Import

Smart migration tool

15-03-2017

Import
Overview

- Migrate data between storage technology
  - transparently for user
  - easy to use and low overhead for administrators
  - efficient to reduce migration time
  - resilient
  - dynamically tunable

Import : Météo France challenge

A seamless transition from SGI DMF …

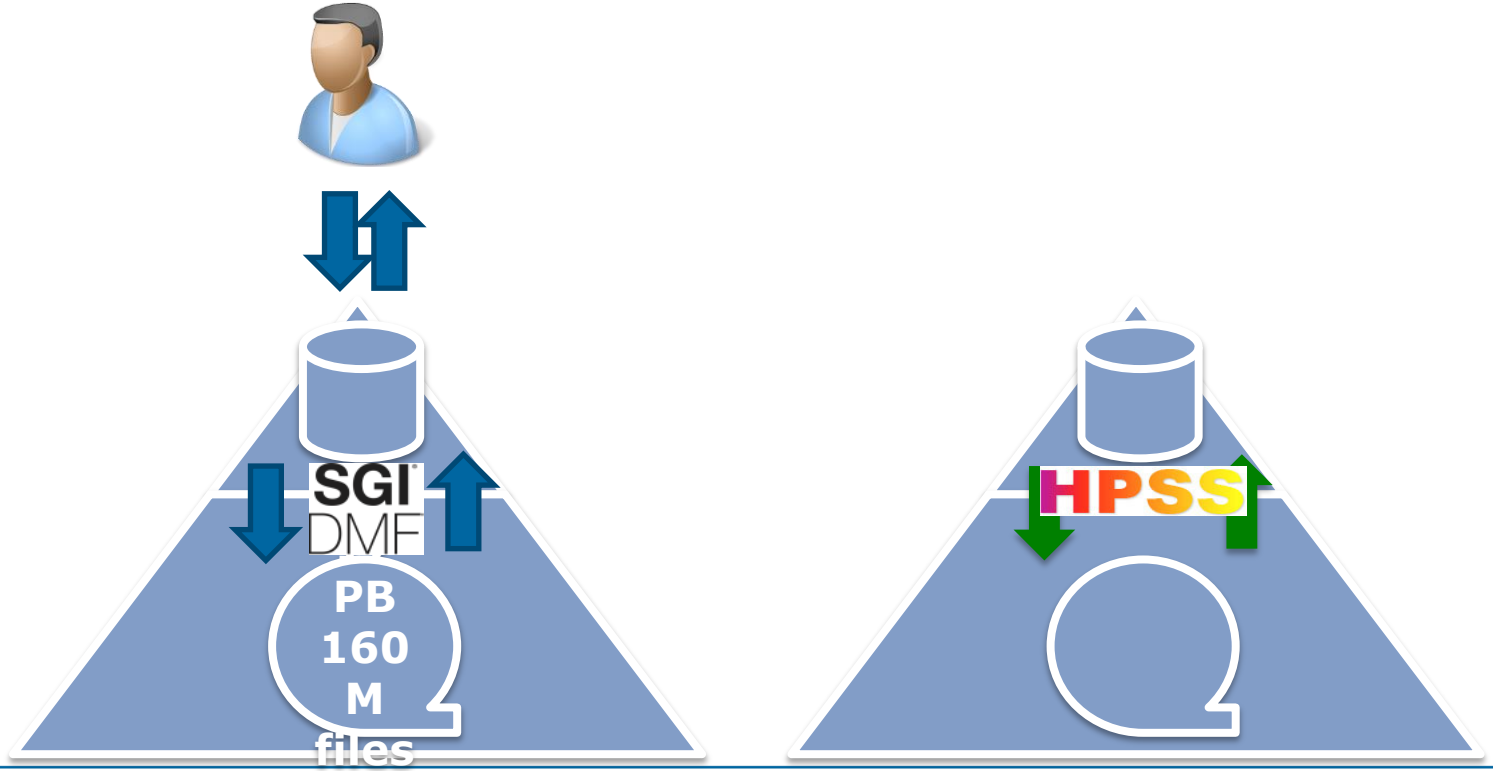

Import : Météo France challenge

… to IBM HPSS

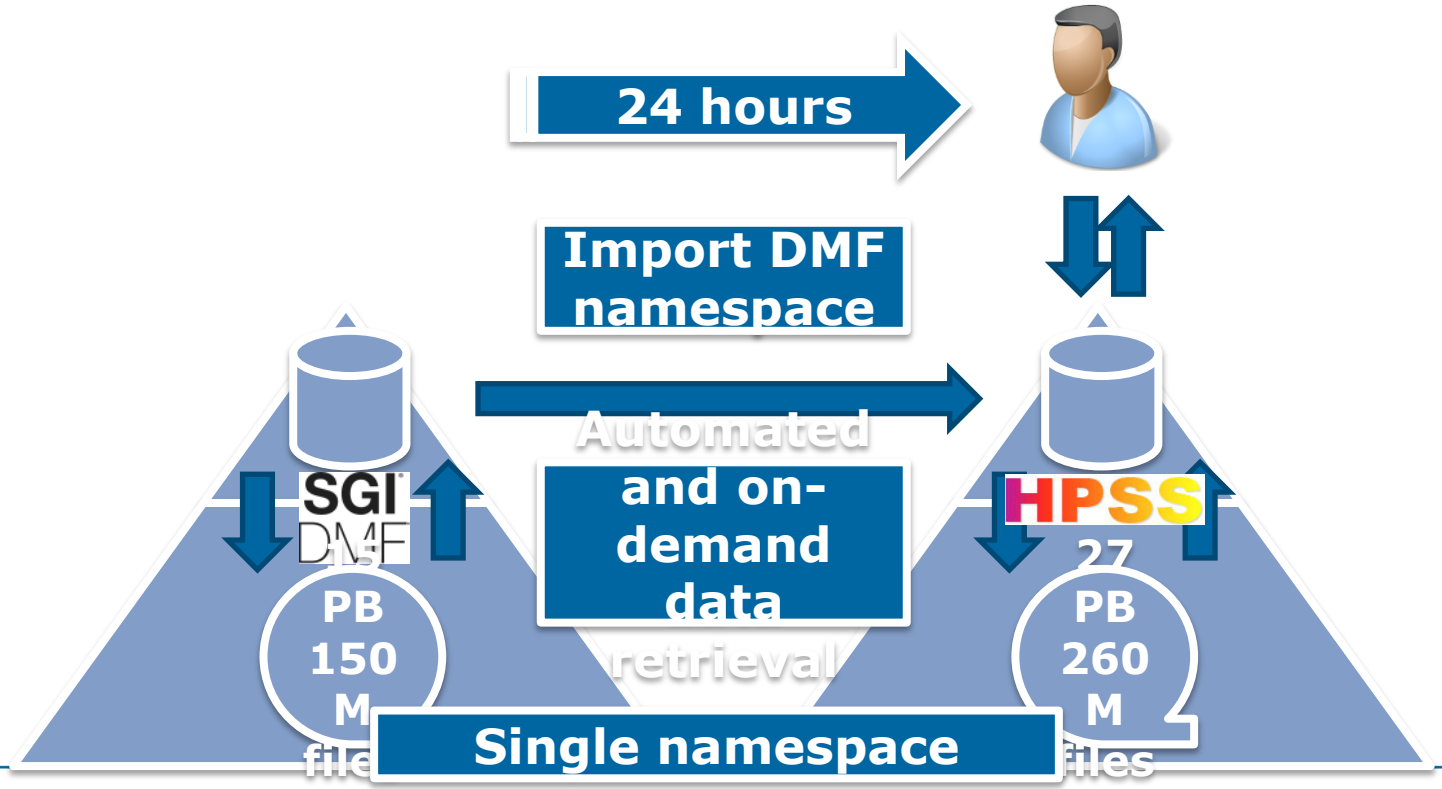

Import : Météo France challenge

Architecture

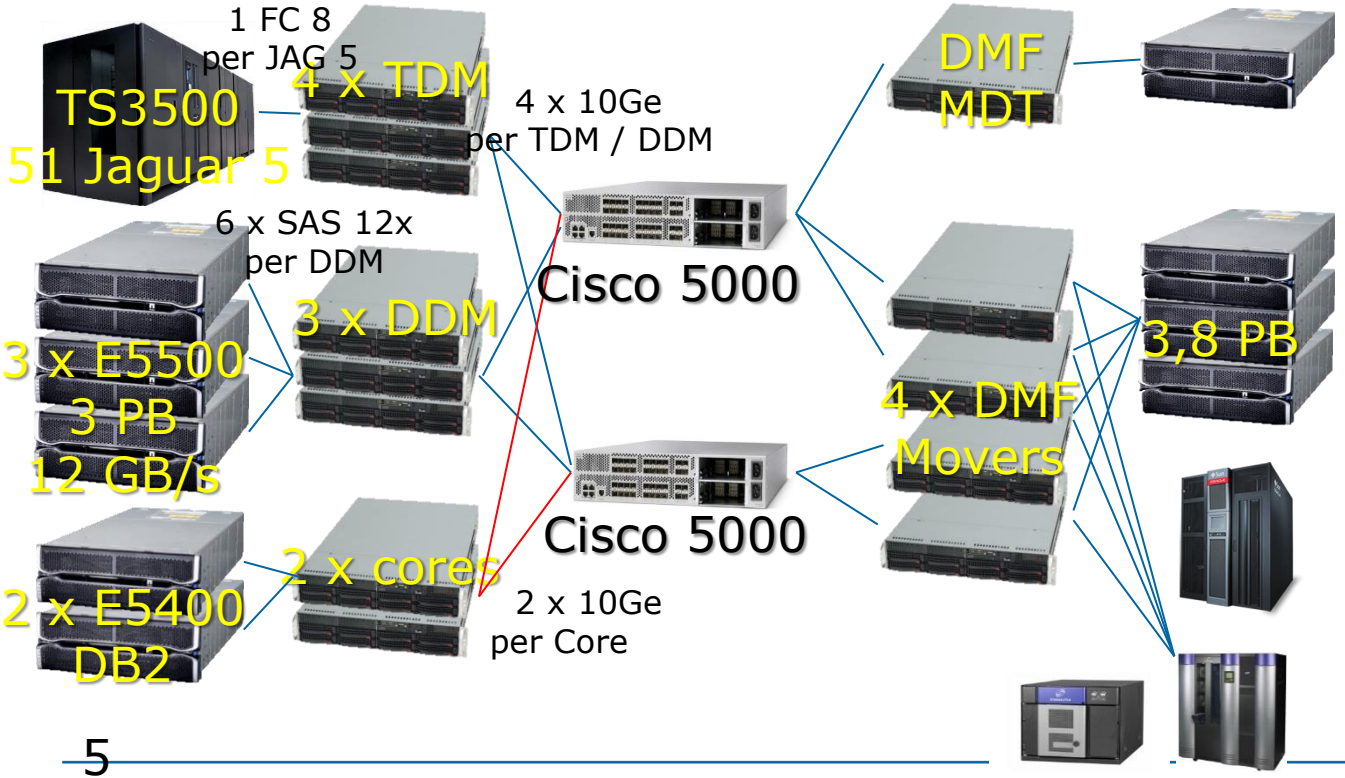

- 18 filesystems
- 4 x NFS servers
  - one filesystem only exported from one server
- 3 libraries
  - sl8500
    - 5 x LTO 4
    - 9 x LTO 5
    - 24 x T10K B
    - 10 x T10K C
  - sl500
    - 3 x LTO 3
  - l1400
    - 4 x LTO 3

Transition steps

Made possible by a joint IBM – Bull effort

- Legacy system is stopped:
  DMF switched to **read-only mode**

Metadata ETL

- DMF **Metadata Extract**: scripts
- **Transform tool**:
  metadata ➔ JSON-format file
- **Metadata Load** in DB2 using the
  Metadata Conversion Utility tool (new table)
- Perform **sanity checks**

Bull atos technologies

IBM

Production downtime: <24 hours

- New system **enters in production**:
  new data is written to HPSS

Data migration

- **Foreground migration**:
  triggered by user requests
- **Background (smart) Migration**:
  optimized to reduce tape mounts
- All client interaction with HPSS are funneled
  through the HPSS Gatekeeper (new site policy)
- **Manual migration** for 0.0007% remaining files

Bull atos technologies

IBM

<1 year

Atos

Transition steps

Details

- Parallel scanner
- collect posix metadata + DMF specifics (BFID, site)
- Météo France run
  - Ran on DMF MDT
    - 2 x Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz, 16 Cores HT
    - 512 GB of memory
  - Started with :

```
# ./scan -t32 -a allfiles.output /chaine /chaine2 /chaine3 /climper /cnrm1_a_mrgu
/cnrm2 /diapason /divmet_previ /divmet /exterieurs /gmap_coope /gmap_obs /gmap_proc
/gmap_recyf /gmgec_eac /gmgec /mercator /preserve
```

    161053555 objects scanned in 33 minutes (576 threads)

Transition steps

Details

Extract DMF Metadata → JSON generator → DB2 Import → Data migration

- Decorate each file with the correct HPSS COS
- Generate a JSON file as specified in the DB2 import tool manual
- Can use data from DMF database to enforce COS decoration
  - objects are stored as they are, not necesseraly how they will based on the dmf.conf
- Météo France run
  - `# ./scan -T -i ./allfiles_cougar.out -s ./dmcatadm_cougar.out -v ./dmvoladm_cougar.out -t 4 -j cougar.json`

    344 seconds (6 minutes) to generate the JSON file of about 40 GB. In this case the performance was limited by the read of input file and write of output file, not processing which is highly paralyzed

Transition steps

Details

- Scan the JSON file format, generate DB2 records compatible with HPSS and import them into HPSS schema
- Generate a dedicated table for gatekeeper (BFID <-> DMF posix location)
- DB2 import miss the capability to import hard links into HPSS
- Météo France run
  - HPSS import tool took 5 hours to ingest those objects into DB2
  - scrub script to recreate hard links (267) took a couple of seconds

Transition steps

Details

- Trigger HPSS mechanism to migrate objects based on the last access time
  - latest accessed files are migrated first
- Know about object location
  - on disk cache
  - on tape, and at which offset

  => only trigger mechanisms for objects stored on disk to avoid HPSS to wait for tape movement
- Purge objects migrated at DMF side to make some space on disk cache
- Purge objects migrated at HPSS side when migrated to avoid standard purge mechanism to remove 'recently' accessed user file in place of 'recently' migrated but 'old' accessed files
- Monitor DMF cache usage and schedule object recall based on
  - access time
  - cache location
  - tape drive availability (limit the number of drives used for automatic migration to let user request have a chance to be proceed)
- Handle shelves tape
- Transactional mechanism for high resiliency (backed by the gatekeeper transactional mechanism)

Transition steps

Details

- Météo France run
  - Metadata import took 3 hours (ran during DB2 import)
  - as simple as running ./import !
    - fine tune thread count per file family to adjust load at HPSS & DMF side
  - require someone to check, on a daily basis
    - that import was still running
    - collect statistics and paste them into an excel spreadsheet
  - started February, 10th 2015, completed September, 9th 2015
    - latest month mostly used to try to recover about 200 files that where located on unreadable tapes.

Import : Météo France challenge

Files migrated over time

average of : 706 611 files / days

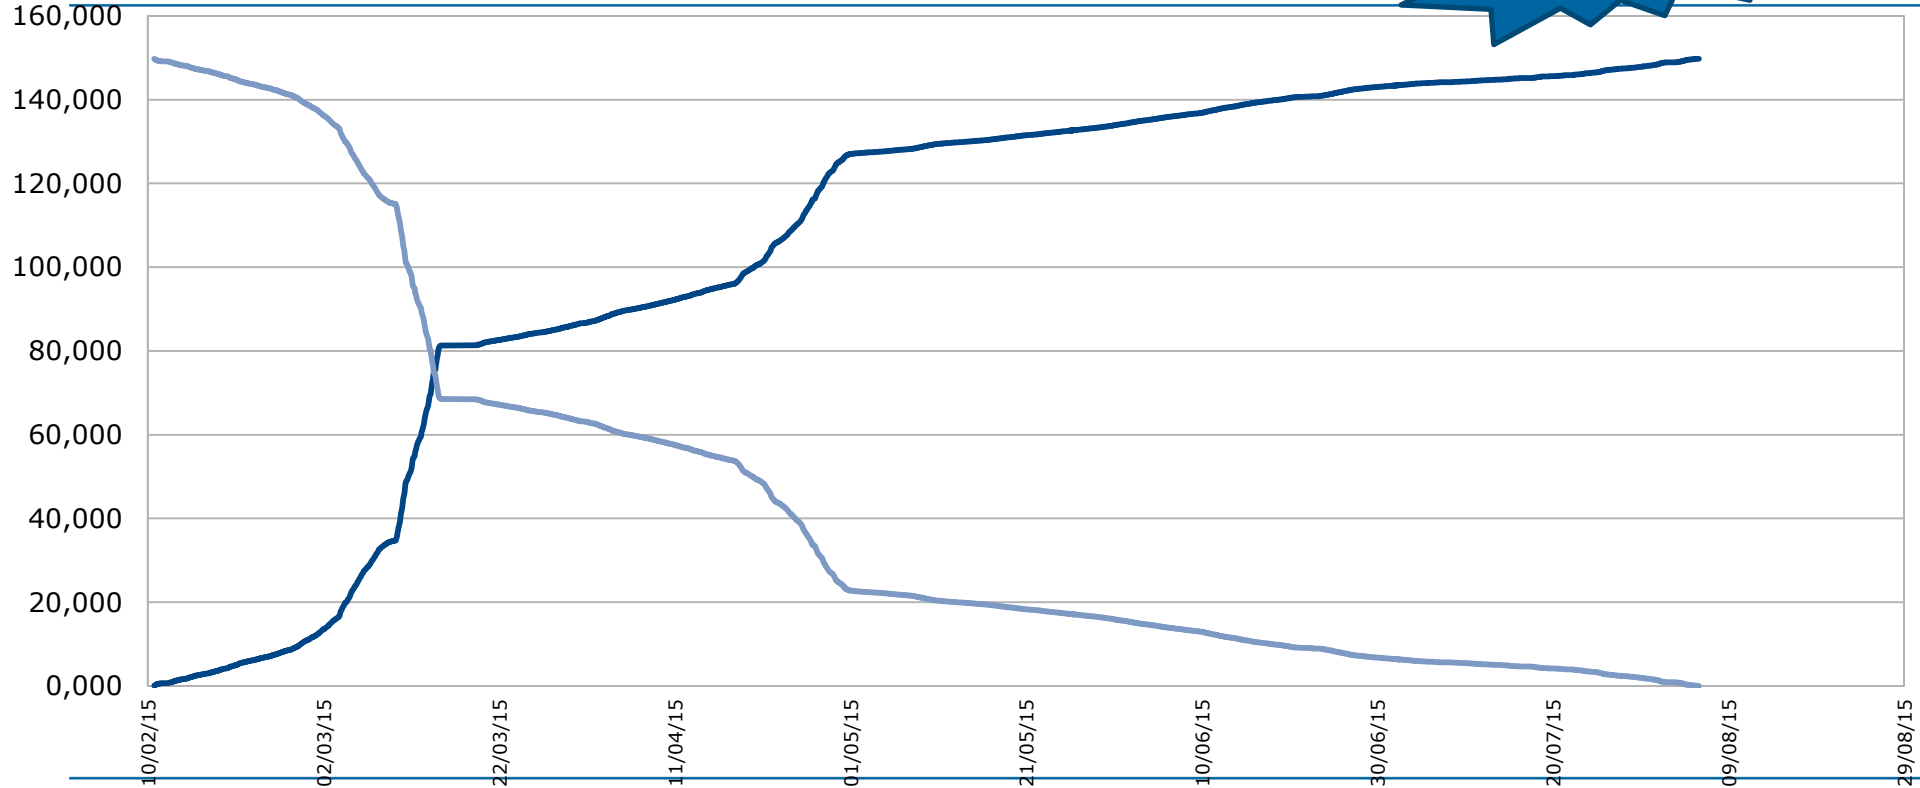

Import : Météo France challenge

Terabytes migrated over time

average of :
- 89 TiB / days
- over 1 GiB / s

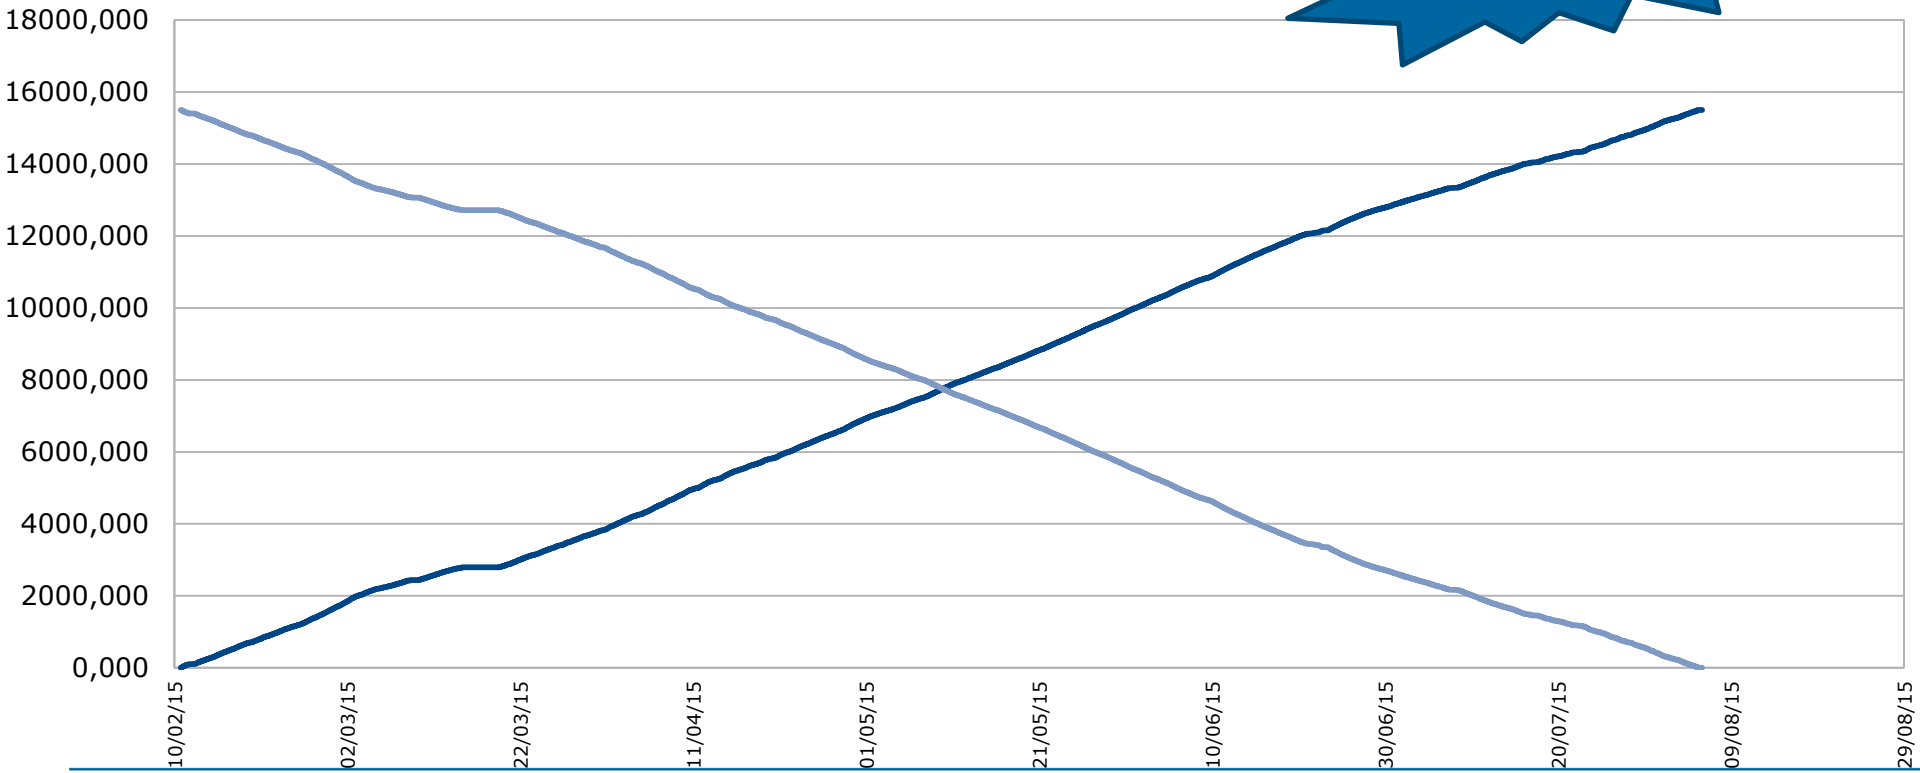

Import : no specific to HPSS

- Import is not specific to HPSS
- Has been used for
  - DMF to HPSS conversion
  - DMF to GPFS conversion (using AFM)
- Can be used for
  - any POSIX or HSM FS to GPFS (using AFM)
- Can be modified for other use.

Thanks

For more information please contact:
T+ 33 1 98765432
F+ 33 1 88888888
M+ 33 6 44445678
firstname.lastname@atos.net

Atos