Storiant - Private Cloud Storage with Public Cloud Economics

Open ZFS Conference - November 2014

Storage is Growing…
… Much Faster than Budgets

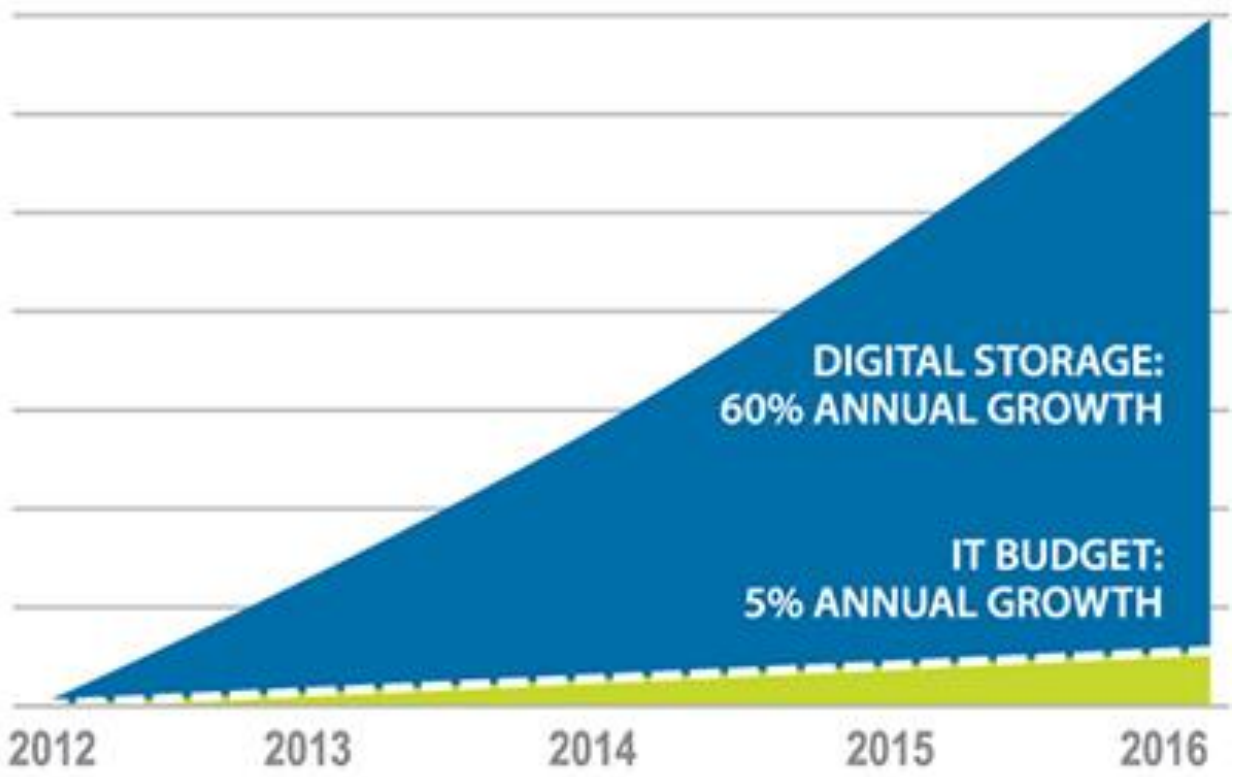

Where Entrepreneurs Thrive

Source: IDC

Storiant Keeps All Your Big Data

- Key Benefits
  - Reliability (in part through OpenZFS)
  - Backup and Archive
  - Compliance - meets SEC rule 17a-4(f) standards
  - High scalability and streaming speeds
  - Lowest TCO
- Economics for your long term data storage
  - Power governor for spin down
  - Commodity hardware and media
- Configure to size and performance
  - Scale out to add storage
  - High throughput supports data analytics

Storiant Market Positioning

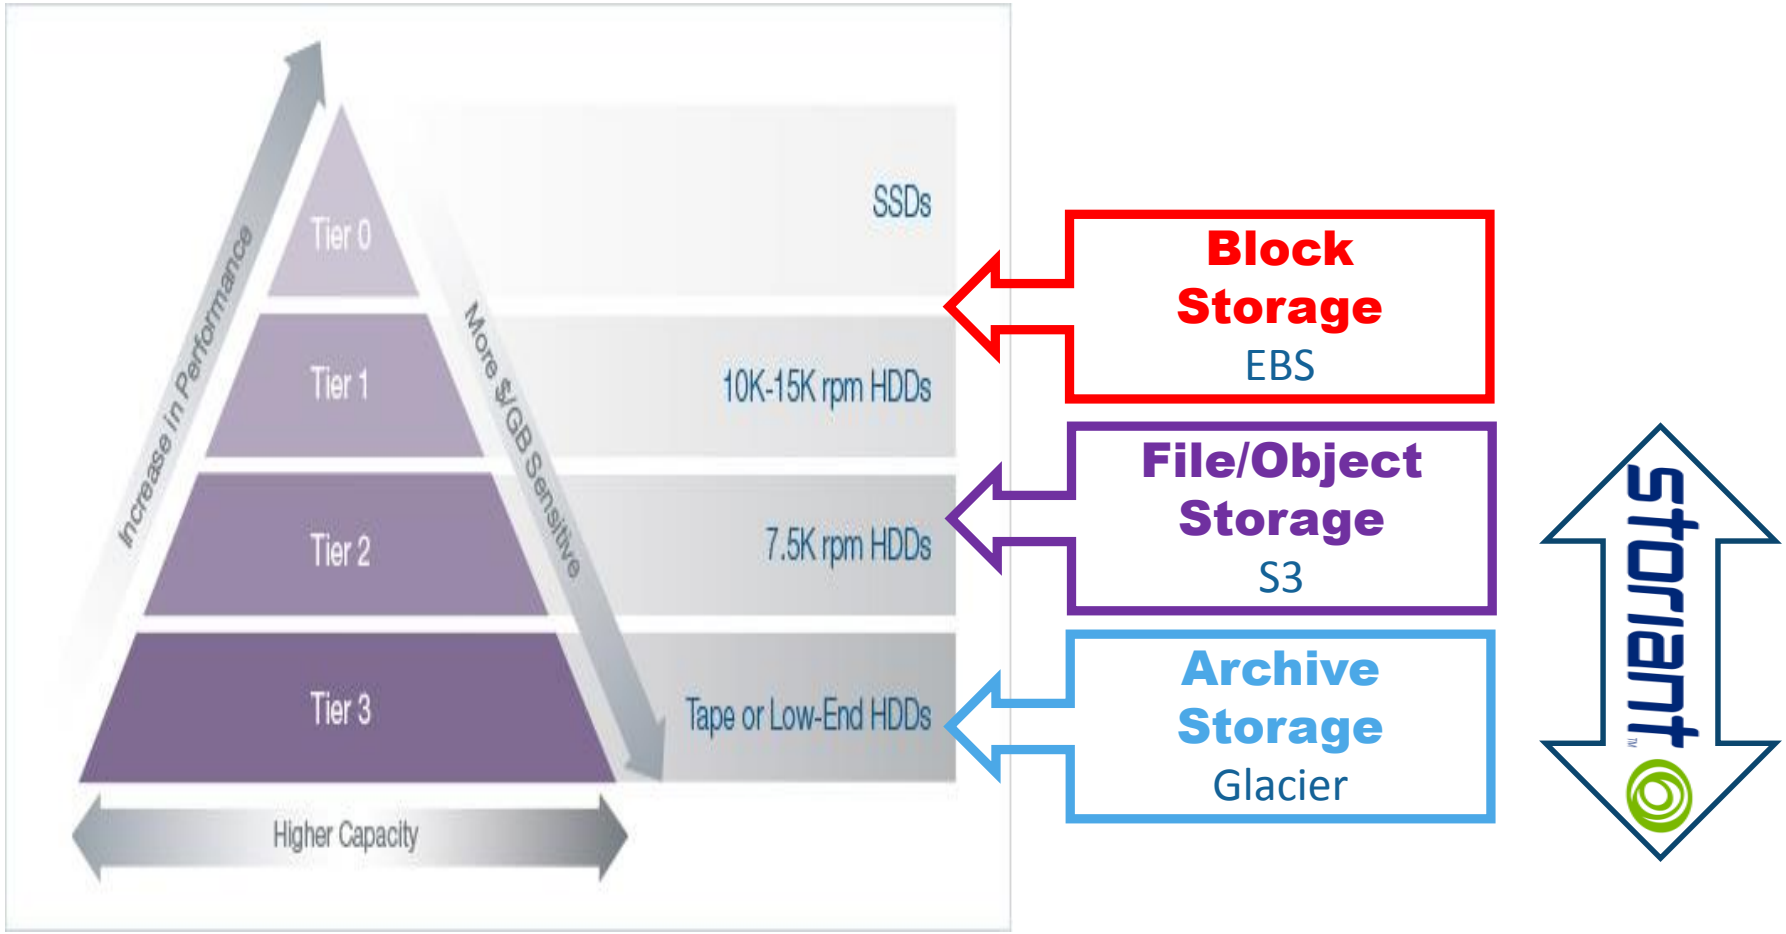

Application Interfaces

Swift Storage
REST Interface

Client Library
C, Java, Python, .NET

Amazon S3
REST Interface

Gateway for
NAS Protocols
NFS and CIFS

Storiant Topology

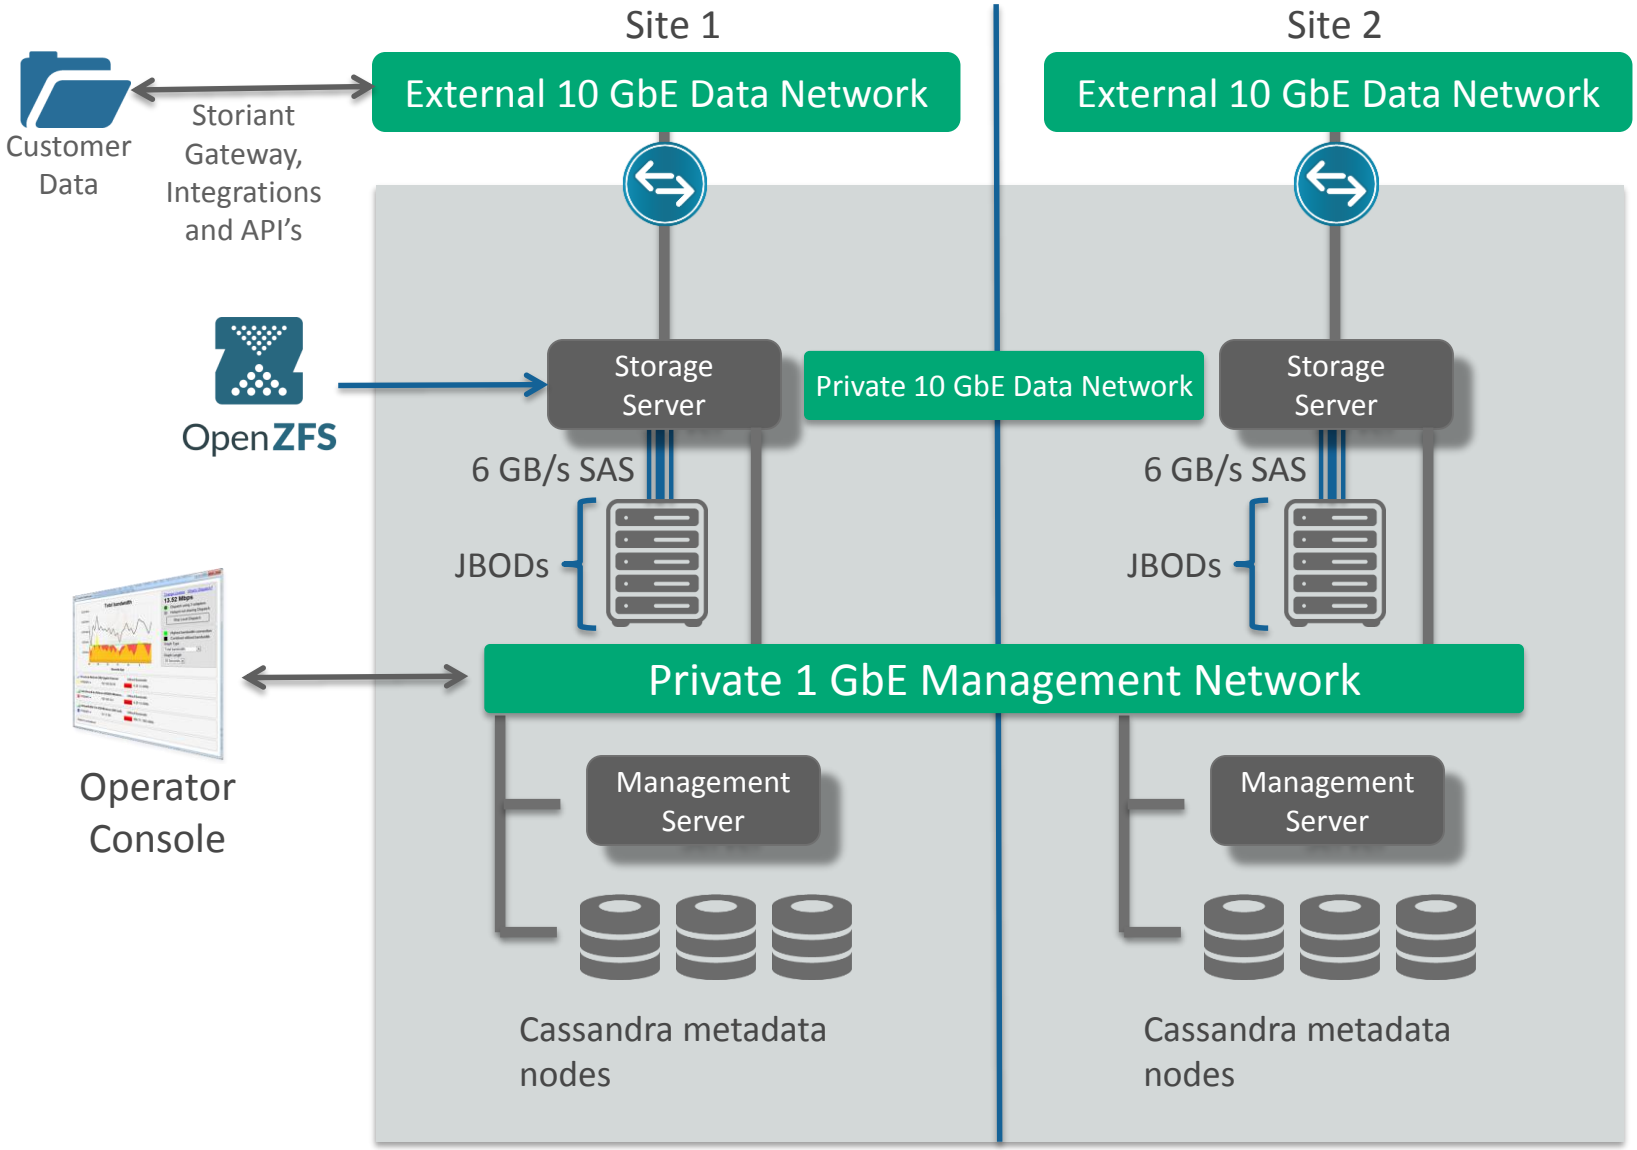

How we use OpenZFS

- We run OpenZFS on Ubuntu independently on each “Storage Server” in our architecture, and manage sets of zpools as the “units” by which we horizontally scale the system
- Each storage server has up to 4 JBOD’s attached
- Our software orchestrates data management across these zpools, managing copies, power control and automatically initiating zpool replace on bad drives
- Use data integrity features, including OpenZFS scrub via maintenance jobs, to find and address bad data, either natively or via multiple copies
- Actively manage power by spinning zpools up and down

How we contribute to OpenZFS

- Advocates for OpenZFS on Linux
  - Pushing the envelope in some areas, e.g. concurrency and disk power management
- Contributed some bug fixes thus far
  - Actively chasing two potential issues right now
- Investigating exposing zpool commands via public API
- Sponsored the beer bash!