Data Logistics

Logistics is the management of the flow of resources from the point of origin to the point of consumption

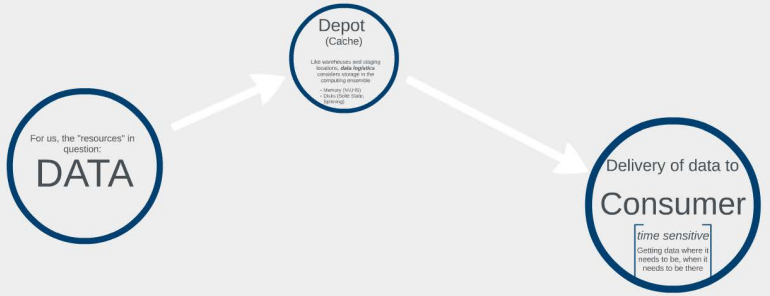

time sensitive

Getting data where it needs to be, when it needs to be there

Depot
(Cache)

Like warehouses and staging locations, ***data logistics*** considers storage in the computing ensemble

- Memory (NV,HS)
- Disks (Solid State, Spinning)

Like warehouses and staging locations, ***data logistics*** considers storage in the computing ensemble

- Memory (NV,HS)
- Disks (Solid State, Spinning)

- Caching is just one form of data logistics
- Prepositioning, prefetching, pipelining

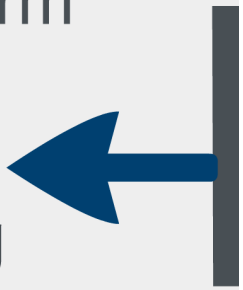

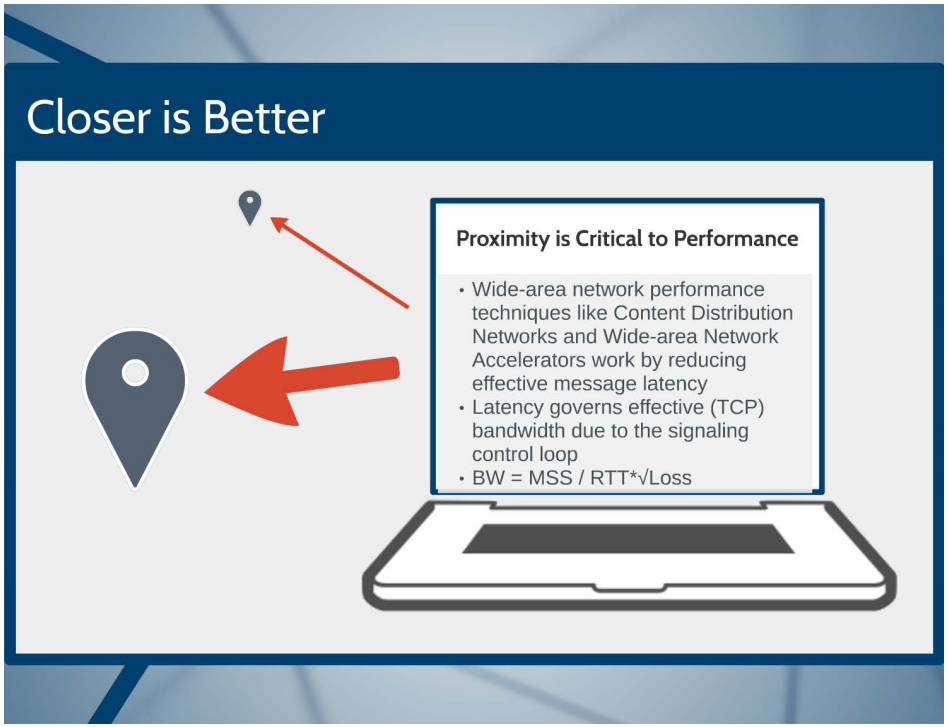

Proximity is Critical to Performance

- Wide-area network performance techniques like Content Distribution Networks and Wide-area Network Accelerators work by reducing effective message latency
- Latency governs effective (TCP) bandwidth due to the signaling control loop
- BW = MSS / RTT*√Loss

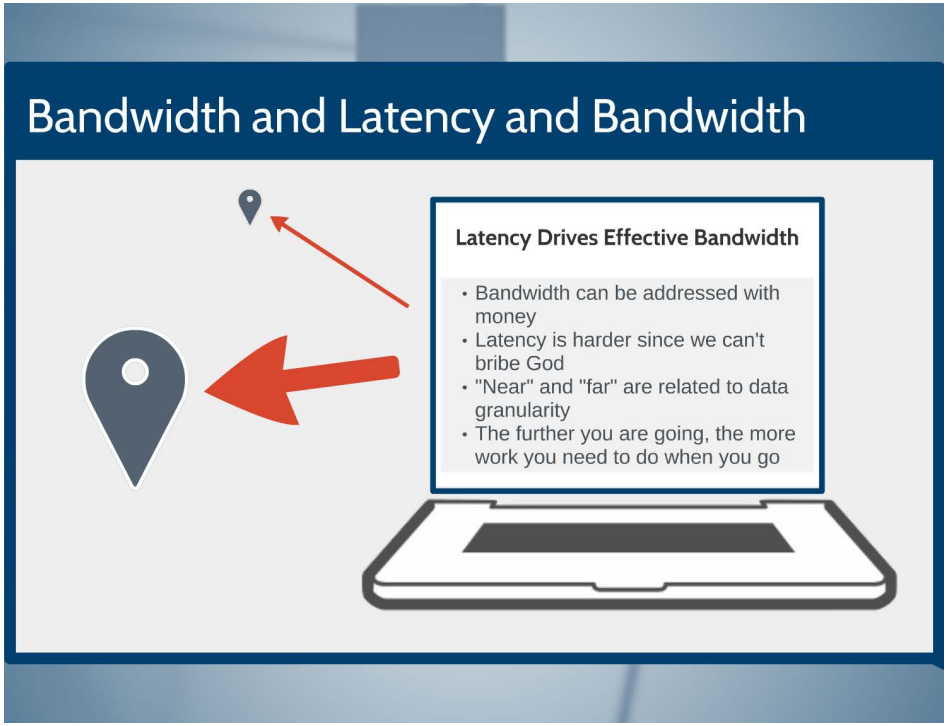

Latency Drives Effective Bandwidth

- Bandwidth can be addressed with money
- Latency is harder since we can't bribe God
- "Near" and "far" are related to data granularity
- The further you are going, the more work you need to do when you go

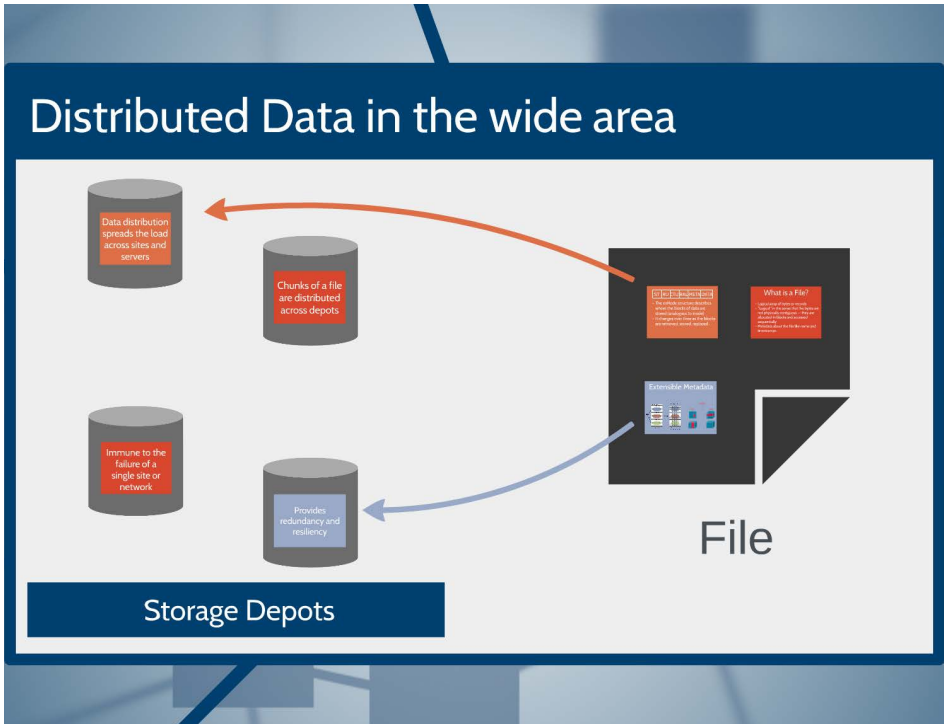

Chunks of a file are distributed across depots

Data distribution spreads the load across sites and servers

Provides redundancy and resiliency

Immune to the failure of a single site or network

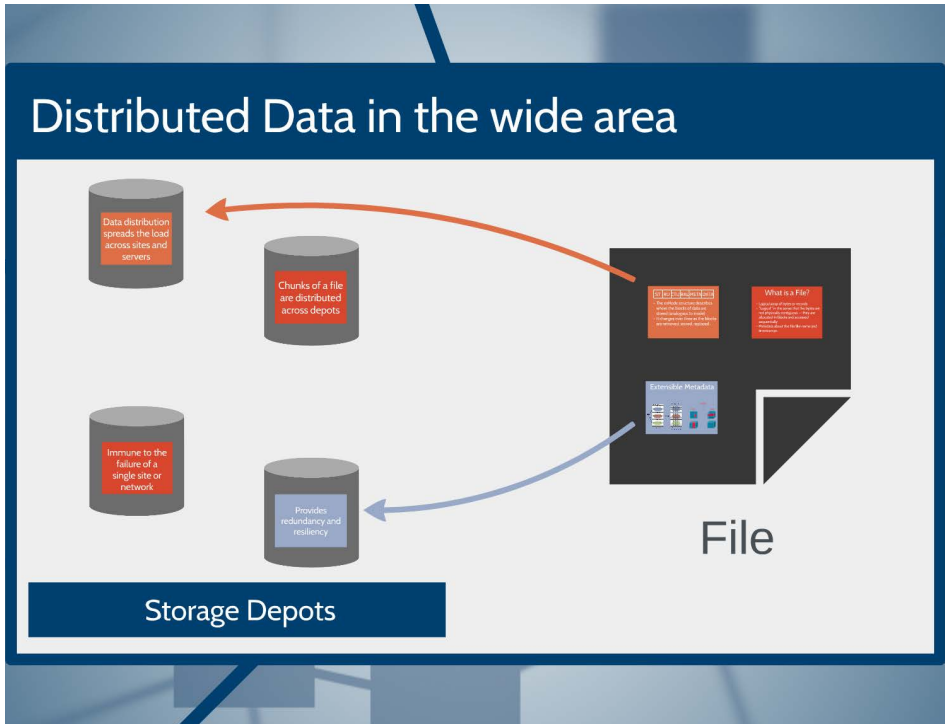

What is a File?

- Logical array of bytes or records
- "Logical" in the sense that the bytes are not physically contiguous -- they are allocated in blocks and accessed sequentially
- Metadata about the file like name and timestamps

ST	RU	CTU	RAL	META	DATA

- The exNode structure describes where the blocks of data are stored (analogous to inode)
- It changes over time as the blocks are retrieved, stored, replaced...

Extensible Metadata

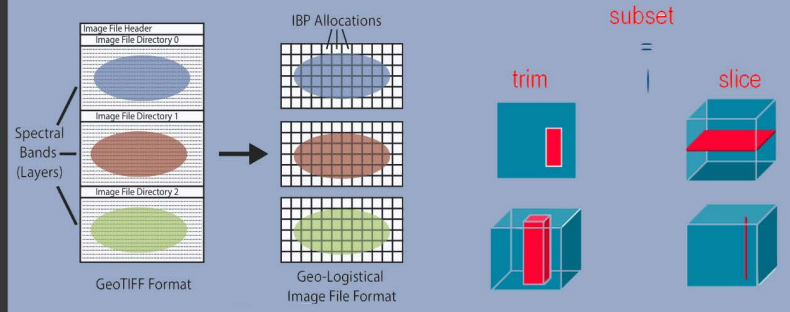

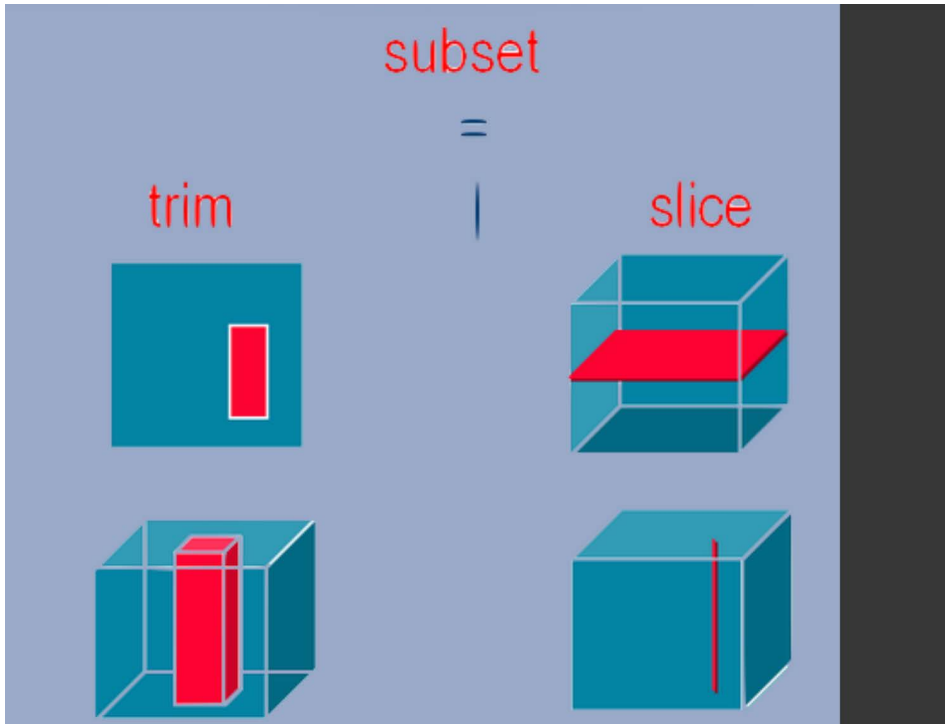

Distributed Data in the wide area

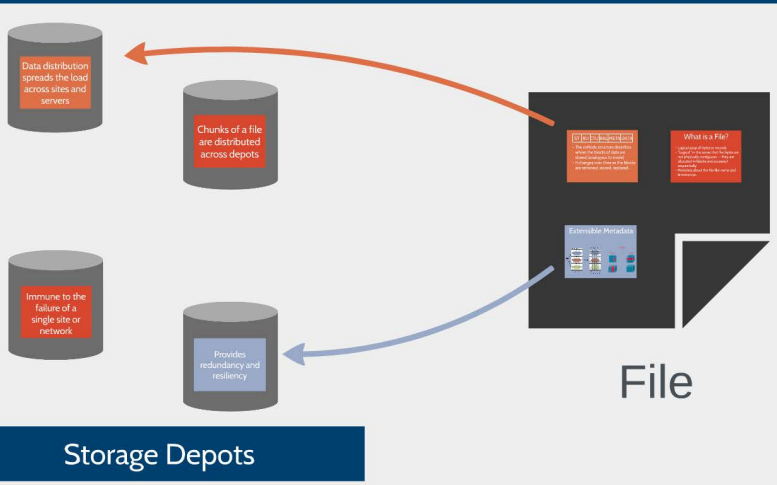

Extensible Storage with the DL Toolkit

The DLT enables simple, edge nodes to serve as storage devices

From small, embeddable systems to large storage-rich servers

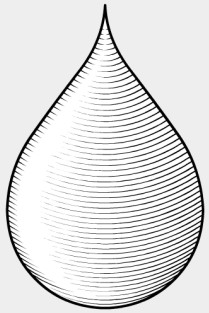

Current data clouds live in a relatively small number of large data centers

Users and sites can add storage to extend the cloud using the DLT

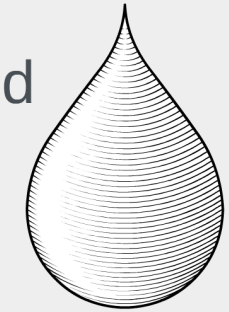

Soft-state Storage

- Time-limited, best-effort storage - times out like a cache
- Distributed storage must be managed this way as reliable reference counting is key to garbage collection

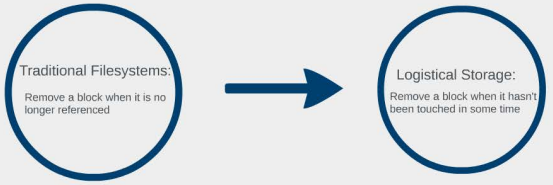

Traditional Filesystems:

Remove a block when it is no longer referenced

Logistical Storage:

Remove a block when it hasn't been touched in some time

The Dispatcher and Data Management

Policy

- What to move? What to refresh?
- Dynamic information based on use (cache)
- Performance policy (first 1MB available within 20ms)
- Local data completely cached (all of Texas cached in Texas)

Mechanism

- exNodes are stored in a directory
- The dispatcher is an active agent
- It updates blocks' time to live ("warming" them)
- It moves or copies blocks based on policy

Policy

- What to move? What to refresh?
- Dynamic information based on use (cache)
- Performance policy (first 1MB available within 20ms)
- Local data completely cached (all of Texas cached in Texas)

Mechanism

- exNodes are stored in a directory
- The dispatcher is an active agent
- It updates blocks' time to live ("warming" them)
- It moves or copies blocks based on policy

Network Topography

- The DLT includes network measurement software based on perfSONAR
- Measure network traffic and distances
- Measure storage performance
- Measure client requests and performance

Like a map with routes and traffic conditions

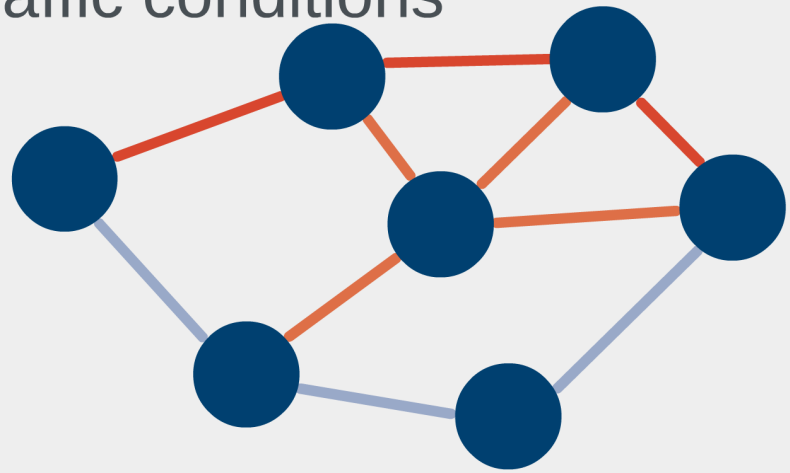

Earth Observation Depot Network

1) An image/view is uploaded to EODN depots and an exNode is created.
2) The exNode is placed in a directory service
3) Using a suitable client (e.g., GloVis, Earth Explorer), the user selects a scene.
4) Finally, the client initiates a parallel download.

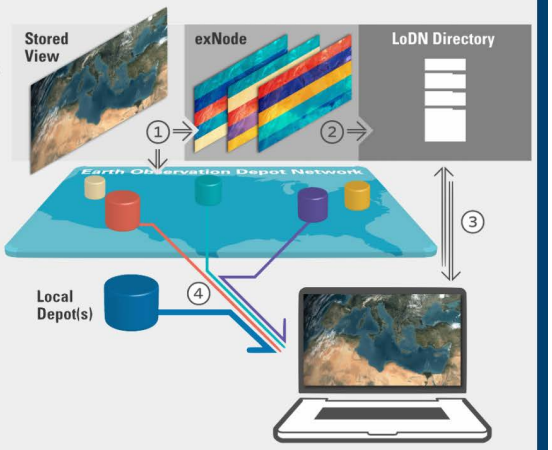

Buffer and Burst

- Accumulate low-rate data into a burst for efficiency of data movement
- Explicit resource allocation with SDN

SESS ID	OFFSET	LENGTH	CRC
s1	0	4	0x6BEF
s2	4	9	0x863B
s3	13	3	0xD329

Session

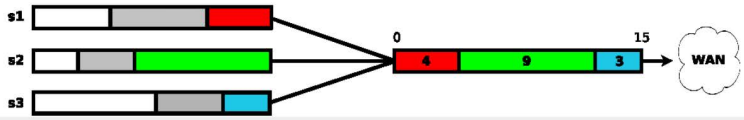

Data Movement

- With the growing diversity of storage, we must consider programming data logistics
- Declarative techniques rather than imperative put or get
- There are silos of optimization that need to be broken down
- Processing elements in various places
- Data movement is the key challenge in both big data and extreme computing