Storage and Services:

Planning and managing cost, quality and risk

Martin Hall-May

IT Innovation Centre

18 November 2011

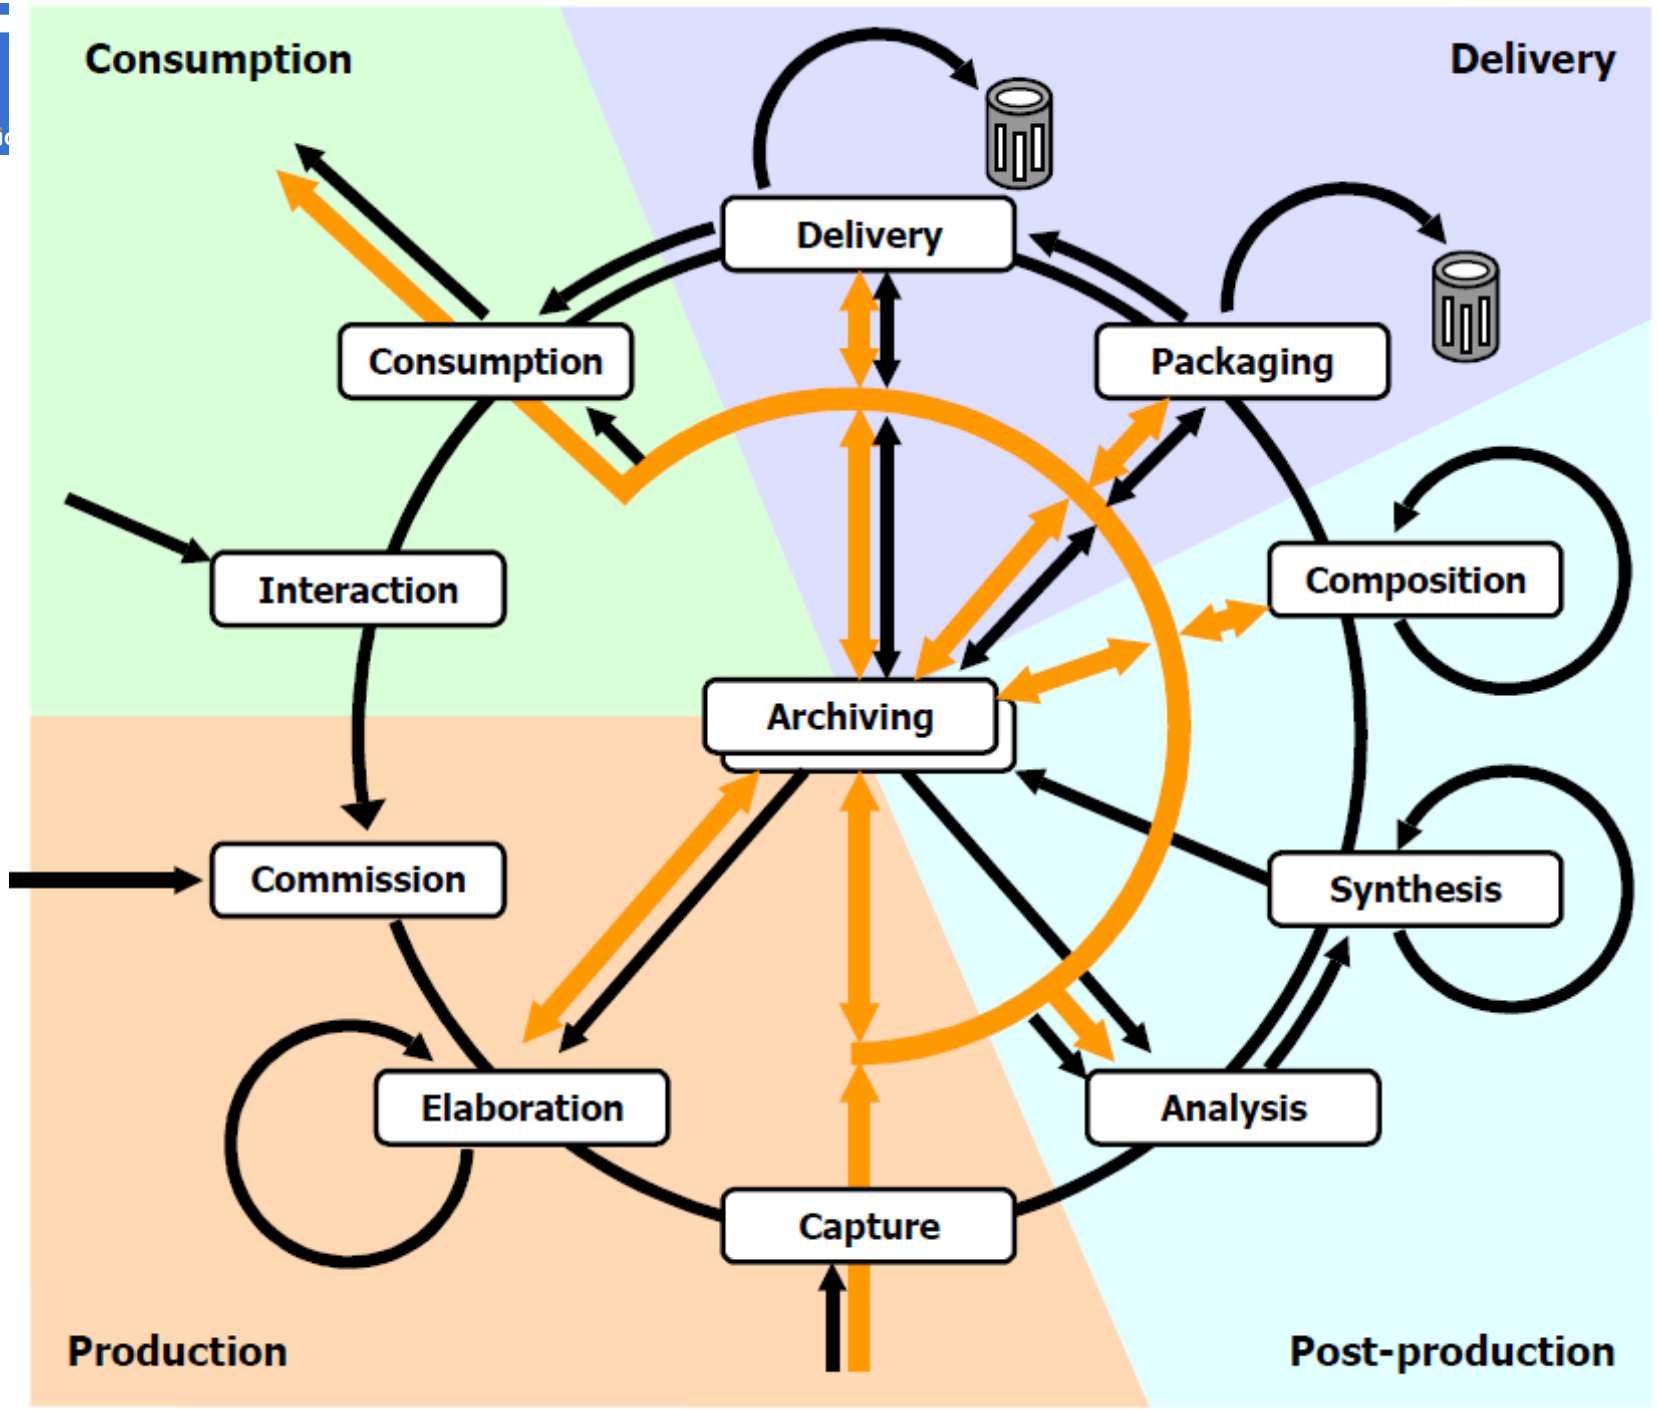
Consumption
Delivery
Delivery
Consumption
Packaging
Interaction
Composition
Archiving
Commission
Synthesis
Elaboration
Analysis
Capture
Production
Post-production

An archive provides *services*

- Ingest, access
- Safe storage
- Formats
- Metadata
- Rights

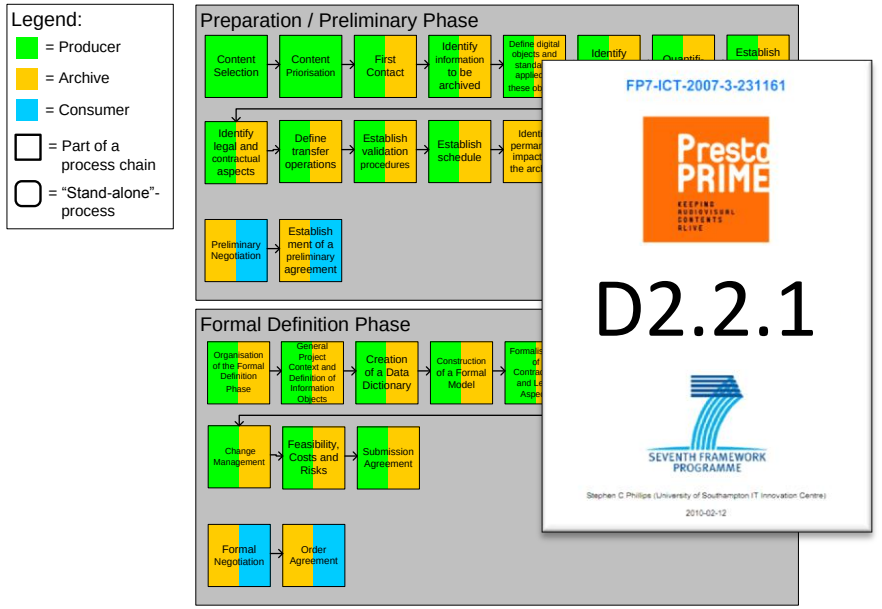

- They all cost money
- They all take time
- Never enough of either!

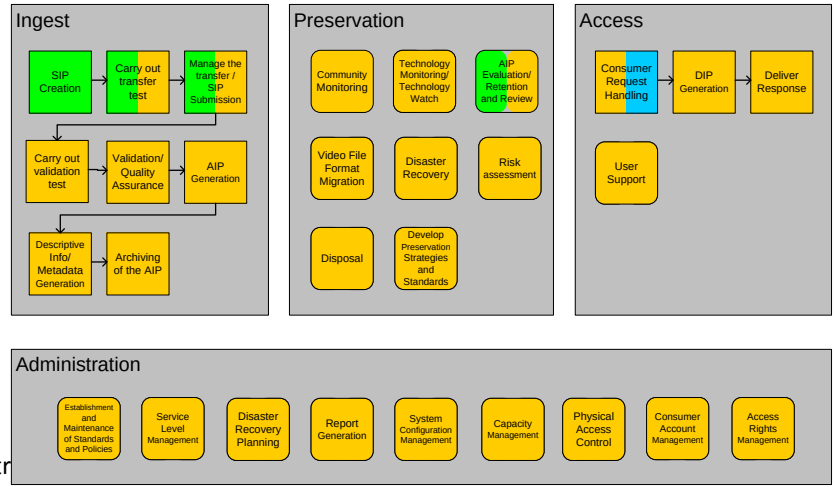

Services have location and responsibility

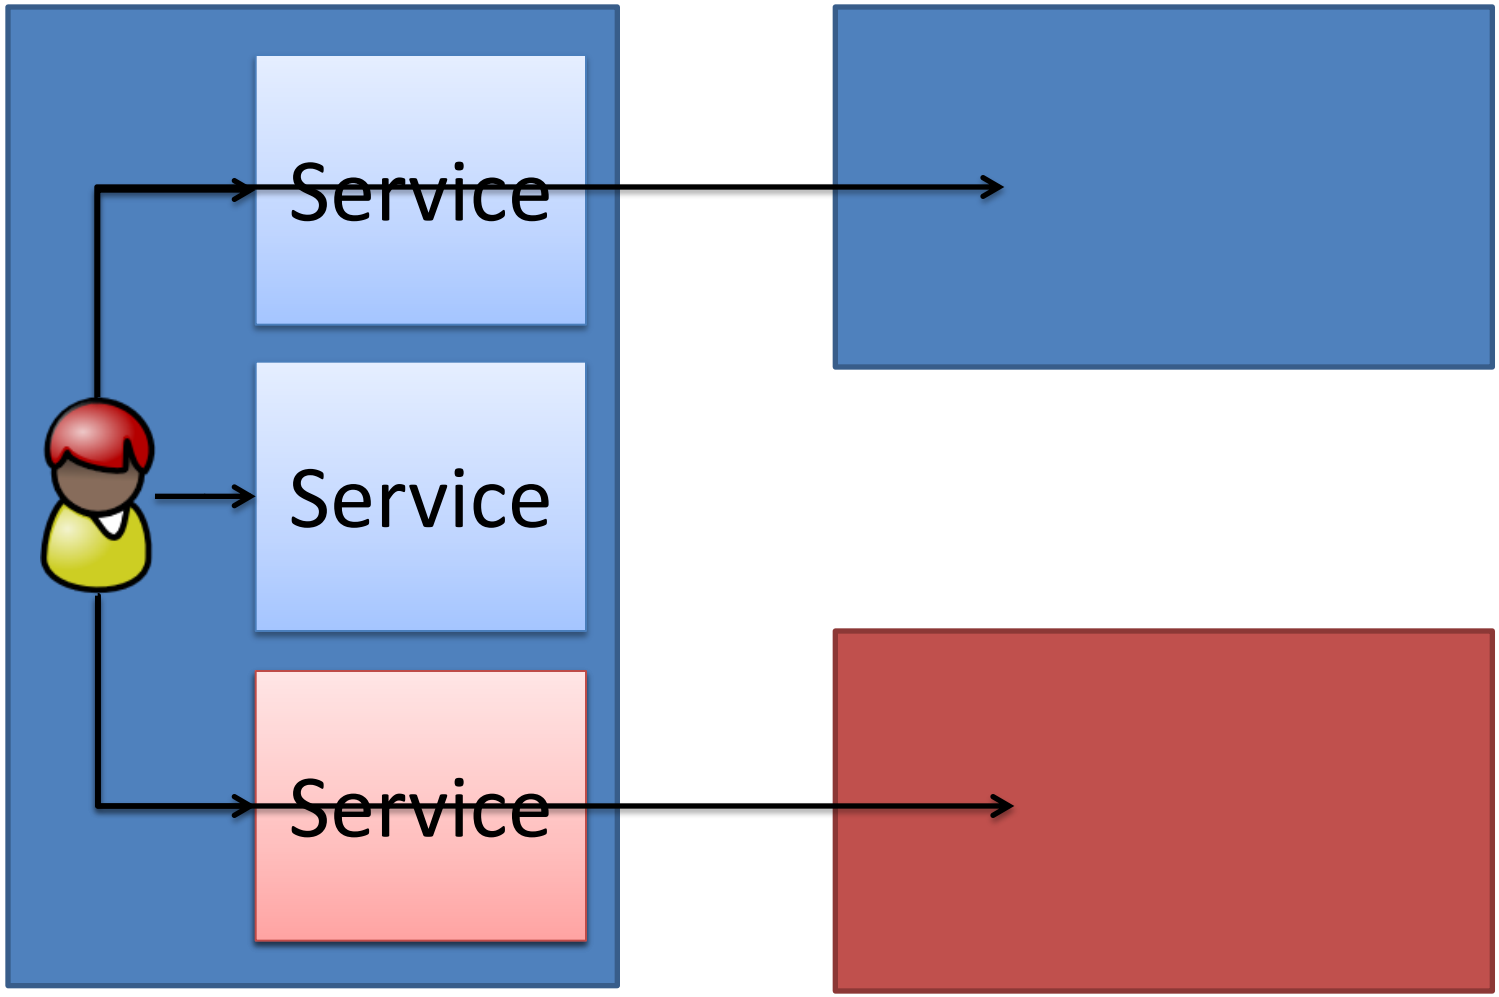

FP7-ICT-2007-3-231161

D2.3.1

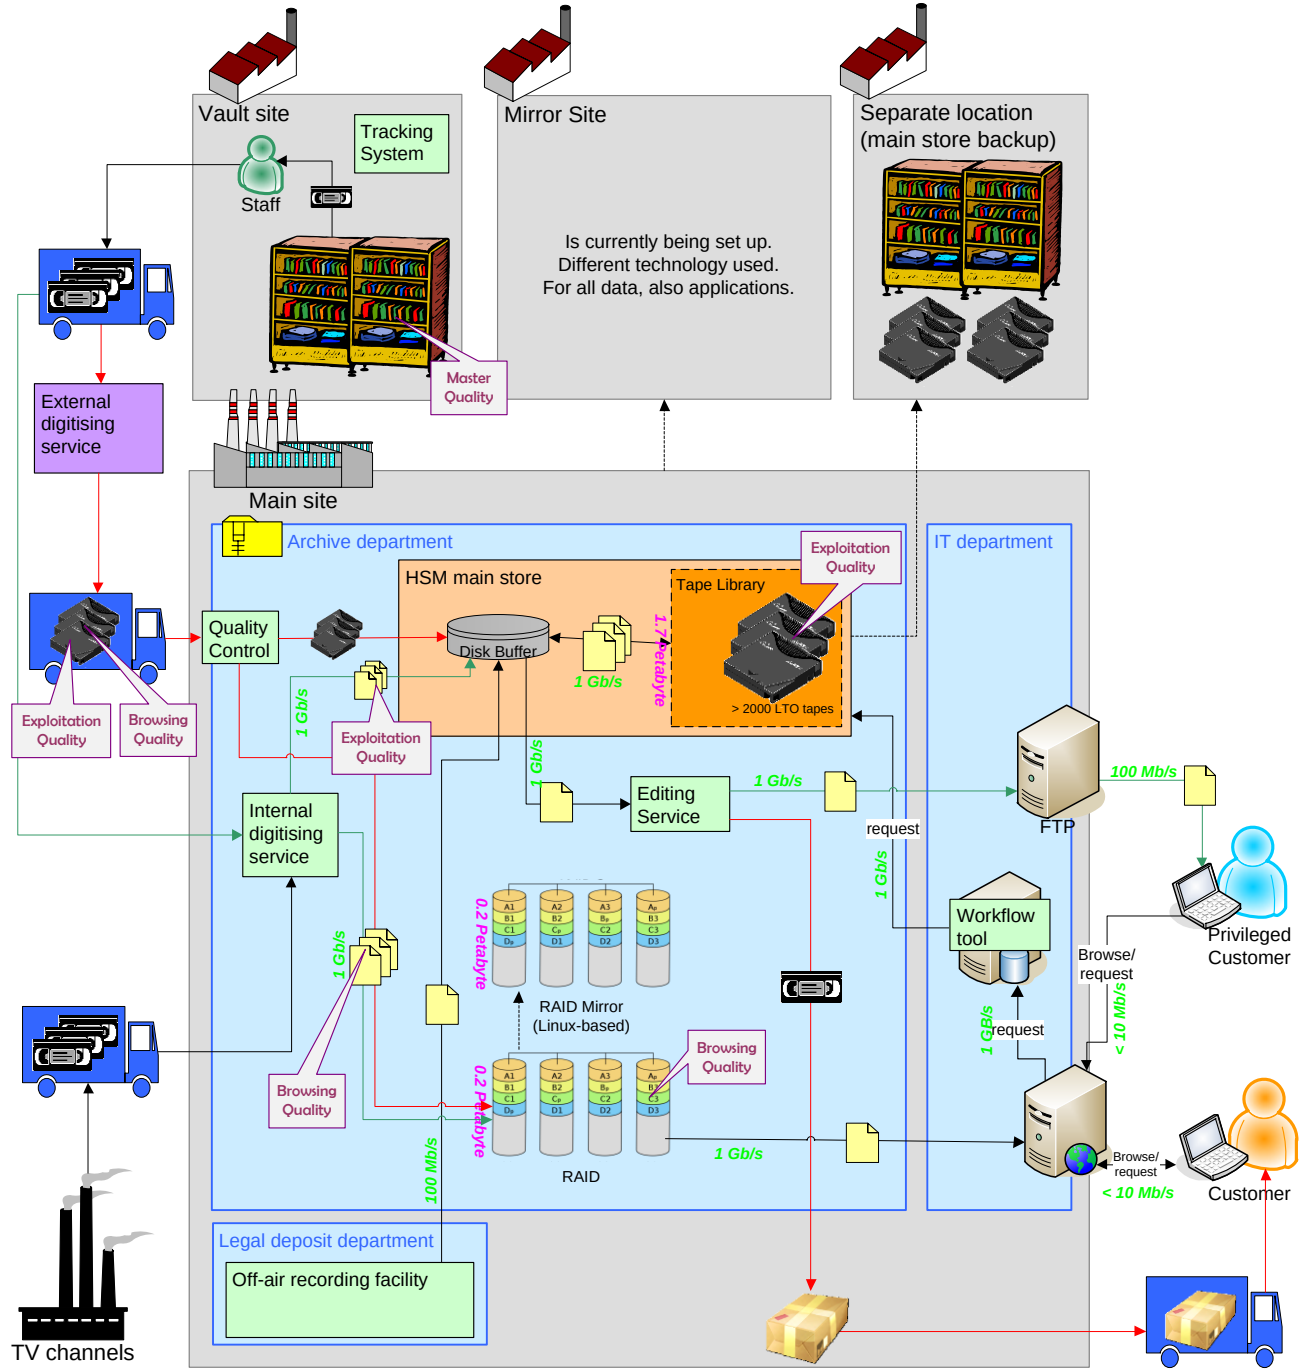

Services need planning and managing

- Service Level Agreements (SLA)
  - What the service does
- Quality of Service (QoS)
  - How well it does it

- If you can't **measure** it then you can't **manage** it
  - Throughput
  - Quality
  - Cost
  - Risk

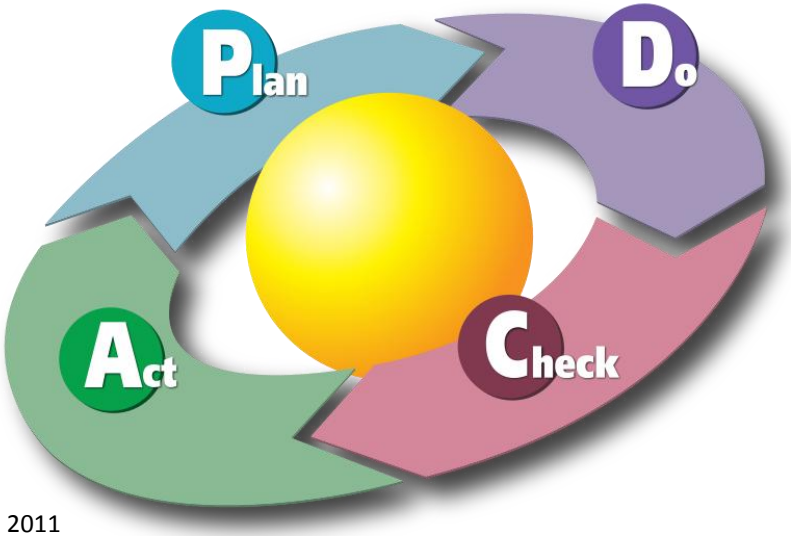

Planning and managing includes compromises

- Volume
- Quality
- Deadline
- Budget

- Digitisation workflows
  - Cost, throughput, quality
  - E.g. QC v.s. automation
- Storage strategies
  - Cost, risk of loss
  - E.g. copies v.s. cost
- Online access services
  - Cost, QoS, Users
  - E.g. KPI v.s. customers

Storage SLA Terms

- Availability
- Integrity / Safety
  - How to measure?
- Ingestion time
  - Indexing, generating access copies
- Search time
- Delivery time
  - From request to start of delivery
- Bandwidth
- Subscription fee
- Charge for data
  - On disc
  - Ingest
  - Access
- Charge for CPU
- Charge per user
- Maximum storage size
- Maximum number of users

Availability

I need the service to be available almost all the time

Can you be more specific?

I need the access service available 99.9% of the time

Is that measured over a day, month or year?

The access service must be available 99.9% of the time each month

That's 43 minutes of downtime each month – what if that's all in one go one afternoon?

When can maintenance be done on the service?

Do you want different uptimes for day and night?

When is "daytime" for an international operation?

Access Time

I want a good response time for delivery

For every delivery, or on average?

The average delivery time must be less than 5 minutes

Average over what period?

The average delivery time each month should be less than 5 minutes

What about the network connection – doesn't that affect it?

Files must be ready for download in less than 5 minutes (on average, per month)

What about big files? Won't they take longer?

Files must be ready for download in less than the size of the file in GB + 15 minutes, on average, per month.

Data Safety

I want everything returned in perfect condition all the time

That's hard: at least three copies and an active management system

OK, maybe I can lose a bit, but not too much

How much risk are you willing to take?

Are we talking about losing files, parts of files or a few individual "bits"?

Or, are you talking about programmes, scenes, shots, frames..?

What about a certain number of programme-minutes at risk per year?

Is losing 1%-3% of the archive over 20 years acceptable?

Maybe, but what's the chance that I could lose more than that?

OK, let's run a model to look at the options

Storage

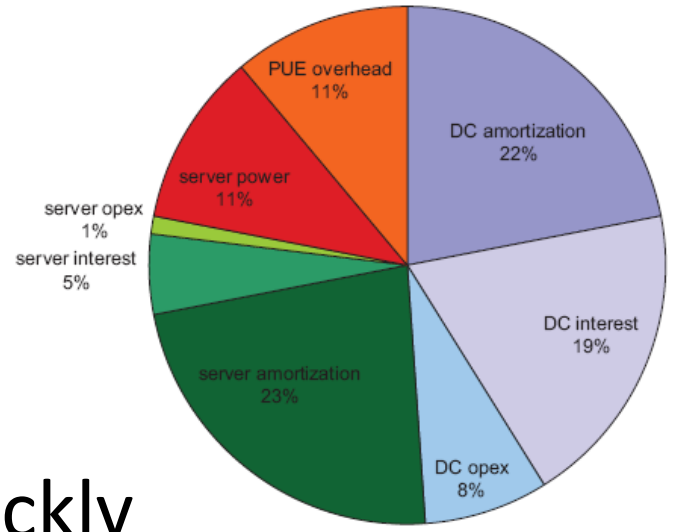

- Is not 100% safe
- Becomes obsolete quickly
- Total cost is high, but falls quickly
- Fast access and safety don't always go together

Medium	Storage Density bits/cm²	Life, years
Stone	10	10000
Paper	10^4	1000
Film	10^7	100
Disc	10^{10}	10

Many storage choices

- **Longer lived** storage technology
  - E.g. Printing bits to film
- **More reliable** storage technology
  - E.g. data tape instead of HDD on shelves
- Make **more copies**
  - E.g. off site deep archiving
- Encode to **make content more resilient**
  - E.g. Graceful behaviour if a few bits and bytes are corrupted
- **Conceal errors**
  - E.g. Interpolation of corrupted frames or blocks
- **Check often** and **fix quickly**
  - E.g. ‘scrubbing’ of HDD servers

Comparing ‘cost of risk of loss’

- Diversity (copies) keeps things safe
- Active management of data integrity
- Migration to address obsolescence
- All activities have a cost, especially **access**

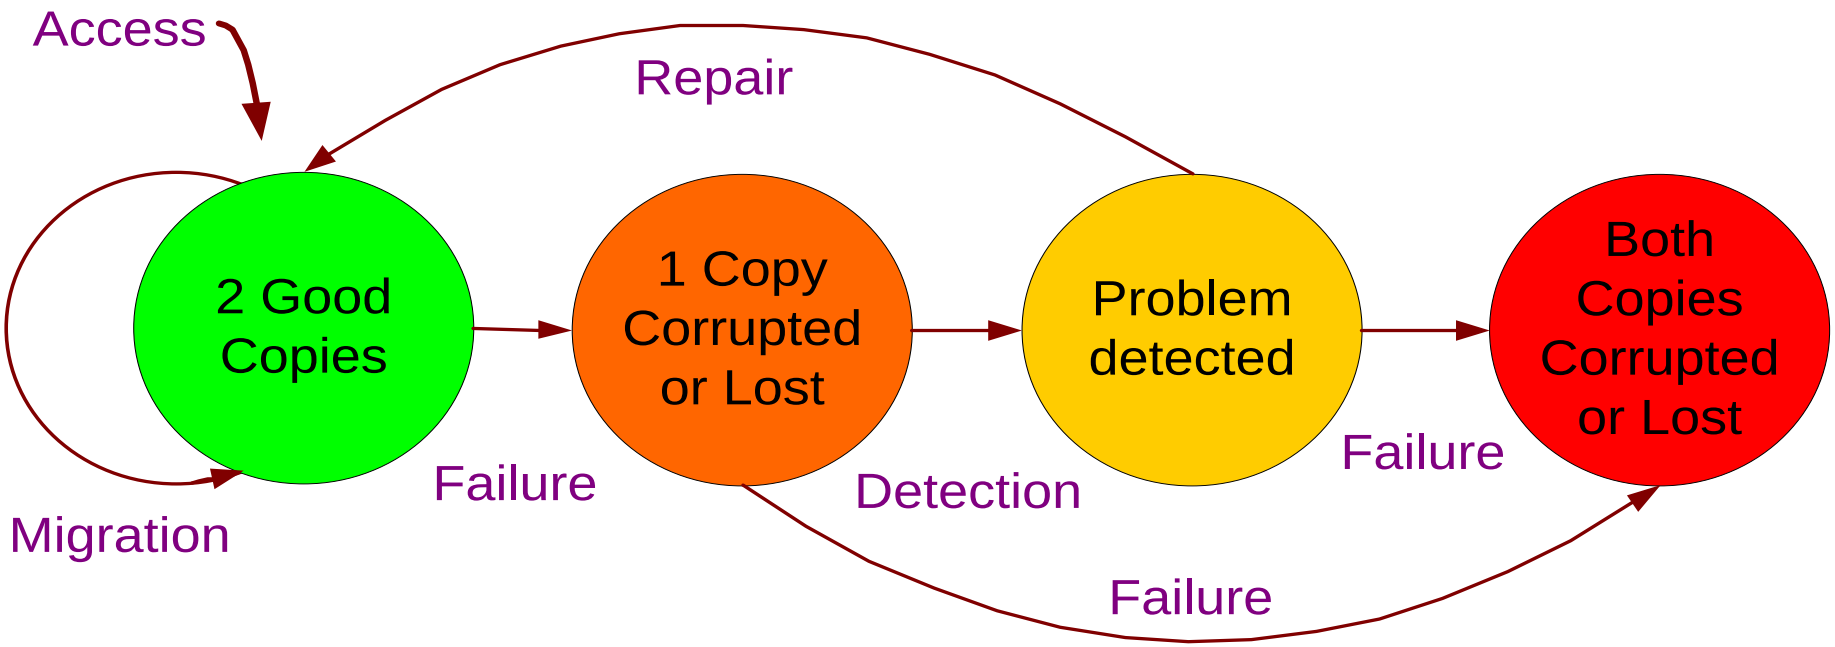

Cost, safety and access: Simple comparison of IT storage

	Data tape on shelves	HDD in servers	Storage as a Service
Storage Cost	**Low** (media, shelves, climate control)	**High** (servers, power, cooling, maintenance)	**High** (fully managed service)
Access Cost	**High** (people retrieve and load media)	**Low** (internal network, automated)	**High** (bandwidth, charges for i/o)
Latent Failures	**Low** (data tape is reliable)	**Med** ('bit rot')	**Low** (replication and monitoring)
Access Failures	**Medium** (drives eat tapes)	**Low/Medium** (depends on system)	**Low** (automated checks)

Two tools that might help

- Long term planning
  - 25 years
  - High level choices
  - Estimates of total cost and loss
  - Narrow down the options
- Short to medium term simulation
  - Simulates actual events
  - Corruption, loss, catastrophes
  - Ingest, access, ‘active preservation’
  - Impact of limited resources

Storage Systems

Found 5 storage systems. Add...

HDD in servers read-only
Migration required every 4 years.
Running Costs
Access: €0.1 per GB
Storage: €1 per GB per year
Corruption Rates
Access: avg. 1 in 500 files
Latent: avg. 1 in 750 files per year

HDD on shelves read-only
Migration required every 4 years.
Running Costs
Access: €1 per GB
Storage: €0.25 per GB per year
Corruption Rates
Access: avg. 1 in 100 files
Latent: avg. 1 in 500 files per year

Data tape in a robot read-only
Migration required every 6 years.
Running Costs
Access: €0.2 per GB
Storage: €0.4 per GB per year
Corruption Rates
Access: avg. 1 in $1x10^4$ files
Latent: avg. 1 in $1x10^5$ files per year

Data tape on shelves read-only
Migration required every 6 years.
Running Costs
Access: €1 per GB
Storage: €0.1 per GB per year
Corruption Rates
Access: avg. 1 in $1x10^4$ files
Latent: avg. 1 in $1x10^5$ files per year

mystorage Edit Delete
Migration required every 5 years.
Running Costs
Access: €1 per GB
Storage: €1 per GB per year
Corruption Rates
Access: avg. 1 in 10 files
Latent: avg. 1 in 10 files per year

Risk and Loss

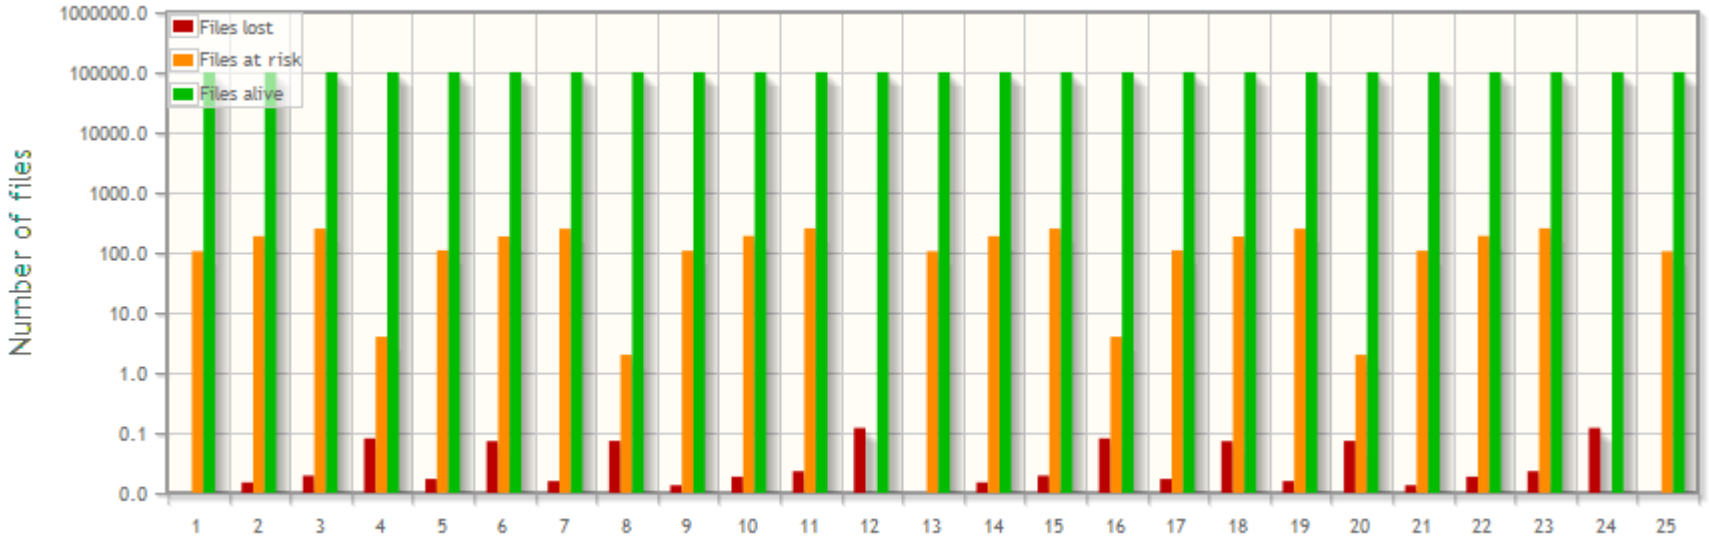

Risk and Loss

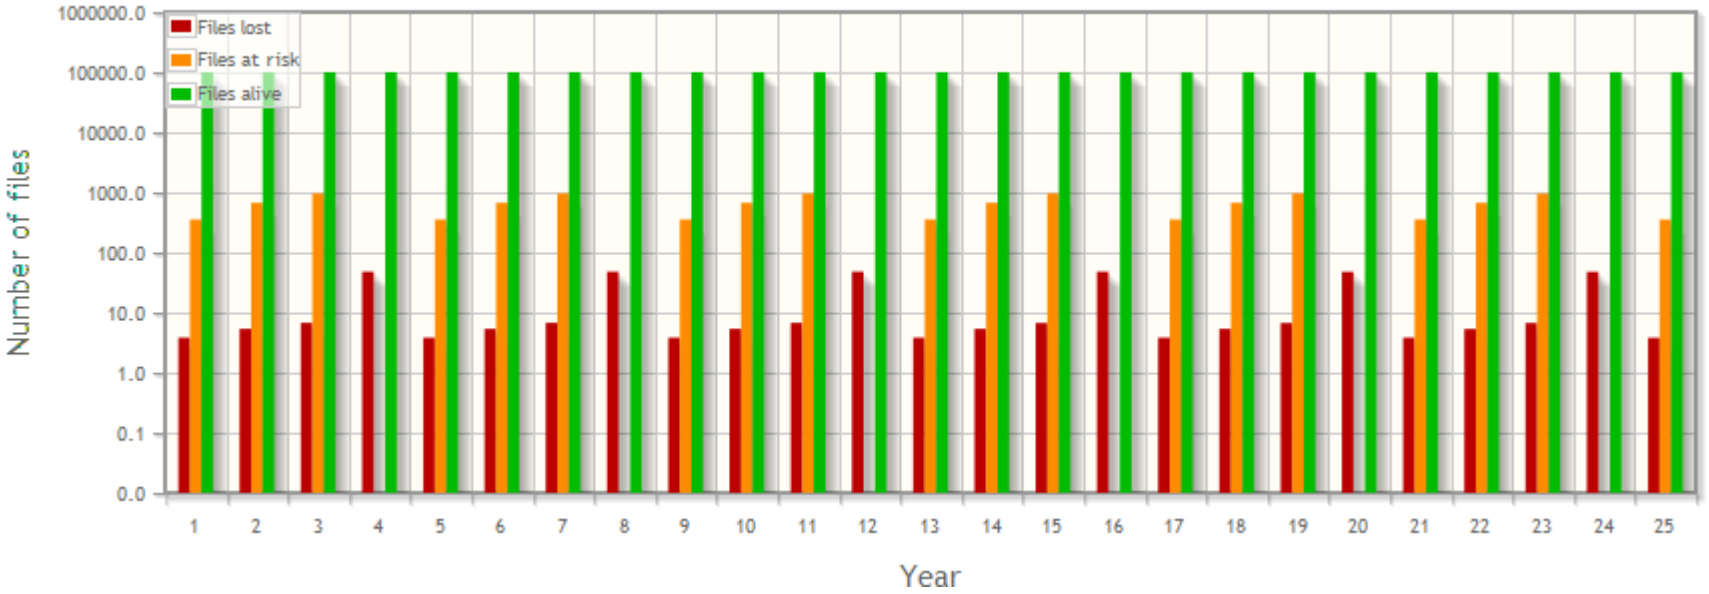

Long term cost

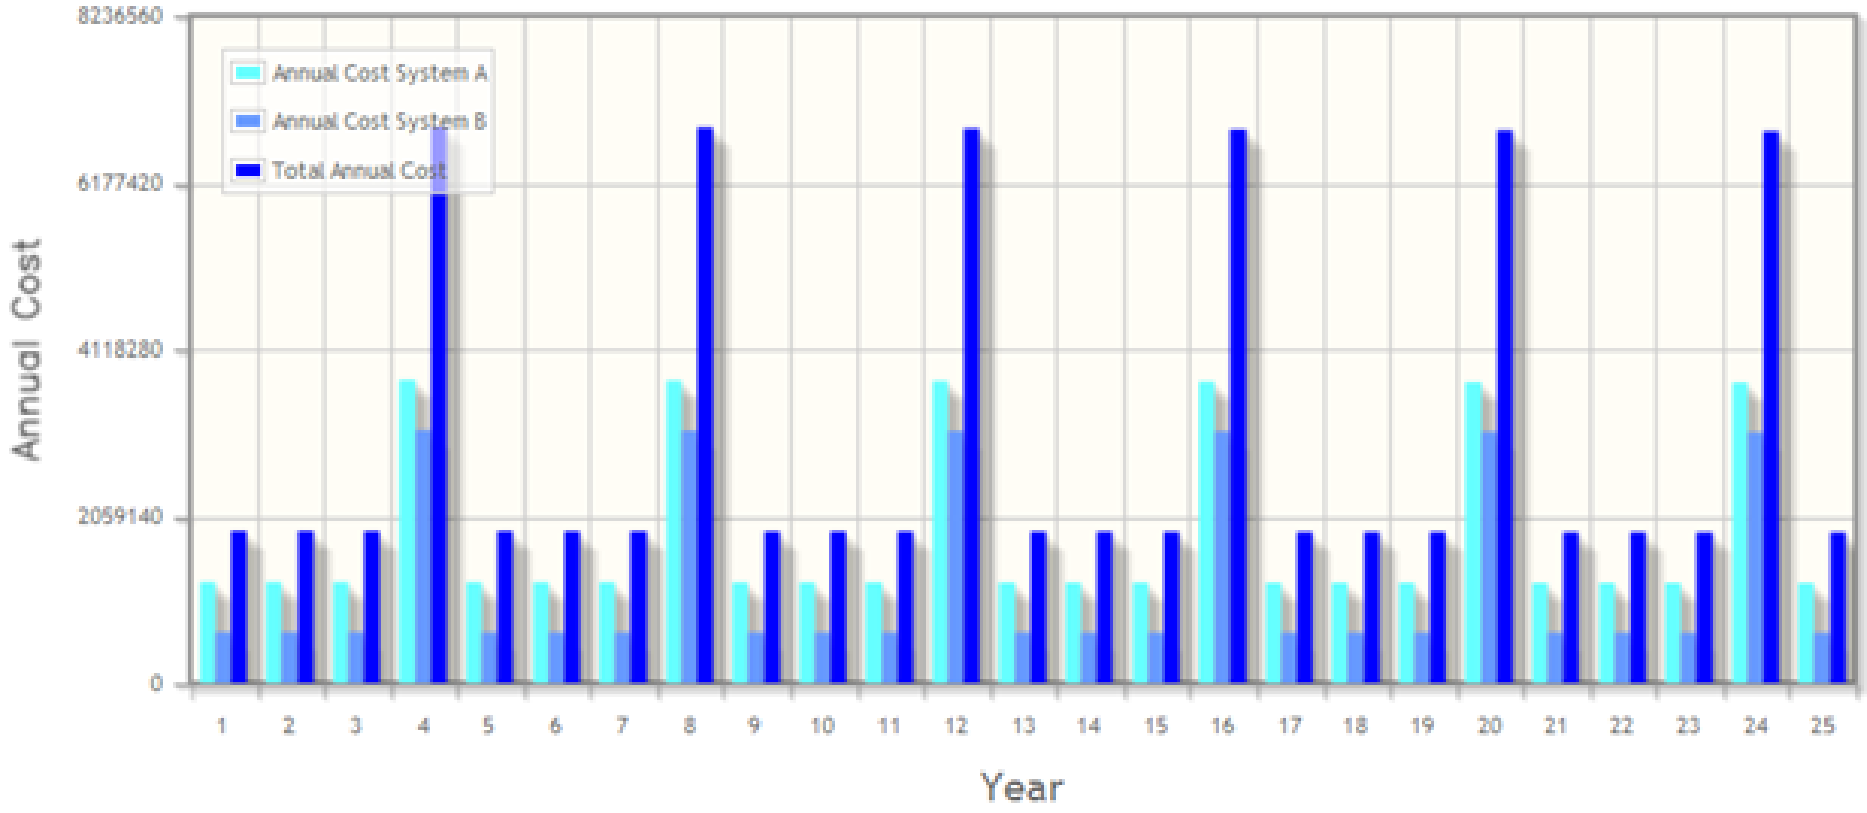

Simulating retention and access

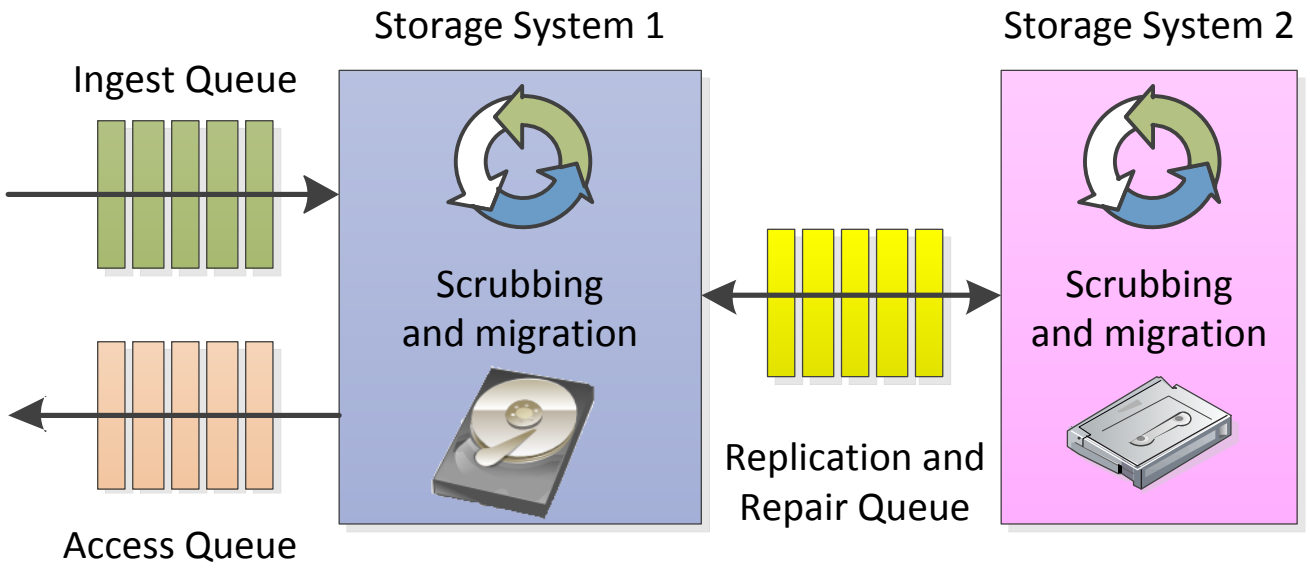

- Resources are often limited
  - People, servers, bandwidth
  - Contention and priorities
- Capacity planning, Disaster simulation, Training

Interactive Storage Simulation Tool v 1.0-beta-1 (IT Innovation Centre)

Global Performance | Ingestion Queue | User Access Queue | StorageSystem 1 | StorageSystem 2

Archive Performance

Simulation time: [Current simulation time: 2016-09-23T25:01Z] Ticks: 251160
Archive running costs (EU): 3.9102938243511766E7
Currently Used Storage (TB): 12455.5
Total assets: 62275 Files: (124555)
Assets at Risk: 779 Lost Assets: 25

Running Cost | Used storage space

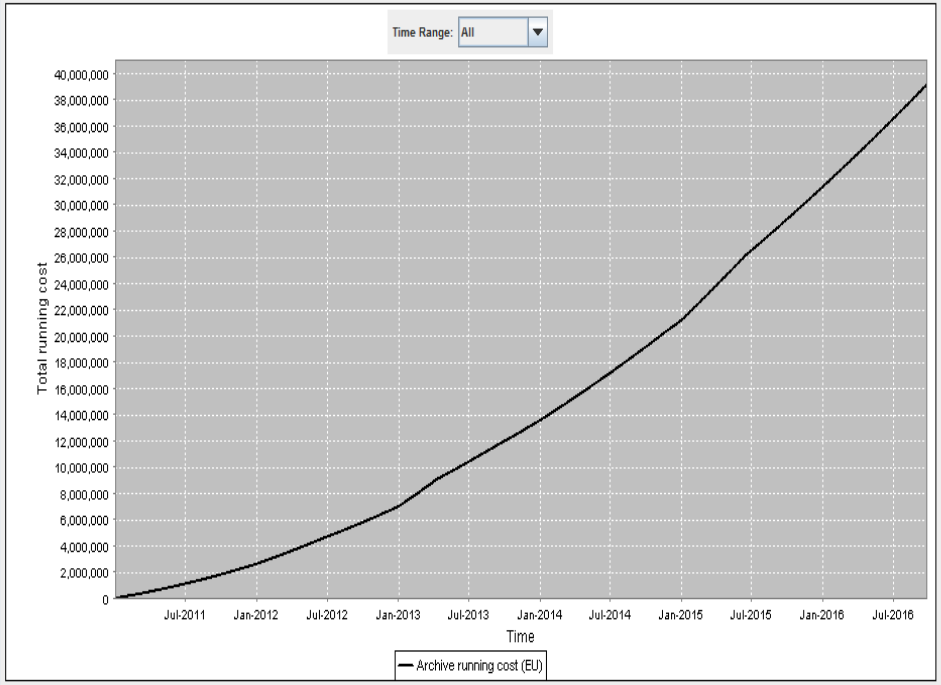

Configuration Template Management

Select Configuration Template to Load

Default 2 Storage Model Configuration | Load Selected Template

Template Management

LOADED TEMPLATE: mja_model1

? | Update Model | Save As New Model | Exit Simulation

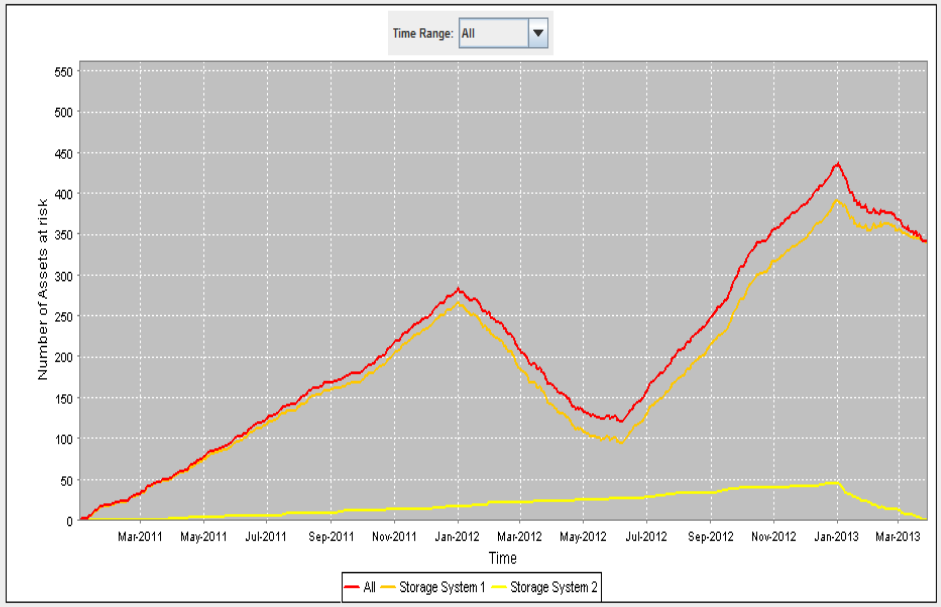

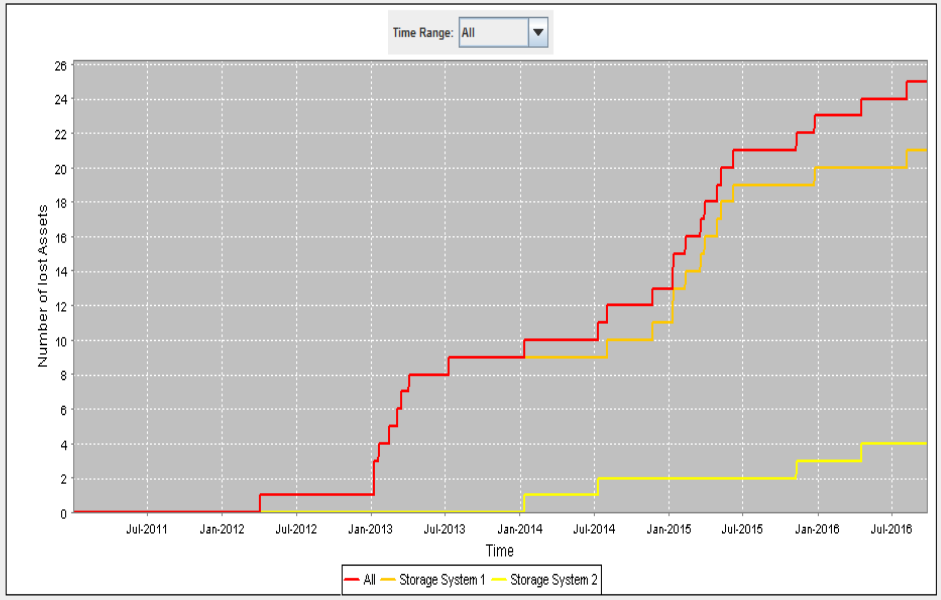

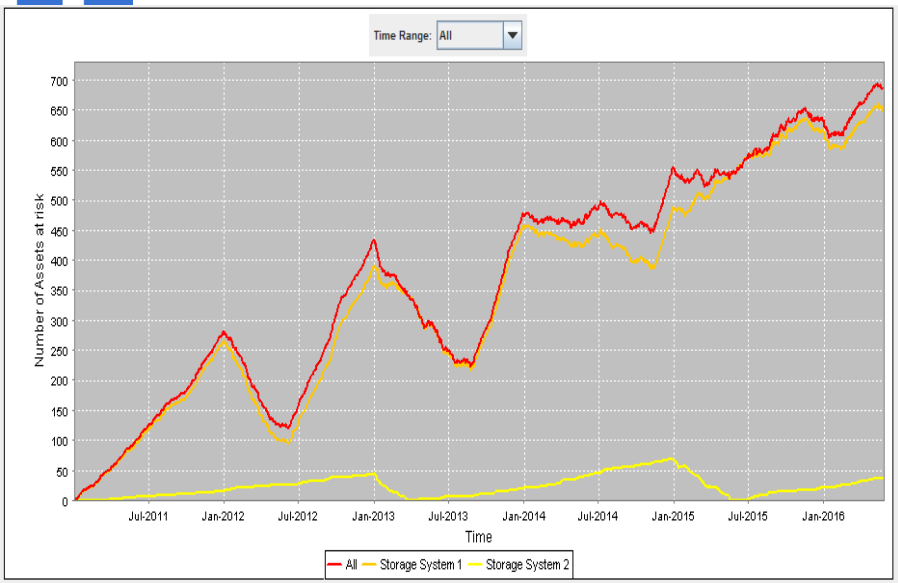
Time Range: All
Number of Assets at risk
Time
All
Storage System 1
Storage System 2

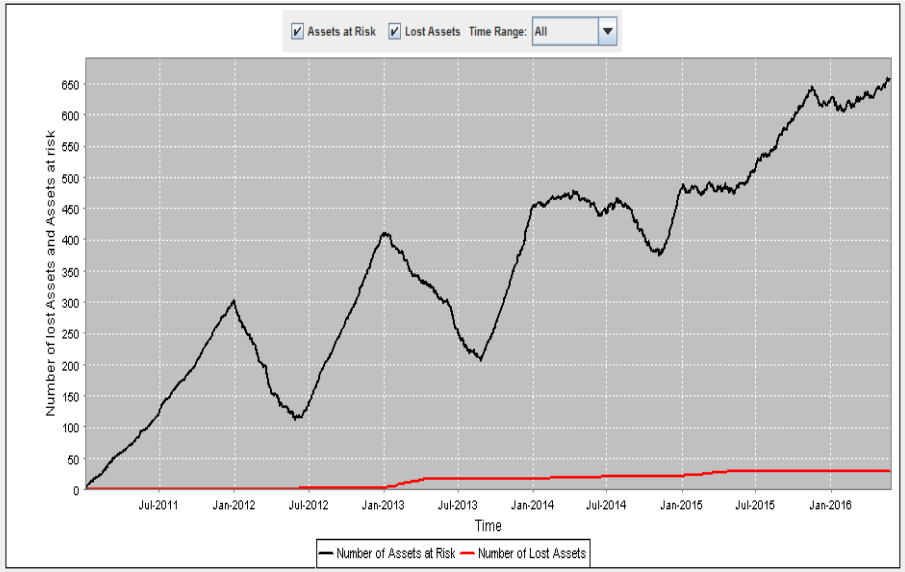
Assets at Risk
Lost Assets
Time Range: All
Number of lost Assets and Assets at risk
Time
Number of Assets at Risk
Number of Lost Assets

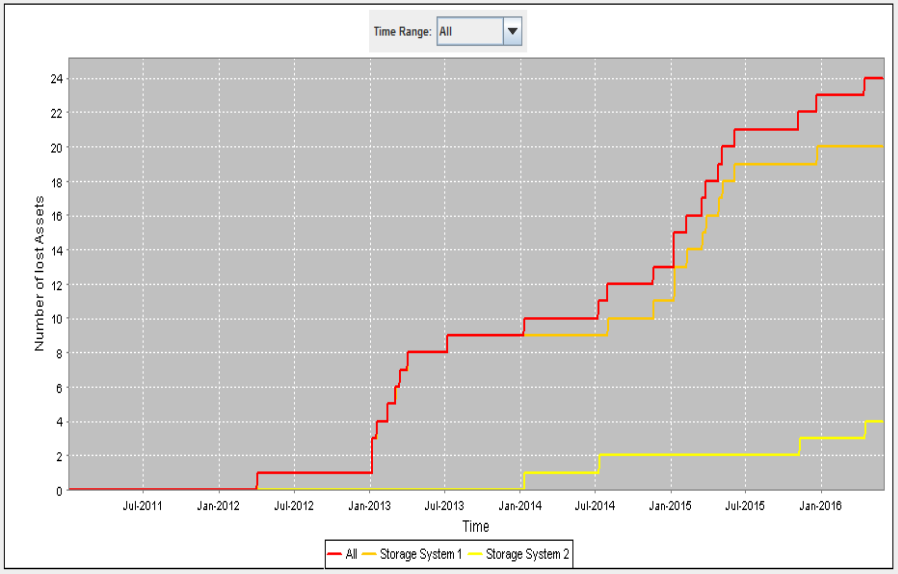
Time Range: All
Number of lost Assets
Time
All
Storage System 1
Storage System 2

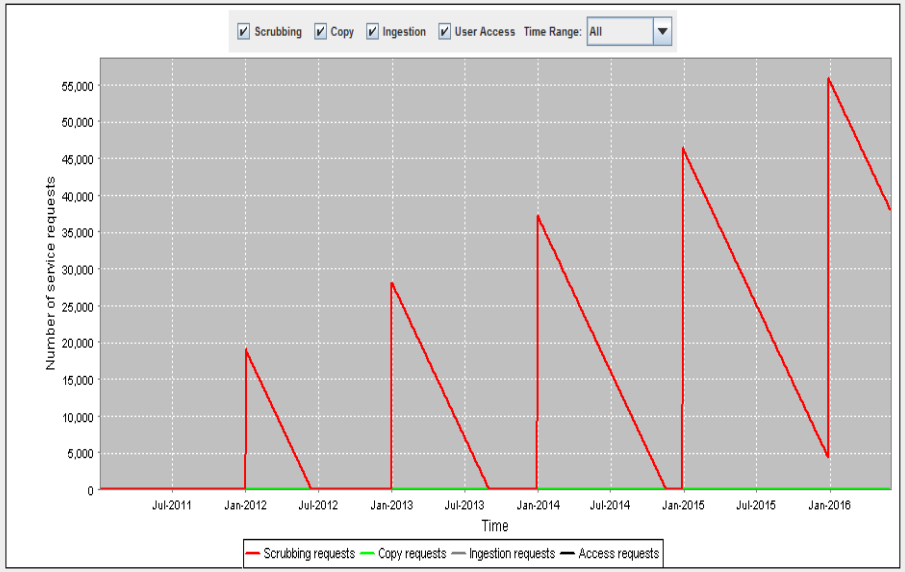
Scrubbing
Copy
Ingestion
User Access
Time Range: All
Number of service requests
Time
Scrubbing requests
Copy requests
Ingestion requests
Access requests

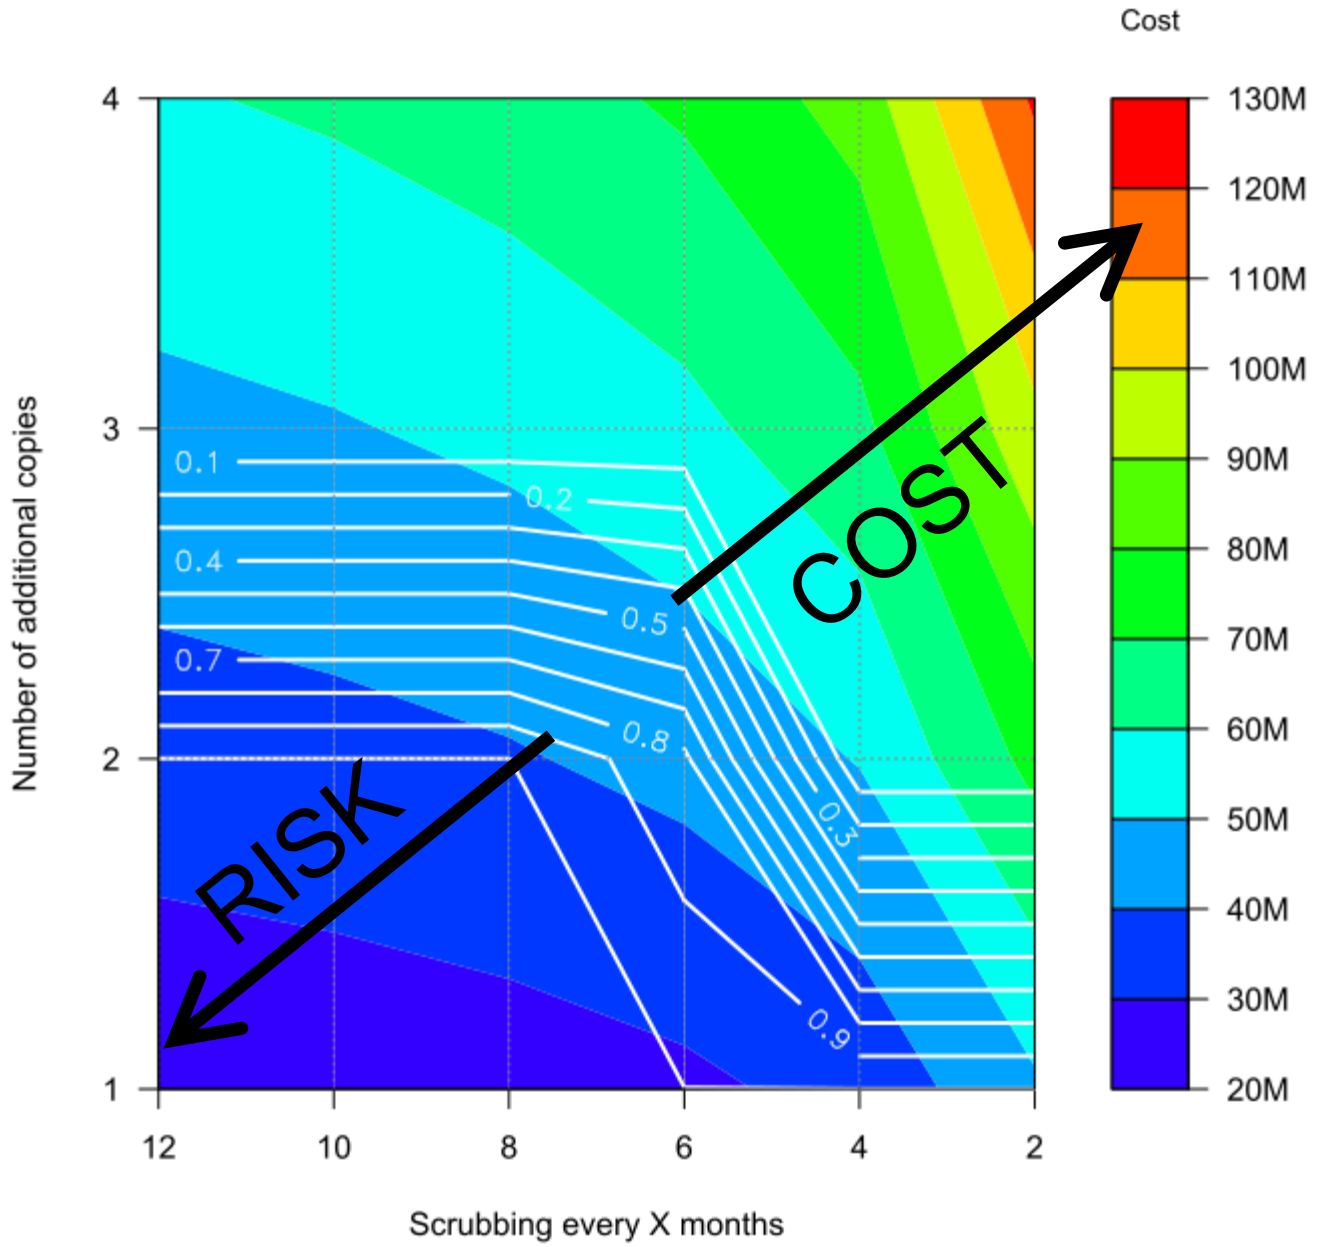
Cost
130M
120M
110M
100M
90M
80M
70M
60M
50M
40M
30M
20M
Number of additional copies
4
3
2
1
Scrubbing every X months
12
10
8
6
4
2
0.1
0.2
0.4
0.5
0.7
0.8
0.3
0.9
COST
RISK

Example: migration workflows

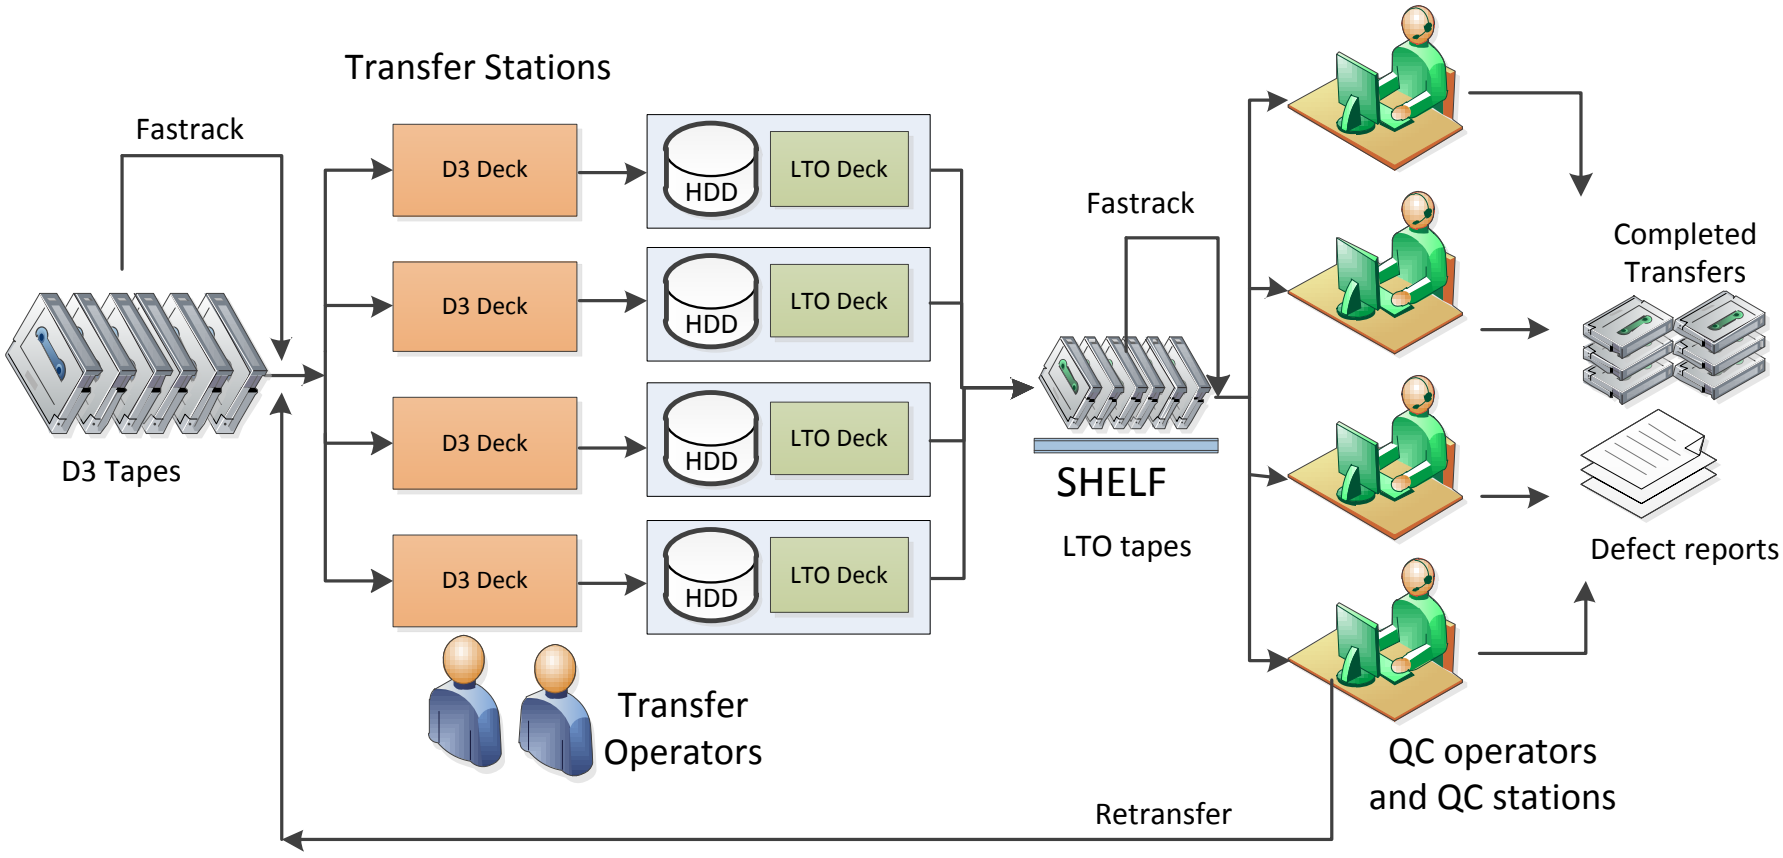

Interactive D3 Migration Simulation Tool v 1.0 (IT Innovation Centre)

Global Performance | D3 Workflow 1

Workflow Configuration

Currently viewing: Quality Control Configuration ?

Number of QC operators and QC stations: 4 Update

Time spent by QC operator reviewing transfer defects: 30 seconds per defect Update

Time spent by QC operator logging transfer defects: 30 seconds per defect Update

Cost per hour of QC operator: 30 Update

Cost per QC station: 20 Update

Amortisation time of QC station: 5 years Update

Number of spot check operators: 1 Update

Spot checked tapes proportion: 100.0 % Update

Percentage of failed tapes: 0.0 % Update

Pre-existing defects detection method: Manual Check (full sweep)

Time spent by QC operator reviewing the full programme for pre-existing defects: 1.2 x length of programme material Update

Time spent by QC operator logging pre-existing defects: 30 seconds per defect Update

Percentage of pre-existing defects not picked-up by QC operator: 5 % Update

Percentage of tapes that will need retransfer: 7 % Update

D3 Tapes | Transfer | Content Throughput | Quality Control | Running Cost

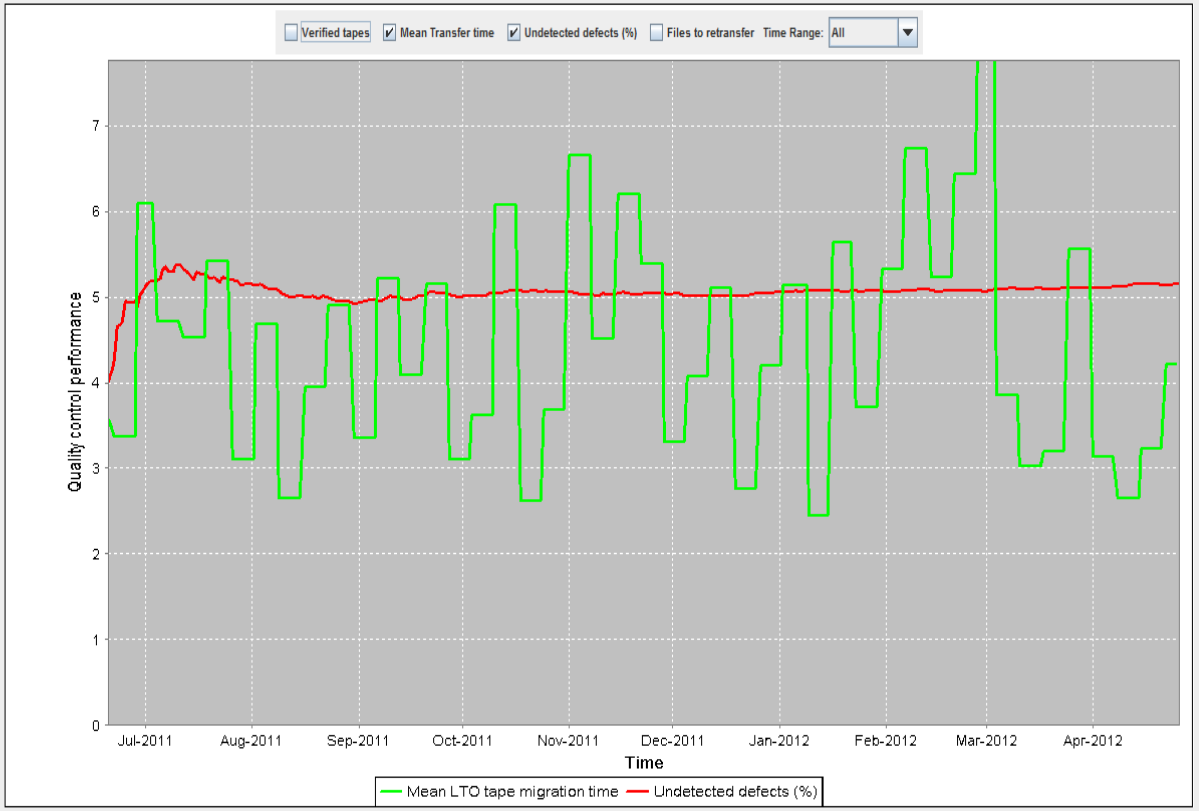

Model Performance Output

Simulation time: [23h/25d/4m/2012y]
D3 Tapes Queue Size: 7298.0 Fast tracked D3 items: 209 Fast tracked LTO tapes: 31
Undetected errors in D3 Tapes: 0.0 Undetected preexisting defects (after QC): 5.145506804953935 %
D3 Decks Number: 1 Headlife used (h): 1528.2585850015912 Headlife remaining: 8471.741414998409 D3Deck staff total work hours: 1672.0
HDD Cache Size (in GB): 97.31835822996912 Files in cache: 2.0 Transferred content size: 137445.95429191386
LTO Tapes Size (in GB) : 126087.78403288871 Produced Tapes: 338.0 LTO staff total work hours: 587.0
Size of the Material Migrated To LTO Tapes: 137445.95429191386
Average time of item to go through the workflow: 4.215686274509804
Number of undetected defects: 234.0 QC staff total work hours: 6660.0

Volume of material awaiting quality control (in GB): 395.8958216205666
Volume of quality checked material (in GB): 136670.884775118
Infrastructure cost: 42129.687699669856
Human resources cost: 351920.0
Annual total cost: 394049.68769966747

SEVENTH FRAMEWORK PROGRAMME

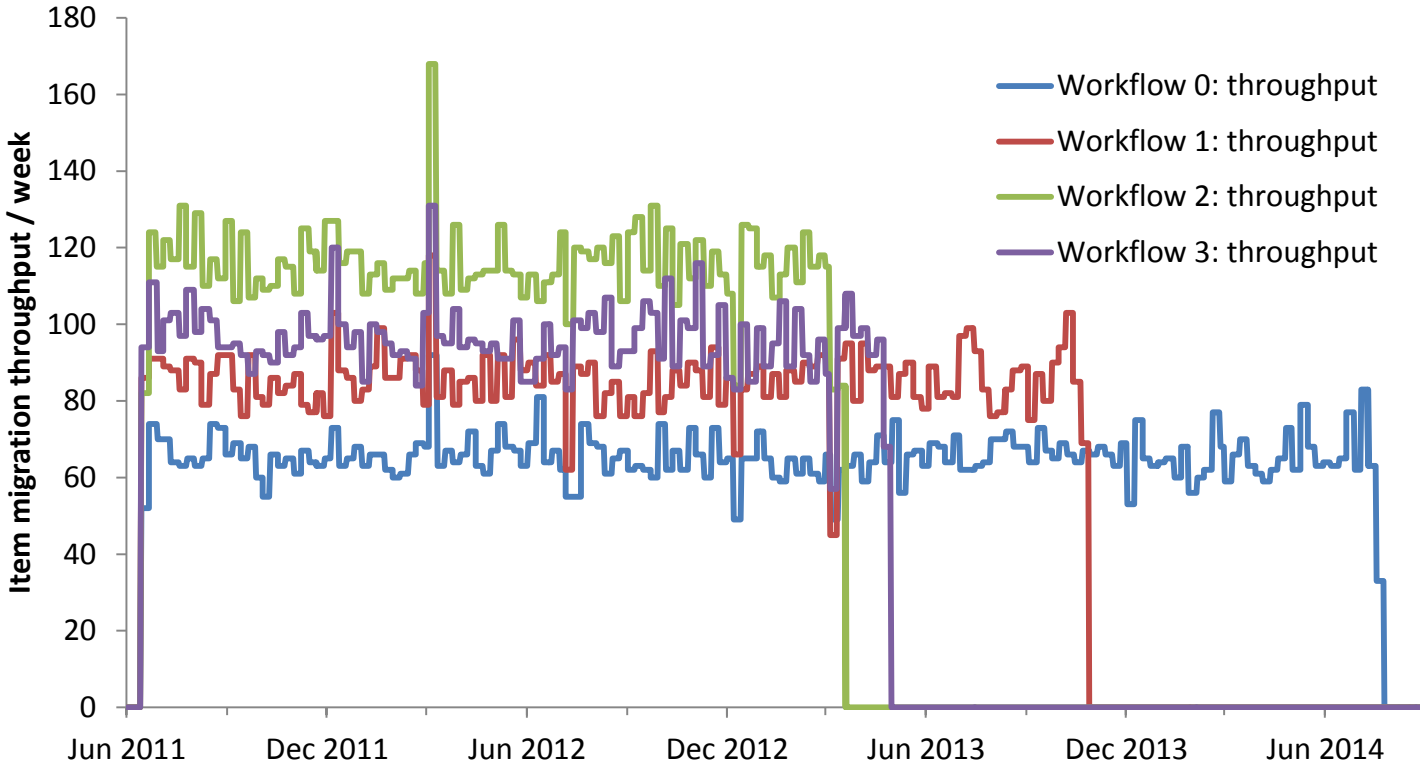
Item migration throughput / week
180
160
140
120
100
80
60
40
20
0
Jun 2011
Dec 2011
Jun 2012
Dec 2012
Jun 2013
Dec 2013
Jun 2014
Workflow 0: throughput
Workflow 1: throughput
Workflow 2: throughput
Workflow 3: throughput

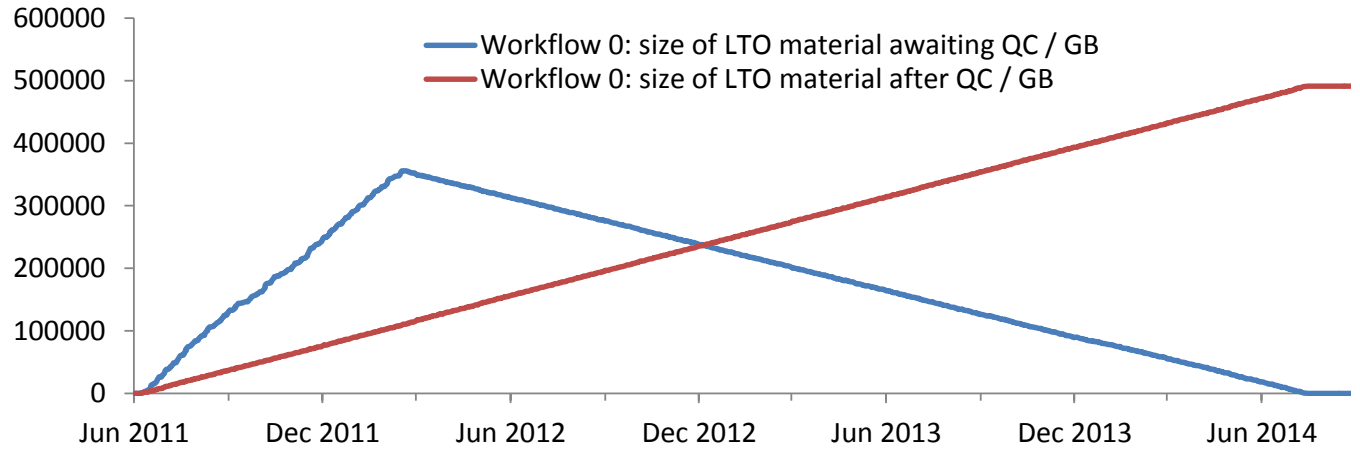
600000
500000
400000
300000
200000
100000
0
Jun 2011
Dec 2011
Jun 2012
Dec 2012
Jun 2013
Dec 2013
Jun 2014
Workflow 0: size of LTO material awaiting QC / GB
Workflow 0: size of LTO material after QC / GB

Decision Loops for Services

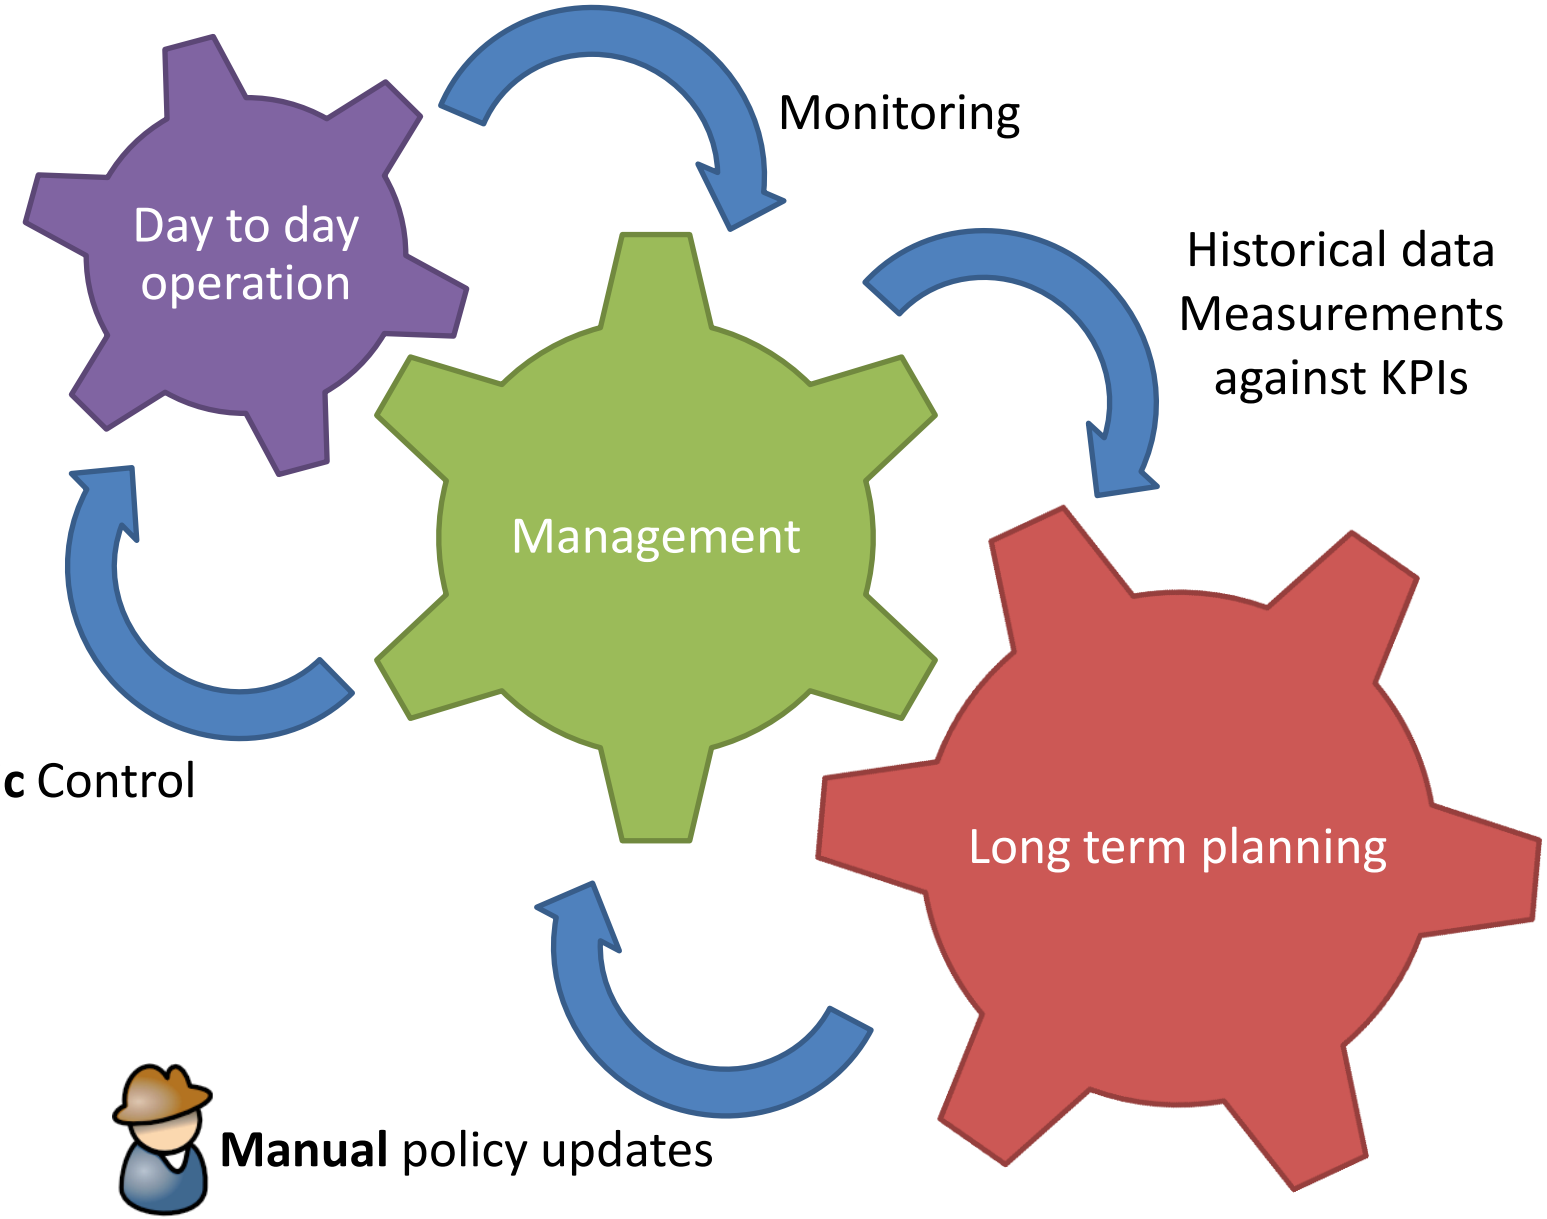

Data Service Management

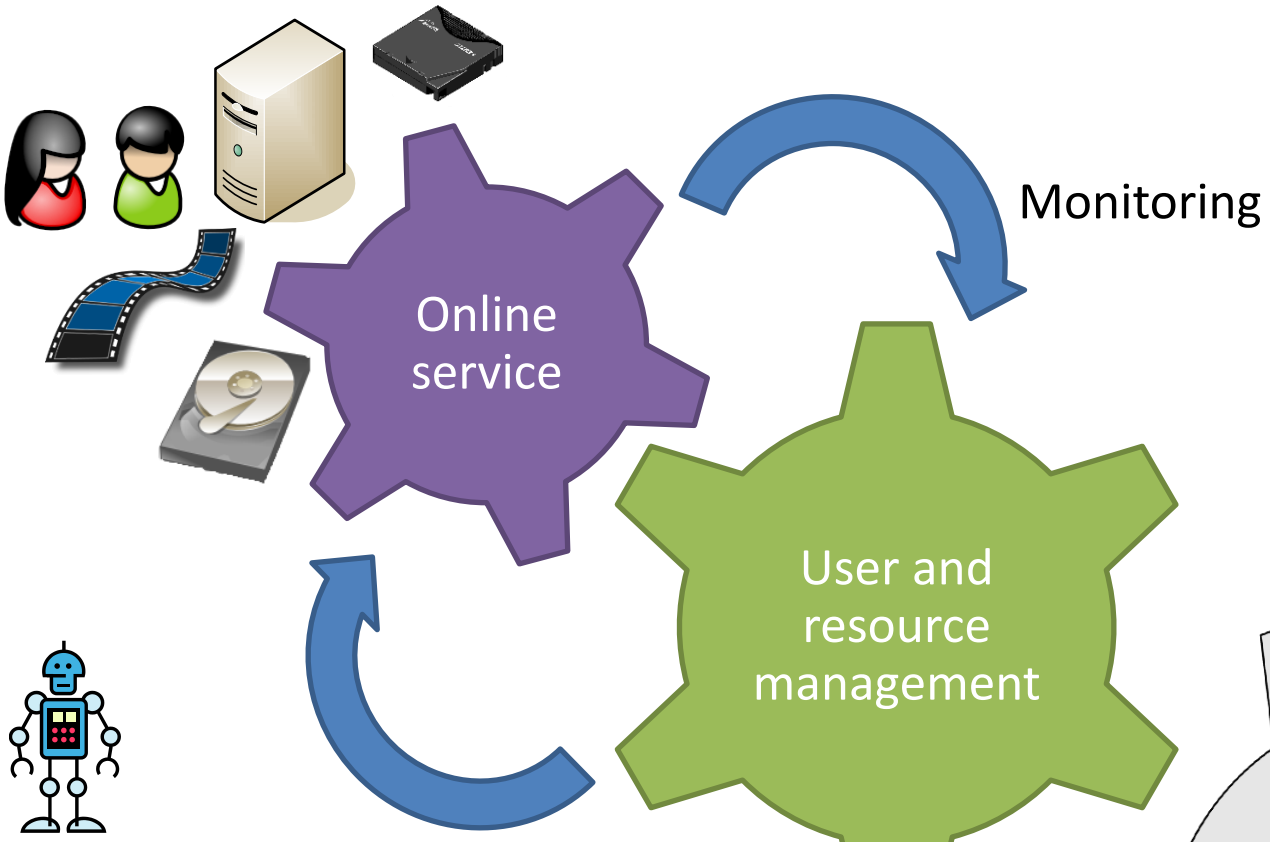

What have we built?

Online service

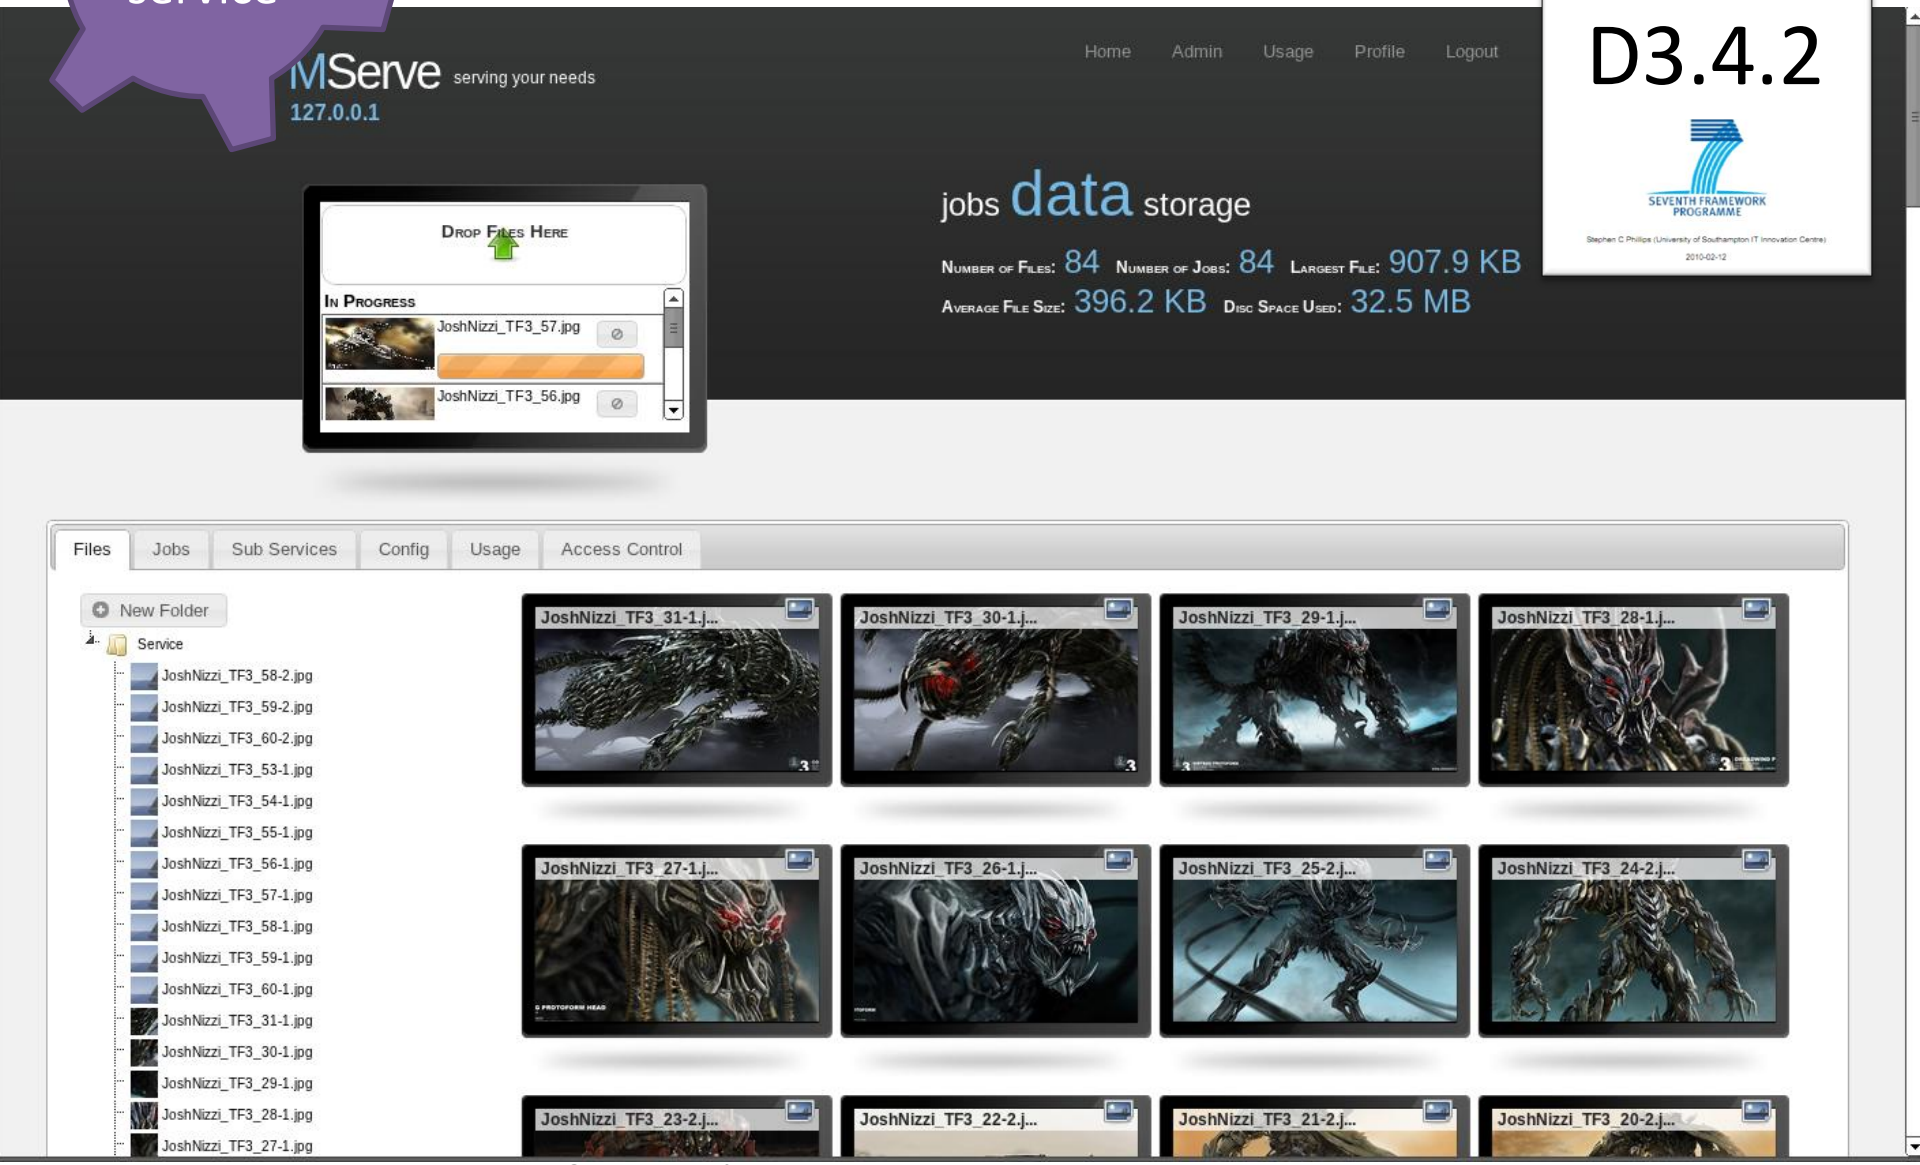

FP7-ICT-2007-3-231161

PrestoPRIME

D3.4.2

SEVENTH FRAMEWORK PROGRAMME

Stephen C Phillips (University of Southampton IT Innovation Centre)

2010-02-12

Usage Summary

Traffic

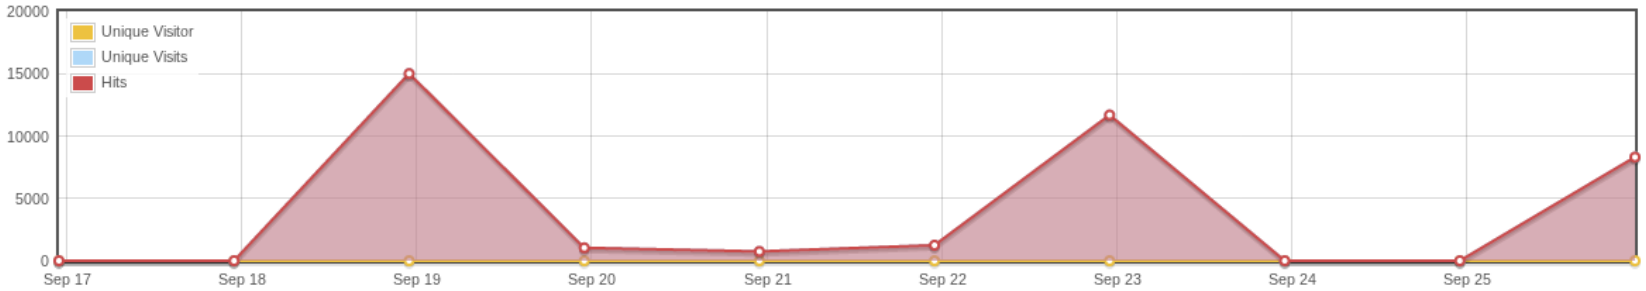

Tasks in last 24 hours

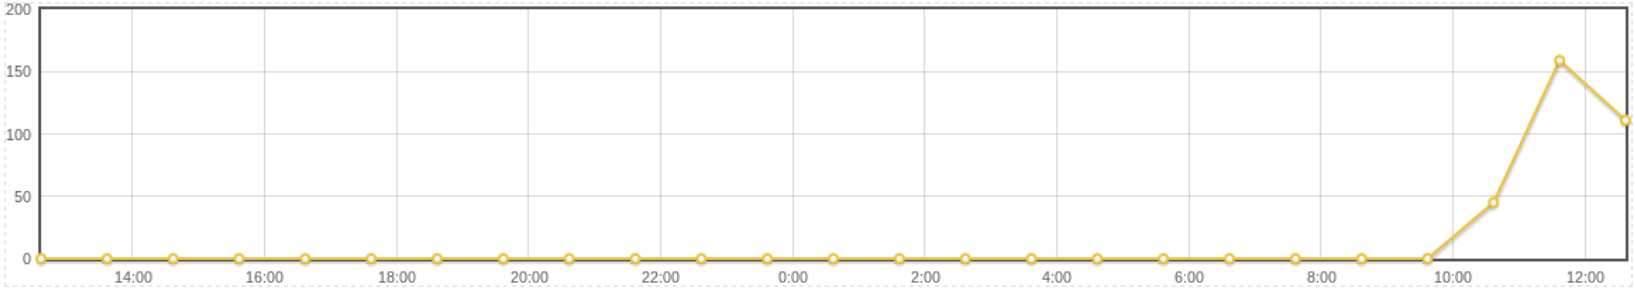

Tasks in last hour

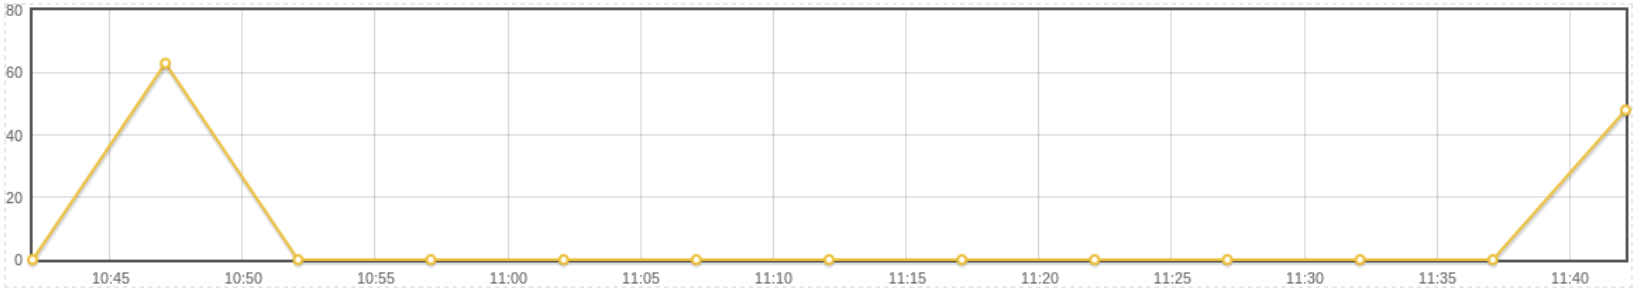

All Jobs

Job md5file

Job backup_mfile

Job posterimage

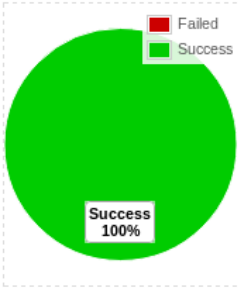

Job thumbimage

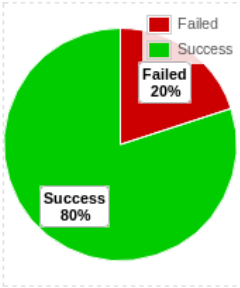

What have we built?

Management

FP7-ICT-2007-3-231161

D3.4.2

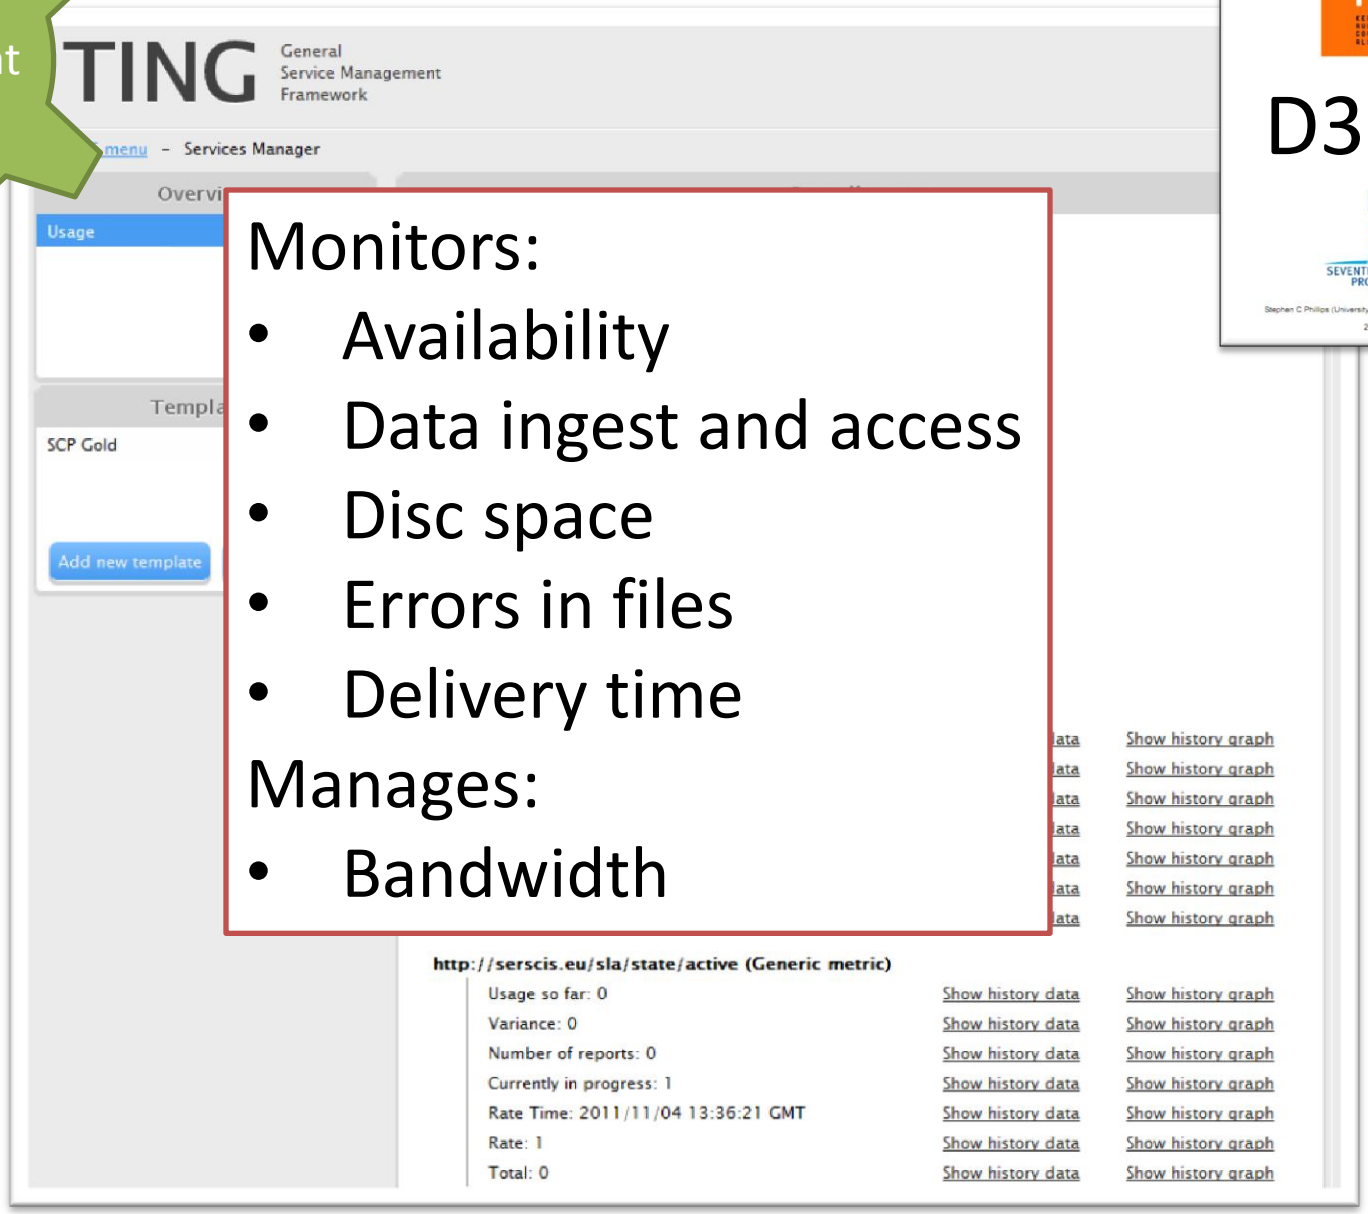

Monitors:

- Availability
- Data ingest and access
- Disc space
- Errors in files
- Delivery time

Manages:

- Bandwidth

Predict Future Trends

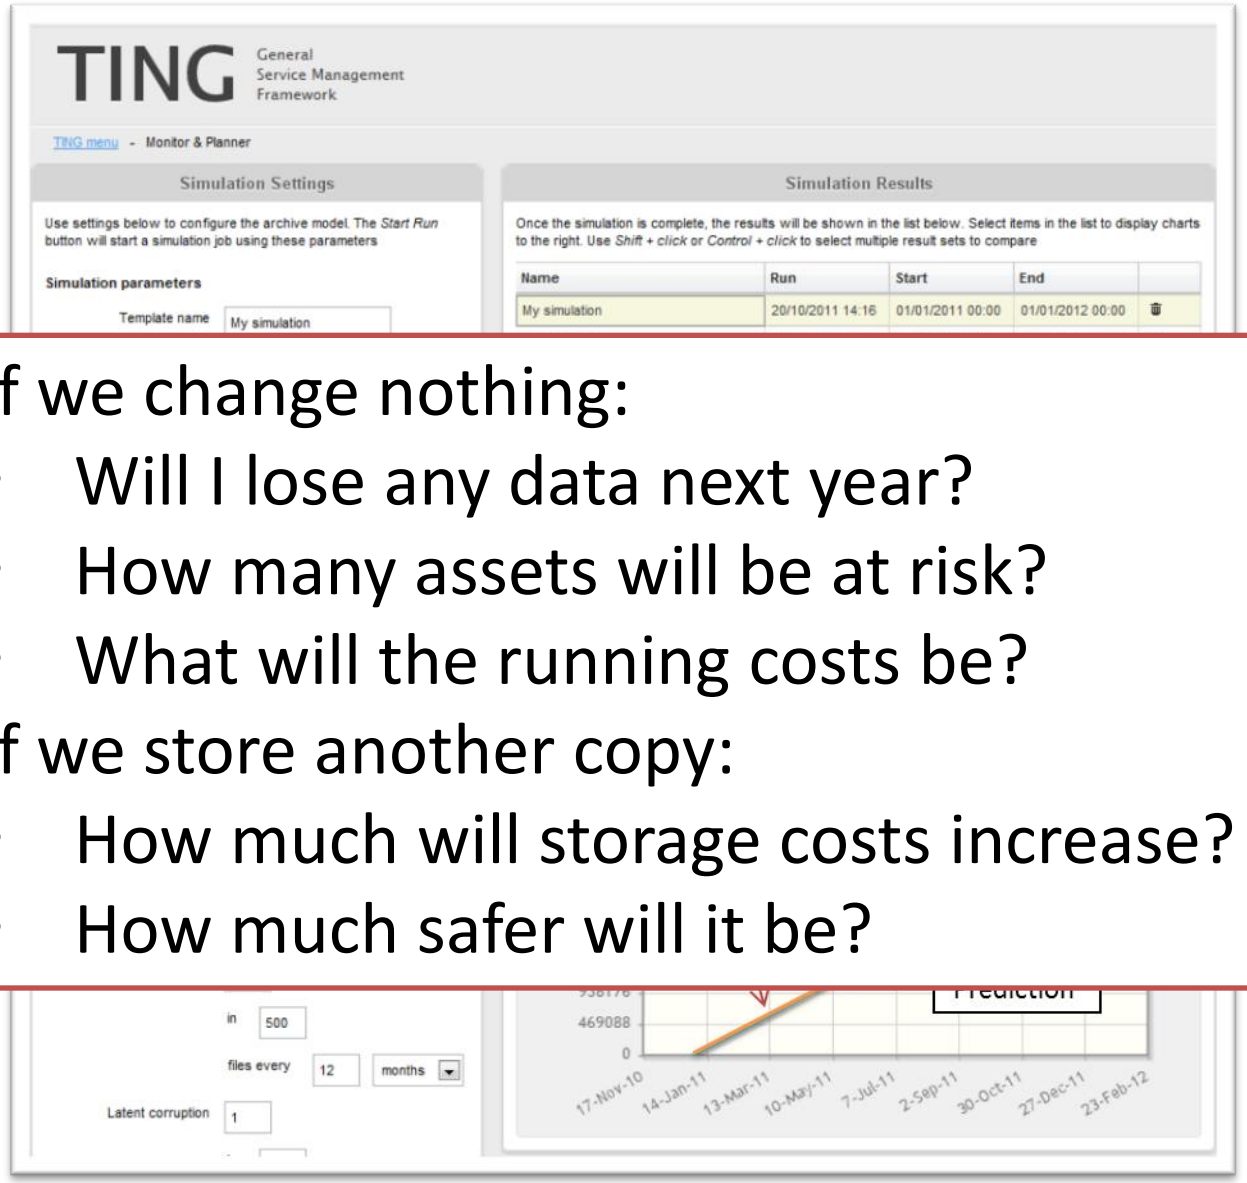

If we change nothing:

- Will I lose any data next year?
- How many assets will be at risk?
- What will the running costs be?

If we store another copy:

- How much will storage costs increase?
- How much safer will it be?

Optimising the System

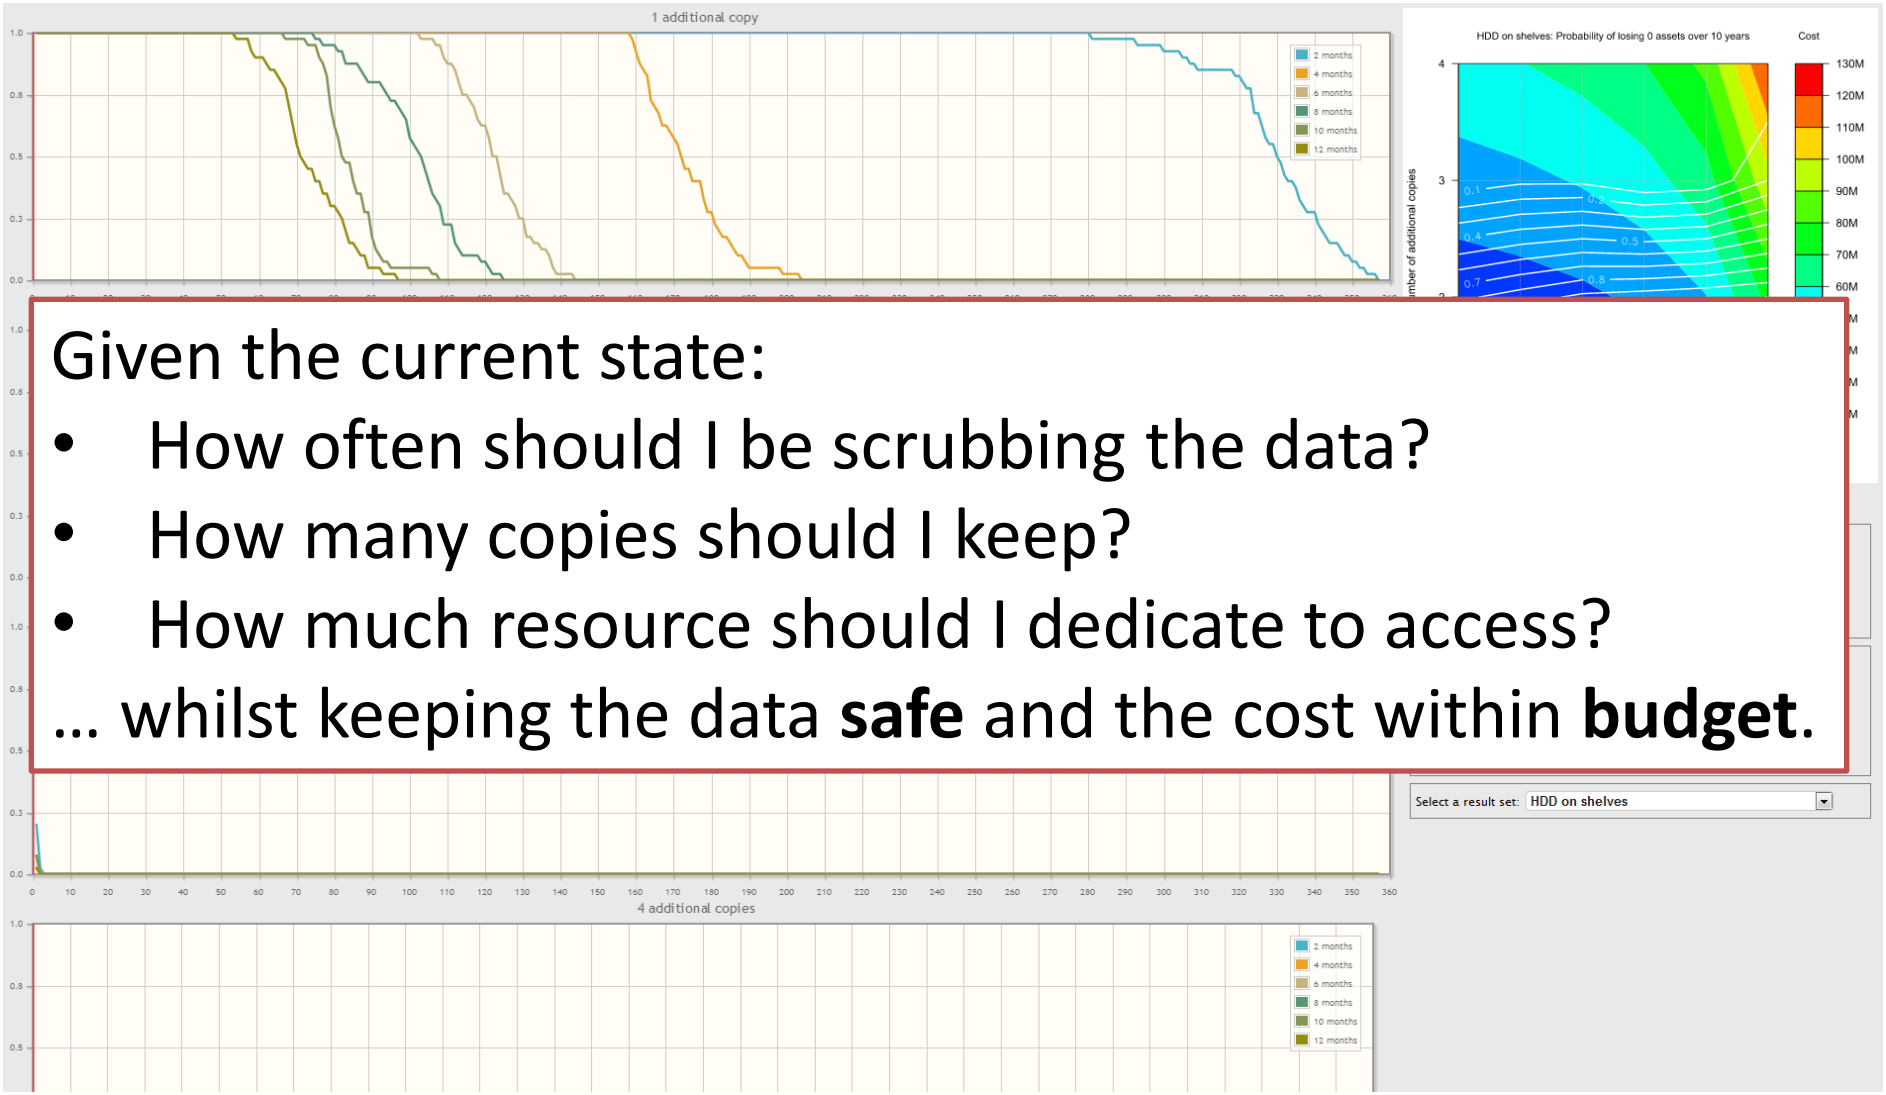

Given the current state:

- How often should I be scrubbing the data?
- How many copies should I keep?
- How much resource should I dedicate to access?

… whilst keeping the data **safe** and the cost within **budget**.

More information

- D2.1.1 Preservation Strategies
- D2.1.2 Preservation Modelling Tools
- D2.2.1 Processes for preservation and access
- D2.3.1 SOA for AV storage
- D3.2.1 Threats from mass storage
- D6.3.1 Financial models and cost calculation
- D7.1.4 Annual AV preservation report(s)

All available from the PrestoCentre

Try out the tools

http://prestoprime.it-innovation.soton.ac.uk