

The Cloud Storage Market

- Cloud storage providers expose simple interfaces to developers. Amazon S3's data model provides flat namespaces ("buckets") into which named objects can be uploaded for later retrieval. Other storage services can be mounted as network filesystems. There is no widely agreed-upon standard interface, but S3's REST API has been adopted by smaller providers and by the open-source Eucalyptus server software.
- These interfaces differ, but are similar enough to be considered interchangeable. Storage providers are forced to compete on price rather than by offering unique services.
- Cloud storage is a highly competitive market. These are simplified pricing schemes for the top two cloud storage providers:

Operation	Amazon S3	Rackspace Cloud Files
put / list request	\$0.01/1000 requests	\$0.01/500 requests*
get / other request	\$0.01/10000 requests	free
delete request	free	free
Data transfer in	\$0.10/GB	\$0.08/GB
Data transfer out	\$0.17/GB	\$0.22/GB
Storage	\$0.15/GB/month	\$0.15/GB/month


*requests are free for files above 250KB in size

Why Should We Diversify?

- Cloud storage providers promise high availability, data persistence, and strict impressive SLAs. So why should we diversify storage?
 - Outages and Operational Failures:**
 - Cloud storage providers can experience transient outages. Sometimes lasting up to several hours. Diversifying storage improves data availability.
 - Technical issues at a provider's site can have unintended consequences. In October of 2009 a failure at a Microsoft data center resulted in data loss for many T-Mobile smart-phone users.
 - Economic Failures:**
 - A change in pricing scheme or the emergence of new competition can render a particular provider unfavorably expensive compared to its alternatives.
 - Clients may not be able to pick an optimal cloud storage provider because the switching cost overrides the desired benefits. Thus, clients experience vendor lock-in if their stored data is large.
 - The fundamental problem is that clients have to make an *all-or-none* decision in switching their data to new providers.
- Main point:** By striping data across multiple providers and adding appropriate redundancy, clients can tolerate outages and operational failures, as well as adapt to changes in the economic landscape.

Error Correcting Codes

- RACS uses Reed-Solomon error correcting codes to tolerate failures without data loss. Starting with m equal-size disks of original data, we fill k additional disks with redundant data. **Any combination of m disks (data or redundant) is sufficient to reconstruct the original data.** We write (m, n) to indicate that there are $n = m + k$ total disks, any m of which are sufficient to reconstruct all original data.

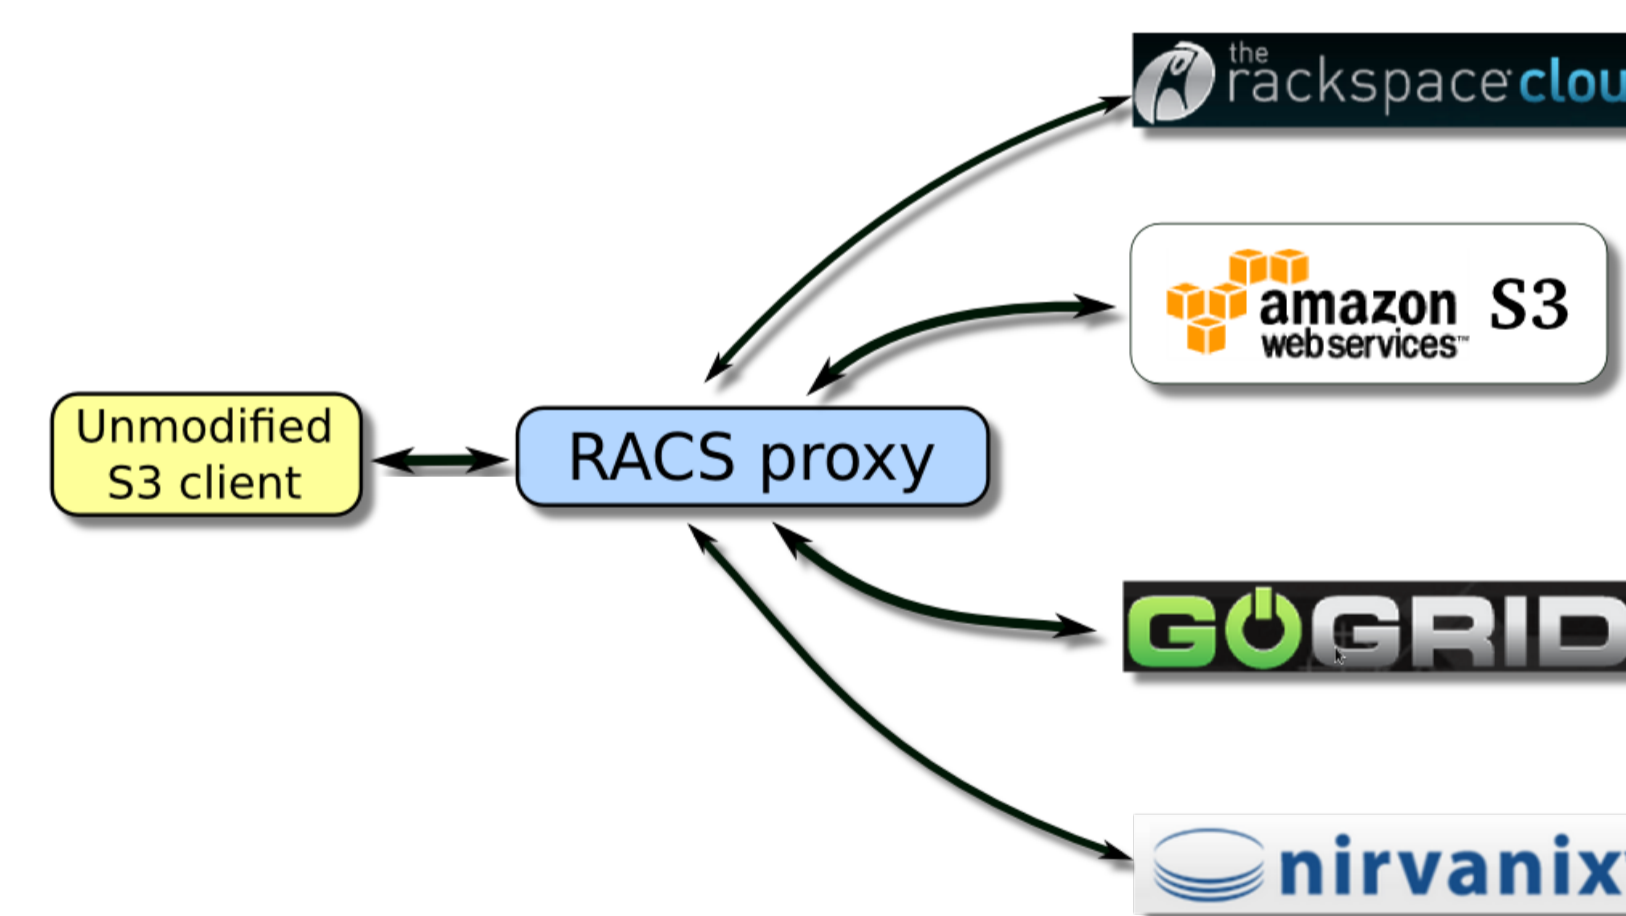


$(4,6) \rightarrow$ "Tolerate up to two failures"

- The choice of parameters m and n is a trade-off: Overhead for data storage and write operations is increased by the ratio $n : m$. Interestingly, read operations are not significantly more expensive, since only m disks must be read under normal operating conditions.
- RAID-5 uses a similar strategy to tolerate up to one failure in an array of hard disks.

RACS: Redundant Array of Cloud Storage


- Redundant Array of Cloud Storage (RACS) operates on the same principle as RAID-5, but rather than using hard disks, it stripes data across cloud storage repositories.




- The goal of RACS is slightly different than RAID-5. Cloud storage is assumed to be much more reliable than hard disks, so data loss prevention is a much less compelling reason to use error correcting codes.
- RACS lowers the cost of switching providers, e.g., as a result of **economic failure**.
 - Only $\frac{1}{m}$ of all data needs to be moved to leave a vendor.
 - By reducing the impact of vendor lock-in, RACS increases the leverage of customers when negotiating contracts with cloud providers.
- RACS is implemented as an HTTP proxy with the same interface as Amazon S3

Evaluation

- Internet Archives Trace:** Trace represents 18 months of activity on the Internet Archive's FTP sites.



- Cost of Hosting on the Cloud:** Simulated cost of hosting the Internet Archive's trace on various cloud storage services.



- Cost of Switching Vendors:** Simulated cost of switching cloud storage vendors.

