

---



# Media for Long-Term Archiving: Flash

Bruce Moxon

Principal Architect, NVM Storage Solutions



**Strategy and Innovation Center**

# **Flash** for Long-term Archiving

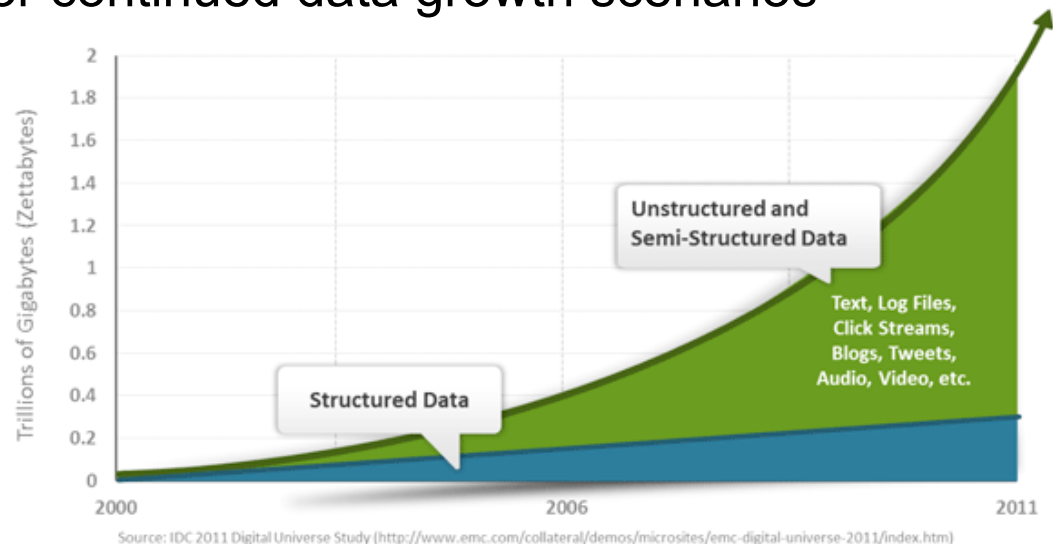
---

## ***Flash ... Archival (Cold) Storage ... Huh?***

- Trends
  - Unstructured data, Flash supply (uses) and pricing
  - “Cheap, crappy Flash” (Facebook); crossover
- All “Archival” (“Cold”) Storage is not equal
  - Long-tail Distributions
  - Access Patterns and Expectations
- The Move to Object Storage
  - “Big” O and “Little” O; application metadata
  - Object caching (keys, k-v, application object metadata)

# Flash “Crossover” ... Do the Stars Align?

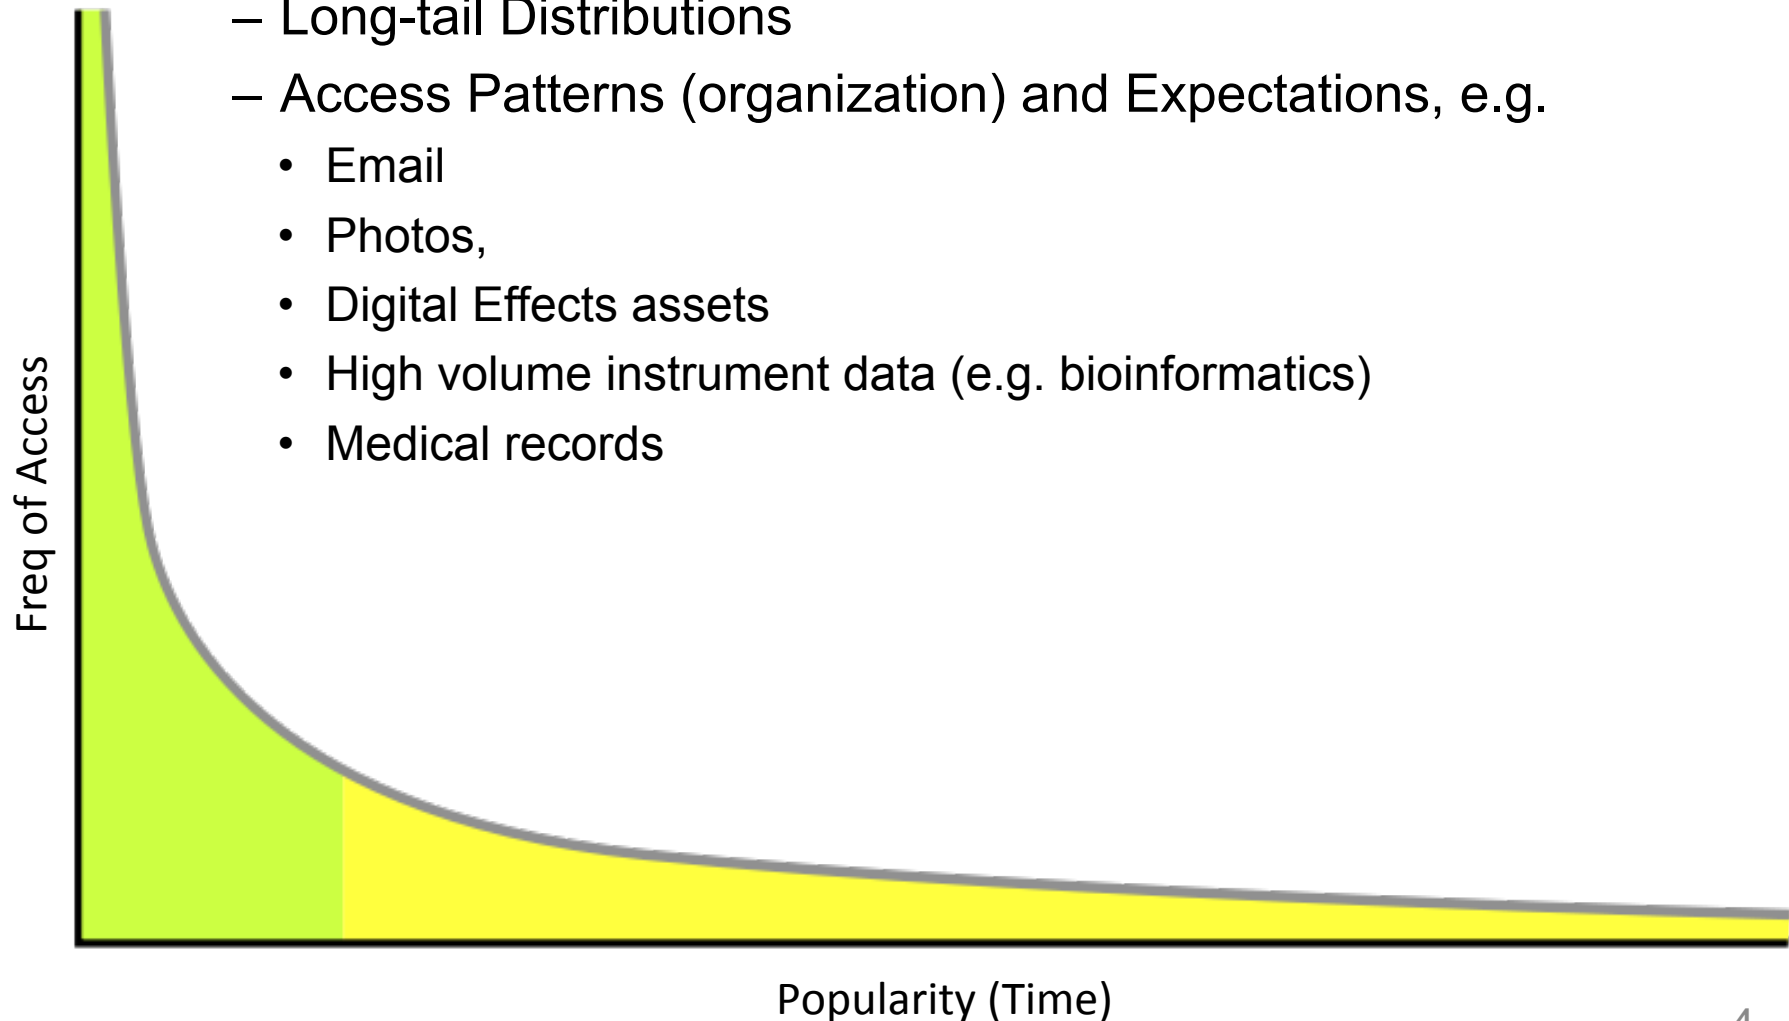
- Facebook ... “cheap, crappy Flash”
- Gartner<sup>1</sup>: Cold storage requires sub-\$0.25 per GB
  - Post-2015 based on 3D-NAND and 8LC projections
- OK, but ...
  - That doesn’t account for the fact that the NAND **supply** will continue to be *relatively* limited for continued data growth scenarios
    - Fabs
    - Competing uses
      - Consumer/mobile
      - Laptop/desktop
      - Enterprise T1/cache



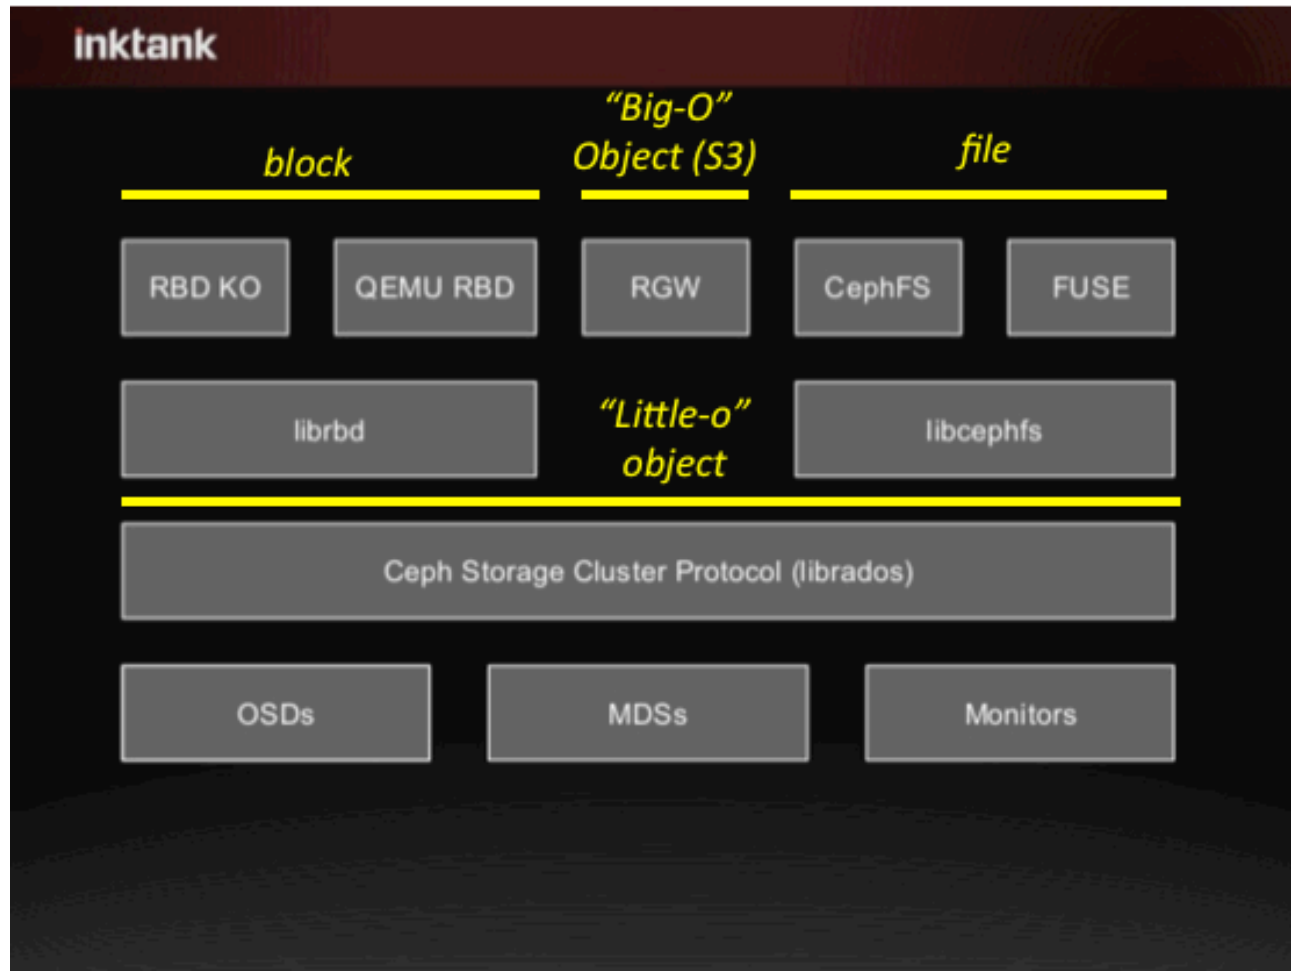
<sup>1</sup> 2014 NAND & SSD Industry The Year of ‘Forceful Restraint’ (Gartner, Feb 2014)

# The Long Tail

- All “Archival” (“Cold”) Storage is not equal
  - Long-tail Distributions
  - Access Patterns (organization) and Expectations, e.g.
    - Email
    - Photos,
    - Digital Effects assets
    - High volume instrument data (e.g. bioinformatics)
    - Medical records



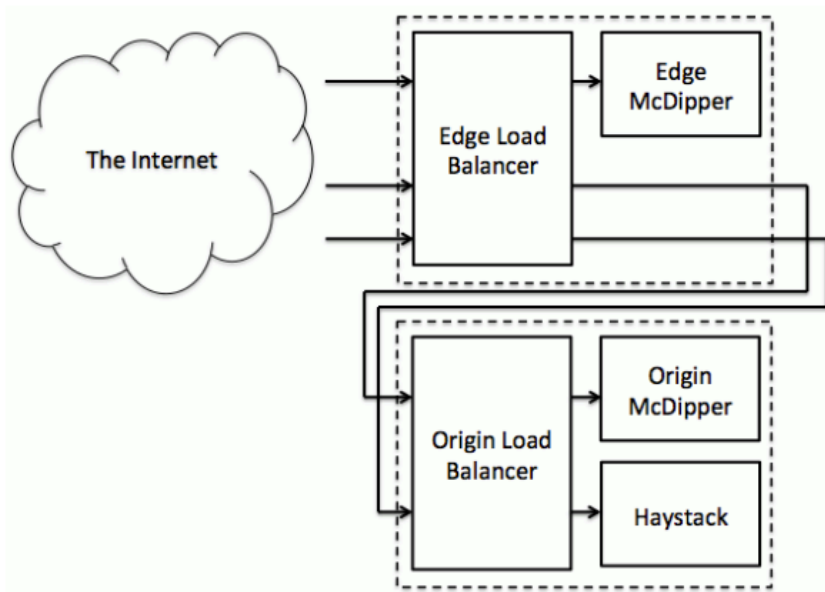
# Object Storage: Big “O”, Little “O”



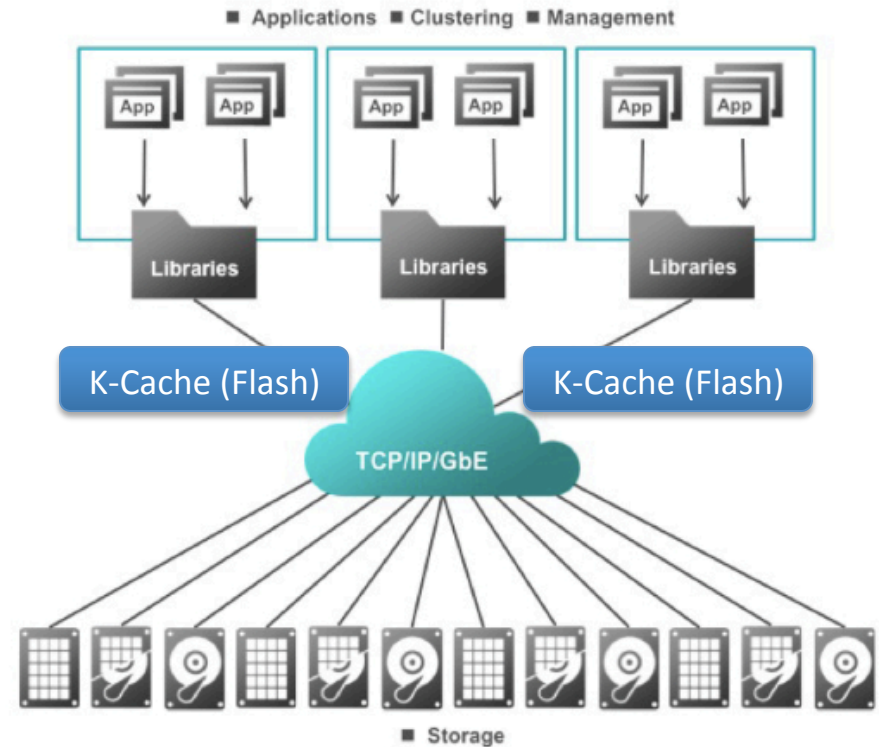
Little “O” => Embedded, key-value, single node (“chunk”)

Big “O” => Application metadata, multi-node (incl. geo-distributed), protected/replicated

# Object Storage and Flash Caching



**Facebook McDipper**



**Kinetic (Application) Caching**