Data Protection for LAMP/J applications

Paddy Sreenivasan
paddy@zmanda.com
April 2006

Agenda

- LAMP/J application data
- Open source data protection tools
- File system data protection
- MySQL data protection
- Amanda: Application data backup and recovery
- LAMP/J application data protection solutions
- More information

LAMP/J Application stack

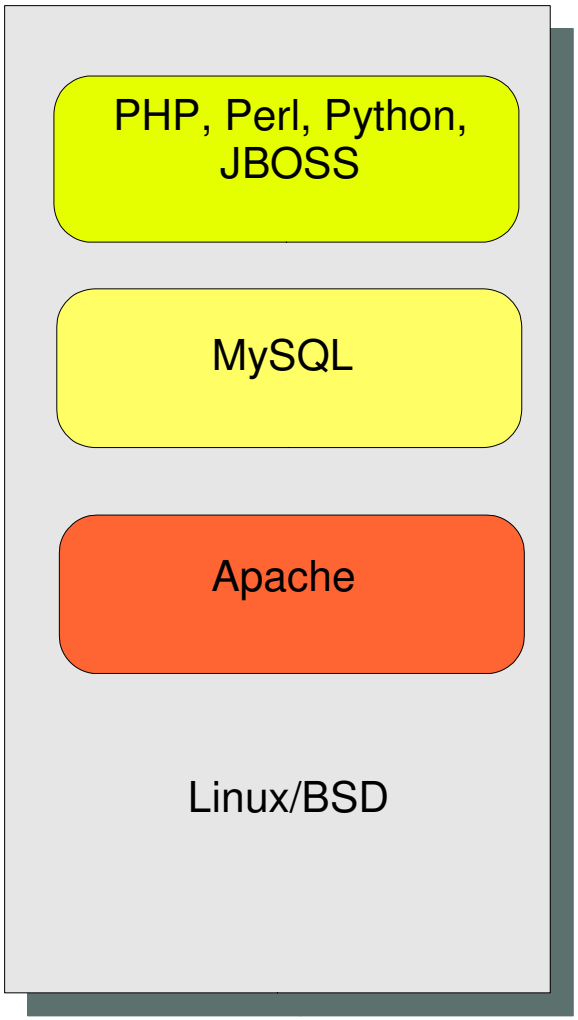

LAMP application
Example: Sugar CRM, MediaWiki

Database

Web server

Operating System

LAMP/J Application data

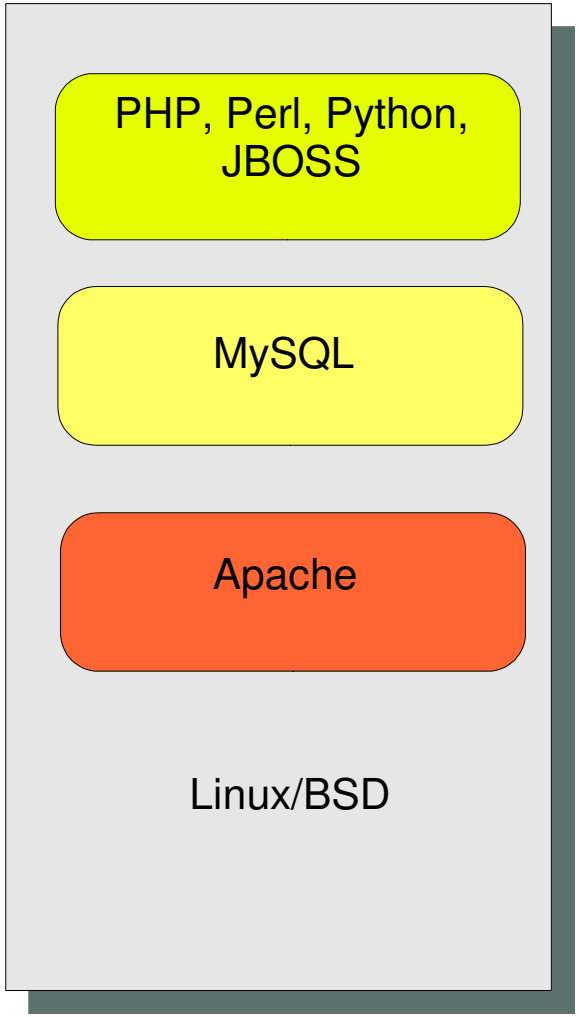

- LAMP application and configuration data

Example: Sugar CRM
Sugar CRM database
Sugar CRM files in /var/www/html

- Perl/PHP/Python/JBOSS configuration files
- MySQL configuration files
- Apache configuration files

LAMP/J Application: Where does data reside?

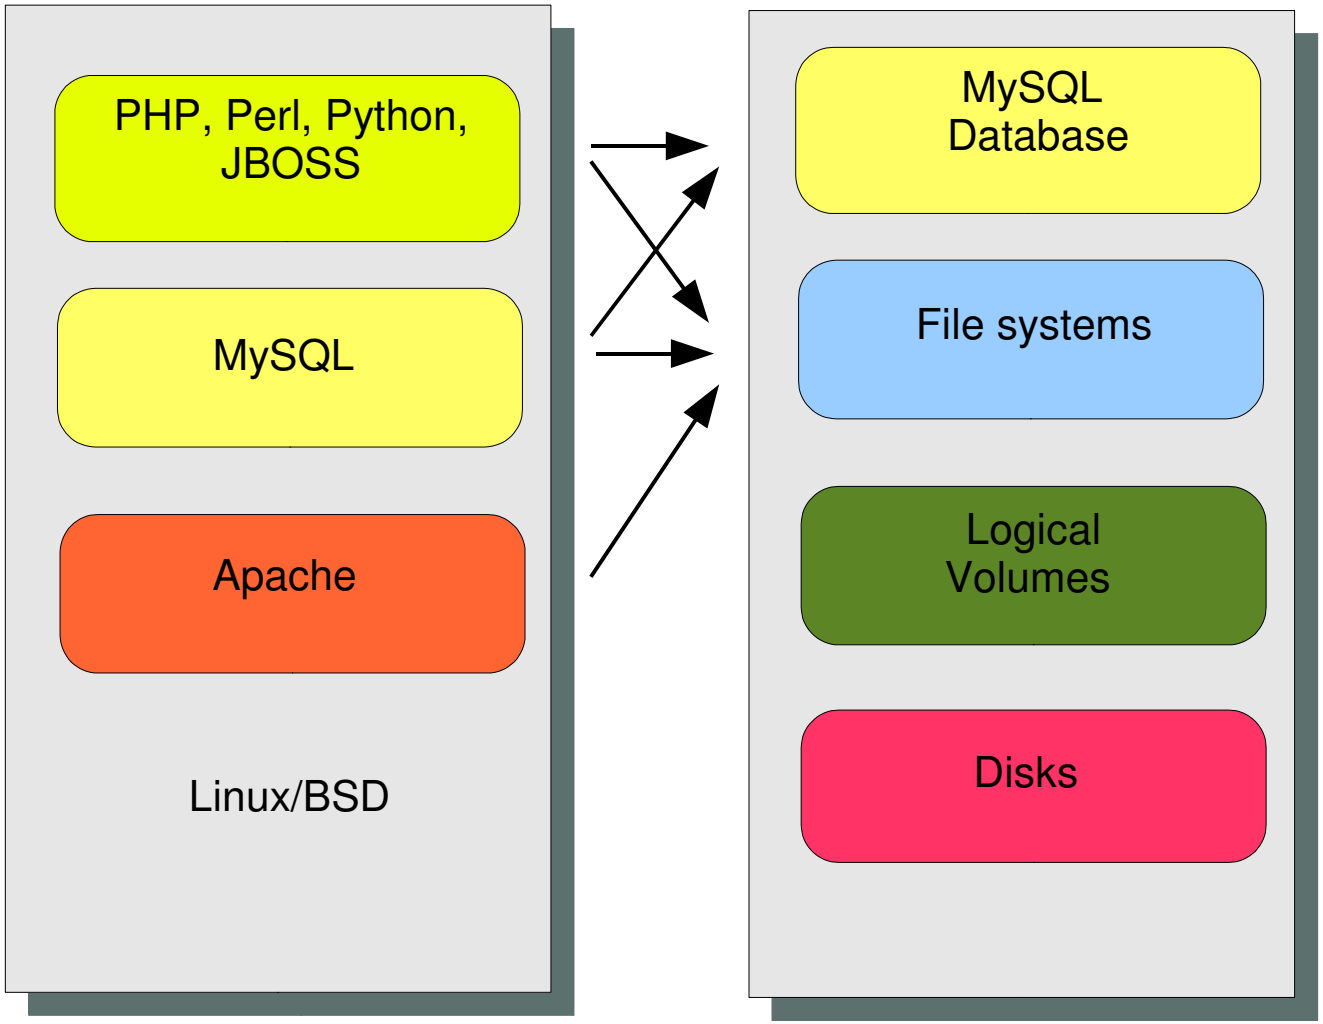

Open source data protection tools

- Use RAID or Volume manager mirroring
- Data snapshots
  - Application consistency
  - Crash consistency
- Replication (synchronous, asynchronous)
- File system backup for recovery/archival needs
- Media: Disks, Tapes, Optical media, Network attached storage

Network based replication (DRBD)

- Block device synchronous replication using LAN
- Requires cluster membership
- R/W access on the primary node only.
  - Disconnect DRBD on secondary node
  - Mount read only
  - Backup
  - Connect DRBD
- DRBD+ provides asynchronous replication.
- Expensive solution for Backup only.
- Operator errors/security flaws also get replicated!!!

Logical Volume Manager (LVM 2)

- Read/Write snapshots
- File system consistent snapshots with 2.6 kernel

Creating a snapshot volume for backups of size 50MB

```
# lvcreate -L50M -s -n mysqlbackup /dev/mysqlvolume
```

Remove logical volume after backup

```
# lvremove /dev/mysqlvolume
```

Enterprise Volume Manager system (EVMS)

- Unifies all Volume Management and Filesystems
- LVM, MD (Software RAID), File system plug ins
- Bad block remapping
- Snapshot expansion - Long backup windows
- Snapshot reinitialization
- Snapshot rollback - Quick backup/restore

File system snapshots

- Journaled file systems (ext3fs, XFS, JFS, ReiserFS, GFS)
  - Data and meta-data journaling does not require snapshots
- File system snapshots
  - XFS (xfs_freeze -f and -u options)
  - GFS (gfs_tool freeze/unfreeze)
- Backup steps
  - Freeze file system
  - Create a volume manager snapshot
  - Thaw file system

MySQL full backup - mysqldump

- Online logical backup of the database.
- Obtains READ LOCK on the tables followed by FLUSH TABLES to flush active indexes to the disk.

mysqldump [options] db_name [tables]

mysqldump [options] --databases db_name1 [db_name2 db_name3...]

mysqldump [options] --all-databases

MySQL full backup – mysqldump options

- **--opt and –extended_insert** : Creates optimized MySQL statements in the dump file. Speeds up restore operation.
- **--lock_all_tables and –flush-logs** : Locks all tables and flushes the data before backup. Pending transactions are committed. Consistent backup for HEAP, MERGE, ARCHIVE and MyISAM tables.
- **--single-transaction** : Creates backup in a single transaction. Consistent backup for transactional tables such as InnoDB and BDB.
- **--routines, --triggers**: Backs up Stored routines and table triggers.
- **--master-data** : Rotates binary logs. Marks the position when the backup was done. Useful for Point-in-time recovery.

MySQL full copy - mysqlhotcopy

- Backs up MyISAM, MERGE, ARCHIVE, FEDERATED tables.
- Faster than mysqldump for backup and restore.
- Obtains Read lock on the tables and does flush tables to create consistent backups.
- Copies the database using cp or scp commands.

```
mysqlhotcopy --flushlog db_name_1 ... db_name_n
                                    /path/to/new_directory
```

MySQL full backups - innoDB

- innoDB Hot Backup (ibbackup)
  - Non-Blocking physical backup of innoDB tables and indexes.
  - Transaction consistent backups.
  - Only Full backups.
  - Does not backup .frm files, MyISAM tables/indexes.
  - All database pages are not backed up at the same time. Apply innoDB log to roll forward the back up to same log sequence number.

MySQL incremental backups

- Enable binary logging (--log-bin option for mysqld)
- Binary log file should on a different disk from MySQL data.
  - Improved performance
  - More availability
- New binary log is created when MySQL is restarted
- Rotate the binary logs using
  - FLUSH LOGS SQL statement

    or
  - mysqladmin flush-logs

MySQL incremental restore

- Restoring incremental backups from log files starting with “binlog”

```
mysqlbinlog binlog.[0-9]* | mysql -u root -p
```

- Point in time recovery using --stop-time or --start-time

```
mysqlbinlog --stop-date="2005-04-20 9:59:59" \
        /var/log/mysql/bin.8989 | mysql -u root -p
```

- Recovery to a position using --start-position or --stop-position

```
mysqlbinlog –start-position="3786" /var/log/mysql/bin.768 \
        | mysql -u root -p
```

MySQL server replication

- Master servers
  - Handles client requests
- Multiple slave database servers
  - Same data as master server
  - Statement based replication
  - Row based replication (MySQL 5.1)
- Use database backup methods on the slave server to backup the database after stopping replication.
- Easy to set up and administer.
- No safeguards against operator errors.

MySQL cluster

- Synchronous replication.
- Integrated with MySQL as NDB tables.
- Write scalability via data fragmentation.
- One of the nodes can be used to take a MySQL database backup.
- Expensive if used only for data protection.

Amanda – file system backup & recovery

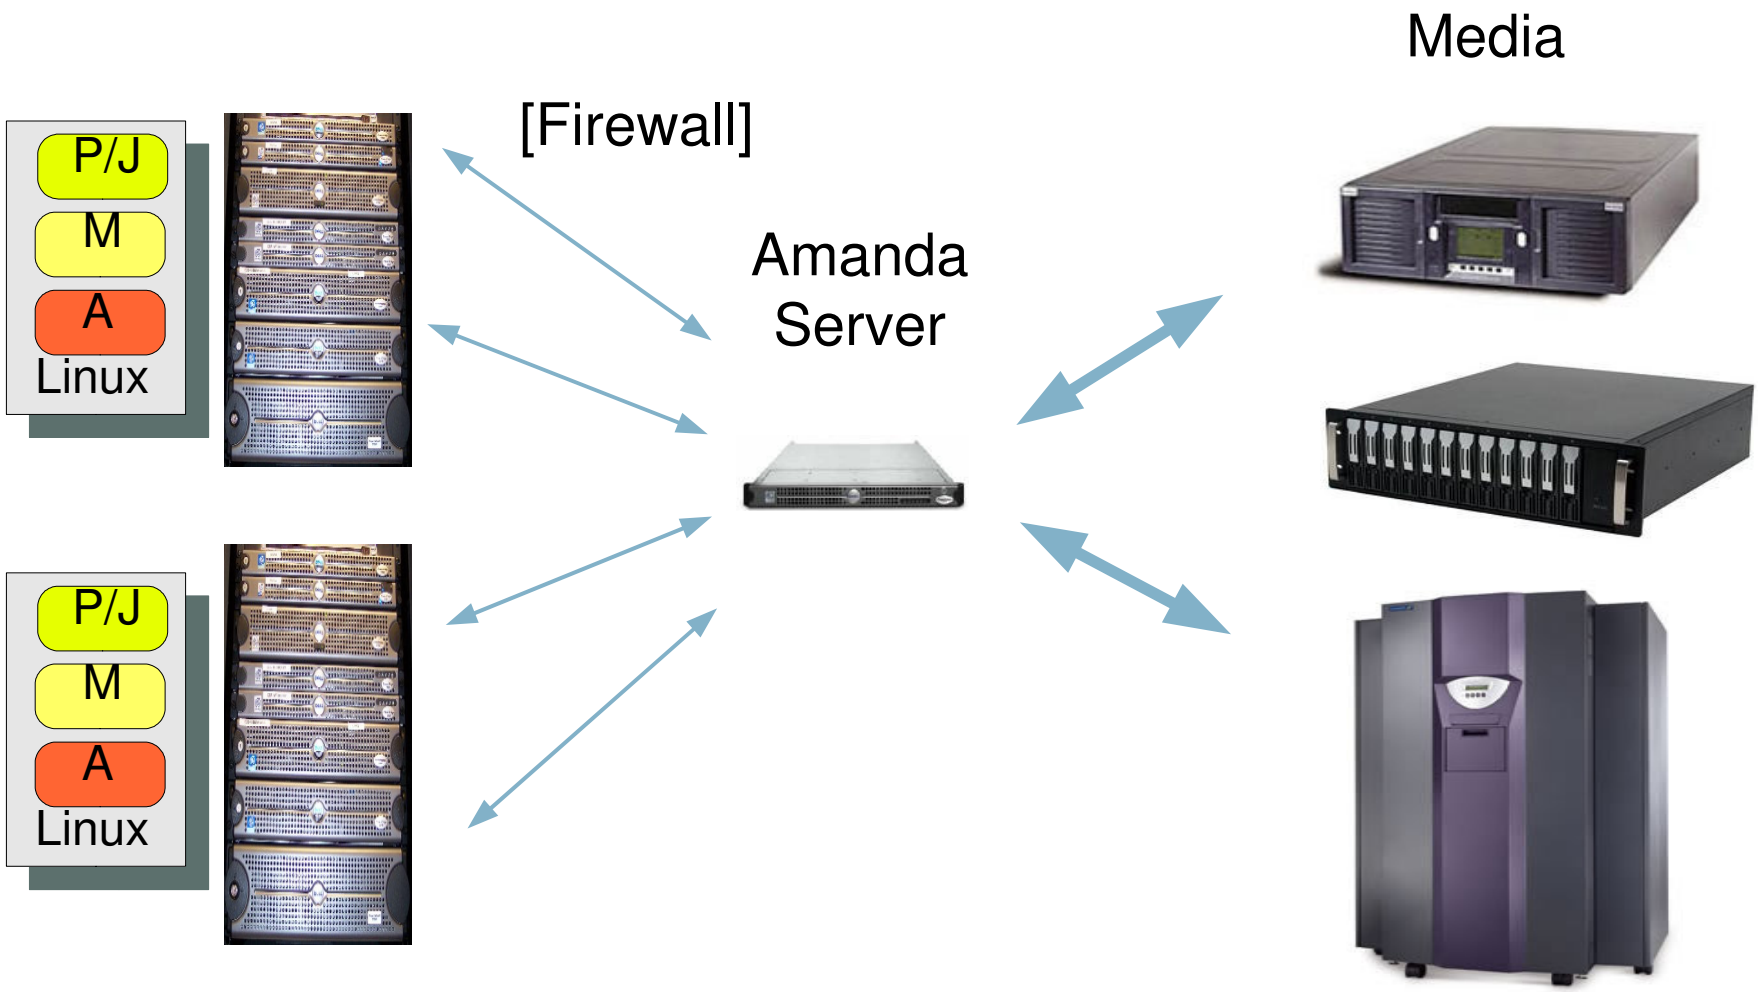

Amanda project

- Active sourceforge.net project with tens of thousands of users and developers.
  - Current release: 2.5
- Active documentation.
  - Amanda wiki http://wiki.zmanda.com
- Supports LAMP operating systems - Linux/BSD
- Distributed as part of Linux and BSD operating systems
- Latest RPMs available at http://www.zmanda.com/downloads.html

Amanda features

- Consistent backup window
  - Tries to backup same amount of data every run
  - Supports multiple backup levels
- Uses native platform tools
  - GNU tar/dump/star for backup
  - gzip/bzip2 for compression
  - gpg/aespipe tools for data encryption
- Media format – IEEE standard pax format
- Supports tapes, disks, changers, optical media, RAIT
- Network load balancing

Recovering data

- Amanda stores all media, file information in a database (MySQL in the future)
- The database can be browsed using Amanda commands
- Data in the media is stored in native format and can be recovered using OS tools without Amanda commands
  - Backup images spanning multiple media can also be restored using OS tools

Amanda media file header


```
# mt -f <media_dev> fsf 1
# dd if=<media_dev> bs=32k count=1
AMANDA: FILE 20060228 natasha /boot  lev 1 comp N program
    /bin/gtar
To restore, position tape at start of file and run:
       dd if=<tape> bs=32k skip=1 |    /bin/gtar -f... -


1+0 records in
1+0 records out
```

Data protection comparison

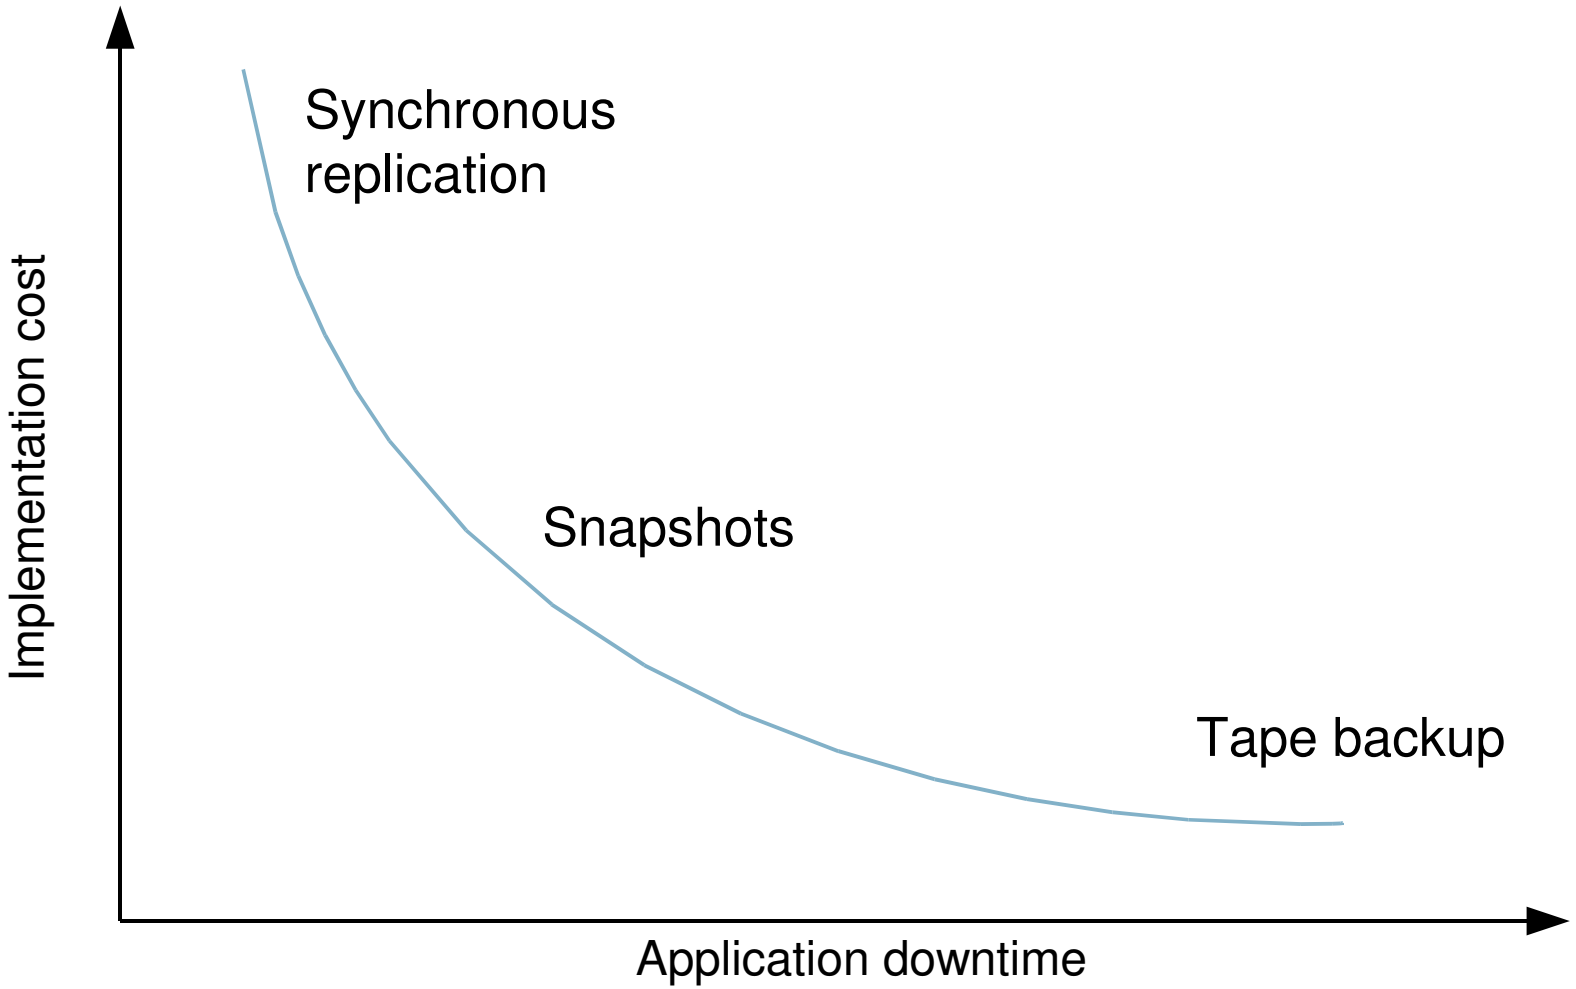

LAMP application data backup (snapshot solution)

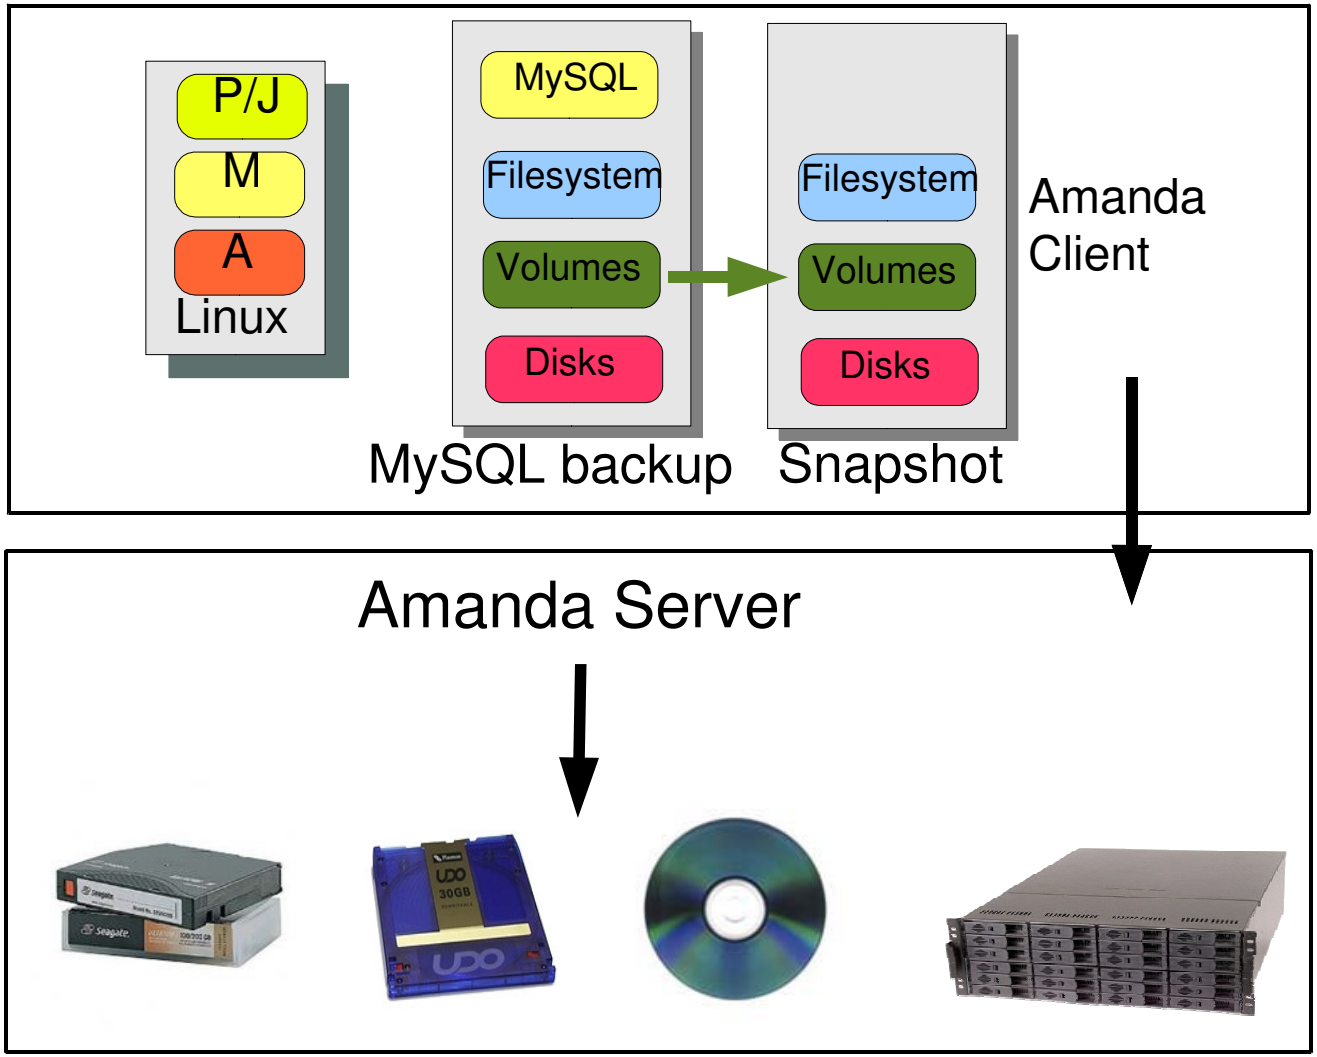

LAMP application data backup (MySQL replication solution)

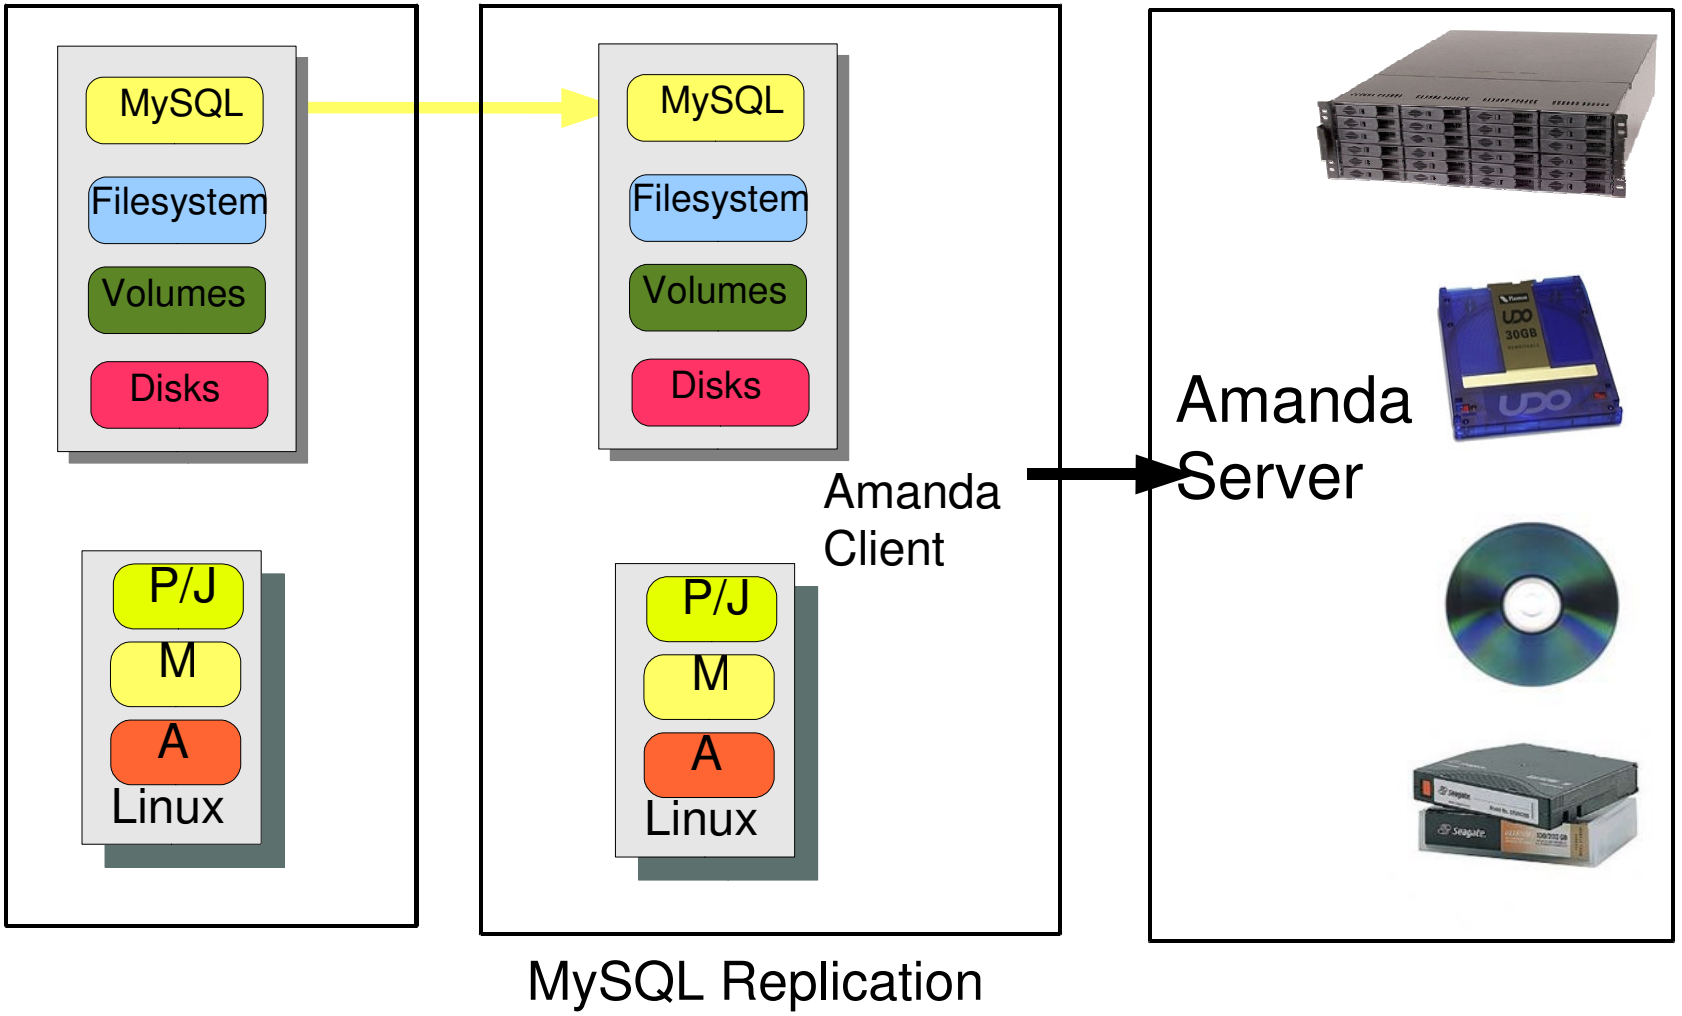

LAMP application data backup

- Create Amanda dump program wrappers
  - Prepare for backup (Data consistency)
    - Backup the MySQL database (full/incremental)
    - Create filesystem snapshots using LVM
  - Post-backup actions
    - Delete LVM snapshot
- Amanda's new Application API will not require program wrappers.
- LAMP application crash consistent backup

LAMP application data restore

- Restore from the media
  - Amanda tools – amrecover command
  - Native platform tools (dd, mt)
- MySQL database restoration from
  - MySQL full backups
  - MySQL binary log files

LAMP application data security

- Backup data security
  - encrypted data in the media
- Communication security
  - Kerberos
  - OpenSSH
- Working in Secure environments
  - SE Linux
  - Apache server virtual hosts

Zmanda – Open source data protection

- Company that provides open source data protection solutions and services.
- Active developer and supporter of open source projects.
- Offers Amanda enterprise edition on an annual subscription.
- Services and Solutions based on Amanda enterprise edition.
- Zmanda professional services provide open source data protection solutions for LAMP applications.

More information

- Zmanda network: http://www.zmanda.com
- Amanda documentation: http://wiki.zmanda.com
- Amanda images: http://www.zmanda.com/downloads.html
- Amanda forums: http://forums.zmanda.com/
- Amanda sourceforge project: http://sf.net/projects/amanda
- MySQL backup and recovery:
  http://dev.mysql.com/doc/refman/5.0/en/disaster-prevention.html

Contact information

- Talk to us at MySQL users conference
  - Zmanda booth 416
  - BOF session on “MySQL backup”

    Wednesday, April 26th, 2006 7:30pm Ballroom A
- Zmanda website
  - http://www.zmanda.com