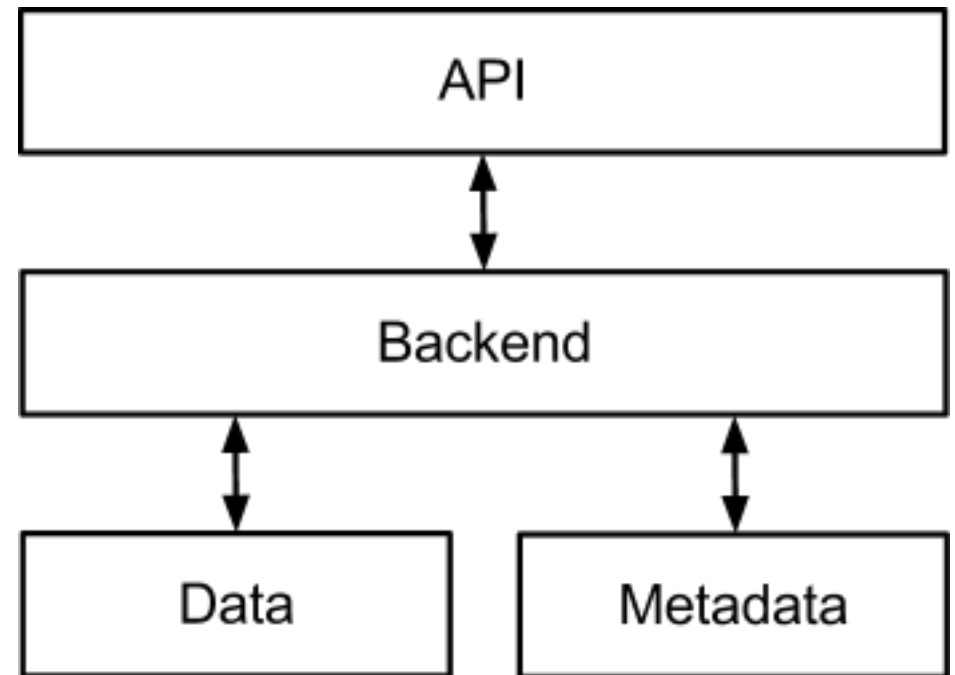


Pithos+

- Based on OpenStack Object Storage API
- Major extensions:
 - Block-based
 - No size limits
 - Deduplication (content hashed)
 - Permissions (sharing/public)
 - Versions (time-variant listings)
 - Metadata-based queries
 - Policies
- Other operations:
 - Partial upload/download
 - Object updates (data may come from other objects)
 - UUIDs
- Layered, modular implementation

Pithos+ layers


- API speaks HTTP (extended OpenStack)
- Backend is a library meant to be used by "internal" code (ex. meta domains)
- Data provides hashmap and block storage
- Metadata is a database of:
 - Nodes (local)
 - Permissions (global)



Now: data goes to filesystem, metadata to SQL
Future: data will be saved in some distributed block storage, metadata to NoSQL

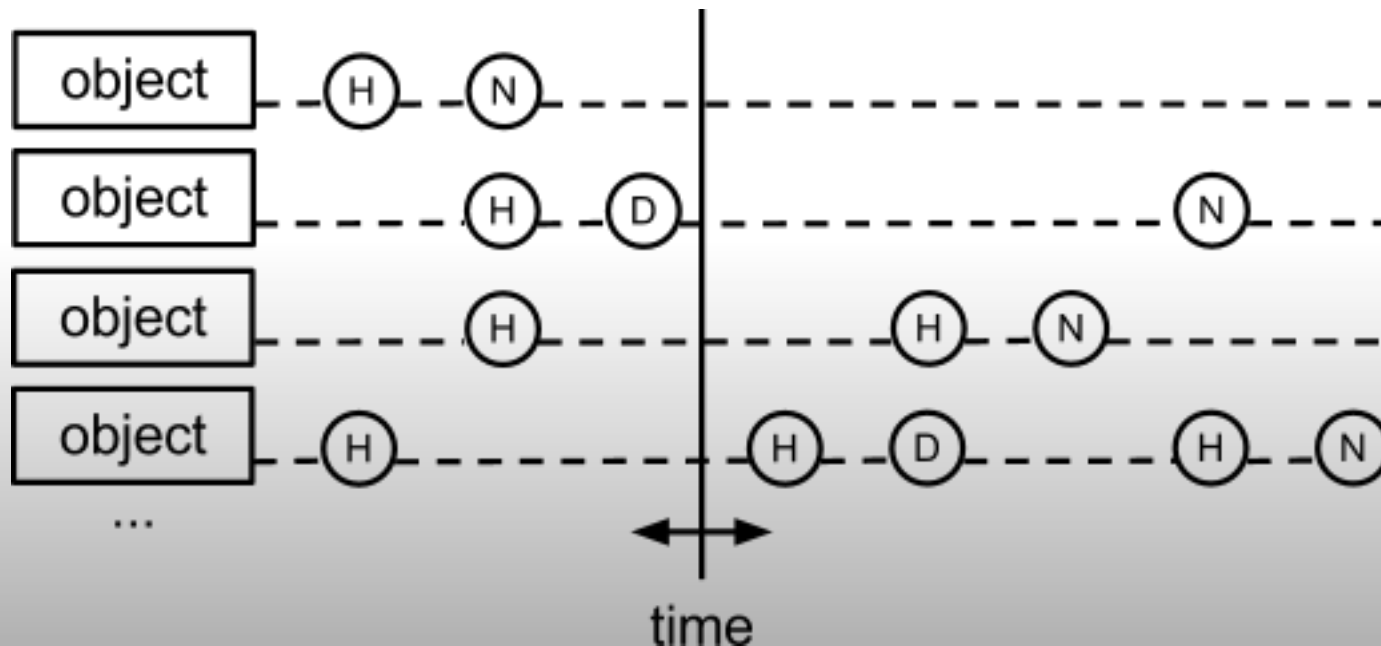
Pithos+ backend nodes

- Nodes are organized in a tree hierarchy (one node per path entry)
- Each node may have multiple versions
- Each version has properties (fixed metadata, like size and mtime) and arbitrary metadata



Pithos+ backend versions

- Each version belongs to a cluster:
 - Normal
 - History
 - Deleted
- For every node, keep statistics (size, population, latest mtime) per cluster



Pithos+ backend permissions

- Permissions supply paths to filter node listings
- The backend decides if the permission path is:
 - Exact
 - Prefix (apply for "path/*")
- Completely independent of nodes module