1. Configure the 2 VLANs on the Ethernet switch.
2. Cable all of the systems (power cables, Ethernet cables, etc.).
3. Install and configure the router.
4. Install and configure the DHCP server. Initially, configure *Option 3* on each DHCP scope.
5. Install and configure the primary RDM server.
6. Begin using the primary RDM server as your production deployment tool.

3.2 Capture an image of the primary RDM server

There are many tools for backing up a system. We chose to use RDM to do this, because it was convenient in our lab environment. In a production environment, this would be an inefficient way to do your backup. You should, instead, use your standard backup tool to do this.

Note: RDM is not a backup/recovery tool.

7. Shut down your primary RDM server.
8. Install another RDM server on the DHCP server. We will use this RDM (in step 11 below) to capture an image of the primary RDM server.
9. Add *Option 60* to each DHCP scope, and restart the DHCP service. We had to do this because the RDM Master Deployment Server and the DHCP server are installed on the same computer. If we had installed RDM on a different computer, we would skip this step.
10. Discover the primary RDM server (that you built in step 5 above). For example, you could power it on and press F12 to force a network boot. Then RDM will run its Basic Scan task on that system. Make sure that you press F12 to force any subsequent boot to do a network boot, too.
11. Run the *Get Donor* task on the primary RDM server. That is, the primary RDM server is the target system for this task. During System/Task Configuration, choose *Reseal* as the Sysprep method (even though you have not run SYSPREP.EXE on that system).

    Important: Do not run Microsoft SYSPREP.EXE on the primary RDM server.

The result of this step is an RDM image of the primary RDM server's drive.

3.3 Deploy an image of the primary RDM server

12. Using the RDM server that you built in step 8 above, create a *Windows Clone Install* task that uses the image that you created in step 11 above. Call the task *Put Donor*. It does not matter how you answer the other wizard questions, because we will modify this task so that it uses none of that data.
13. Edit the Put Donor task, and modify the command list as in this picture. Notice that we merely deleted all of the commands related to SYSPREP.EXE and mini setup.

    Important: If your image is particularly large, you may have to increase the timeout value in the command list. The default value of 240 minutes is shown here.

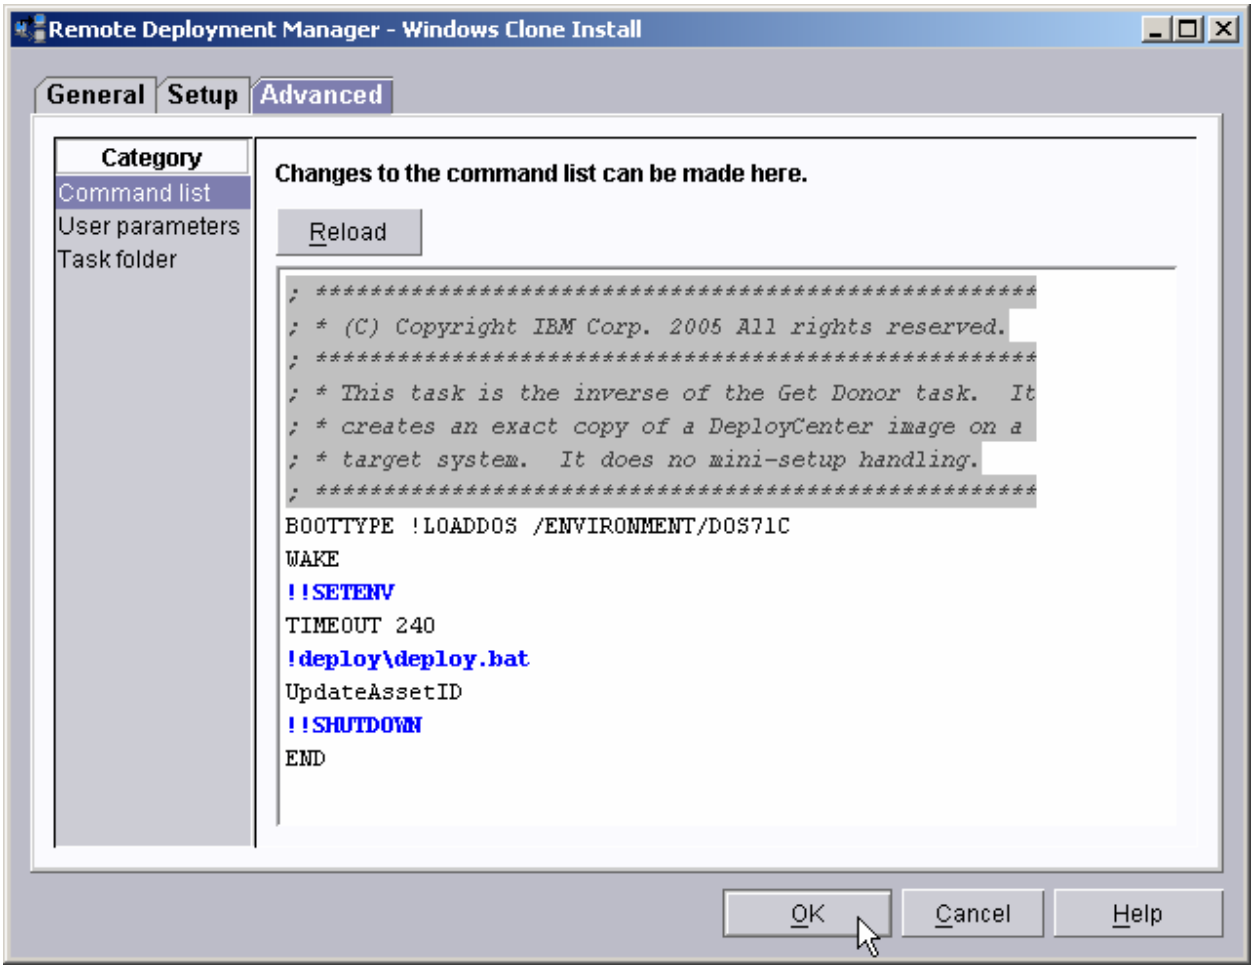

14. Discover the standby RDM server. For example, you could power it on and press F12 to force a network boot. Then RDM will run its Basic Scan task on that system. Make sure that you press F12 to force any subsequent boot to do a network boot, too.
15. If the standby RDM server uses RAID, run an appropriate RDM task (e.g., *Express RAID Configuration*) to configure the RAID.
16. Run the *Put Donor* task on the standby RDM server. That is, the standby RDM server is the target system for this task.

The result of this step is that the standby RDM server's drive looks exactly like the primary RDM server's drive.

3.4 Return the environment to production state

17. Stop the RDM services on your RDM server (the one installed on the DHCP server in step 8 above).
    - RDM 4.20: Stop the *IBM RDM D-Server* service.
    - RDM 4.11: Stop the *IBM RDM D-Server Service, IBM MTFTP Service,* and the *IBM PXE Service*.

    We cannot just shut down this server, because we need its DHCP server to be up and running.
18. Remove *Option 60* from each DHCP scope, and restart the DHCP service. We had to do this because the backup RDM Master Deployment Server and the DHCP server are installed on different computers. In a typical production environment, your DHCP server would be on a different computer, and you would not have *Option 60* configured at all.

The result of this step is that there is no RDM server running on the network and the DHCP server is returned to its production configuration.

3.5 Reconfigure the backup RDM server

19. Power on the backup RDM server. You should get many "New hardware found" messages, followed by the following *System Settings Change* message. After allowing your IBM Director server to completely start (i.e., the green circle appears in the system tray), select the Yes button on the message to restart the backup RDM server.

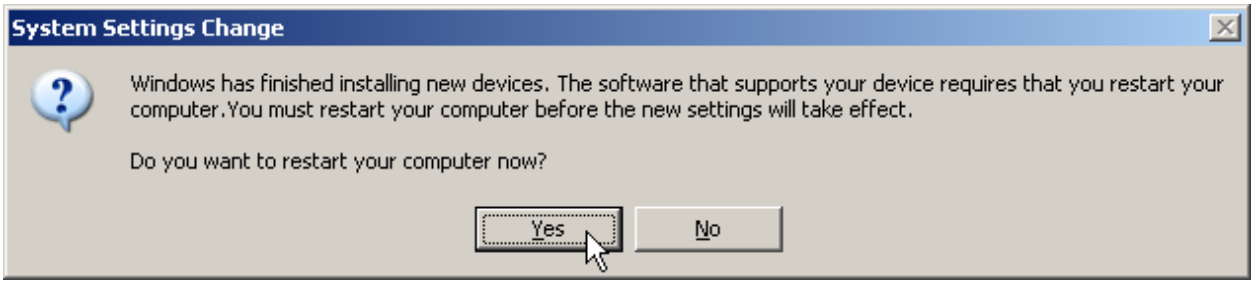

20. Change the backup RDM server's IP address from DHCP to the correct static IP address (10.4.0.5).

    Because the original static IP address is invalid on the other subnet, Windows changed the network configuration from using static to using DHCP.
21. Change the backup RDM server's computer name, as appropriate. (We changed RDMSERVER1 to RDMSERVER3.)
22. Reboot the backup RDM server, and wait until your IBM Director server completely starts (i.e., the green circle appears in the system tray).
23. Stop IBM Director's and RDM's services with the following commands:
    a. NET STOP TWGIPC
    b. NET STOP DSERVER

       You will probably get a message that says this service is not started. Just ignore that message.

24. Edit the C:\Program Files\IBM\RDM\local\DSERVER.INI file.

    a. Change both occurrences of the old IP address to the new one (10.2.0.5 to 10.4.0.5).
    b. Change the old hostname (RDMSERVER1) to the new computer name (RDMSERVER3).

25. Change the IBM Director configuration:

    a. TWGIPCCF
    b. Change the old system name to the new computer name (RDMSERVER3).
    c. Check the *Enable driver* box.

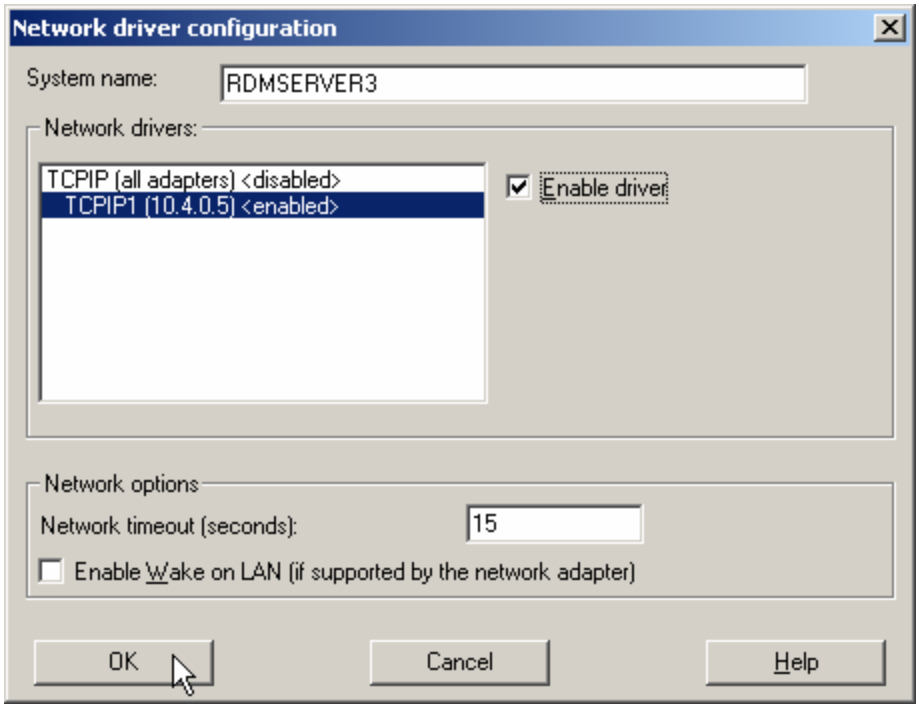

26. Start IBM Director's and RDM's services with the following commands:

    a. NET START DSERVER
    b. NET START TWGIPC

27. Open the IBM Director console. You will probably have to update the login information with the new computer name.

28. If the console contains an (off-line) IBM Director native managed object for the primary RDM server, it is an orphan. The easiest procedure is to delete that object now, and then to do a Director discovery; alternatively, you could just rename that object (we renamed it from RDMSERVER1 to RDMSERVER3).

The result of this step is that the backup RDM server will now contain exactly the same configuration and data that the primary RDM server had.

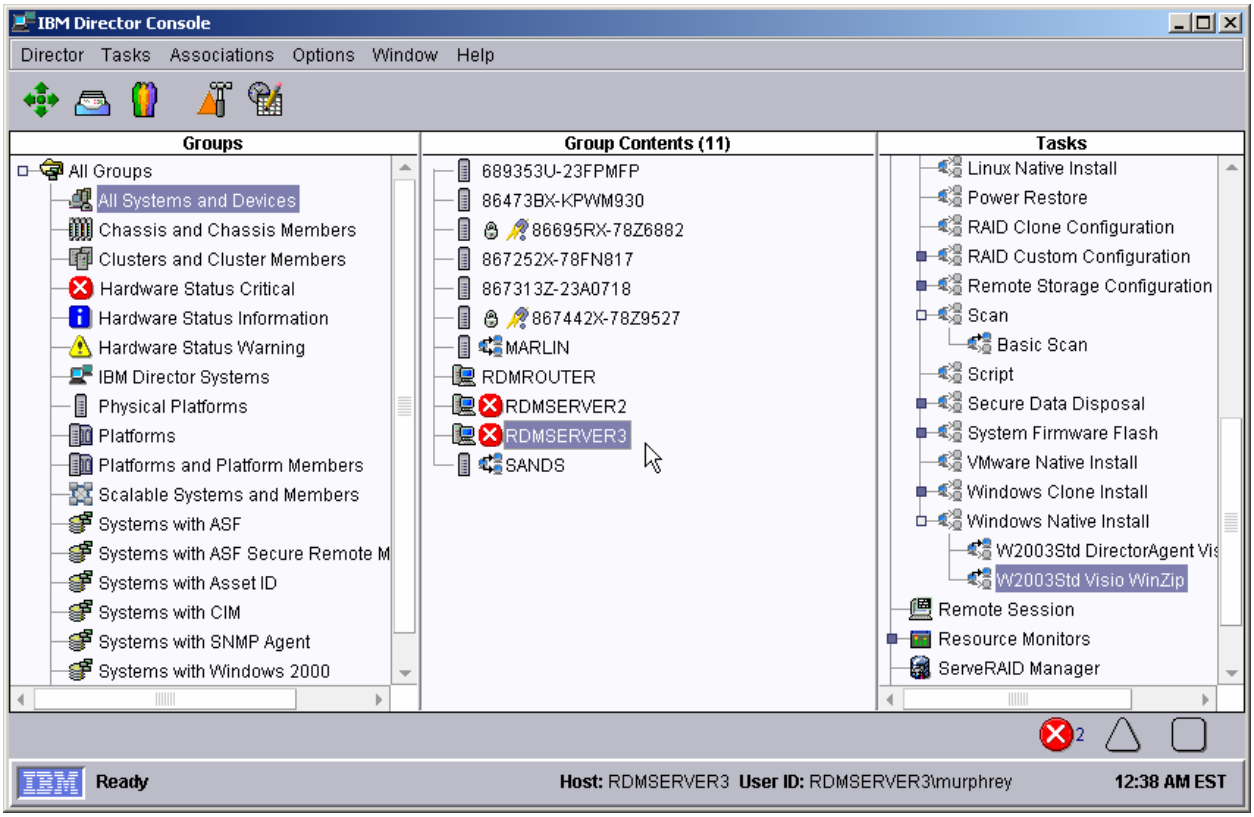

3.6 Test the backup RDM server

29. You should run a few RDM tests to confirm that the backup RDM server is working properly. For example, run a *Basic Scan* task on target systems on both subnets.

The result of this step is that the backup RDM server will now be in production use. It will behave exactly as the primary RMD server did.

4 Modifications

The procedure described in section 3 above is an example. You may choose to do it differently.

4.1 Disk cloning

The basic premise of that procedure is that the backup RDM server becomes an exact copy of your primary RDM server. The steps above used a disk cloning tool (the Symantec™ DeployCenter (light version) tool that is part of RDM) to do this. You can accomplish the same thing in other ways:

- Use a different cloning tool with RDM. For example, you could use DeployCenter 5.6 (full version) or Symantec Ghost™ 8.2. See section 1.2 above to find White Papers that describe how to do this.
- Use a backup/recovery or cloning tool outside RDM. Consult your tool's documentation for instructions on how to do this.
- Move the disk from the primary RDM server to the backup RDM server.
- If the disk is a Fibre (or iSCSI) LUN, map it to the backup RDM server.

4.2 Automation

It may be possible to automate some of the procedures described herein.

Reconfiguring the backup RDM server

You could use a batch file like the following to accomplish the reconfiguration described in section 3.5 above:

```
NET STOP TWGIPC
NET STOP DSERVER
DEDITD /r /n0 DSERVER.INI 10.2.0.5 10.4.0.5
DEDITD /r /n0 DSERVER.INI RDMSERVER1 RDMSERVER3
START /WAIT TWGIPCCF
NET START DSERVER
NET START TWGIPC
```

See the RDM 4.20 User's Reference for the DEDITD.EXE syntax. This example replaces (/r parameter) all occurrences (/n0 parameter) of 10.2.0.5 with 10.4.0.5.

4.3 Using remote D-servers

The purpose for remote D-servers is to have a local server from which RDM target systems obtain all of the medium or large files. The idea is that you replicate the files from the master D-server, one time, to each remote D-server; and then you use the files many times from the remote D-server that is local to the target systems. This keeps the RDM network traffic over the Wide Area Network to a minimum, restricting all of the heavy traffic to the Local Area Network.

The example given in sections 2 and 3 above uses no remote D-servers. However, adding a remote D-server to each subnet would require only a slight modification of the procedure. Essentially, you would have to modify each D-server's DSERVER.INI file in a similar way to that outlined in step 23 through 25 in section 3.5 above.

4.4 Redundant remote D-servers

It is also possible to have redundant remote D-servers. For example, you can just install 2 RDM D-servers on the same subnet. As long as you ensure that any file that gets replicated to one of them always gets replicated to the other, the 2 D-servers will always contain the same data (i.e., they will be synchronized) and all RDM tasks will work properly.

Typically, it is an error to have multiple PXE servers serving a particular subnet. This is because there is no control over which PXE server responds to any particular request. So, for example, you cannot run RDM and Microsoft RIS on the same subnet.

Although each RDM D-server contains its own (internal) PXE service, this redundancy does not generate any errors. That's because regardless of which D-server responds to a PXE request, it gives the correct response.

5 Notices

This information was developed for products and services offered in the U.S.A.

IBM might not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service might be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right might be used instead. However, it is the user responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM might have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM might make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM might use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Some software might differ from its retail version (if available) and might not include all user manuals or all program functionality.

IBM makes no representations or warranties regarding third-party products or services.

5.1 Edition notice

5.2 Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM

IBM (logo)

Asset ID

IntelliStation

LANClient Control Manager

Netfinity

ServeRAID

ThinkPad

Wake on LAN

xSeries

Adaptec is a trademark of Adaptec Inc. in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

6 Glossary

BAT file. A file that contains a batch program (that is, a set of commands).

bind. Associating one or more systems to a task. This causes all information to be verified (by one of the STC modules) and a resulting job to be scheduled to run.

console, or RDM Console. The group of programs that make up the user interface to RDM. RDM is client/server in nature so that the Console might run on any computer and not necessarily be running on the same computer as the RDM server or other RDM components. The RDM Console is actually an IBM Director Console on which the RDM Console component is installed.

image. An image is the software stored on a deployment server that is downloaded to a system during an operation. Images vary in size and in the type of software they provide to the system. The purpose and content of each image depends on the task to be accomplished, as well as the method used to download the image from the deployment server to the system. A *native* image is built off a product installation CD. A *clone* image is copied from a donor system.

job. An object managed by the scheduler and created by STC. A job is a binding of one task and one or more systems. A job can be scheduled to run once or to recur. Sometimes a job is called by a different name (Scheduled Task, Running Task), to emphasize some aspect of the job.

managed system. The IBM Director term for its system. Mentioned here only for clarity; the term *system* is preferred when referring to an RDM system.

preboot DOS agent. The preboot DOS agent is a DOS operating system with a communications stack that is booted from the network by the bootstrap agent. The preboot DOS agent performs actions on a system as directed by the RDM server.

Preboot Execution Environment (PXE). PXE is an industry standard client/server interface that allows networked computers that are not yet loaded with an operating system to be configured and booted remotely. PXE is based on Dynamic Host Configuration Protocol (DHCP). Using the PXE protocol, clients can request configuration parameter values and startable images from the server.

The PXE process consists of the system initiating the protocol by broadcasting a DHCPREQUEST containing an extension that identifies the request as coming from a client that uses PXE. The server sends the client a list of boot servers that contain the operating systems available. The client then selects and discovers a boot server and receives the name of the executable file on the chosen boot server. The client downloads the file using Trivial File Transfer Protocol (TFTP) and executes it, which loads the operating system.

Redundant Array of Independent Disks (RAID). RAID is way of storing the same data in different places (thus, redundantly) on multiple hard disks. By placing data on multiple disks, I/O operations can overlap in a balanced way, improving performance. Multiple disks increase the mean time between failure (MTBF) and storing data redundantly increases fault-tolerance.

system. An individual, target system being deployed or managed by RDM. In IBM Director terminology, an RDM system is always a platform managed object. These can represent any of the supported-by-RDM systems. They cannot represent an IBM Director object that RDM does not process, such as a chassis or an SNMP object.

task. An already defined and configured unit of work that is available to be applied to a system or a group (of systems). You create a task by clicking on the applicable task template from the RDM main console. RDM is installed with predefined tasks, such as data disposal and scan.

task template. A prototype of a specific kind of RDM task. This is a term used to describe the different kinds of tasks shown on the task pane in the main window of the RDM console. Each task template has its own characteristics and attributes. RDM comes with a set of task templates.

Wake on LAN. Technology developed by IBM that allows LAN administrators to remotely power up systems. The following components are essential for the Wake on LAN setup:

- Wake on LAN-enabled network interface card (NIC).
- Power supply that is Wake on LAN-enabled.
- Cable which connects NIC and power supply.
- Software that can send a magic packet to the system.

If the system has the first three of the previous components, the system is called a Wake on LAN-enabled system. Even though a system might be powered off, the NIC keeps receiving power from the system power supply to keep it alive. A network administrator sends a magic packet to the system through some software, for example, RDM or Netfinity IBM Director. The NIC on the system detects the magic packet and sends a signal to the power supply to turn it on. This process is also called *waking up the system*. Using RDM, this process can be scheduled for individual systems. The Wake on LAN feature and RDM together make it very easy for you to deploy software on individual systems on a scheduled basis.

Printed in U.S.A.