



Scale-out DC: Storage trends

Manoj Wadekar
June 3, 2015

Data Center Trends

Explosive growth in Storage

- Big Data/Analytics, Object Store and NoSQL

Disaggregation of compute and storage

- Independent scaling/refresh
- Application-driven tiering

Network bandwidth increasing

- Full bi-sectional bandwidth
- Helping disaggregation solutions

Scale-out solutions driven by TCO

- Storage density, Cost, Performance



EBay Storage Infrastructure

Mission Critical


Balanced Performance

High Capacity

| | OLTP | NoSQL | Big Data | Cloud |
|-----------------------|-----------------------|--------------------------|-----------------------------|----------------------------|
| Topology | Centralized | Local → Disaggregated | Local → Disaggregated | Local → Disaggregated |
| Scale | Petabytes | Petabytes | 100's Petabytes | 10's Petabytes |
| Wire Protocols | FC, iSCSI, NFS | DAS, iSCSI | DAS, iSCSI, Swift, Ceph, | DAS, iSCSI, Swift, Ceph |
| Management | SAN/NAS Management | Cinder, REST APIs | Cinder, REST APIs | Cinder, REST APIs |




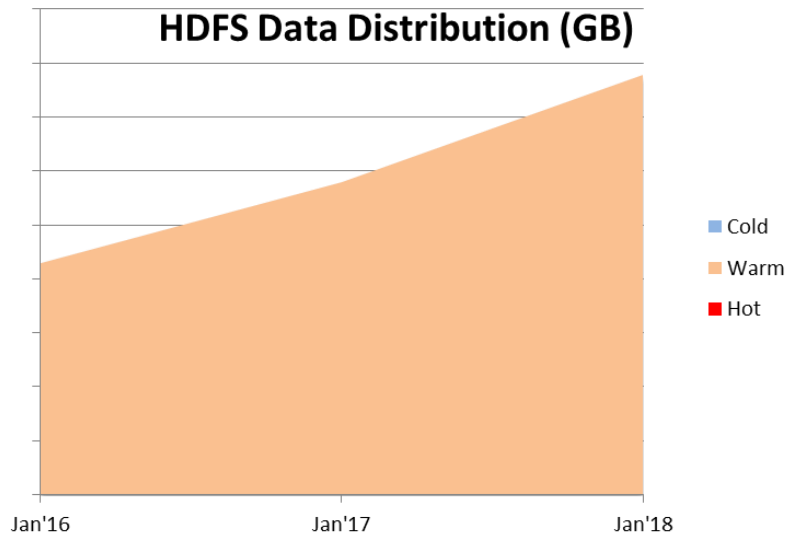
EBay Storage Infrastructure



| OLTP | |
|-----------------------|--------------------|
| Topology | Centralized |
| Scale | Petabytes |
| Growth | Low |
| Wire Protocols | FC, iSCSI, NFS |
| Management | SAN/NAS Management |

| NoSQL | Big Data | Cloud |
|-----------------------|---|-----------------------|
| Local → Disaggregated | Local → Disaggregated | Local → Disaggregated |
| Petabytes | 100's Petabytes | 10's Petabytes |
| Medium | Scale-out Server-Storage High | High |
| DAS, iSCSI | DAS, iSCSI, Object | DAS, iSCSI, Object |
| Cinder | Cinder, Swift, Ceph | Cinder, Swift, Ceph |

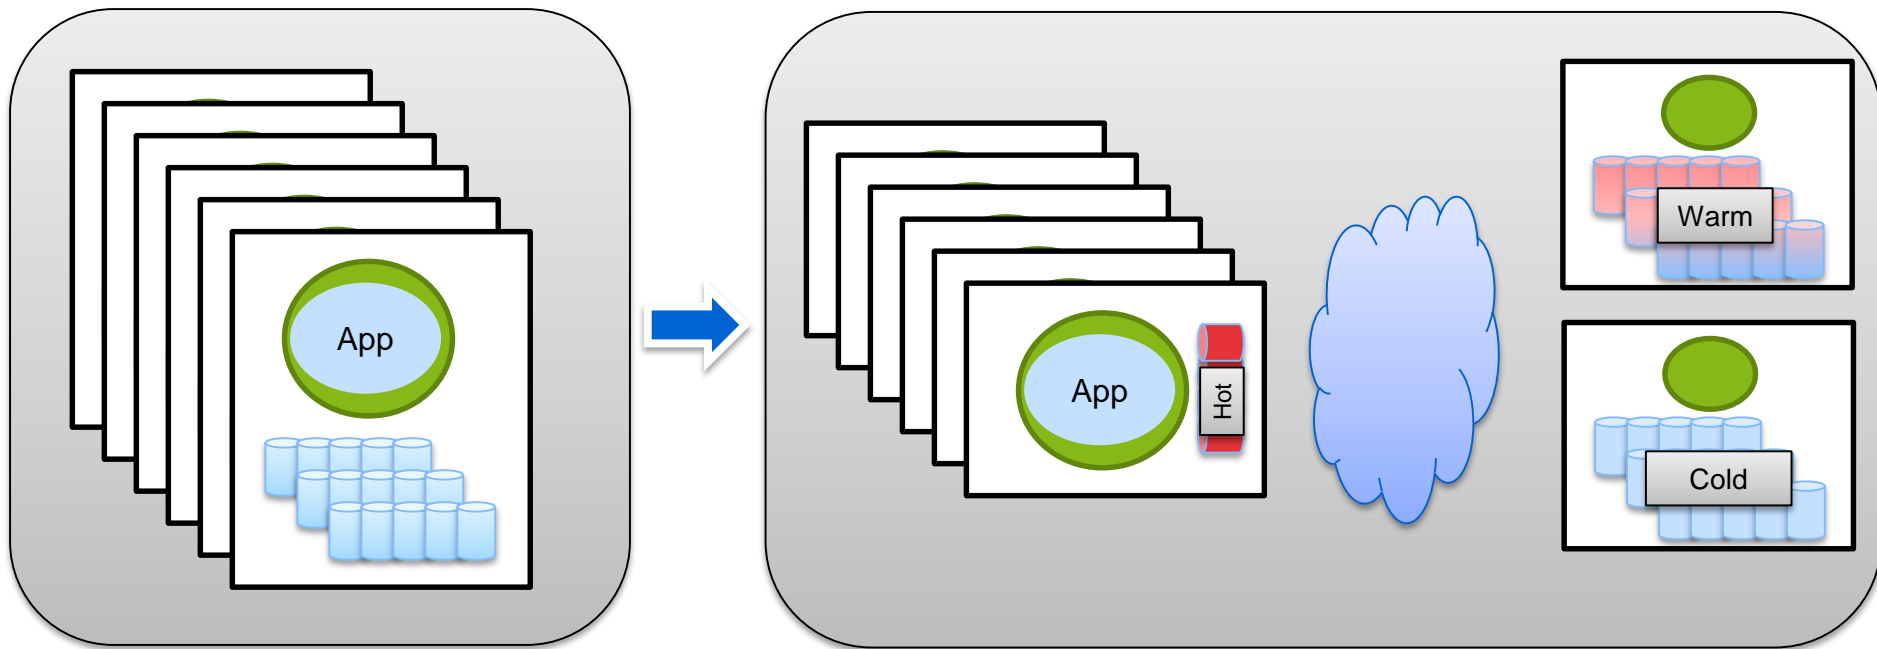
Tiering in Hadoop



Ref: [HDFS Storage Efficiency Using Tiered Storage](#)



Tiered and Disaggregated



Server-Storage Tiers: Characterization

