



FleSSR Business Models

Or: is Amazon beatable?

Matt Johnson, Head of Research



Overview



Overview of Amazon Web Services

- Delivering Cloud solutions since 2006
 - Wide range of services (compute, storage...)
 - Global reach (US, EU, Asia)
- High demand
 - S3 peaks at 200,000 requests per second
 - Hundreds of thousands of servers
- Economies of scale



Not all good news...

- Beware of “hidden” costs
 - Particularly around I/O throughput and bandwidth
- “Low-value” service levels
 - NO guarantees on data durability, only availability
- Possible business risks
 - Lack of long-term service stability (monthly contracts only)
 - Continuing concern around data and the US Patriot Act
 - No (auditable) access to infrastructure or operations



Why try to compete?

- To deliver a Cloud service focussed on education
 - Customised service models
 - Fully UK-based infrastructure
- To build on sector infrastructure
 - JANET network, UK Access Federation
- Long-term planning and stability
- What about costs?



S3 Storage Costs



AWS S3 – key features

- 99.999999999% (“eleven 9’s”) durability
 - Files are stored across 3 devices
 - Copy-on-Write protection
 - Regular analysis of stored data integrity
- 99.99% availability
 - Clustered front-end API access
 - Multiple availability zones



AWS S3 – pricing model (February 2011)

- Amazon pricing model reflects its infrastructure costs:
 - **Data stored:** disks, storage servers;
 - **Network usage:** bandwidth and Internet transit;
 - **Requests:** front-end web/API servers;
- Sliding-scale pricing
 - **Data stored:** \$0.14 p/GB to \$0.055 p/GB
 - **Network usage:** \$0.15 p/GB to \$0.08 p/GB
 - **Requests:** \$0.01 per 1,000 writes, \$0.01 per 10,000 reads



AWS S3 – price example (for 1 TB of data)

- Assumes an upload of full dataset at start of year
 - Includes updates (uploads) of 1% stored per month
 - Includes usage (downloads) of 1% stored per month
 - Assumes a 1 MB average file size
- Assumes EU (Dublin) datacentre



AWS S3 – price example (for 1 TB of data)

- **Data stored: \$1,720 p/a**
- **Network usage: \$130.56 p/a**
 - 1 TB upload of dataset: \$102.36
 - 1% of 1 TB (p/m) change: \$28.20 p/a
- **Request volumes: \$11.32 p/a**
 - 1.12m writes: \$11.20
 - 0.12m reads: \$0.12
- **TOTAL: \$1,861.88 p/a (~ £1,160)**



AWS S3 – assumed storage architecture

- Storage service
 - Commodity mass storage servers
 - no RAID within storage units
 - Resilience derived from replicated copies of data
- Suggests that 30% of raw storage is usable
 - ~90% storage efficiency per single storage device
 - Data is stored across 3 separate devices $(90\% / 3) = 30\%$



Possible FleSSR Cloud storage infrastructure

- Broadberry 4U Storage Server
 - 95 TB “raw” storage for \$30,000 (assume 3-year lifespan)
- Annual costs: \$5,000
 - Power (PUE of 2.0): 35c per hour, \$3,000 p/a
 - Support costs: \$2,000 p/a
- Three-year costs: \$45,000 for 95 TB
 - \$160 per “raw” TB per year
 - **\$530 per “usable” TB per year**



Some considerations

- FleSSR pricing model is comparable to AWS:
 - **Data stored:** as shown previously
 - **Network usage:** JANET connectivity benefits
 - **Request volumes:** front-end web/API servers



What's not included in the FleSSR example?

- Overheads (firewalls, routers, security)
 - Would be factored across storage and compute usage
 - Shared services offer economies of scale
- Management software
 - Amazon have invested significantly in this area...
 - ...but OpenStack Storage is OSS and production-ready
- Sales/marketing/training overheads
 - When did you last speak to an AWS salesman?



Conclusions

- Amazon have a great business model
 - First-mover advantage, economies of scale, high margins
- There is a justified case for academic storage services
 - Specific requirements and integration with existing HE services
 - Potential for flexible business models
 - Costs can be comparable with S3
- We anticipate similar findings around compute services



Thank you!

Matt Johnson

Head of Research

email: Matt.Johnson@eduserv.org.uk

twitter: @mhj_work

web: <http://labs.eduserv.org.uk/>

Slideshare: <http://www.slideshare.net/eduserv/flessr-business-models>