ARCITECTA™
www.arcitecta.com

Optimising DMF Storage Workflows with LiveArc

NOVEMBER 2011

Peter Cross
Arcitecta
Business Development

MEDIAFLUX™
OPERATING SYSTEMS FOR META+DATA

LIVEARC

- Commercial Software Platform
  - Also known as Mediaflux
- Developed by Arcitecta
  - Based in Melbourne, Australia
- Global customer base in diverse industries
  - Research, Government, Defence
- SGI is global reseller
- Ocean Informatics OEM for e-health
- Customers since 2003 – AAHL & Weta soon after

CLIENT SERVER TOPOLOGY

- LiveArc Server(s)
  - Close to the data
- Client application, e.g. LiveArc Desktop
  - Ajax based Web 2.0 application
  - Dynamically constructs menus, panes, pop-ups etc. based on data in the associated LiveArc repository
  - Bundles specific purpose applications in desktop metaphor
- Also Aterm, MACDDAP, CKM, DaRIS, Clinical Viewer, Clinical Hub, DINS, nQuisitor

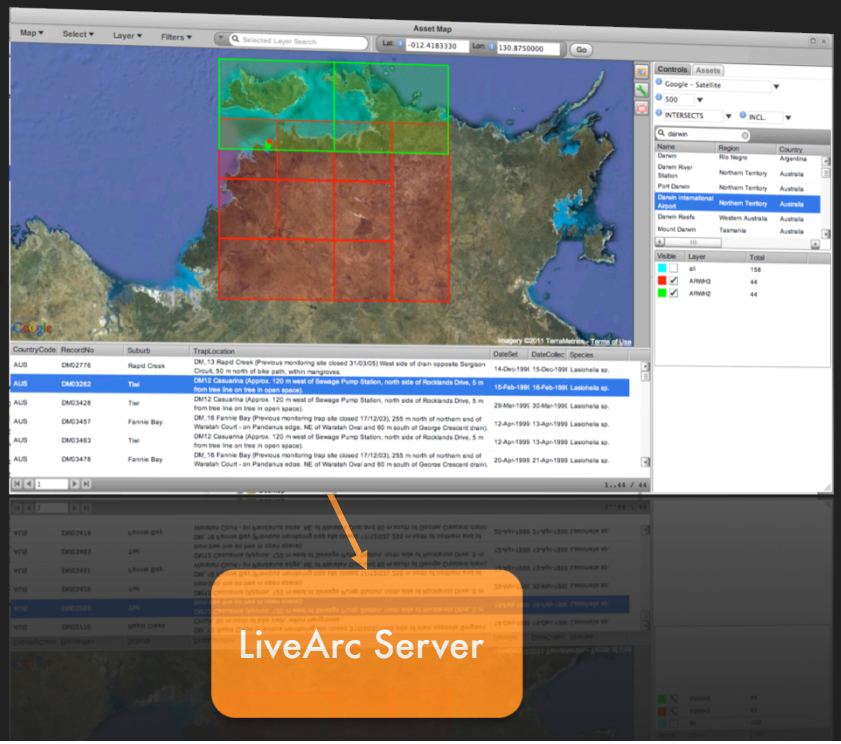

DMF STORAGE WORKFLOW

- DMF
  - Typically tape based
  - Functions optimally with larger files rather than many small files
- LiveArc is designed to deal with large data:
  - Large files
  - Large numbers of files
- Optimise DMF workflow with:
  - File compilation profiles
  - Arcitecta archives
  - Parallel I/O
  - DMF integration
  - Putting it all together

FILE COMPILATION PROFILE

- Simple language to define the layout and content of an external file system tree structure to be processed and ingested
- Patterns define which files to process or ignore, which files to coalesce, and which files are related
- Can assign MIME types and extract metadata, including geo-spatial extents and sidecar XML files
- Groups into an archive file with selectable compression

EXAMPLE PROFILE SPECIFICATION:

```
  construct esri.shape {
            match {
                    /* A 'group' is any set of files with the same name, but
different extensions: */
                  group {
                        file extension "shp"
                        file extension "dbf"
                        file extension "shx"
                        optional file extension "prj"
                }
           }
           logical type "application/esri-shape”
               consume yes
    }
```

ARCITECTA ARCHIVE

- Coalesce many files into a single archive (AAR) – up to 2^63 bytes
- Automatically re-inflate on extraction
- Significantly reduces the number (and potentially size) of files managed by DMF
- Parallel compression and decompression, which drastically improves compression and decompression times on multi-CPU machines
- Can extract individual files from an AAR and transmit to a client without intermediate decompression
- TOC can be extracted and stored locally
- Archives can be split and merged without decompression/recompression

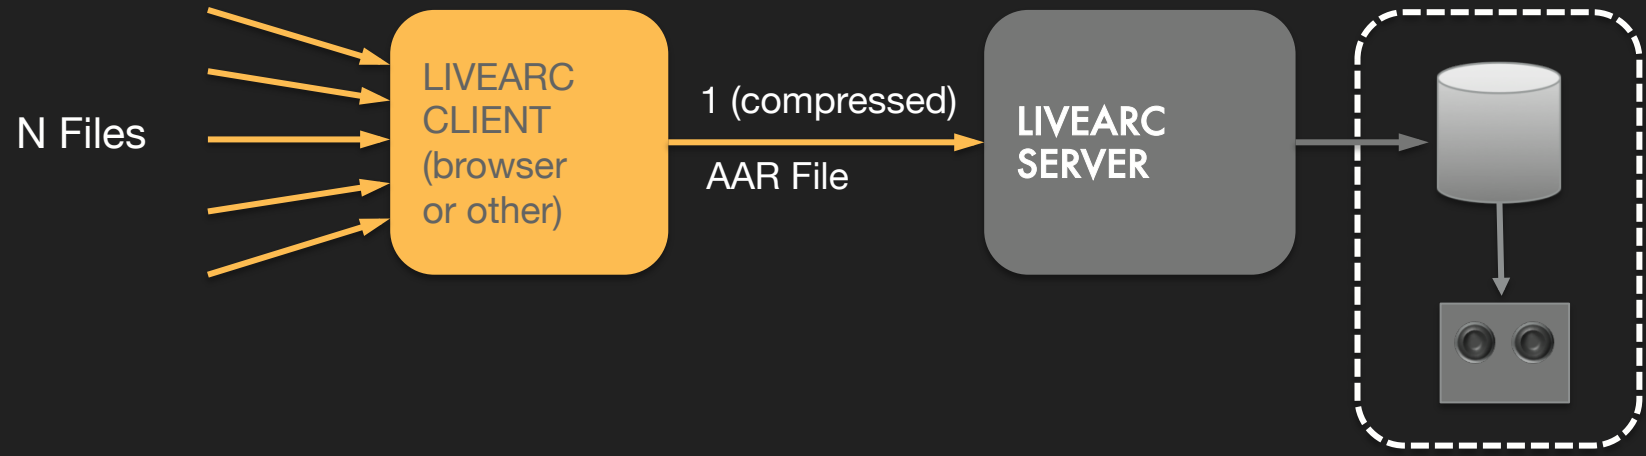

SERVER SIDE ARCHIVES

- Support for zip, tar.gz, ISO images
- However with AAR:
  - TOC can be browsed without decompressing archive
  - TOC can be extracted and stored outside the archive enabling access without recalling the data
  - Specific files can be extracted without decompression and transmitted to the client
  - Check integrity with asset.content.archive.check
  - Convert from one format to another with asset.archive.convert
  - Partial file recall

PARALLEL I/O

ARCITECTA™

- Parallel I/O employed for upload
  - Can specify the number of concurrent packet transmissions
  - Increases performance for wider area network transmission, but also improves performance for local area network transmission by ensuring the network is fully utilised
  - During upload to the server, data can be transmitted and arrive out of sequence – it is properly ordered within LiveArc before passing to the service to receive/process the data
- Both the DTI (browser-based upload) and Aterm utilise parallel I/O for transmitting the data to the server.

DMF

ARCITECTA™

- LiveArc has specific support for DMF:
  - Retrieving status to:
    - Inform end-users of potential delays
  - Migrating data between tiers
- These capabilities are exposed via services:

```
> asset.content.status :id 1
    :asset -id "1"
        :state "online”
> asset.content.migrate :id 1 :destination offline
```

- The following query sends all assets in a particular area off-line

```
asset.query :where \
“geoshape intersects rectangle [(16,143),(16.56,144.89)]” \
:action pipe \
:service -name asset.content.migrate
< :destination offline >
```

POSIX FILE SYSTEM

- LiveArc has a new POSIX File System interface which allows the server to be mounted as a POSIX compliant file system
- Linux and Mac OS/X supported
- Allows direct access to namespaces and assets as though they are directories and files within a file system
- The “shape” of the file system can be determined per user by linking to asset views. For example, for one user, the file system may only present files (“assets”) that have been labelled as PUBLISHED.

PUTTING IT ALL TOGETHER

INGEST

1. File compilation bundles files into an archive file
2. Compress (multi threaded)
3. Transmit with parallel I/O
4. Extract TOC to XODB
5. Commit to DMF managed store

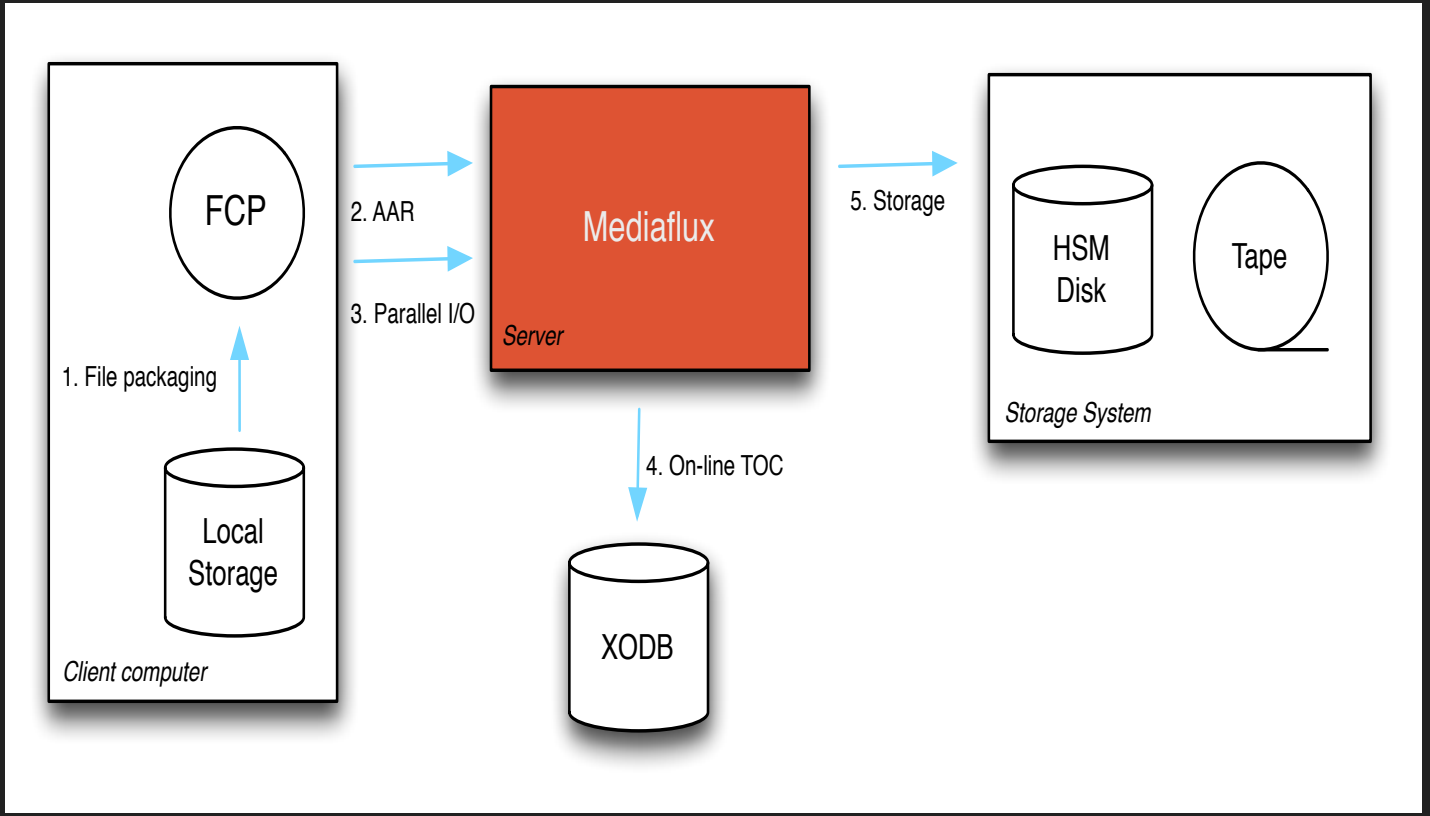

PUTTING IT ALL TOGETHER

EGEST

1. Search for assets - finds archive
2. User requests specific file from the archive
3. Server requests byte range from DMF
4. File(s) extracted and transmitted to client
5. Re-inflate at network boundary
6. Alternatively, LiveArc mounted using AAFS and files accessed directly

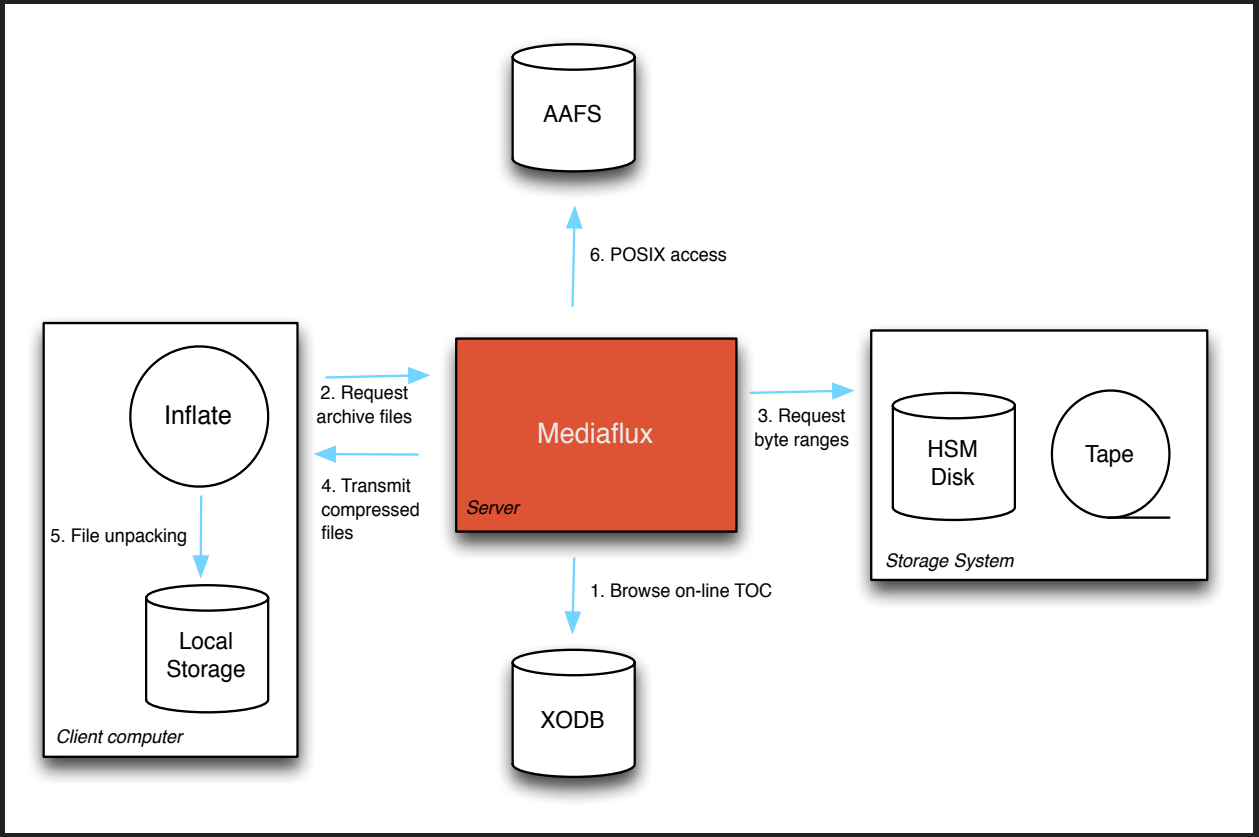

QCMG (@UQ IMB)

INGESTION

ARCITECTA™

Notes:
1. Compression to ~25%, files up to 1TB each
2. All data has a CRC which is checked on replication
3. Tape barcodes recorded in metadata
4. Tape operator notified if tapes running out

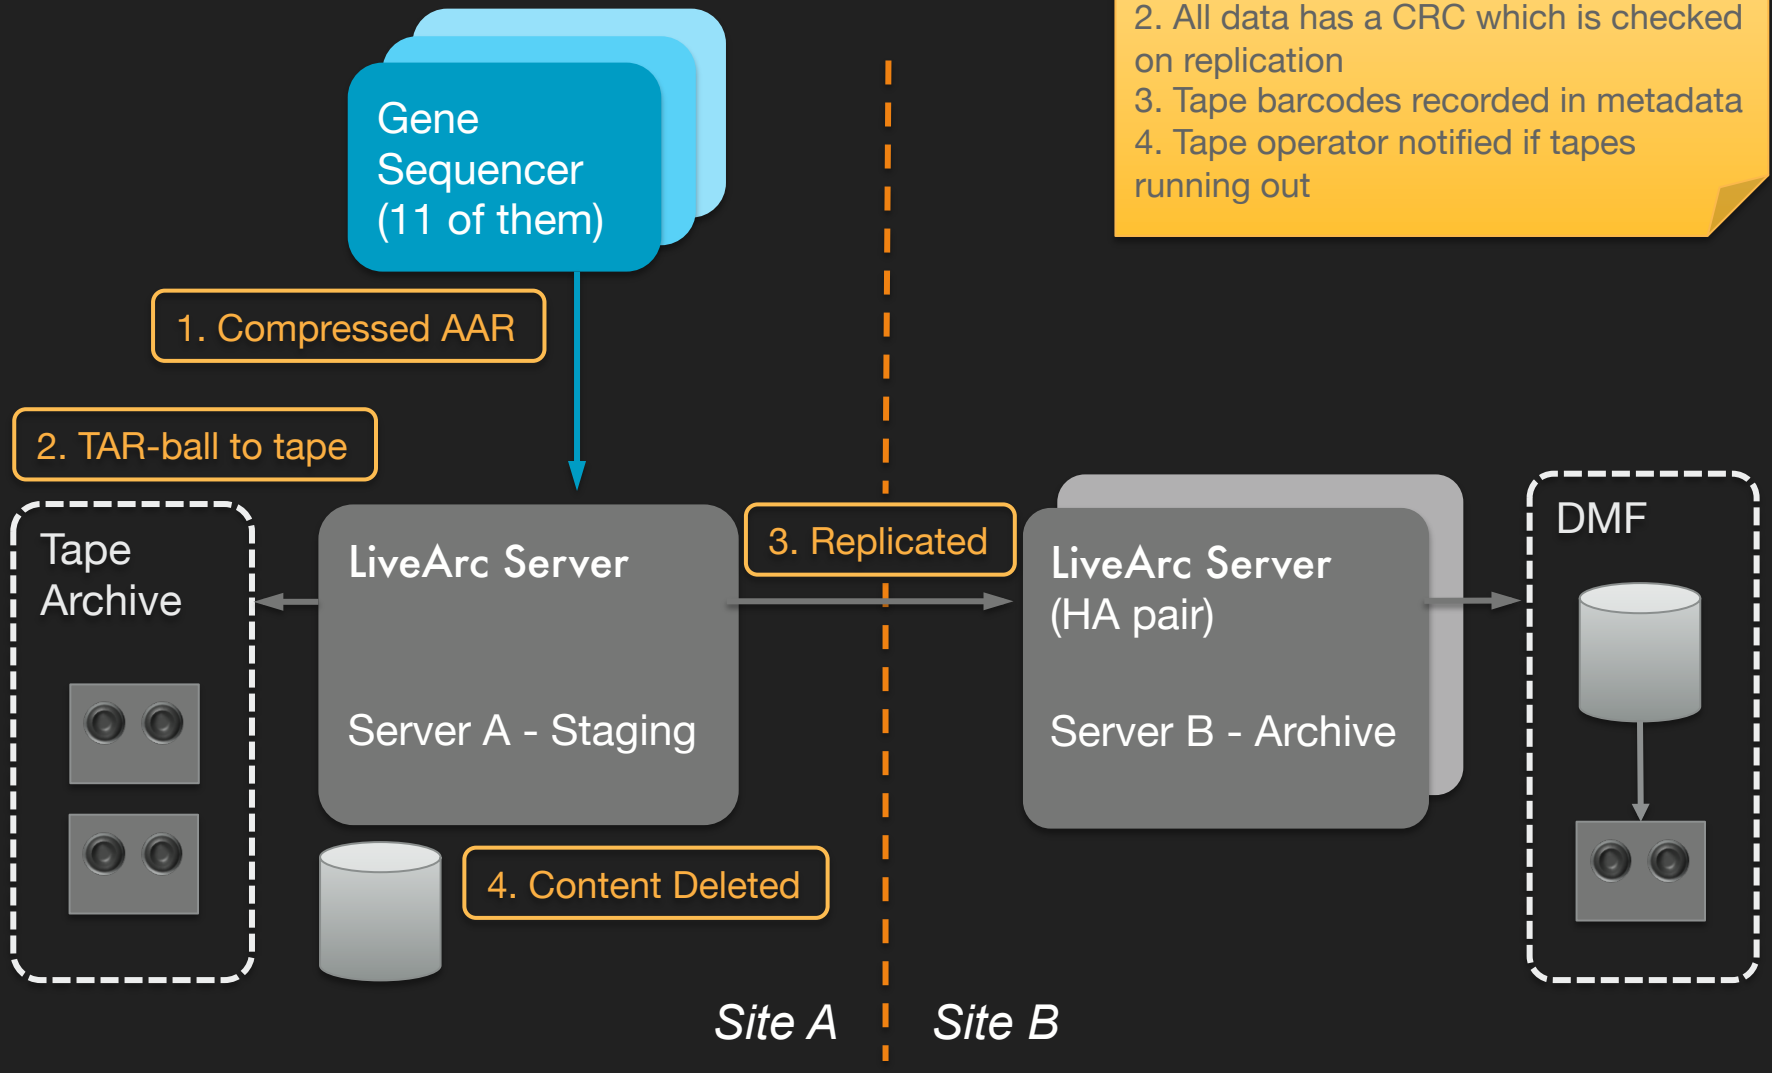

ARCITECTA™