



# HyperTape

from BridgeHead Software

## Overview White Paper

---

Document M8015 Version 3.1 rev.3 April 2002

<b>Document Reference Version</b>	<b>HyperTape Overview White Paper M8015 3.1 rev.3 April 2002</b>
-----------------------------------	--

The information in this document is subject to change without notice.

BridgeHead Software makes no warranty of any kind with regard to this material, including, but not limited to the implied warranties of merchantability, performance, or fitness for a particular purpose. BridgeHead Software shall not be liable for errors contained herein or for any direct, incidental, or consequential damages resulting from the use of this material.

© Copyright 2002 BridgeHead Systems Ltd. and its licensors. All rights reserved.

You may view and print this publication solely for personal, informational, internal, non-commercial purposes. You must not change it in any way or remove any copyright or other proprietary notices. You may not reproduce, distribute or translate any part of the publication without the written permission of BridgeHead Software. These rights constitute a license to use and not a transfer of title.

All brand names, trademarks and registered trademarks are the properties of their respective owners.

## Contents

Contents .....	2
What is HyperTape? .....	3
System architecture.....	3
HyperTape files .....	7
Performing backup operations.....	8
Automating backups .....	9
Customizing backups.....	9
Media management.....	9
Further reading .....	9
For more information.....	10

### For information on any BridgeHead products contact any of our offices ...

<b>BridgeHead / UK</b>	<b>BridgeHead / USA</b>	<b>BridgeHead / Germany</b>
Bailey House 215 Barnett Wood Lane Ashted Surrey KT21 2DF UK	400 West Cummings Park Suite 6600 Woburn MA 01801 USA	Spanierstrasse 69 76879 Essingen Germany
Tel: +44 (0)1372 221950	Tel: (001) 781 939 0780	Tel: +49 (0)700 38400000
Fax: +44 (0)1372 221977	Fax: (001) 781 939 5607	Fax: +49 (0)700 38400001
<a href="mailto:Sales.UK@BridgeHeadSoftware.com">Sales.UK@BridgeHeadSoftware.com</a>	<a href="mailto:Sales.US@BridgeHeadSoftware.com">Sales.US@BridgeHeadSoftware.com</a>	<a href="mailto:Sales.DE@BridgeHeadSoftware.com">Sales.DE@BridgeHeadSoftware.com</a>

[www.BridgeHeadSoftware.com](http://www.BridgeHeadSoftware.com)

## What is HyperTape?

HyperTape is a leading storage management solution that automates the backing up, archiving and restoring of critical data, databases and applications across networks. It has a cross-platform, multi-protocol, highly scalable architecture, which conforms to industry standards and supports centralized, distributed and mixed backup technologies.

Because it uses your network's existing hardware and software components, for example by providing interfaces to standard backup and recovery utilities, HyperTape protects your investment. It lets you use existing storage more effectively and helps you adapt easily to organizational and operational changes.

HyperTape allows unattended backup of PCs, workstations, departmental servers and enterprise servers across networks and supports a variety of magnetic storage media. Its user interfaces enable you to monitor your backup and restore operations at a glance. If an error does occur, detailed log files are on hand to help you determine the cause of any problem.

With the right operating system and enough bandwidth to send data to the mainframe or server, HyperTape automates the management of LAN/WAN-based data. It is primarily designed for high-speed backup, across a network, to automated, high-capacity peripheral systems controlled by a server or mainframe backup station. As such, HyperTape handles an unlimited number of concurrent sessions between any number of network nodes, while, at the same time, its security and access procedures protect the integrity of your data.

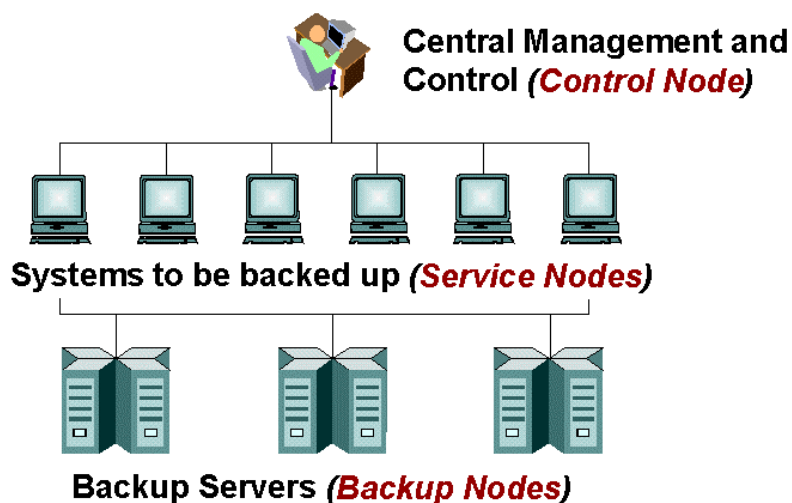
OpenMedia is HyperTape's multi-platform enterprise media manager. Its facilities include choice of tape and device, offsite (vaulting) policies, tape recycling and retirement, and automated drive cleaning. In a Storage Area Network (SAN), OpenMedia allows tape drive sharing between different hosts and operating systems. Its design makes this process simple to achieve.

HyperTape Robot Manager completes the family, providing control of automated tape libraries and automatic integration with OpenMedia.

## System architecture

HyperTape uses a five-layer model. The first three layers relate to the backup operation:


- The **Control Node** is the central point from which HyperTape manages backup operations
- The **Service Nodes** are individual systems that hold live data
- The **Backup Nodes** are the systems that migrate backups to secondary storage



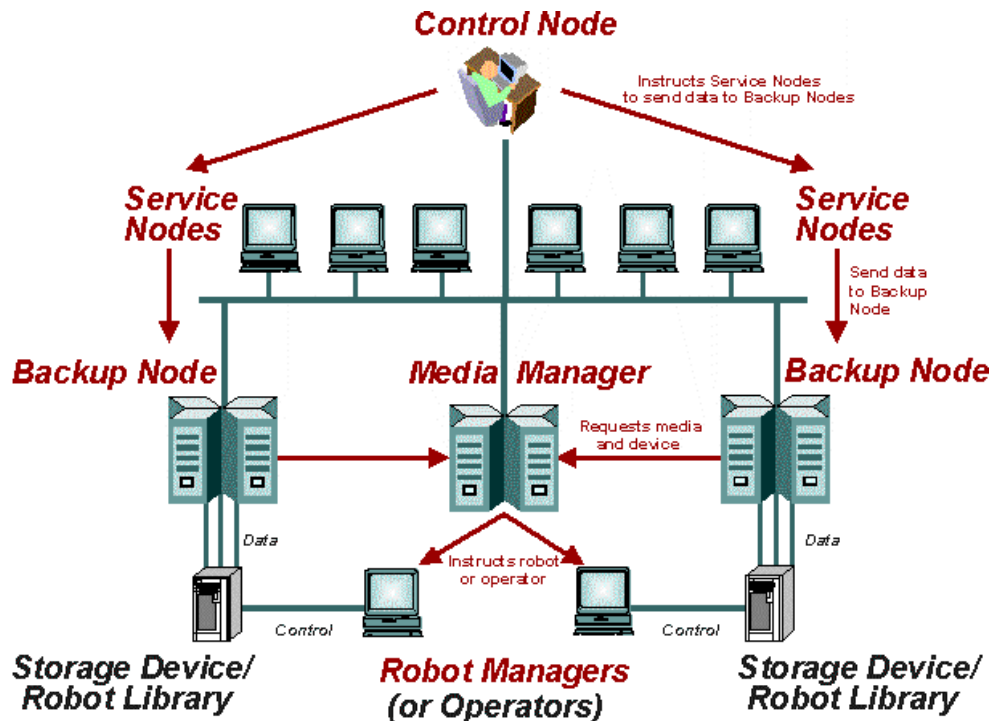
Then there are two optional layers:

- Media management
- Robotic storage management

**Control Node**



The following diagram shows how HyperTape controls the flow of data from the Service Node to the backup media.



## HyperTape nodes

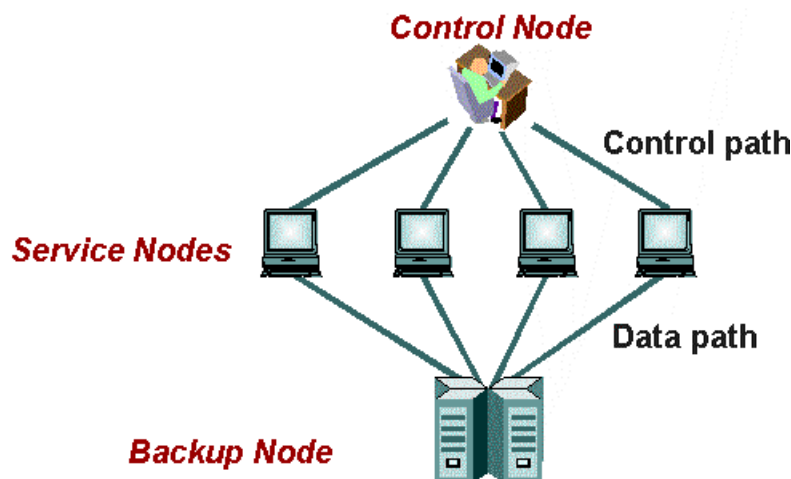
The **Control Node** initiates unattended backups of data on the Service Nodes. It accesses its central *Object Database* to find out what to back up, when to do it and where to store the data. The Control Node performs all the management functions related to backup and recovery operations.

The **Service Node** performs the actual backup operation, using a standard backup utility where available. It moves the backup saveset to the designated Backup Node using standard File Transfer Protocol (FTP) or its own *Saveset Transfer Utility (STU)*, which is based on FTP enhanced to include media management and direct to tape storage.

The **Backup Node** carries out any tape management functions necessary to save the data to secondary storage, usually with the help of a media manager.

Once the backup operation is completed, the Service Node notifies the Control Node of the completion status, so that it can update its databases as appropriate.

An important advantage of HyperTape is the ability to segregate control and management functions in a separate server. The Control Node uses the *control path* between itself and its Service Nodes to initiate backup operations. A separate *data path* transfers the backup between the Service Node and Backup Node, and on to the attached storage device.



Of course, one system may be any combination of Nodes. For example, in a SAN environment, many systems will be both Service Nodes **and** Backup Nodes.

The three types of node work together to ensure reliable network data transfer. You can choose from a number of supported network protocols, including TCP/IP and DECnet, to establish sessions between the nodes. You can select one protocol for the control path and another one, geared toward higher throughput, for the data path.

HyperTape fully supports the most common backup-related tasks via three *backup operations*:

- **SAVE** - saves backup objects. It creates a saveset containing backups of designated file systems or databases.
- **RESTORE** - restores a selected file or system from a remote saveset.
- **LIST** - lists the date and time a saveset was created and the names of the files it contains. You can list from the saveset itself, or from a journal kept on the Service Node.

HyperTape supports full, incremental and periodic (monthly) backup types, which you can schedule in a variety of ways. You can also specify the number of monthly and full backups to be retained. For incremental backups, however, you specify the number of *backup generations* to be retained, to ensure that a full set of related incremental backups is kept for each full backup.

HyperTape also supports facilities for executing user-written routines on the Control Node and Service Node, both before and after a backup is performed, so you can customize the backups to suit your needs.

Information about all these options is held in the Object Database on the Control Node, which initiates all the backup operations and provides uniform control and administration through its easy-to-use interfaces.

## HyperTape OpenMedia

OpenMedia provides the media management layer of HyperTape. It has facilities for:

- Volume, saveset and device management, including onsite/offsite rotations
- Robotic or human operator control
- User and host administration
- System policy application

OpenMedia can manage multiple robots, allowing them to be shared between applications. It sends instructions to robots via an open interface to the HyperTape Robot Manager or a third-party robot manager, thus providing the interface between the Backup Node and the tape library.

The Backup Node uses OpenMedia to allocate tapes and devices, and to store information about where backups are saved. The parameters that OpenMedia uses to select volumes on the Backup Node are held in the Object Database on the Control Node and passed across to OpenMedia when the backup operation takes place.

OpenMedia is flexible and adaptable, which makes it suitable for use in a variety of organizations. It can stand alone as a media manager, work in concert with other HyperTape components, or integrate with other storage applications. It is designed to be easy to integrate with both robot managers and tape writing products such as backup or hierarchical storage manager (HSM) applications.

When integrated with a robot manager, OpenMedia automatically chooses suitable tapes and devices, then calls the robot manager to issue appropriate load and unload commands. It monitors and records the location and status of the robot's tapes and devices when the robot loads, unloads and relocates the tapes.

## HyperTape Robot Manager

The HyperTape Robot Manager provides software to control a storage robot or automated tape library. The Robot Manager is called by OpenMedia to load or unload tapes. It provides a Graphical User Interface (GUI), which allows you to:

- Load and unload tapes
- Move tapes from any source to any destination
- Load tapes using a bar code label or slot number
- View the contents and status of all robots

Because of the Robot Manager's client/server architecture, you can view and manage robots throughout the network from a central point. This allows you to deploy remote storage robots and so optimize your network backup throughput.

## HyperTape files

HyperTape maintains the following databases:

- Object Database on the Control Node – holds definitions of all your HyperTape backups
- Queue Database on the Control Node – holds information about your scheduled HyperTape backup jobs
- History Database on the Control Node – holds information about completed HyperTape save operations
- Journal Database on the Service Node – holds information about completed HyperTape save operations run on that node
- Log databases on the Control, Service and Backup Nodes – hold information, warning and error messages produced by HyperTape jobs

All platforms maintain similar databases, but the names and structure of individual files may differ according to the platform.

### Object Database (Control Node)

HyperTape helps you to automate your regular backups by maintaining a central database of backup definitions for individual files, directories, databases and disks. See the *HyperTape System Manager Guide* for a discussion of the concepts involved in designing a backup strategy.

The Object Database, which resides on the Control Node, holds your backup definitions as *backup objects*. It is a core component of the HyperTape system, as the backup objects specify what to back up, when to do it, and where to put the saved data. HyperTape's Object Manager is a User Interface that allows you to specify and manage your backup objects.

Each backup object contains all the information needed to carry out a save or restore operation. This includes the names of the Service Node and Backup Node to be used, identification of the data to be saved or restored, the backup utility that will save the data, and so on.

Depending on the Control Node platform, HyperTape provides either backup object templates or default backup objects, which are designed to help you create your own backup objects. They contain some preset values for defining save and restore operations on specific operating systems and databases.

### Queue Database (Control Node)

Every day, at a time you can specify, a Control Node job called the AutoSave Job searches the Object Database to determine which backups are scheduled to run during the day, and submits them to a job queue. At the same time, the AutoSave Job tidies up the job queues, removing any expired jobs.

Most HyperTape Control Nodes maintain a queue database and corresponding Queue Manager User Interface, which allow you to control and manage your job queues. You can use the Queue Manager to view logs of completed jobs, ascertain what happened and reschedule jobs when necessary.

### History Database (Control Node)

The HyperTape History Database records the names of savesets, journal files, backup times and saveset sizes. It is located on the Control Node and is used by all HyperTape utilities.

When a backup object is created, the Control Node software automatically creates a separate history file for it. HyperTape updates the history file after every successful save operation on the backup object. It can then use the backup object's history data to determine which backups ran when and with what parameters, to enable restoration of the data when it is needed.

## Journal Database (Service Node)

The Journal Database is maintained on the Service Node, although references to journal files are kept in the Control Node history files. Service Node software creates a new journal file for a backup object after every successful full backup, and provides a catalog of the files saved in the backup. Journal information from incremental backups is automatically appended to the current journal for the backup object.

## Log files

All HyperTape components maintain log files, usually in the log subdirectory of the component's root directory.

The Control Node software maintains a detailed log file for every job that runs. In addition, the *Master Log File* on the Control Node records the completion status of every job and enough information to locate the detailed log files for the job on both the Control Node and Service Node. There are facilities for selecting and viewing log files, so that you can analyze the events that occur.

## Performing backup operations

From the Control Node, you can initiate save and restore operations using any Service Node and Backup Node on the network.

The backup operation itself is performed by the Service Node, using a *backup utility* for a save operation, or a *restore utility* for a restore operation. The backup and restore utilities are usually standard software native to the Service Node platform. The Service Node then directs the output from the backup utility to the Backup Node, using a *transfer utility*.

The backup, restore and transfer utilities are specified in the backup object, and any messages arising from them are automatically sent to the Control Node.

## Updating HyperTape files

When the backup operation completes, the Service Node returns its progress and completion status to the Control Node. If the operation has been successful, the Service Node also updates the backup object's journal file and reports to the Control Node on the total number of kilobytes saved. The Control Node then updates the Object Database and History Database. You can monitor the progress and status of backup operations on the Control Node via the Queue Manager interface.

## Performing restore operations

The HyperTape restore operation returns data from a saveset to its original directory or to a new target location that you specify. Files, directories or volumes are restored if they have been deleted, or if the disk containing the original data has been corrupted. You can initiate restore operations from the Control Node, or independently from the Service Node.

Control Node restore operations are best suited to restoring entire file systems or directory trees. They rely on information held in the Control Node history files and user-entered patterns to locate the savesets and files to be restored.



Service Node restore operations may be used for selective restores of individual files and groups of files. They rely on information held in the Service Node journals to locate the files to be restored.

The restore operation, which can be initiated from either the Control Node or the Service Node, supports a selective or full restore of a backup object. By default, the restore operation restores the data to the original location, but you can specify an alternative location. When restoring a set of full and incremental backups, HyperTape restores the full backup of the specified object first, followed by all the incremental backups in chronological order.

## Automating backups

You can automate your backups with HyperTape's own internal scheduler, which allows you to specify set days of the week or selected dates on which various types of backups will be taken. You can also specify dates on which no backups should be performed.

For example, you could specify incremental backups every Monday to Friday, full backups every Saturday and monthly full backups (with a longer retention period) on the fourth Saturday of every month. In addition, you could specify the dates of public holidays on which no backups should be taken.

## Customizing backups

You can customize individual backups by automatically executing your own routines on the Control Node and Service Node before starting a backup operation (*preprocessing*) and after the backup operation finishes (*postprocessing*).

The preprocessing and postprocessing routines on the Control Node address issues concerning the backup job; those on the Service Node are used for activities restricted to the Service Node.

The availability of Service Node preprocessing and postprocessing depends on the Service Node platform, as does the method of implementation.

## Media management

The Backup Node needs to perform various media management functions, to ensure that HyperTape runs efficiently and to enable migration of saved data to secondary storage. These functions include:

- Allocating scratch tapes
- Managing the pool of used and free tapes
- Requesting the next tape at the end of the current tape
- Associating saveset names with tape serial number(s)
- Issuing mount messages to the operator or robot

When you use the Backup Node with OpenMedia, these functions are automatically performed by OpenMedia.

It is also possible to link HyperTape with other media managers, in which case these functions may be achieved via Backup Node preprocessing and postprocessing routines.

## Further reading

HyperTape is distributed on CD-ROM. The CD booklet has a quick start guide that explains how to transfer HyperTape components from the CD to the systems on which they are to be installed. There is a ReadMe text file for each component. This file contains the Release Notes, which provide information-specific to the version of the component, any late information that is not covered in the product manual and detailed installation instructions.

HyperTape documentation is provided in electronic form on the HyperTape CD, depending on your license. The manual set includes the following:

*HyperTape Enterprise Backup White Paper*

Introduces the features and facilities of HyperTape to a general audience and explains how HyperTape provides a comprehensive solution to current storage management requirements.

*HyperTape System Manager Guide*

Introduces HyperTape to system managers and administrators, describes its components and explains how to implement a network backup strategy using HyperTape.

*HyperTape Control Node Guides*

There is a *HyperTape Control Node Guide* for each supported platform. Each guide explains how to set up and operate the Control Node software on the relevant platform.

*HyperTape Backup Node Guides*

There is at least one *HyperTape Backup Node Guide* for each supported platform. Each guide explains how to set up and operate the Backup Node software on the relevant platform using the relevant media manager.

*HyperTape Service Node Guides*

There is a *HyperTape Service Node Guide* for each supported platform. Each guide explains how to set up and operate the Service Node software on the relevant platform. The guides also explain how to define parts of the backup object that are specific to the Service Node.

*HyperTape Service Node Restore Utility*

This manual is intended for end users. It explains how to use HyperTape's Service Node Restore Utility to connect to a HyperTape History Database on a HyperTape Service Node and restore files to that Service Node.

*HyperTape Database Backup Reference Manual*

Provides a reference manual for system managers, explaining how to set up HyperTape Database Agents for various databases on supported Service Node platforms, and how to define the backup objects that use the Database Agents.

*HyperTape OpenMedia Concepts and Facilities*

Provides an overview of HyperTape OpenMedia, explaining what it is and what it does. This manual is intended for anyone who needs an introduction to the product, from technical decision makers and system managers who are assessing HyperTape OpenMedia, to software engineers who are developing applications that will integrate with it.

*HyperTape OpenMedia System Manager Guide*

Intended for use by system managers and administrators, this manual explains how to set up HyperTape OpenMedia, and how to use the Graphical User Interface (GUI) and Command Line Interface (CLI) to monitor and control storage systems.

Most software components also have comprehensive online help.

## **For more information**

For more information on HyperTape and OpenMedia, contact BridgeHead by referring to page 2 of this paper, or to our corporate WEB site:

[www.BridgeHeadSoftware.com](http://www.BridgeHeadSoftware.com)