



Protocol Specification Summary

Document Version: ***1.7.1 S***

Last Update: ***08/26/96***

1. OVERVIEW.....	3
1.1 MOTIVATION	3
1.2 AUDIENCE	3
1.3 TERMINOLOGY	3
2. ARCHITECTURE.....	4
2.1 ARCHITECTURAL MODEL.....	4
2.2 COMPARISON ARCHITECTURES.....	6
2.3 STATE DESCRIPTION	7
2.4 INTERFACES.....	8
2.5 MESSAGES FROM BACKUP TO THE NDMP SERVER	8
2.6 MESSAGES FROM NDMP SERVER TO BACKUP	8
2.7 MESSAGING PROTOCOL.....	9
2.8 MESSAGE DEFINITIONS	9
2.9 HEADER.....	9
2.10 ERROR.....	9
3. INTERFACES.....	10
3.1 CONNECT INTERFACE	10
3.2 CONFIG INTERFACE.....	10
3.3 SCSI INTERFACE	10
3.4 TAPE INTERFACE.....	10
3.5 DATA INTERFACE	10
4. BACKUP INTERFACES.....	11
4.1 NOTIFY INTERFACE.....	11
4.2 LOGGING INTERFACE.....	11
4.3 FILE HISTORY INTERFACE.....	11
4.4 REMOTE INTERFACE	11
 FIGURE 1. SIMPLE CONFIGURATION	 4
FIGURE 2. TWO DRIVE CONFIGURATION.....	5
FIGURE 3. JUKEBOX CONFIGURATION	5
FIGURE 4. BACKING UP TO BACKUP HOST.....	6
FIGURE 5. BACKUP STATE DIAGRAM.....	7



1. Overview

1.1 Motivation

The purpose of this protocol is to allow a backup of a host without requiring a full port of the backup software product. The backup and restore solution is partitioned in a way that minimizes the amount of software required on the host with the tape drive attached.

The system vendors need only be concerned with maintaining compatibility with one, well-defined protocol. The backup vendors can place their primary focus on the sophisticated central backup administration software.

This protocol is specifically intended to support tape drives. This protocol is targeted towards backup software and there are extensive references to the tasks of backup and restore. However, the protocol may be used for other applications in the future.

1.2 Audience

This document is intended for use by software developers considering implementation of Network Data Management Protocol. The reader is assumed to be familiar with network protocol specifications and with the general operation of backup software. The user is not expected to have knowledge of internal backup software behavior.

1.3 Terminology

Backup software host

The host on which the backup software, daemons and the databases exists. The backup software host may or may not have a tape drive attached

NDMP host

The host which has a tape drive physically attached and can perform local backups to that tape drive using the NDMP.

NDMP server

The virtual state machine on the NDMP host that is controlled using the NDMP protocol. There is one of these for each connection to the NDMP host. This term is used independent of implementation.

2. Architecture

2.1 Architectural Model

The architecture is a client server model and the backup software is considered a client to the NDMP server. For every connection between the client on the backup software host and the NDMP host there is a virtual state machine on the NDMP host that is controlled using the NDMP. This virtual state machine is referred to as the NDMP server. Each state machine controls at most one device used to perform backups. The protocol is a set of XDR encoded messages that are exchanged over a bi-directional TCP/IP connection and are used to control and monitor the state of the NDMP server and to collect detail information about the data that is backed up.

In the most simple configuration, the backup software will back up the data from the NDMP host to a tape drive connected to the NDMP host.

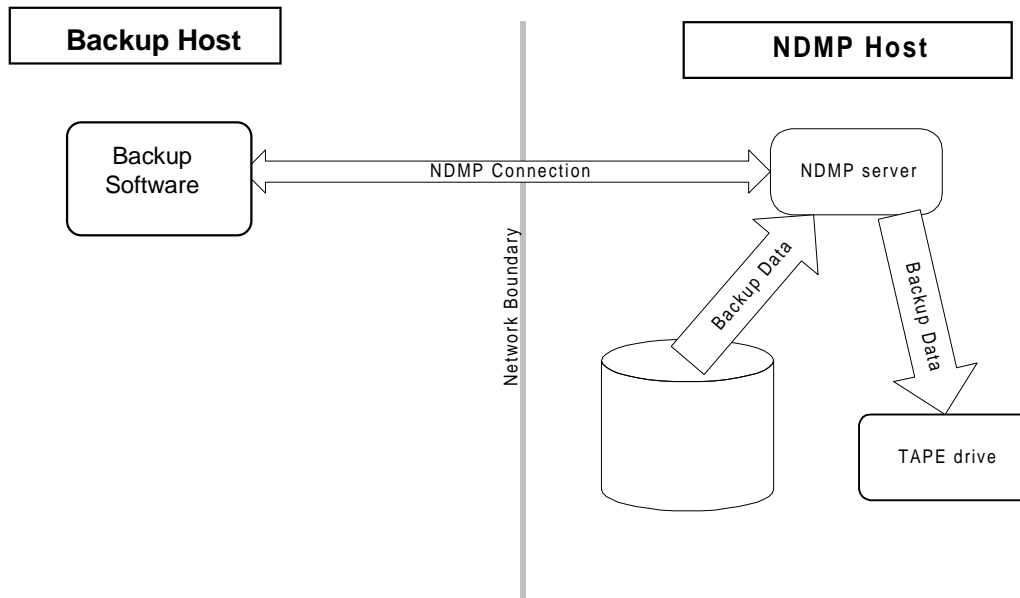


Figure 1. Simple configuration

It is also possible to use the NDMP to simultaneously backup to two tape drives physically attached to the NDMP host. In this configuration there are two instances of the NDMP server on the NDMP host.

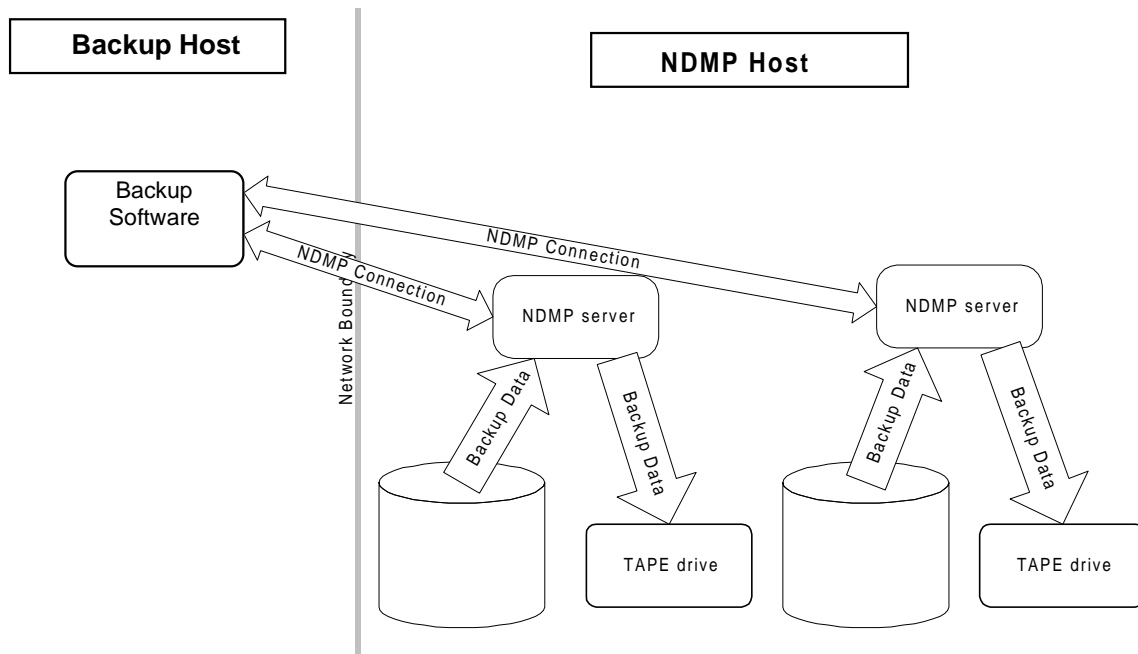


Figure 2. Two drive configuration

The NDMP can be used to backup data to a tape drive in a jukebox that is physically attached to the NDMP host. In this configuration, there is a separate instance of the NDMP server to control the robotic arm in the jukebox.

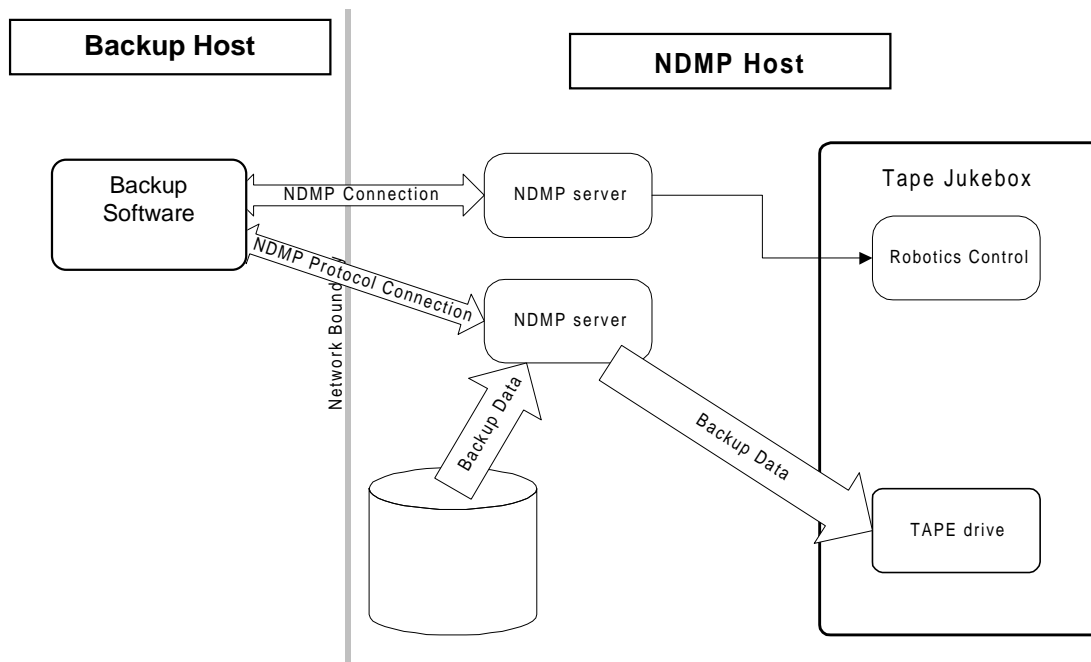


Figure 3. Jukebox configuration

As a special case it is possible to backup a host that supports the NDMP but does not have a locally attached tape drive by sending the data to the backup software host as messages.

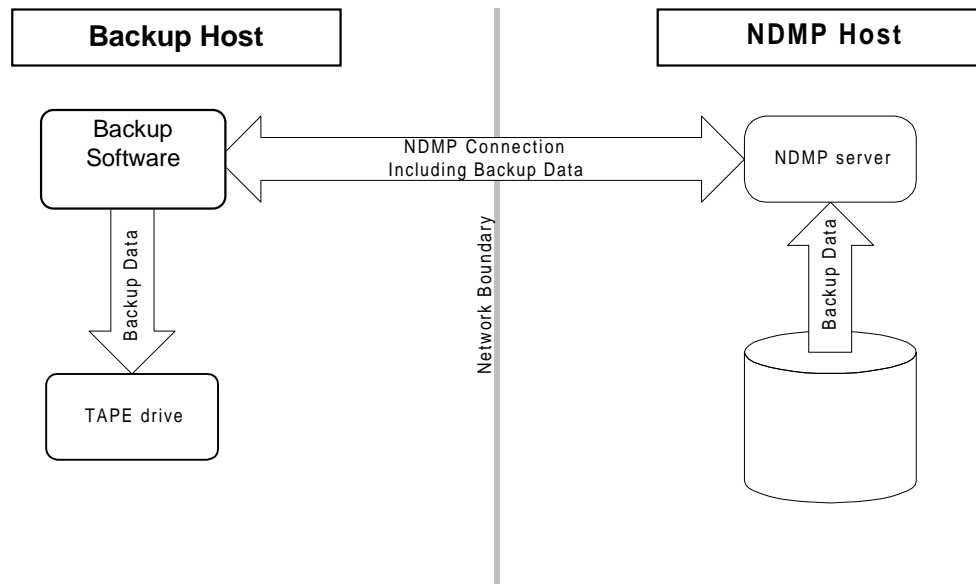


Figure 4. Backing up to backup software host

2.2 Comparison Architectures

It is useful to compare the NDMP architecture to other architectures and note the similarities and differences.

rmt

The architecture is similar to the rmt architecture in that connection is made to a generic server and the server is instructed to open a specific tape drive device. The NDMP differs in that it uses a TCP/IP connection to a dedicated port whereas rmt uses the rsh demon to launch a server.

X11

The architecture is similar to the X11 architecture in that it uses a single connection to a TCP/IP port, however it differs in that the NDMP server is not assigned to a device until the client opens a device and that there is only one client per NDMPserver, whereas X11 is assigned to a display device before the first client connects and accepts connections from many clients.

RPC

The NDMP architecture is similar to RPC in that it uses XDR encoding. NDMP differs in that it is only defined for a TCP/IP connection and that it is not a call-return model, but rather a bi-directional asynchronous messaging model.



2.3 State Description

The following defines the DATA state diagram

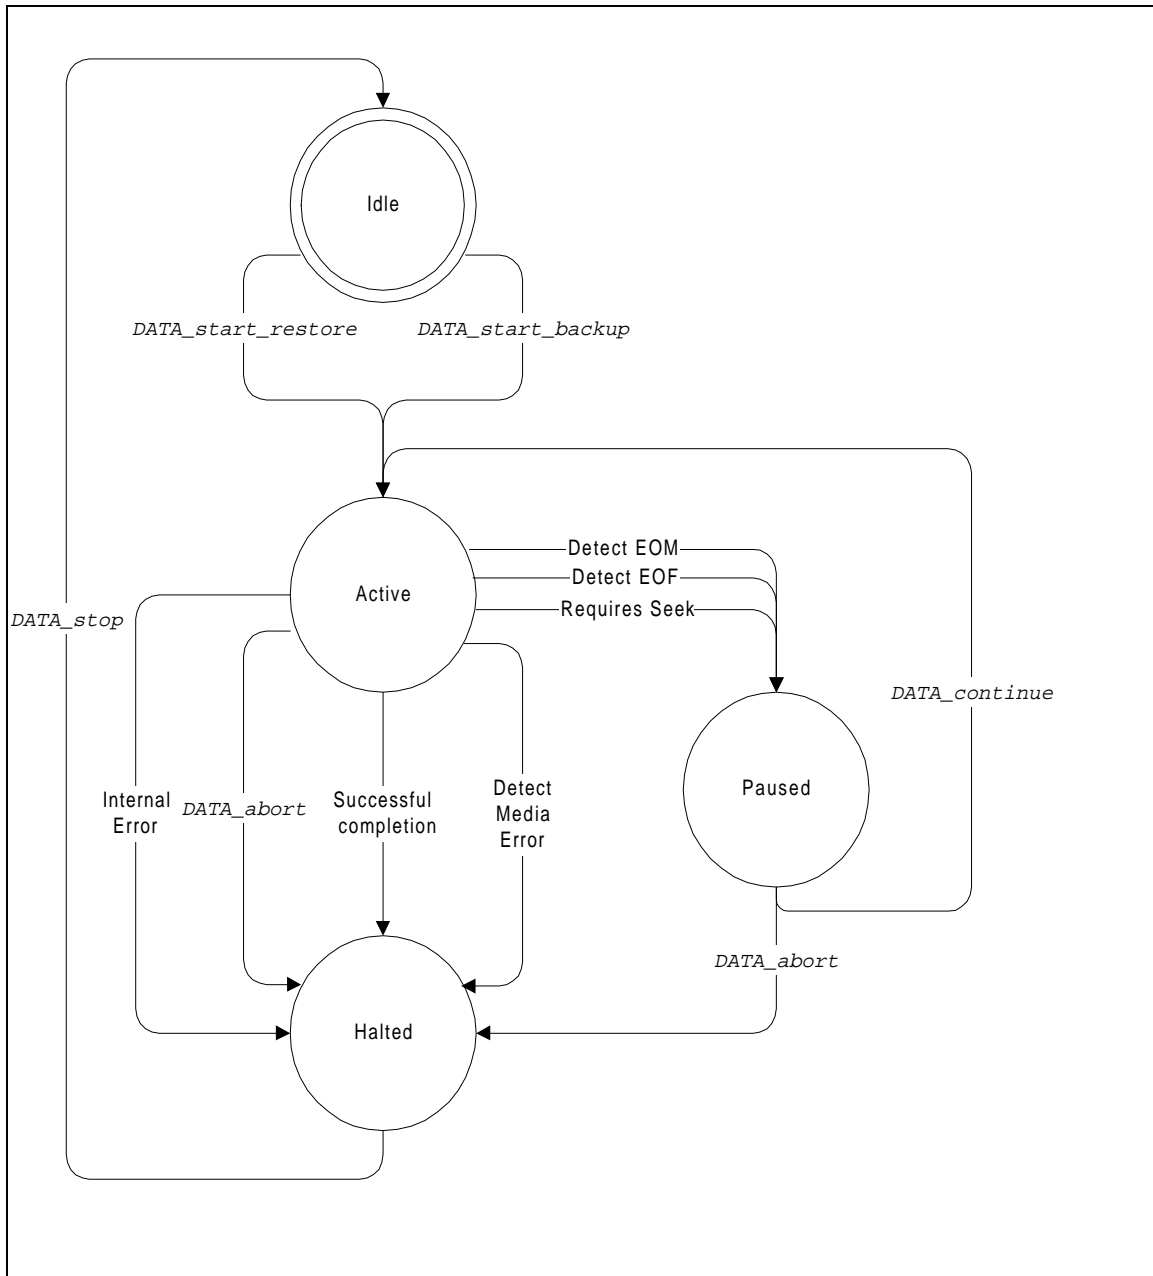


Figure 5 - Backup state diagram



2.4 Interfaces

Messages are group together by functionality into several “interfaces”.

2.5 Messages From Backup Software To The NDMP Server

The NDMP server must implement the following interfaces:

- CONNECT interface

This interface will be used when a client opens the communication to a NDMP server. This interface allows the NDMP server to authenticate the client and negotiate the version of protocol used.

- CONFIG interface

This interface allows backup software to discover the configuration of the NDMP server. It can be used to discover tape drives and jukeboxes as well as file systems and databases. Backup software will use this interface to build request and media server databases automatically.

- SCSI interface

This interface simply passes SCSI CDBs through to the SCSI device and returns the SCSI status. The backup software will use this interface to control a locally attached jukebox. Software on the backup software host will construct SCSI CDBs and will interpret the returned status and data. This interface can also be used to exploit special features of SCSI tape drives.

- TAPE interface

The TAPE interface will support both tape positioning and tape read/write operations. The backup software will use this interface to control the labeling and format of the tape. The backup software will also use this interface for positioning of the tape during backups and restores.

- DATA interface

This is the interface that actually deals with the format of the backup data. The backup software will initiate backups and restores using this interface. The backup software provides all of the parameters that may affect the backup or restore using this interface. The backup software does not place any constraints on the format of the backup data other than it must be a stream of data that can be written to the tape device.

2.6 Messages From NDMP Server To The Backup Software

The NDMP server’s implementation may send the following messages to the backup software host. All of the messages that the backup software host accepts are asynchronous. None of these messages will generate a reply message.

- NOTIFY interface

This message is used by the NDMP server to notify the backup software that the NDMP server requires attention.

- FILE HISTORY interface

These messages allows the NDMP server to make entries in the file history for the current backup. The File History will be used by the backup software to select files for retrieval.

- LOGGING interface

These messages allows the NDMP server to make entries in the backup log. This is used by the operator to monitor the progress and successful completion of the backup. It is also used to diagnose problems.

- REMOTE interface

This interface is used to support backing up and restoring the NDMP host using a tape drive on the backup software host.

2.7 Messaging Protocol

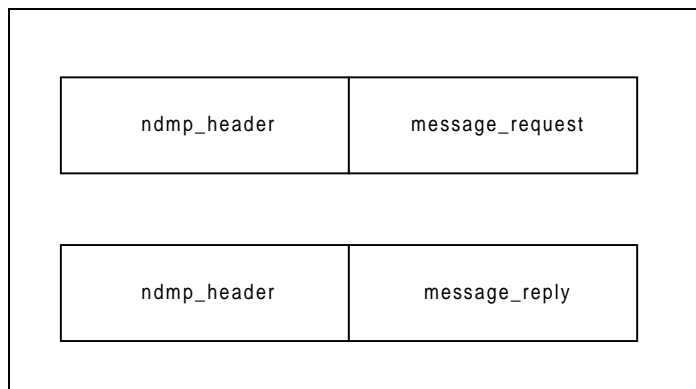
The NDMP uses asynchronous messages and a message does not require a reply, however, many of the messages may result in a reply message.

2.8 Message Definitions

Each message is described using a block of XDR specification. The body is a textual description of what the message is intended to do. If an error status is returned, then some of the reply arguments may be meaningless. If the request has no arguments, then the request structure is omitted. If the request does not expect a reply message, then the reply structure is omitted.

2.9 Header

Each message will be preceded by a message header. The header will be used to identify the message, how to de-serialize the arguments and to dispatch the message.



2.10 Error

Errors are defined.

3. Interfaces

This section will define the interfaces.

3.1 *CONNECT Interface*

This interface allows NDMP server to authenticate the client and negotiate the version of protocol which will be used.

3.2 *CONFIG Interface*

This interface allows the backup software to discover the configuration of the NDMP server.

3.3 *SCSI Interface*

The SCSI interface allows low level control of SCSI devices such as jukeboxes.

3.4 *TAPE Interface*

Provide complete control of a tape drive. If the tape drive is a SCSI tape drive, then this interface also provides low level CDB access to the tape drive. This interface is analogous to the rmt protocol. The physical device is assigned when the server is started.

3.5 *Data Interface*

Selects and formats data for backup. Extracts files from backup for retrieval.



4. Backup Software Interfaces

These interfaces are available to the NDMP server. Only the requests that are intended for use by the NDMP server are available. Many of the interfaces require initialization which will be performed by the backup software before the NDMP server is called.

4.1 NOTIFY Interface

This interface is used by the NDMP server to let the backup software know that the NDMP server requires attention.

4.2 LOGGING Interface

Make an entry into a backup specific log.

4.3 FILE HISTORY Interface

The backup software maintains a file by file record of every file that is backed up. These records are called the file history. This interface provides the ability to add entries to the file history. Backup software currently only maintains file history in a format that is compatible with a UNIX file system, however the interface is defined in an extensible manner.

4.4 REMOTE Interface

These messages support backing up and restoring the NDMP host using a tape drive on the backup software host.