



PERMABIT

The Evolution of a Distributed Storage System

Norman Margolus

nhm@permabit.com


NASA/IEEE MSST 2004

12th NASA Goddard/21st IEEE Conference on
Mass Storage Systems & Technologies

The Inn and Conference Center
University of Maryland University College

Adelphi MD USA

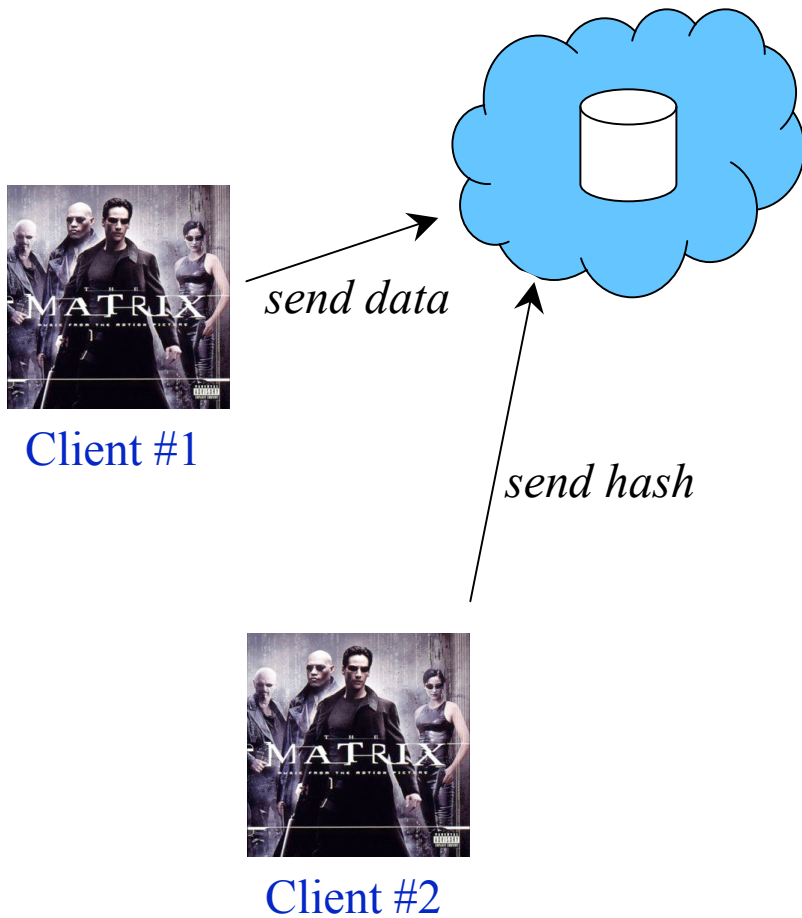
April 13-16, 2004



The Evolution of a Distributed Storage System

- PermaBit makes clustered storage software
- Started with SSP “disk in the sky” model
 - save bandwidth and storage (CAS)
 - legal and societal issues (privacy, IP piracy)
- Near term focus is bulk storage
 - enforced data retention (backup, compliance)
 - geographical distribution (disaster tolerance)
 - low cost (ATA+clustering, scalable, autonomic)
 - familiar file-sharing protocols (portals)
- Long term: bulk storage \Rightarrow “disk in the sky”

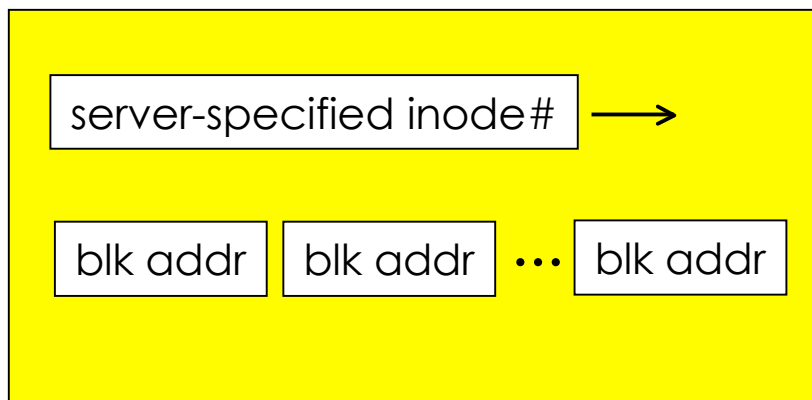
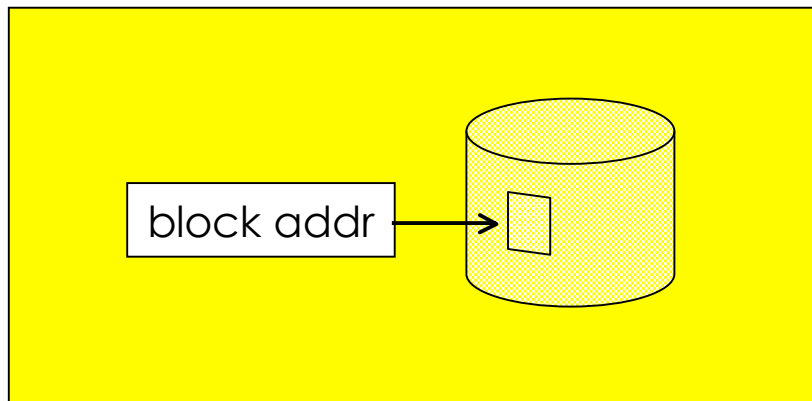
Widely Shared CAS





- *Content Addressed Storage*: shared content is only sent and stored once
- *Self-Encryption*: enables private sharing
- *Non-shared Metadata*: enables access control
- *Deposit Challenge*: must have the data (not just its address)

Object Storage

Unix Filesystem



Permabit CAS

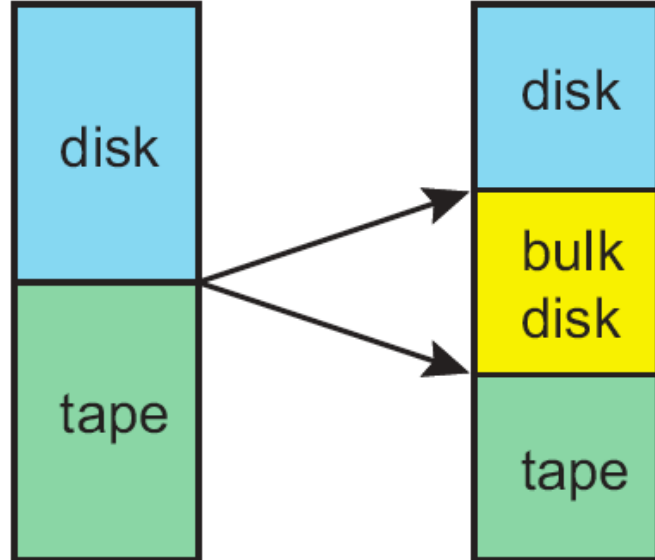


Secure, Self-Named Archive Protocol



Deposit block	<i>(explicit or by hash)</i>
Deposit object version	<i>(client-named)</i>
Extend version retention	<i>(server enforced)</i>
List an object's versions	<i>(specified time range)</i>
Read from object version	<i>(access is controlled)</i>
Delete object version	<i>(if allowed)</i>

Bulk Disk Storage



- Local clustered servers
 - cheap LAN bandwidth
 - striping, cooperative caching
 - scalable, autonomic, immortal
 - NFS/CIFS portals
- Geographical redundancy
 - disaster, maliciousness
- Enforced retention
 - backup snapshots, WORM
- Future: Disk in the Sky!
 - permanence, “main memory”
 - secure/private protocol