Cost model study on cache size / network trade off

Andrea Sciabà

DOMA Access meeting, 10/12/2019

Context

- Site cache simulated by using file access records for CMS sites
  - Data is stored and analysed in the CERN analytix Spark+HDFS cluster
    - File name, size, site, access time, user, type, …
  - Already presented many times
- Showing results for CMS T2 SoCal sites
  - Big T2, it has even an actual production Xcache
- Looking at MINIAOD(SIM) just because it is the most popular format for analysis

Optimal cache size

- From a cost perspective:
  - Too large cache → expensive storage
  - Too small cache → too much WAN traffic → need to buy more network bandwidth
- Different cache management strategies:
  - High/low watermarks to free up space, e.g. according to LRU criteria, or
  - Remove files not accessed since more than days, or
  - More sophisticated strategies (might even use ML)
  - The first two are mostly equivalent

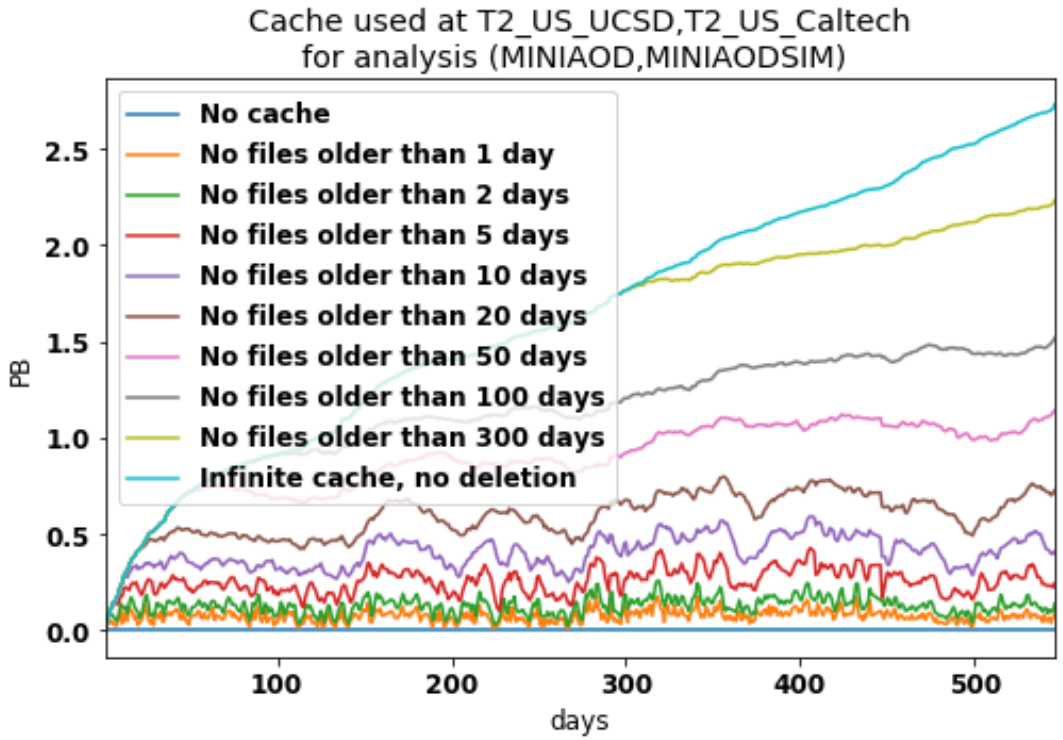

WAN traffic vs. used cache

- WAN traffic is generated for files not cached
- Cache utilisation is generated by files cached ☺
- There must be an optimal value for the maximum file age that minimizes cost
- A “cost” function can be defined

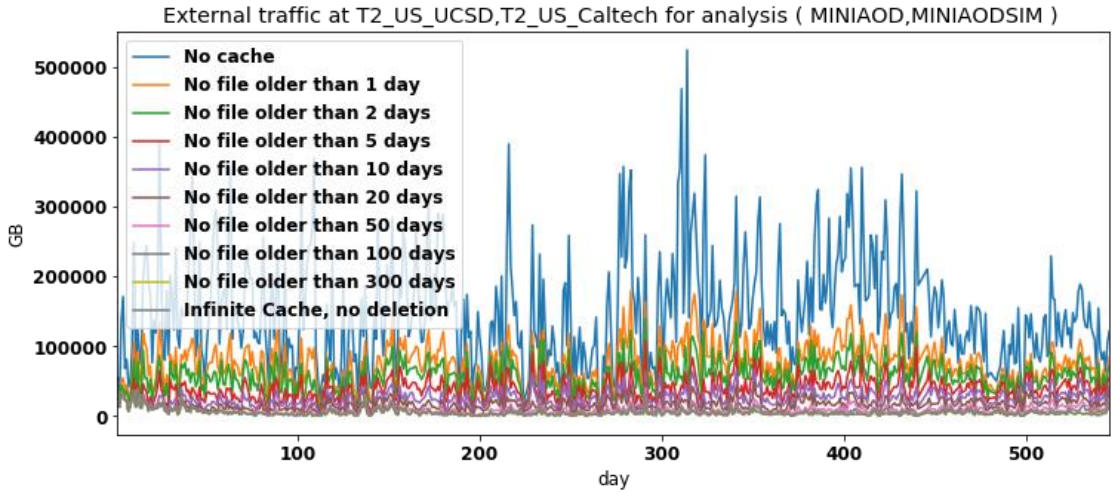

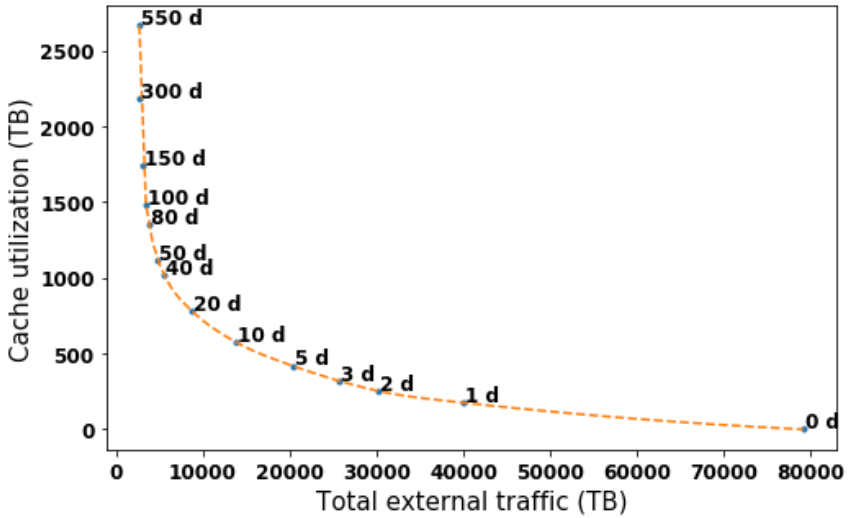

Cost function

- Total cost = network cost + storage cost
- Storage cost = max(cache occupancy) × cost / unit of disk storage
  - Relatively straightforward, caches have low QoS, so cheap HDDs in JBOD configuration are sufficient
- Network cost = avg(external traffic / time) × cost / unit of bandwidth
  - Much more difficult to estimate, as it is not proportional to usage

Cost estimates

- Disk
  - Cost estimated in the WLCG/HSF cost model working group
    - 1 HDD: 8 TB, 400 EUR, 4 years lifetime
    - Disk server cost / TB ~ twice disk cost
    - ⇒ cache cost ~ 25 EUR/TB/year
  - Baseline HDD scenario: 25 EUR/TB/year
  - Pessimistic HDD scenario: 50 EUR/TB/year
  - SSD scenario: 100 EUR/TB/year
- Network
  - Cost estimated very roughly from a couple of WLCG entities (and from my internet provider: they are all in the same ballpark!)
  - Baseline: 1 EUR/TB
  - Pessimistic: 10 EUR/TB
- These estimates can surely be improved and anyway they can be very different at different sites, so take them just as arbitrary but meaningful references

Cost optimisation results

		Disk cost (EUR/TB)		
		40	100	150
Network cost (EUR/TB)	1	8	2	1
	10	68	36	26

Optimal file age (days)

		Disk cost (EUR/TB)		
		40	100	150
Network cost (EUR/TB)	1	522	252	175
	10	1262	980	870

Optimal cache size (TB)

These numbers are for 1.5 years

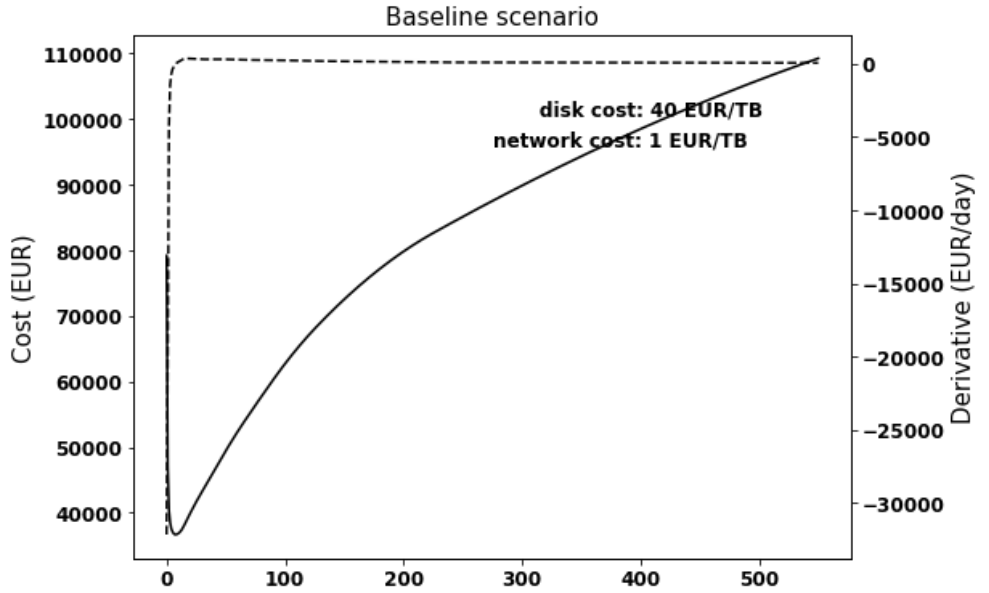

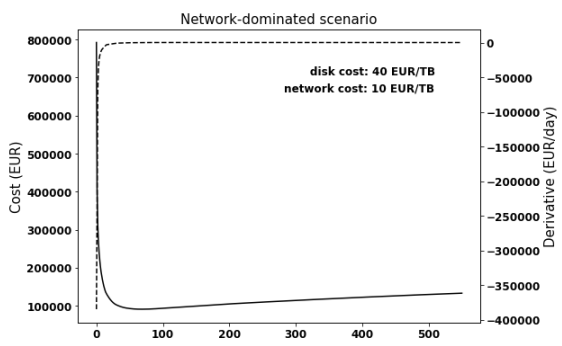

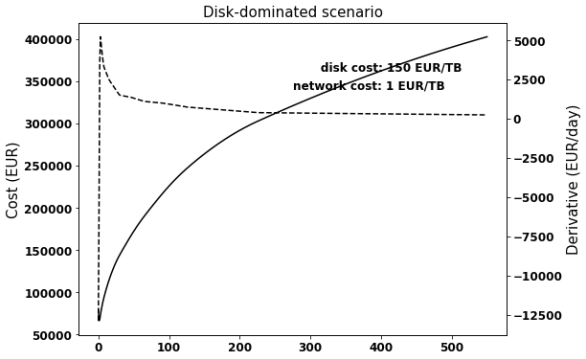

Conclusions

- A simple exercise to show how to choose the best cache size
- The optimal point critically depends
  - on the access patterns and the scale of the dominant workloads
  - on the cost scenarios at the site
- It would be interesting to compare with the actual costs for the actual production cache at SoCal!