



Data Domain™

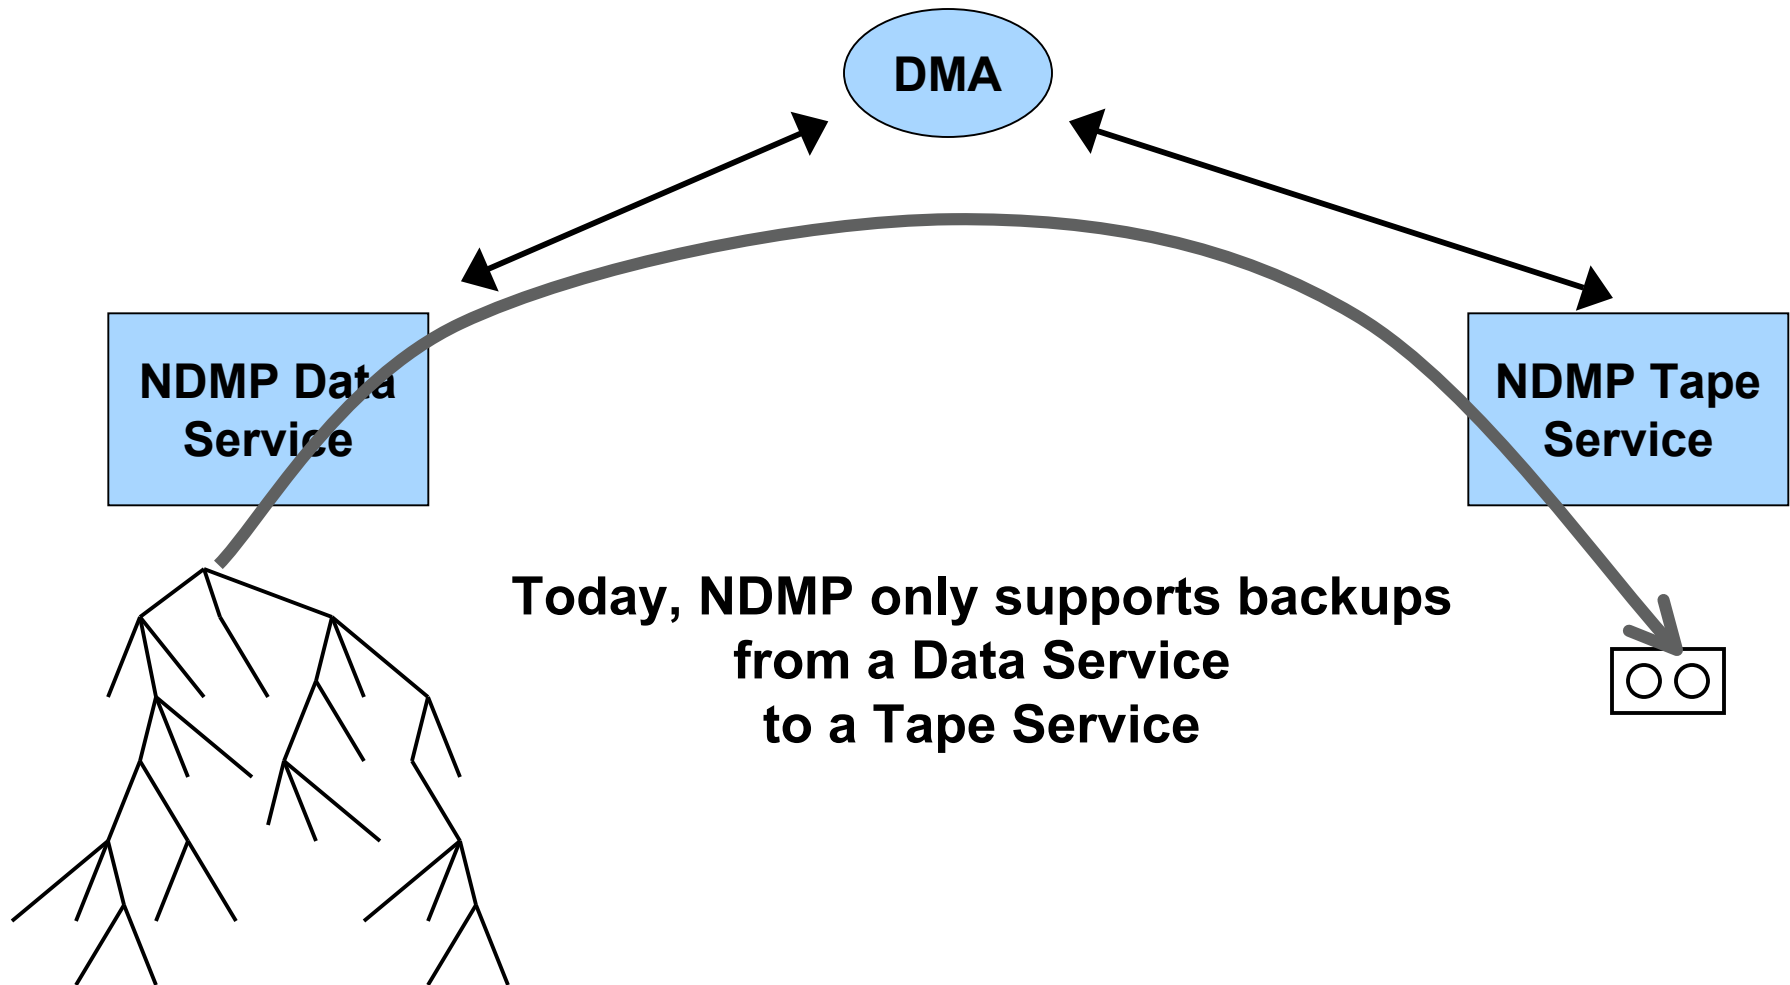
NDMP for D2D Backup

Hugo Patterson

Bob Fozard

Dave Manley

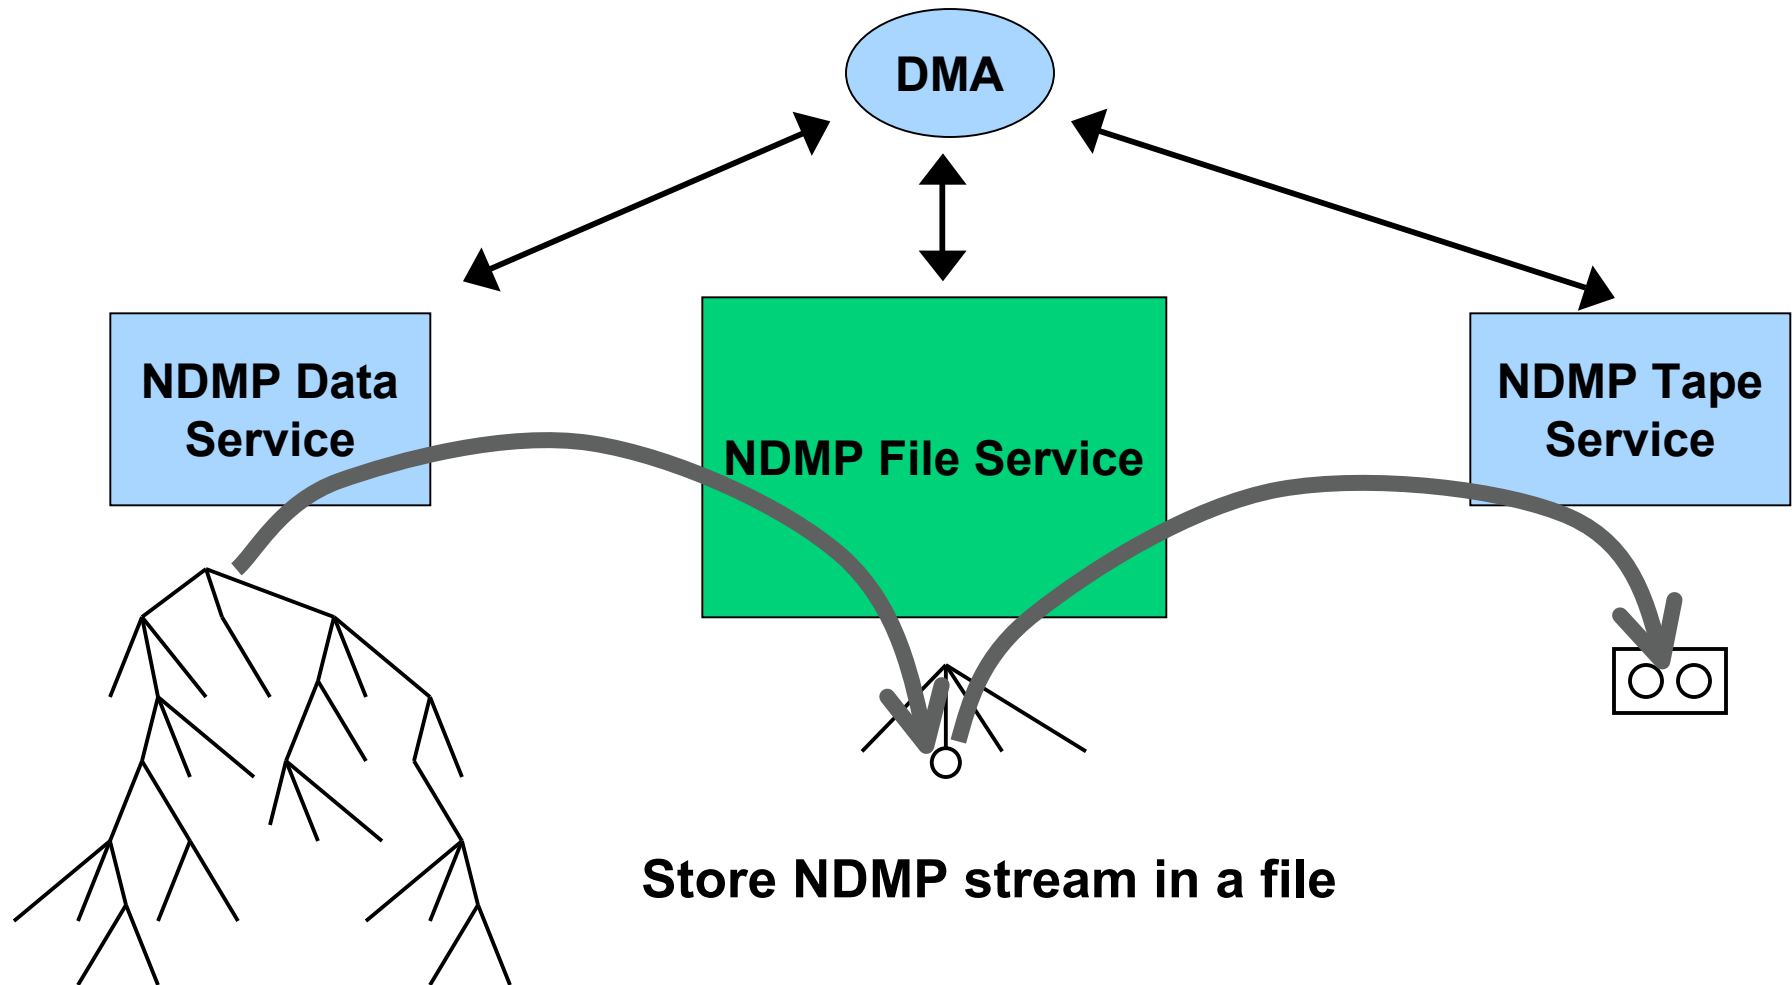
No D2D backup in NDMP today



Disk to disk backup is hot

- **Staging to disk then tape**
 - More reliable backup
 - Less prone to mechanical failures and human errors
 - Lower labor costs
 - Stream fastest tape drives without multiplexing
 - Verifiable recoverability
- **Online recovery copies**
 - More reliable restores
 - Faster restores (especially single file)
 - Lower TCO than tape (if files are compressed)

NDMP File Service for D2D backup



Benefits of NDMP D2D backup

- **Remove media manager overhead**
 - Transfer data directly to data protection storage
 - Avoid 2 hops on network
 - Spend less on media manager hardware
 - More \$\$ available for software
- **Higher performance NDMP backups**
- **Take advantage of new networking infrastructures**
 - FC-IP, RDMA, etc.

Proposal: NDMP File Service

- **An NDMP V4 extension for saving backups in a file on disk**
 - Takes NDMP data stream and stores it in a file on disk
 - Can send contents of file out as an NDMP data stream
- **Interoperates with existing data and tape services**
 - Can stage data to a file and from a file to tape
 - Can restore data directly from tape without staging back to file
- **Supports MOVER interface for data transfer to/from the file**
 - DMAs already know how to use the MOVER
- **New: NDMP_FILE_* messages for managing files/directories**
 - Use these instead of NDMP_TAPE_* messages
 - Supports directory hierarchies

Messages

- **NDMP_FILE_CONFIG_GET_INFO**
 - Equivalent to NDMP_CONFIG_GET_TAPE_INFO
 - Should we just use NDMP_CONFIG_GET_FS_INFO?
- **NDMP_FILE_GET_STATE**
 - Reports things like current size of file, offset in file, etc.
- **NDMP_FILE_OPEN**
 - Opens file for use with MOVER
- **NDMP_FILE_CLOSE**
- **NDMP_FILE_DELETE**
 - Delete file
 - Key difference from tape interface
 - DMA should proactively delete files when they expire
- **NDMP_FILE_TRUNCATE**
 - Truncate file to some size
 - Useful for restartable backups

Messages (cont'd)

- **NDMP_FILE_WRITE**
 - Write data from DMA to file
- **NDMP_FILE_READ**
 - Read data from file to DMA
- **NDMP_FILE_SEEK**
 - Seek to particular offset in file
 - Should be followed by NDMP_MOVER_SET_WINDOW call

Messages (cont'd)

- **NDMP_FILE_MAKE_DIR**
 - Make a directory
 - Takes full path as argument
- **NDMP_FILE_REMOVE_DIR**
 - Remove directory at path
- **NDMP_FILE_LIST_DIR**
 - List files in directory and return attributes of the files
- **NDMP_FILE_RENAME**
 - Rename a file or directory

Issues

- **Risks of using NDMP_MOVER_***
 - Spec language assumes MOVER is associated with tape
 - We want to transfer to file
 - Spec says must have a `_tape_` open before connect
 - We want to open a file
 - Amend spec or add NDMP_FILE_MOVER_* ???
- **NDMP_CONFIG_GET_FS_INFO**
 - Intended use is for data service
 - Has most of what a file service needs
 - Reuse or add NDMP_FILE_CONFIG_GET_INFO ???
- **Don't think any of these tweaks would break existing implementations**
 - E.g. ndmjob

Optional “Advanced” messages add flexibility

- **MOVER treats disk like tape**
 - Easy for DMAs
 - But doesn’t take full advantage of disk
- **“Advanced” messages leverage random access**
 - Open multiple files
 - Read/write arbitrary length records at arbitrary offsets
 - Overwrite data
- **Could enable new DMA capabilities**
 - Create synthetic fulls (if DMA knows data format)
 - Multiplexing/Demultiplexing
 - ???
- **DMAs do not need to use advanced messages**
 - Current functionality available via MOVER interface

Optional advanced messages

- **NDMP_FILE_OPEN_FID**
 - Open file and get a file ID handle for future operations
- **NDMP_FILE_CLOSE_FID**
- **NDMP_FILE_WRITE_FID**
 - write from DMA to file
- **NDMP_FILE_READ_FID**
 - read from file to DMA
- **NDMP_FILE_RECEIVE_FID**
 - write from stream to file
- **NDMP_FILE_SEND_FID**
 - read from file to stream
- **NDMP_FILE_TRUNCATE_FID**
- **NDMP_FILE_SEEK_FID**
- **NDMP_FILE_WRITE_METADATA**
 - Metadata gives DMA a place to store labels etc. separately from data
- **NDMP_FILE_READ_METADATA**
- **NDMP_FILE_TRUNCATE_METADATA**
- **NDMP_FILE_SEEK_METADATA**

Conclusion

- **NDMP should support disk-to-disk backup**
 - Faster, more reliable backup and recovery
- **Proposal: File Service extension to V4**
 - Functionality similar to Tape Service but stores NDMP stream in files instead of on tape
 - Interoperates with Data and Tape Services
 - Leverages familiar MOVER interface
 - New NDMP_FILE_* messages manage files and directories
- **Advanced messages expose random access to DMAs**
 - Offers DMAs greater flexibility for new features
 - Use by DMAs is optional