UTS Global Tape Services Suite Technical Overview

Phil Glossop

February 2003

Phil@utsglobal.com Linux@utsglobal.com

TSS Components 1

Tape Management Subsystem (TSS-TMS) - the core component, which provides device management through kernel drivers and the Tape Daemon (tpdaemon)

Tape Silo Feature (TSS-TSF) - provides access to tape drives housed in StorageTek® tape silos

Distributed Tape Service (TSS-DTS) - provides tape management services on the Linux/390 server to tapeless client systems

Virtual Tape Service (TSS-VTS) - provides virtual tape services (backed by a disk cache) on Linux/s390 and other platforms

TSS Components 2

Media Manager (TSS-MGR) - provides a centralised repository to control and administer the usage of the Linux/390 magnetic tape library.

Backup and Restore (TSS-BAR) - a native application to provide both filesystem and device image backups to tape

Oracle Backup (TSS-OBR) - provides the ability to backup and restore your Oracle database(s) using the UTS Global tape product suite.

DB2 Backup (TSS-DBR) - provides the ability to backup and restore your DB2 database(s) using the UTS Global tape product suite.

TSS Components – Availability Schedule

- **Tape Management Subsystem (TSS-TMS)** – Available Now
- **Tape Silo Feature (TSS-TSF)** – Available Now
- **Distributed Tape Service (TSS-DTS)** - Available Now
- **Virtual Tape Service (TSS-VTS)** - Preliminary version now. Full function in 3Q2003
- **Media Manager (TSS-MMG)** – Available 3Q2003
- **Backup and Restore (TSS-BAR)** – Available Now
- **Oracle Backup** – Available 3Q2003
- **DB2 Backup** - Available 3Q2003

TSS Components

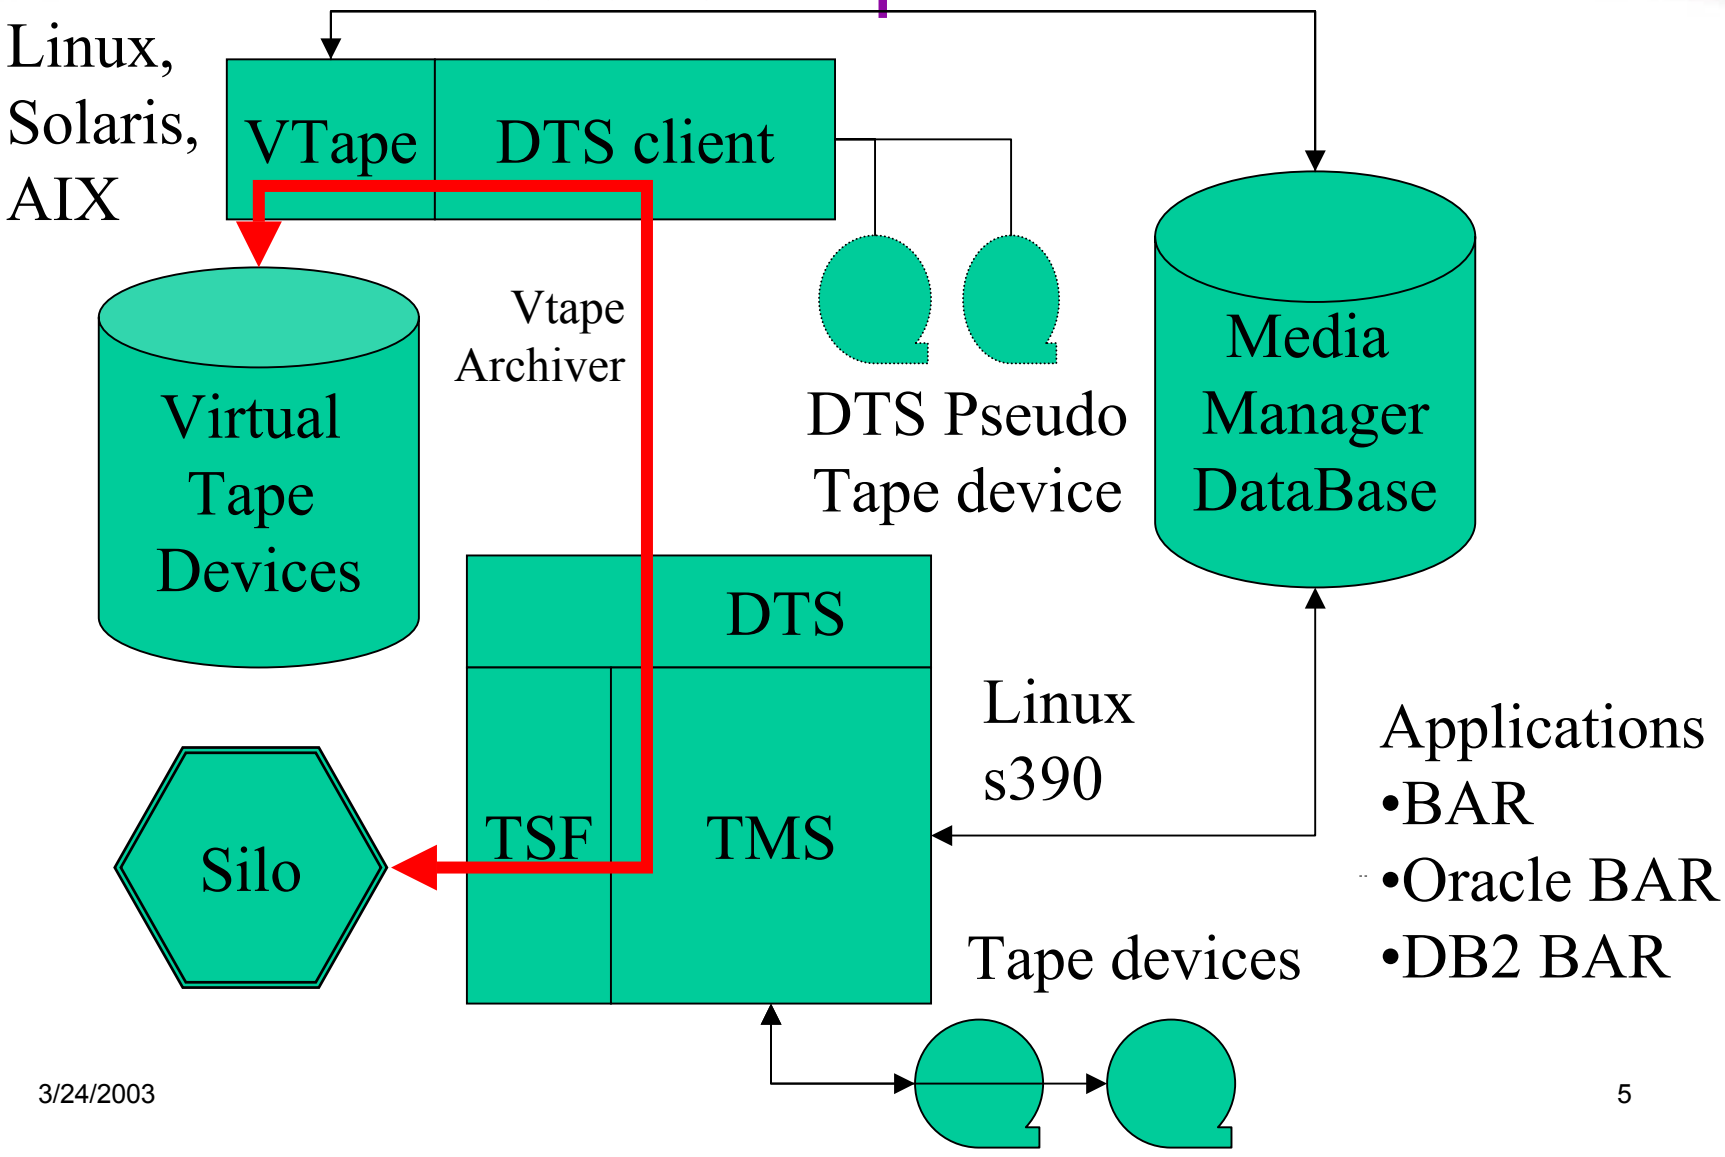

TSS-TMS the Tape Daemon

- The core of the tape subsystem
- tpdaemon plus a sophisticated driver, lots of ioctls
- Command line interface to mount/unmount tapes
- API to mount/unmount tapes
- Add/remove tape drives dynamically

Tape Access Philosophy

- Access tapes by volume name not by device
  - /dev/tape/<volname> not /dev/rmt etc.
- Device class – group of similar devices
  - C.F. JCL unit=180 unit=3590 unit=tape
- Can co-exist with the system provided driver
  - Assign devices via config files in /etc

TSS-TMS User Commands

- tape – mount/unmount a tape
  - e.g. tape –m –l IBM <volser>
  - Mounts IBM Standard Labeled tape <volser>
  - Makes device node /dev/tape/<volser>
  - Lots more options (see below)

- tapemt – display status of tape mounts

- tapevary – tape drive:enable/disable/display status

More TSS-TMS User Commands

- label – read/write tape labels
- tm – manipulate a tape … or ….
- mt – standard UNIX command to manipulate a tape
- tapersv – reserve a tape device (privileged)

Tape command options 1

Mount, Unmount (cancel) or Requeue request

Mount tape readonly or writable
Sequence number of target file
Physical (label) name of target file
Expiration date or Retention period

Bypass Label Processing
Non-labeled tape
Label type of target volume
Standard (IBM) labeled tape
Internal volser and external volser
A user specified file name in /dev/tape

Tape command options 2

Mount scratch tape
Silo scratch pool to use
Target tape drive address
Key if device is reserved
type of tape drive (3590 etc.)
Density of tape IDRC|NONIDRC
Tapeclass that a drive belongs to
Target host and optional port
Quiet mode, no messages to terminal
Target request number

Command executed when mount completes

TSS-TMS Admin Commands

- tapeadm – start, stop, query for tpdaemon
- tapeclass – groups devices by attrib e.g. location
- tapevary – tape drive:enable/disable/display status
- tpdconfig – parse and print tpdaemon config file

/etc/tpdaemon.conf

- Configuration file for tape daemon
- Many parameters, see sample at end of presentation
- Examples :
  - tapeoperator tapeoper
  - tape 3490 “V49*” # valid 3490 volsers
  - filemode 600 # dflt permission /dev/tape/<vol>
  - tape_log_file /var/log/tpdaemon

Silo feature

- Mount/Unmount STK silo tapes
- Silo option for <tape> command
- or C language programming interface
- Silo administration commands

Silo Feature Administration

- /etc/tapesilo.conf
- Two daemons
  - ssi
  - mini_el
- Assorted admin commands
  - E.g. audit the library, define/delete tape pools,
  - Lock drives or volumes, on/offline devices

/etc/tapesilo.conf

```
CSI_TCP_RPCSERVICE     TRUE
CSI_UDP_RPCSERVICE     TRUE
CSI_CONNECT_AGETIME    172800
CSI_RETRY_TIMEOUT      4
CSI_RETRY_TRIES        5
CSI_HOSTNAME           tapesilo
TRACE_VALUE            00000000
LOG_SIZE               64
ACSAPI_SSI_SOCKET      50004
EVENT_FILE             /var/log/silo_event.log
TRACE_FILE             /var/log/silo_trace.log
```

Distributed Tape Service

- Intended for use by tape-less clients
- Network connection to Linux/s390 tape devices
- tape –m –l IBM –H <hostname> <volser>
- Pseudo-tape driver ……..
  - Creates /dev/tape/<volser> on local machine
- Platforms : Linux (s390 and i386), Solaris, AIX(3Q03)

Distributed Tape Components

- dtape_cli : The client daemon
  - -c : cold start, ignore established connections
  - -t secs : heartbeat timer, default 60 secs
- dtape_svr : The server daemon
  - lives with tpdaemon
- /dev/ptap/ - psedotape driver

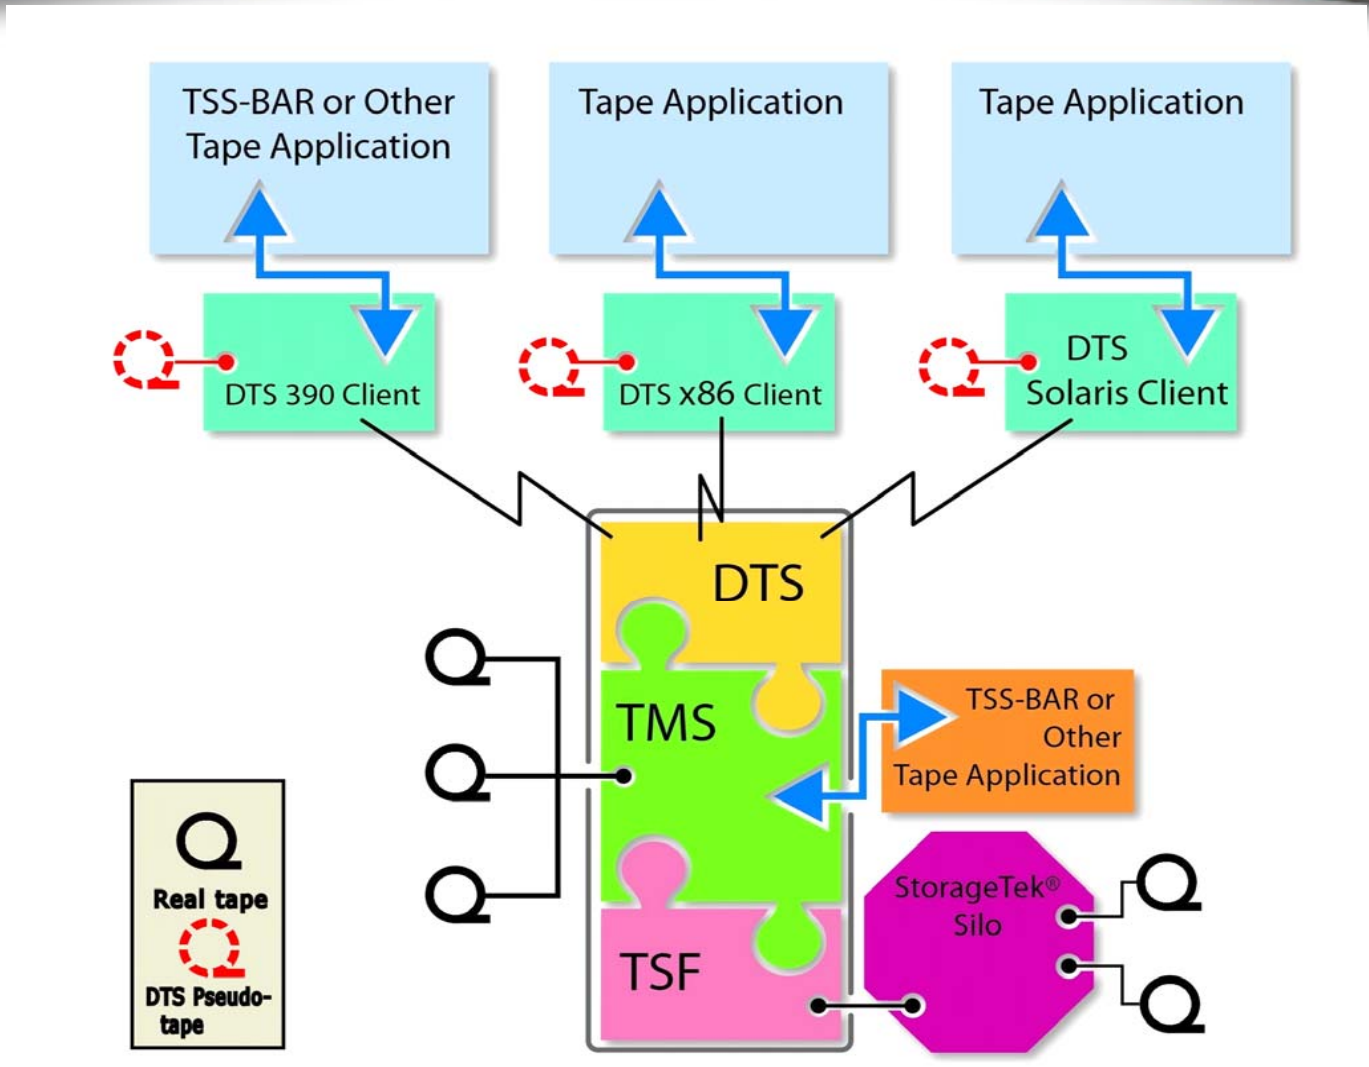
TSS-BAR or Other Tape Application
Tape Application
Tape Application
DTS 390 Client
DTS x86 Client
DTS Solaris Client
DTS
TMS
TSF
TSS-BAR or Other Tape Application
StorageTek® Silo
Real tape
DTS Pseudo-tape

Virtual Tape Service

- Tape devices backed by disk cache
- Software solution, no extra hardware required
- Uses standard filesystem input/output functions
- tape –m –t virtual <volser>
- Linux/s390, Linux/i386, Solaris, AIX(3Q03)

Virtual Tape Features

- Any utility that can use real tapes can use virtual tapes.
- Automatic scratch of expired tapes.
- Unix like permissions.
- Multiple 'files' per tape.
- Use commands or API to mount, to unmount,

  to scratch tapes or to search the tape database.
- Data can be imported from, or exported to real tapes

Virtual Tape Commands

- TSS-VTS can co-exist with TSS-TMS or run alone
- vtape – mount/unmount virtual tapes
  Same as <tape –t virtual> if TSS-TMS is installed
- vtape_cp - copy tape to tape, real or virtual,
  (takes care of mount/unmount)

Virtual Tape Commands II

- vtape_ls – list tape attributes
- vtape_rm – return a tape to the free pool
- vtape_init – initialize the cache
- vtape_db – vtape admin, add, remove, list, stats
- All commands have C language equivalents

Virtual Tape Archiver

- Planned availability early 3Q 2003.
- Automated import /export between real/virtual tapes
- Options :
  - Archive now (real time)
  - Archive soon (background)
  - Archive default (based on disk cache thresholds)

Media Manager Database

- Tapes are grouped into named pools
- Tapes are in <free> or <allocated> state
- Implements ownership, access permissions, expiry
- Tracks physical location of volumes
- Audit trail of tape usage
- Fully integrated with other Tape Subsystem products

TSS Applications

- UTS Global
  - Backup and Restore
  - Oracle Backup
  - DB2 Backup

- Or ……
  - Any standard operating system utility
  - Any 3rd party application

Backup and Restore (TSS –BAR)

- Native Backup and Restore for Linux/s390
- Integrated with UTSG tape management system
- Works with Silo tapes for automated backup
- Backup Types
  - File level
  - Physical region (disk partition)
  - Disk Image

Incremental and Differential Backup

- Full Backup – Complete “Object” e.g. all of /home
- Incremental – Everything since last “Full” backup
- Differential – Everything since last backup of any kind

TSS–BAR Commands

- Command line or GUI interface

- backup <parmfile>
  - Physregion - e.g. /dev/dasda1
  - Object - e.g. /usr
  - Image – very fast but needs kernel changes
- recover <options>
  - Recover has many options (see man page)
  - Simple example : recover ~/mail ~/'test*'

    Will recover my mail directory plus any files or directories beginning with the string "test"

More TSS–BAR Commands

- bktapelib – maintain BAR tape library
  - Add, delete tape volumes
  - Mark volumes as used or unused
  - Change expiry dates etc.

- bkuprpt <logfile>
  - Report on objects backed up

TSS-BAR Parameter File

- Catalog – specify backup catalog directory
  - Tape library (created by bktapelib)
  - Index – what has been backed up and where
  - Log directory (audit trail)
- Backup – List objects to backup
  - E.g. /usr, '/home/*/stuff', '/dev/pty*' (inode)
  - Physregion – e.g /dev/dasda1
- Other stuff – Examples : Device – 3490, 3590 etc.
  - Expiration – days, Silomount – silo or nonsilo

TSS Applications – Oracle BAR

- Backup and Restore for Oracle on Linux/s390
- Uses Oracle RMAN API
- Planned for early 3Q 2003
- Integrated with UTSG tape management system

TSS Applications – DB2 BAR

- Backup and Restore for DB2 on Linux/s390
- Planned for early 3Q 2003
- Uses DB2 XBSA interface
- Integrated with UTSG tape management system

TSS Components

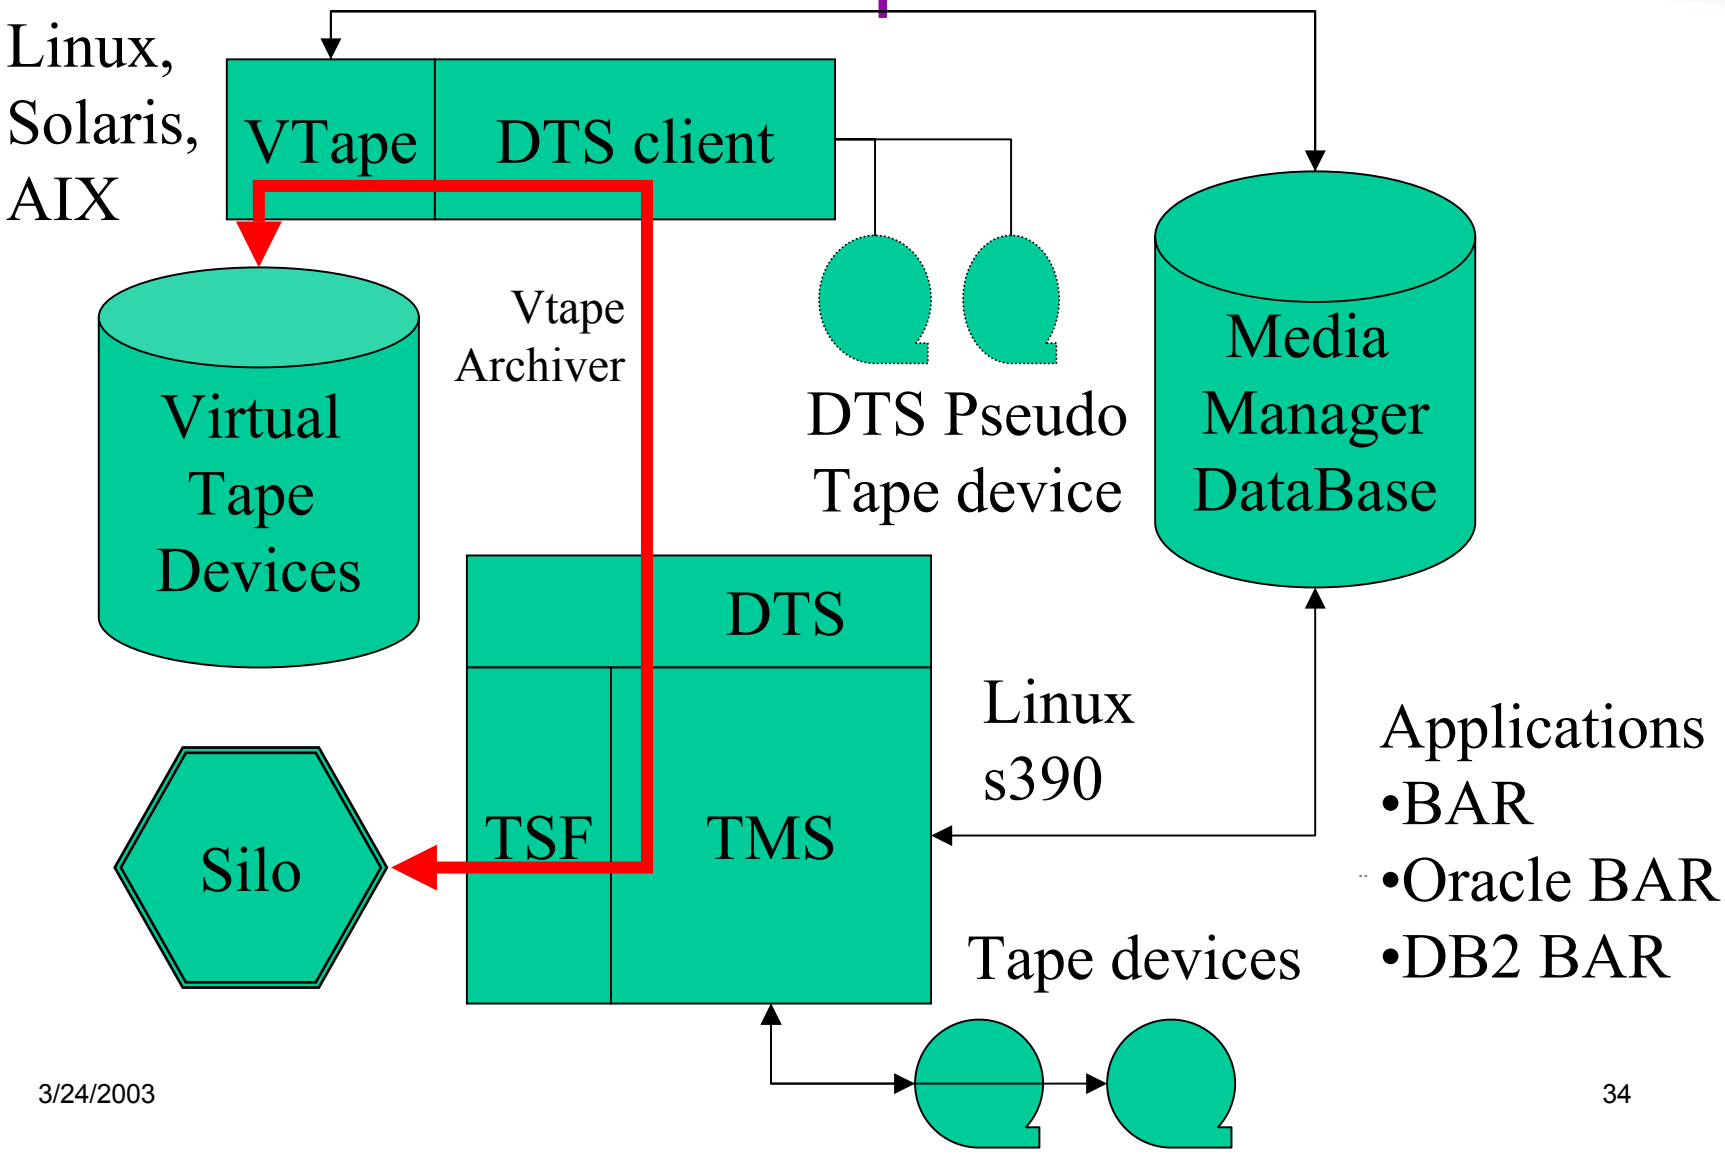

UTS Global Tape Services Suite Technical Overview

Phil Glossop

February 2003

Phil@utsglobal.com Linux@utsglobal.com

Sample tpdaemon.conf

```
# Valid volser ranges are defined using POSIX 1003.2 extended regular
# expressions which are discussed in detail in regex(7).
# The regular expressions are those used by regcomp(5) NOT
# shell regular expressions.
#

tape    any    ".*" # valid external volsers for any media type
tape    3590   ".*" # valid external volsers for 3590 media type
tape    3490E  ".*" # valid external volsers for 3490E media type
tape    3490   ".*" # valid external volsers for 3490 media type
tape    3480   ".*" # valid external volsers for 3480 media type

#
# The following are internal tape range definitions.
# The default for internal volser ranges is the
# definitions for valid external volser ranges.
#

# internal_tape    any     ".*"   # valid internal volsers for any media type
# internal_tape    3590    ".*"   # valid internal volsers for 3590 media type
# internal_tape    3490E   ".*"   # valid internal volsers for 3490E media type
# internal_tape    3490    ".*"   # valid internal volsers for 3490 media type
# internal_tape    3480    ".*"   # valid internal volsers for 3480 media type
```

Sample tpdaemon.conf - continued

```
#
# Definitions for system defaults
#

# devtype        any               # set device to 3480,3490,3490E/3590 or any
# tape_class     backup            # set tape class.  If not specified
                                                    # in this file, the tape daemon will
                                                    # have no default value defined for
                                                    # tape class.
# density        3480   NONIDRC    # set tape density for 3480
# density        3490   IDRC       # set tape density for 3490
# density        3490E  IDRC       # set tape density for 3490E
# density        3590   IDRC       # set tape density for 3590
# labeltype      IBM               # other options are VOL ANSI BLP and NL

#
# Definitions for tape operator and tape administrators
#

# tapeoperator                  tapeoper  # tape operator
# tapeadmin                     tapeoper  # tape administrator
# tapeadmingroup                adm       # authorized group
# groups_adm_and_bin_are_admin  off       # groups "bin" "adm" administrators ?

#
# Set up defaults for tpdaemon(1m) request queue size, file retention
# period and permission mode bits for /dev/tape/<volser> files.
#

# request_queue_entries     100 # size of request queue
# retpd                 forever # default file retention period
# filemode                  600 # default permission bits for /dev/tape/<volser>
```

Sample tpdaemon.conf - continued

```
#
# Set up default for number of dynamic drives which can be online
# simultaneously
#

# dynamic_drive_entries     16  # default number of dynamic drives

#
# Set up directory and file structure for tpdaemon(1m), home and log
# directories
#

# tape_homedir                  /var/run/tpdaemon       # home directory
# tape_lock_file                daemonlock              # name of lock file
# tape_shared_memory_keyfile    tpd_shmkey              # share memory key file
# tape_log_file                 /var/log/tpdaemon       # log file name

#
# Following flags effect tpdaemon(1m) log, and messages
#

# send_log_output_to_stdout            off   # send log to standard output ?
# do_not_send_messages_to_operator     off   # message sent to operator ?
# timestamp_log_entries                 on   # time stamp log entries ?

#
# All the flags for the tape daemon
# Default for all of these flags is "off"
#
```

Sample tpdaemon.conf - continued

```
# enforce_requester_only_open        off   # requester ONLY can open tape ?
# block_user_after_io_error          on    # cancel request if I/O error ?
# allow_case_sensitive_labels        off   # ok to use lower case labels ?

#
# Tapes label types non-administrators can mount. Non-administrative users
# can always mount the tapes with the default label type.
#

# allow_nonadmin_blp                   on    # can non-admin mount BLP ?
# allow_nonadmin_IBM_label_spec        on    # can non-admin mount IBM tape ?
# allow_nonadmin_ANSI_label_spec       on    # can non-admin mount ANSI tape ?
# allow_nonadmin_VOL_label_spec        on    # can non-admin mount VOL tape ?
# allow_nonadmin_NL_label_spec         on    # can non-admin mount NL tape ?

#
# Time interval and number of reminder messages to the operator
#

# oper_msg_limit          0   # Default: do not reissue messages to operator
# oper_msg_interval      10   # 10 minutes: interval between messages

#
# Ask the tape operator to confirm non-scratch non-labeled (NL) tape mounts
# and non-scratch tapes mounted with by-pass label proecessing enabled
#

# verify_non_scratch_NL_label_spec     on    # confirm non-scratch NL mounts ?
# verify_non_scratch_BLP_label_spec    on    # confirm non-scratch BLP mounts ?
```

Sample tpdaemon.conf - continued

```
#
# Set flag to determine if non-administrators can skip volser format check.
# Administrators can always skip volser format check.
#

# nonadmin_skip_label_format_check     off   # can non-admin skip volser check ?

#
# Set flag to determine if a ring-check is made when a tape is mounted
#

# enforce_ring_check     on         # should ring-check be made for mounts ?

#
# Specify which users are authorized to reserve tape devices. If the
# username does not appear here, even administrative users cannot reserve
# devices. Note that values below are not defaults, just examples.
#

# enable_device_reservation ADDR  580  larry  # user "larry" can reserve dev 580
# enable_device_reservation TYPE  3480 curly  # user "curly" can reserve a 3480
# enable_device_reservation CLASS ipt  moe    # user "moe" can reserve a dev in
                                                                    #     tapeclass "ipt"

#
# Define action to take if an unexpired tape volume is mounted writable
#
```

Sample tpdaemon.conf - continued

```
#
# Define action to take if an unexpired tape volume is mounted writable
#

# if_mounted_unexpired_volume_for_write nonsilo ask   # ask tape operator whether to mnt unexpired volume on
non-silo drive
# if_mounted_unexpired_volume_for_write silo cancel   # cancel request if unexpired volume mounted on silo driv

# Specify the default mode for mounts.  The default is extended
# mode.  This can be overridden by inividual mount requests (tape [-U|-X]).
#

# compat_21_mode off                # by default do not mount tapes in compatibility mode
# check_scratch_expiration off      # by default all tapes mounted as scratch are expired
# allow_premount_scratch off        # by default any premounted tape will be rejected as a scratch tape.

#
# Silo related parameters.
#

# The location of the STK SSI start script is configurable to cater for
# differences between Linux distributions.
# stk_ssi_path                /etc/rc.d/init.d/ssi    # STK SSI start script

# Define drives that are in the silo
#

# silo_devices   0  0  9  0 580-583    # values are examples only, not defaults
```

Sample tpdaemon.conf - continued

```
# silo_lock_protocol   lockvry         # can be lockvry, lockdrv, lockvol, lockall
                                                  # lockall2 or nolock

#
# Define silo scratch pool id
#

# silo_scratch_pool    0               # scratch pool id of 0 is default

#
# Define whether daemon should come up in "silo" mode or "non-silo" mode.
# In silo mode, first try to assign silo drives. In non-silo mode, first
# try to assign non-silo drives.
#

# assign_silo_drives_first    off  # by default, non-silo drives assigned first
# ignore_case_for_silo_volsers off # by default, do not ignore case for silo volsers
# ignore_hdrs_for_compat_mode_mount off # by default, do not ignore hdrs for compatibility mode mounts

#
# Specify flag which determines whether non-admin users are allowed access
# to silo drives. Admin users are always allowed access to silo drives
#

# allow_nonadmin_silo_access  off  # by default, no silo access to non-admin uid
```