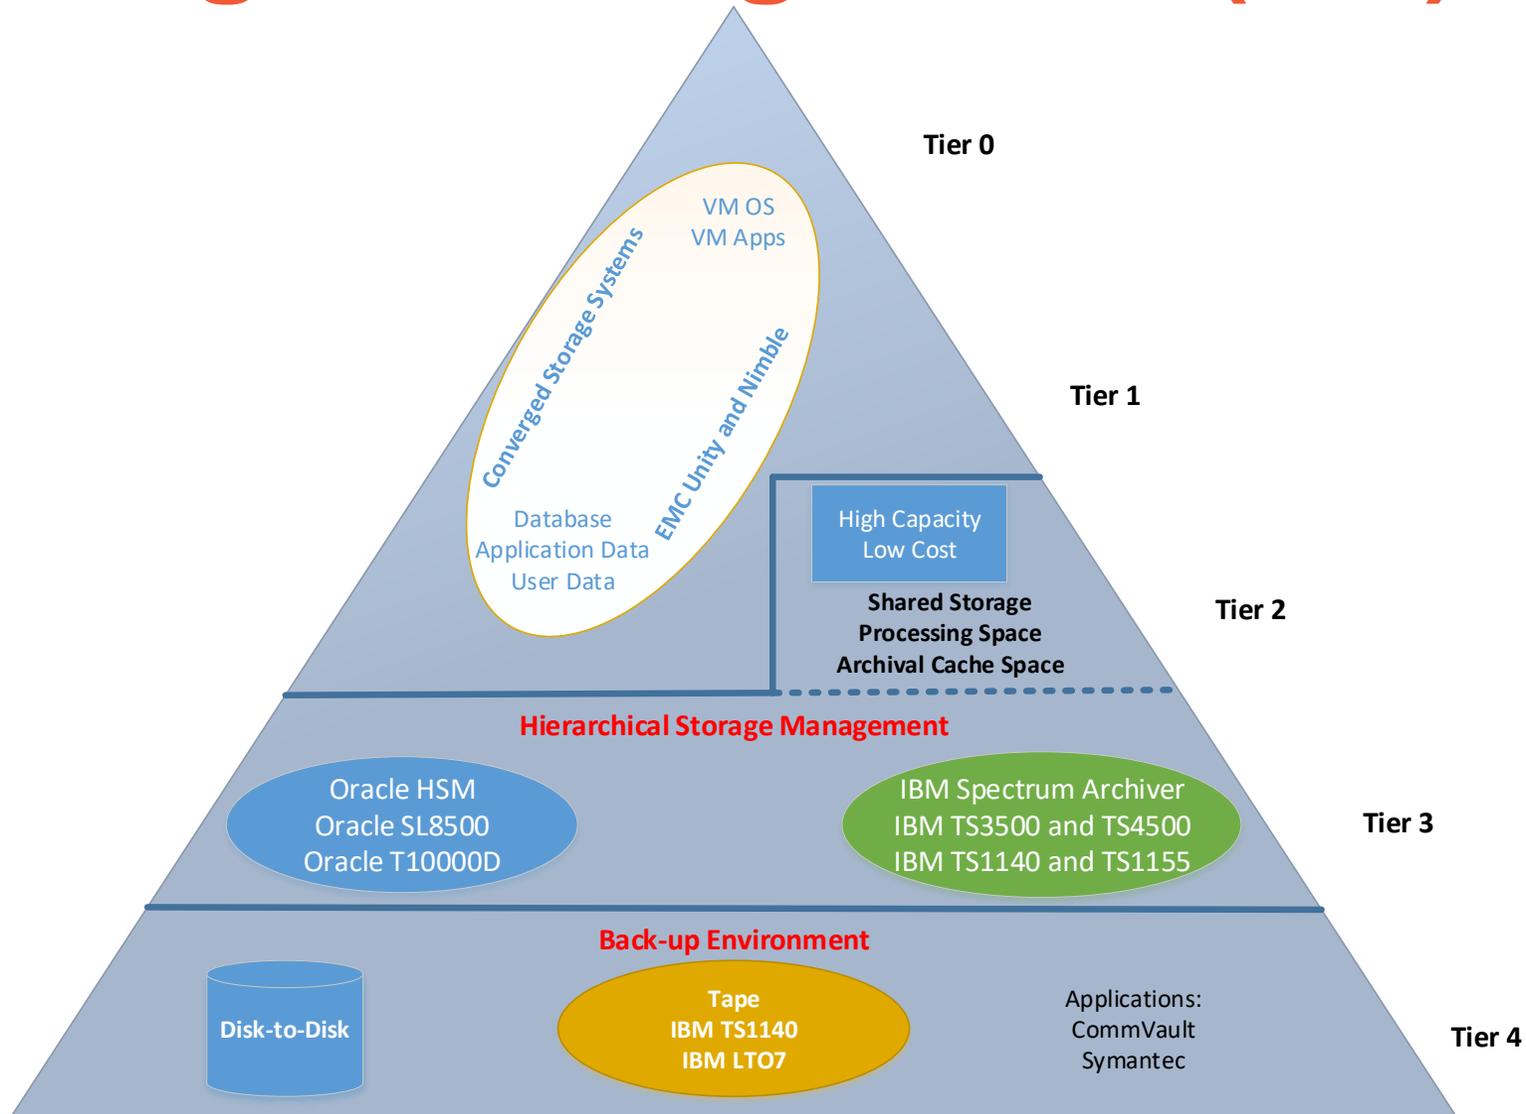

Library of Congress Storage Environment

Update 2019

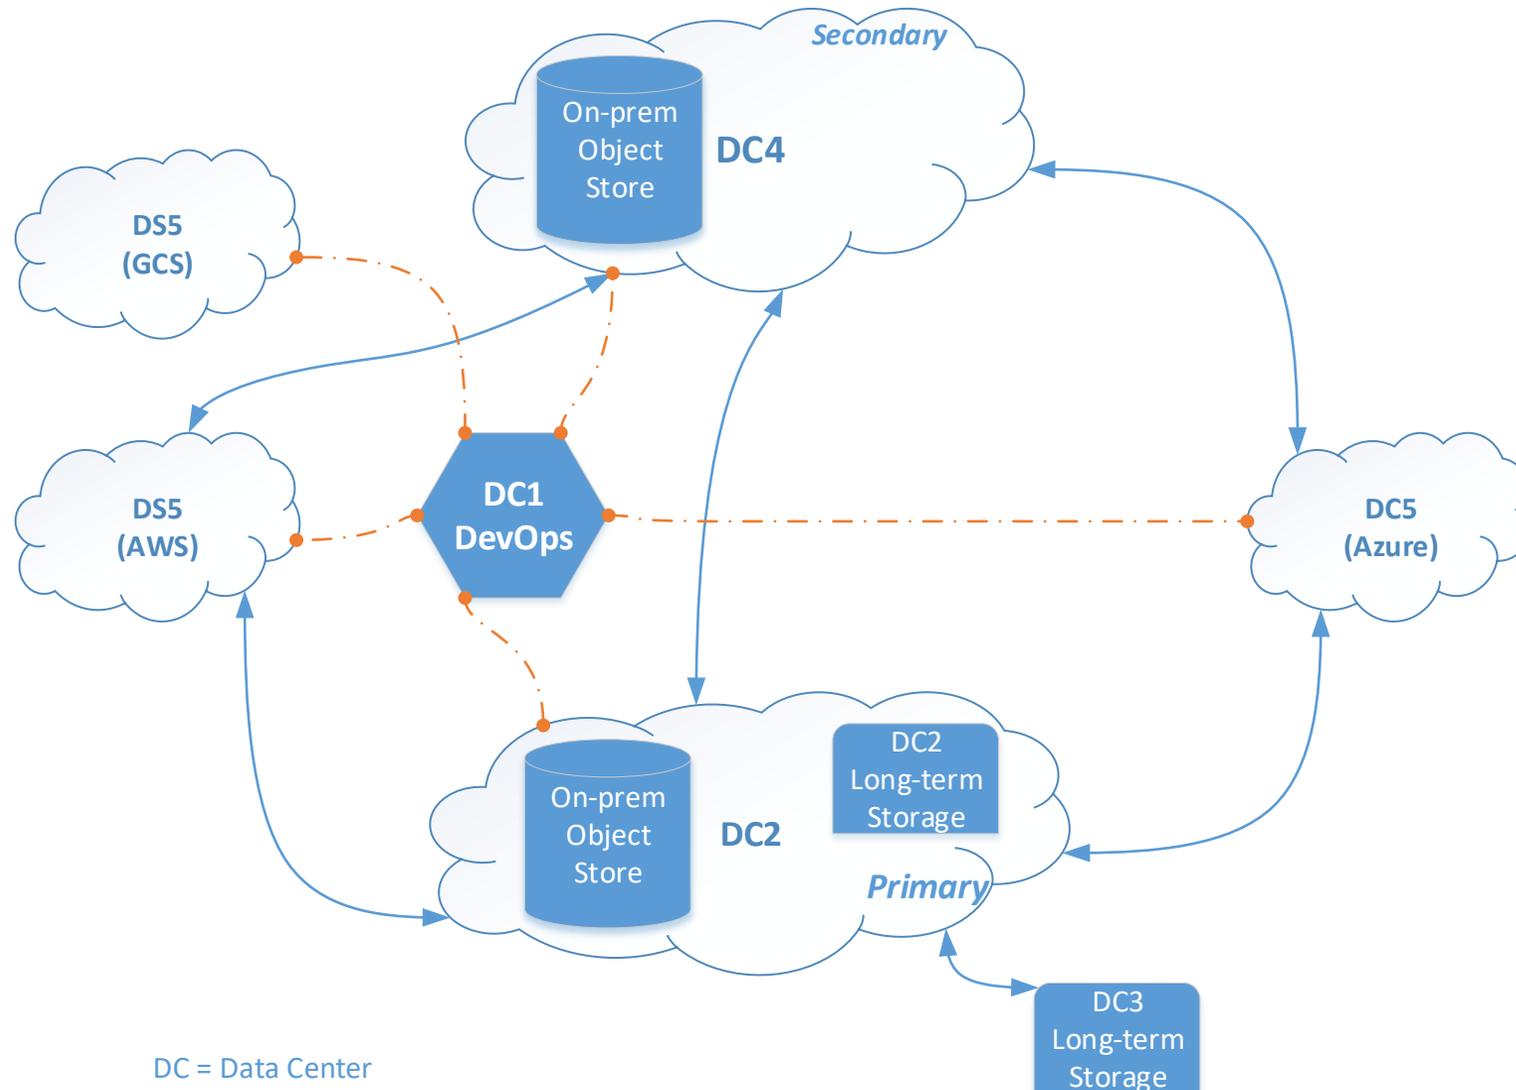


Carl Watts
Information Technology Specialist
IT Services Operations / Operations and Maintenance / Unix Systems

Converged Storage Tiers (old)

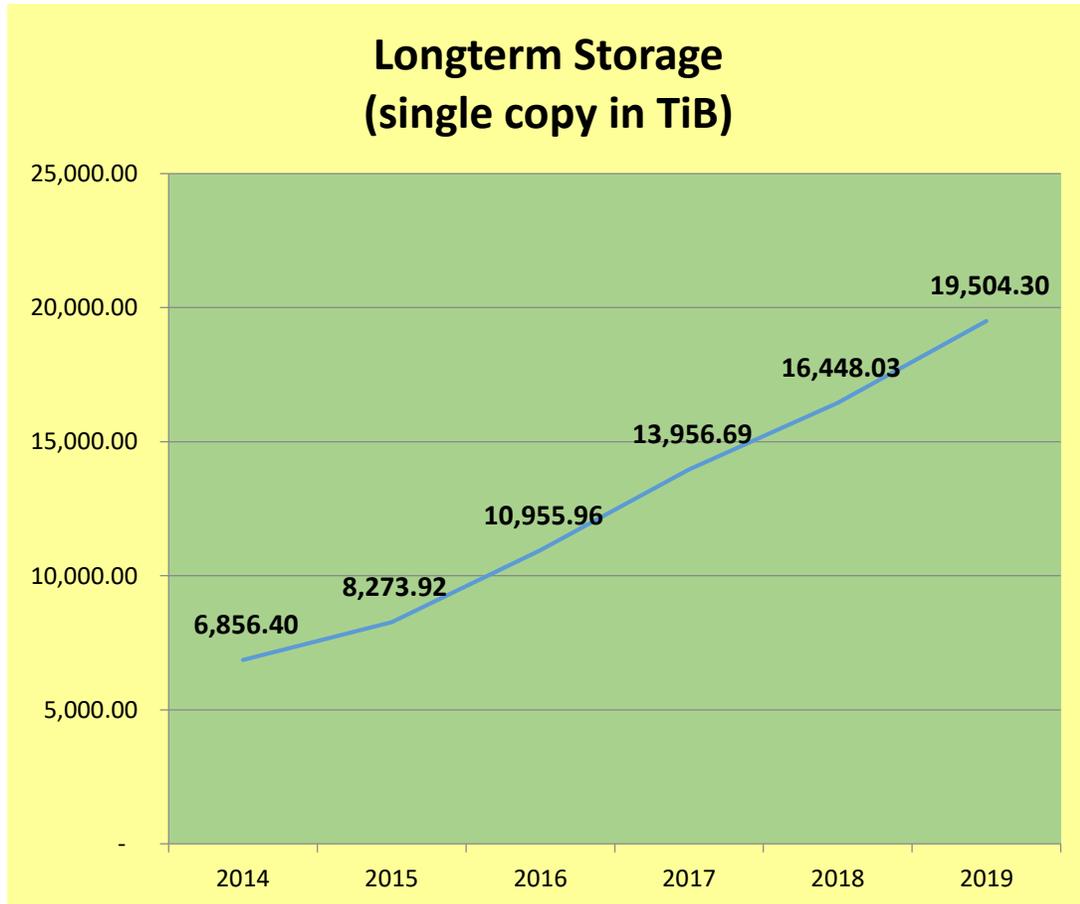


Converged Data Center



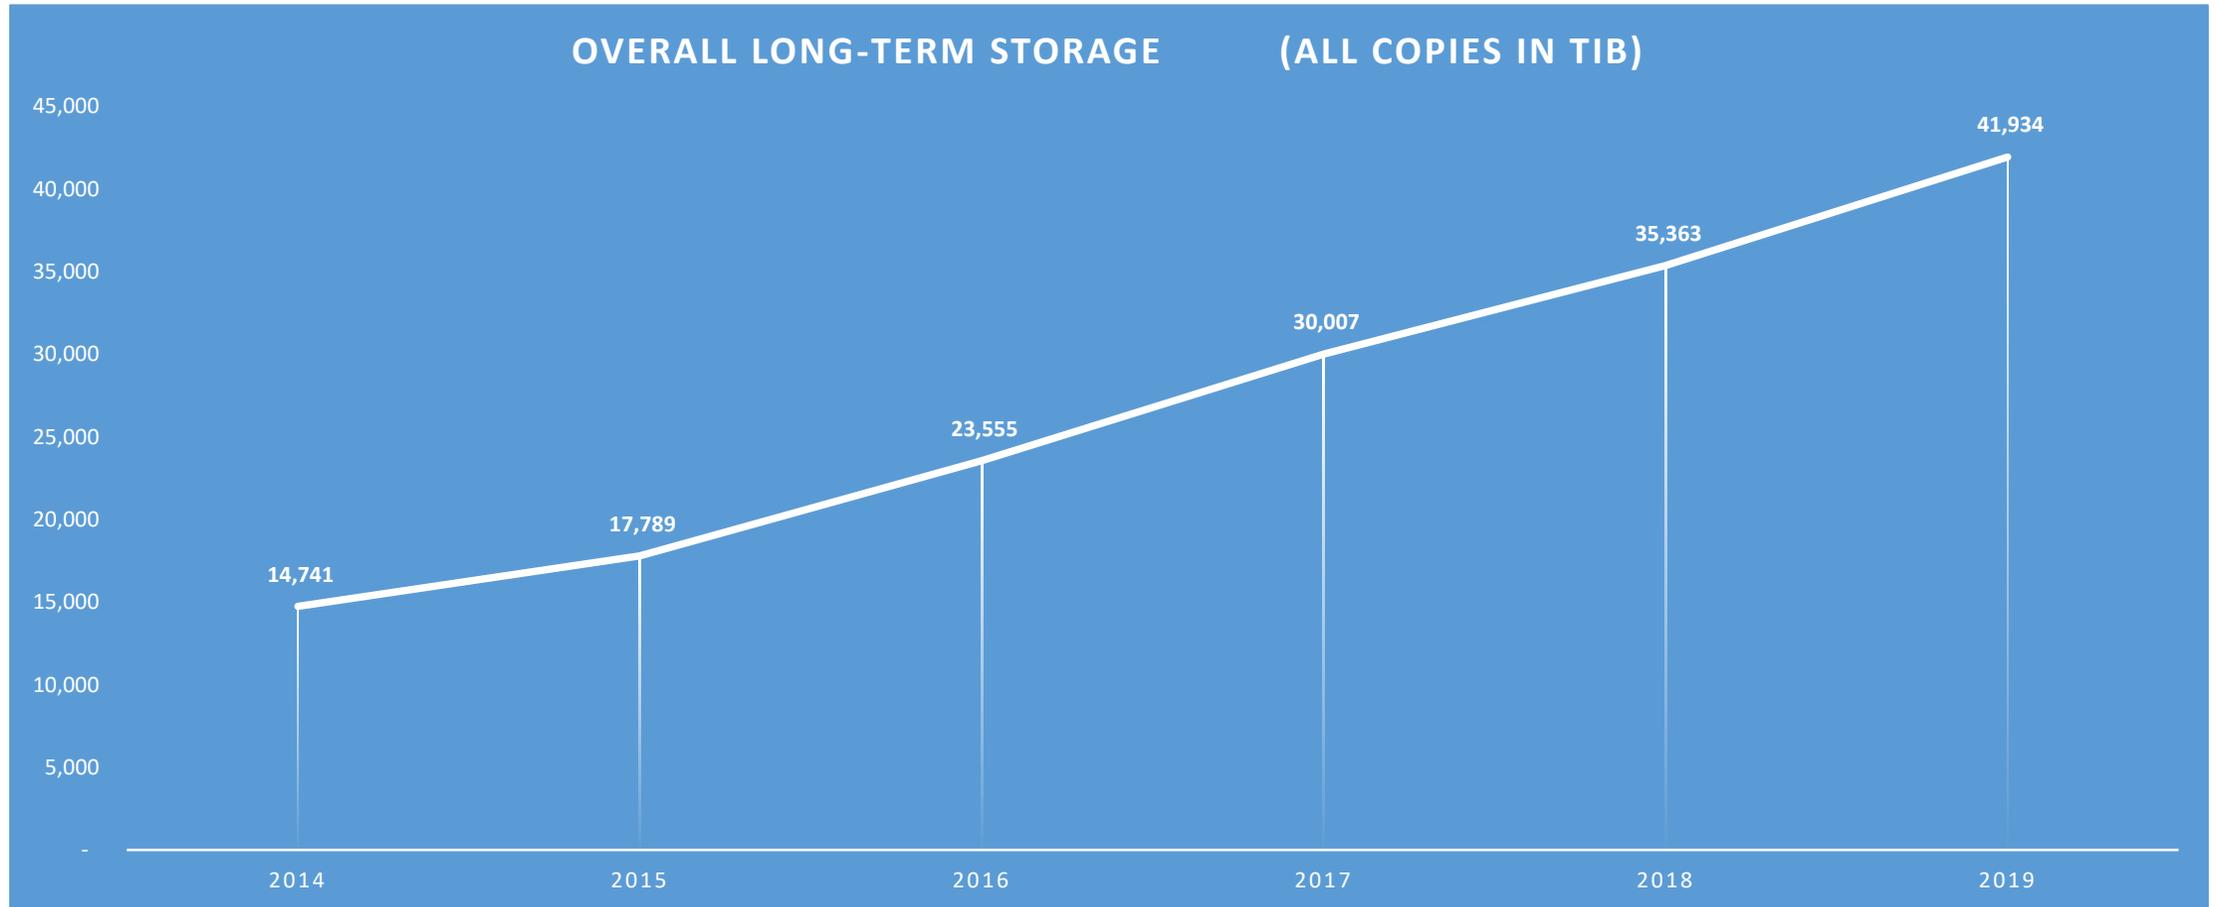
Content Growth – Preservation

Unique File Count:
536M Total Files



Content Growth – Preservation

About 18% annual growth



Content Growth – Presentation

Unique File Count:
374M Total Files

About 18% annual
growth

Access Storage (in TiB)



Annual Terabyte Growth



Migrations Continue

- Data Migrations and propagation are now at a constant churn
- Completed the Consolidation of Preservation Storage
 - Combined resource of three data center to two to reduce cost
- Preparing to Propagate Data Center 2 to Data Center 4
 - Preparing to replicate data to new data center once high-bandwidth network is established
- Migrating Tape Technology
 - Migrated one copy of IBM TS1140 tape to TS1155 tape (2019)
- Propagating Access Storage to AWS
 - Completed 24 of 48 AWS Snowball transfers (over 1.4PB moved in two months)
 - Project is ongoing and should be completed in six weeks
- Propagating Preservation Storage to AWS and Azure
 - Project to start early December to propagate all preservation content to AWS via Snowmobile and Azure via Data Box Heavies

Building out Content Abstraction Layer

- Installation of StrongLink which will become the Content Abstraction Layer
- Content Abstraction Layer (CAL) will provide:
 - Provide a persistent namespace and access method to data
 - Management of the curated data
 - Manage the file fixity and fixity checking
 - Manage the automation of content processing
 - Manage the movement / orchestration of data across multiple
 - Systems
 - Data centers
 - Cloud providers
 - External entities
 - Manage the data migration between old and new storage platforms

Adding On-Prem Cloud Type Storage (STaaS)

- Acquiring a Storage-as-a-Service (STaaS)
 - Replaces:
 - active archive Oracle HSM and Spectrum Archive
 - access storage Spectrum Scale
 - Cost is equal to cloud storage vendor cool/cold (about \$2/TB per month)
 - Access is about the same current NAS and better than HSM products in use
 - No egress fees for access
 - Managed solutions, which lowers staff administration requirements

Content Storage

- **Content** is equal to single copy of a digital object and its associated derivative(s)
- Preservation Copies (currently)
 - Standard Collections – two (2) copies distributed across two (2) datacenters
 - Special Collections – two (2) different platforms holding two (2) copies distributed across two (2) datacenters
- Presentation Copies
 - Currently single online copy
 - Near future – two (2) copies across (2) datacenters
 - Future – multiple copies across datacenters and “cloud” providers

Quad 'P' Dataflow (Proposed)

