

TFDx - ClearSky Data

<http://www.clearskydata.com>

Laz Vekiarides - CTO and Co-founder

Overview

Application mobility and DR
machine data analytics and spunk

“ClearSky’s global storage network delivers enterprise storage, spanning the entire data lifecycle, as a fully-managed service”

Cloud + on-premises
economics
scalability
reliability
security
performance

[TFDx - ClearSky Data - Fragmented Hybrid.tiff]

[TFDx - ClearSky Data - Global Storage Network.tiff]

<http://www.clearskydata.com/clearsky-global-storage-network-white-paper>

<http://www.clearskydata.com/clearsky-data-sheet>

[TFDx - ClearSky Data at a glance.tiff]

Changing how enterprises access data
eliminate storage silos
pay only for what you use - up to 100% useable storage only
guaranteed 100% uptime
multi-site data access without replication
maximum of 30minute response time for Sev 1 and 2 tickets

delivered via consumption-based model
you get charged for only the capacity you use, but apps have all the performance they need
if you delete data, you give back the space ClearSky and are no longer billed for any of it

“Smart Tiered Caching”

ClearSky solves the latency barrier

[TFDx - ClearSky Data - Latency Barrier.tiff]

Edge appliance - 2RU appliance / 24 slots - about 6TB of flash cache
cache is available (on resilient storage), but not copied
SSD failure results in a partial re-warm

ClearSky POPs - distributed and optimised storage
multiple copies to cloud

ClearSky manages the entire data lifecycle
active data - next to your applications, in the metro area near your applications
cold data , backup and DR - multiple copies of data geographically dispersed in the network

Support iSCSI or FC today
Write back cache processed every 10 minutes and pushed to metro cache or cloud

[TFDx - ClearSky Data - Architecture.tiff]

Application Mobility and DR
VMworld 2016 keynote focused on x-cloud mobility, with required data replication
data in the Clearsky network can be accessed from multiple locations without replication

Multi-site availability
Load balancing and disaster recovery

Workload mobility
In-metro and cross-metro
application data can be accessed from other meters
use cases - DR, DC migration, load balancing

Cloud
low-latency / high performance access to CSD service from within customer's virtual DC in the public cloud
application level migration between physical and virtual DCs

Machine Data Analytics Challenge
by 2020, 42% of all data will be machine generated (IDC)
data includes semi- or unstructured log data from
- web/app/ecommerce servers
- mobile apps/devices
- internet enabled things
- network traffic
enterprises want to analyse this machine data to find
- IT security threats
- new revenue opportunities
- ways to increase operational efficiency

Companies use Splunk to analyse machine data
Splunk captures, indexes and correlates machine data in a searchable repository from which it can generate graphs, reports, alerts, and visualisations

Splunk demands high performance and agile storage -
<http://www.clearskydata.com/reinventing-high-performance-storage-for-splunk>

indexing and searching information is data access intensive - requires responsive, scalable storage

Analytics projects: 100s GBs or TBs of new data per day

- need ~23 times more storage, depending on data retention needs
- often tiered (hot, warm, cold, based on age of data)
- storage composes 50% of the cost of a Splunk implementation

Storage needs to scale

- Gartner: "At Splunk customers, data volumes grow gradually for 1st 18months of page, then increase sharply; a hockey stick curve"

Splunk Reference Architecture -

<http://www.clearskydata.com/clearsky-data-for-splunk-enterprise-download>

performance vs cost vs capacity challenge resulted in several storage vendor

"reference architectures"

satisfied Splunk requirements but at expense of introducing technical/operational complexity, and high costs storage related cost can be up to 60% of total Splunk deployment costs

ClearSky for Splunk Enterprise

Indexers simply write to CSD Edge Cache & CSD manages index migration through cache and storage layers - greatly simplifying the solution

Highly consistent ingest performance, cloud capacity, and integrated backup using CSD snapshot technology