


# Object Store


Alain Azagury, Michael Factor, Ealan Henis, Dalit Naor, Ohad Rodeh,  
Noam Rinetzky, Julian Satran, Ami Tavory, Lena Yerushalmi

November 2002

**IBM Haifa Labs**




- **Why is an Object Store Needed?**
- **What is an Object Store?**
- **Security**
- **Status**



## Today's SAN Architecture

- **SANs promise non-mediated, shared access to storage**
- **But the use of block storage in SAN's raises several issues**
  - Security (and Protection)
  - Scalability, in particular for allocation
  - Ability to manage at a meaningful level end-to-end



# Security in Today's SAN with Block Storage

- **Security vs. Protection**

- Protection: buggy clients, inadvertent access, etc.
  - Useful inside and outside glass house
- Security: intentional attempts at unauthorized access
  - Essential outside the glass house

- **SAN Security Today:**


- Essentially doesn't exist
- Assume only trusted clients on the SAN

- **Work arounds**

- Zoning/Fencing
  - Hard to use
  - Physical level
- LUN Masking/Logical Unit (LU) access controls

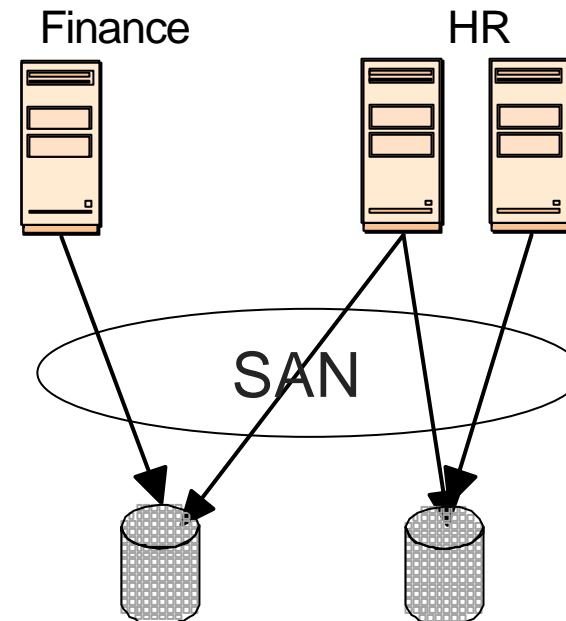

- **At best, provide all or nothing for an LU**

- Too many actively used blocks to provide block-level security
  - Control unit cannot effectively take part in enforcing security decisions
- Too many layers responsible for allocation/management/security



## Security in Today's SAN: No Defense for Human Error

- Storage partitioned between Finance and HR
  - Each division's hosts see only their storage
- Partitioning protected by LUN masking
- Reassign Host from Finance to HR
  - Install new applications
  - Run as different users
- Forget to update the LUN masking on Finance's volumes
- **Oops: HR can access Finance's data**






## What is an Object Store (ObS)?

- Why is an Object Store Needed?
- **What is an Object Store?**
- Security
- Status




## What is an Object Store (ObS)?

- **Raise storage abstraction level**
  - From array of 512 byte blocks to a collection of objects
- **Analogous to a Logical Unit**
- **Allows access to data via *storage-objects***
  - A ***storage-object*** is a virtual entity that groups data a client considers related
    - Similar to a byte-stream file in a flat file-system
    - Size is conceptually unlimited
  - The collection of *storage-objects* are essentially a primitive flat file system
    - No name space - just a flat ID space
    - Security enforcement -- but not management
- **Provides the following basic functionality:**
  - Create or delete an object
  - Read from or write to a byte range within an object
    - Object store manages space allocation within an object
- **Security credentials provided on all operations**
  - ObS validates credentials allow the requested operation on the given object
- **An element of a scalable, networked storage infrastructure**
  - Halfway between a NAS and SAN interface
    - May be implemented on a block device

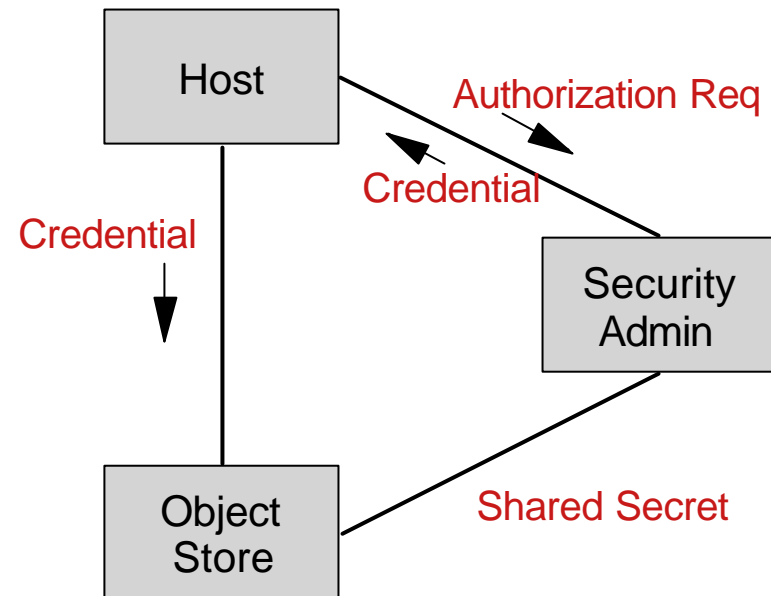



- Why is an Object Store Needed?
- What is an Object Store?
- **Security**
- Status



## Object Store Security


- **All operations are secured by a credential**
- **Security achieved by cooperation of:**
  - Admin - authenticates, authorizes and generates credentials.
  - ObS - validates credential that a host presents.
  - Host - gets credentials from Admin and presents to ObS
- **Credential is cryptographically hardened**
  - ObS and Admin share a secret
- **Goals of Object Store security are:**
  - Increased protection/security
    - At level of objects rather than whole LUs
    - Hosts do not access metadata directly
  - Allow non-trusted hosts to sit in the SAN
  - Allow shared access to storage without giving hosts access to all data on volume






# Security Design Principles

- **Maximal utilization of standards and off-the-shelf mechanisms**
  - Use standard protocols for network security, *e.g.*, authentication, message-integrity, and encryption
  - Use a proprietary protocol for authorization
- **Trust model**
  - Users trust their host's OS
    - Users do not trust other hosts
  - Security Admin and Object Stores are trusted
- **Credential-based access control system**
  - Object access requires obtaining proper credentials
  - No credential revocation
    - However, credentials are short-lived
  - Credentials are cryptographically verified by the object store
    - Credentials contain ID of object store, object ID, rights and time to live
    - Credentials protected with a Message Authentication Code (MAC)
      - *i.e.*, credentials cannot be forged or altered



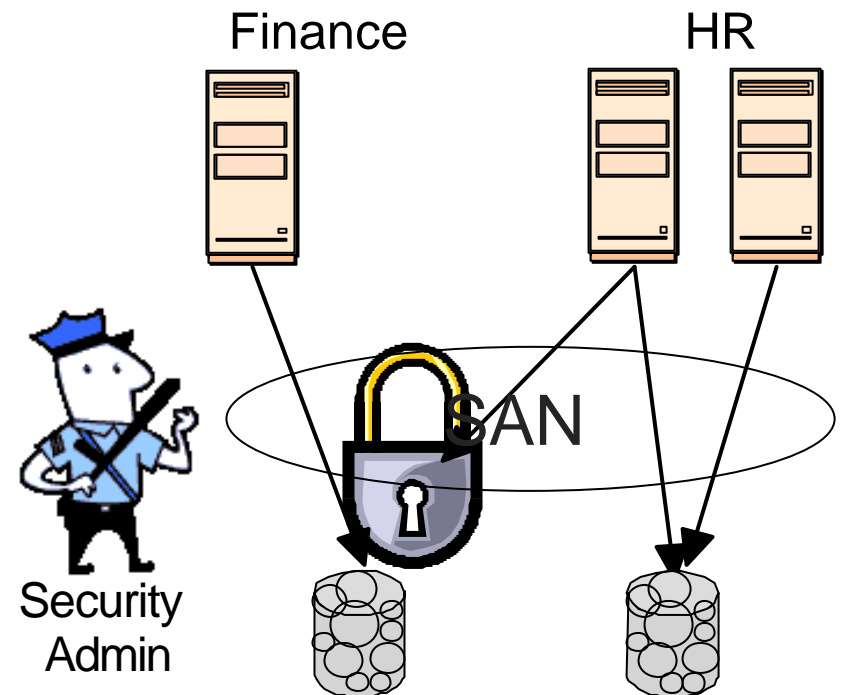
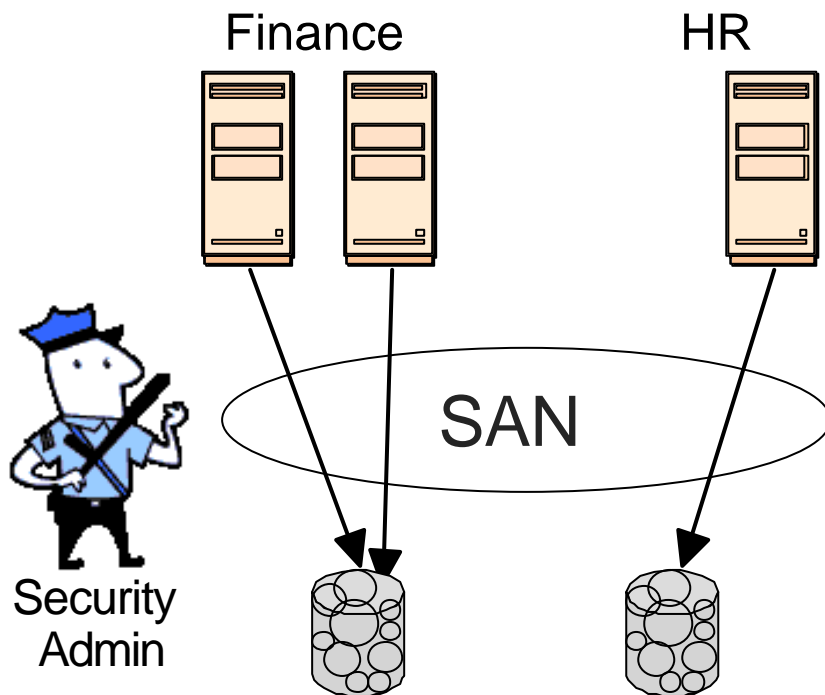
## Security Flow


- **Preliminary**
  - Client authenticates to Security Admin, *e.g.*,
    - Third party: Kerberos in AFS, Verisign, ...
  - ObS authenticates to Security Admin over a secure channel
- **Basic Flow**
  - Client presents authenticator and requests credential from Admin
  - Send requests to ObS with appropriate credential
  - ObS verifies credential and performs operation
- **To improve CPU and bandwidth utilization we use credential caches**
  - At the client: saves accesses to the Security Admin
  - At the ObS per client: saves repeating verification of credentials
    - Only a comparison is required
- **Summary: the critical path remains short**




## Security in a SAN with Object Storage

- Storage partitioned between Finance and HR
  - Each divisions hosts see only their storage
- Partitioning protected by Object Store credentials and security admin
- Reassign Host from Finance to HR
  - Install new applications
  - Run as different users
- Forget to update the LUN masking on Finances volumes
- **No problem: Cannot access data since no valid credential**





- Why is an Object Store Needed?
- What is an Object Store?
- Security
- **Status**



## How Real Are Object Stores?


- **Object Stores**

- First big push
  - Garth Gibson, et al., NASD -- CMU, Panasas
- EMC Centera
  - Claims to be an object store
- Lustre and Object Store Target
- DSF Storage Manager
- . . .

- **Standardization: T10/SNIA**

- **Drivers**

- iSCSI
  - IP access to storage exasperates SAN security problems
- Data Sharing Facility (DSF)
  - A highly scalable research file system which incorporated an object-store like component to ensure local space allocation
- Storage Tank and other SAN file systems
  - Shared access requires SAN security (or trusted clients!)



## Status and Conclusions

- **An ongoing project in IBM Haifa Research Lab (HRL)**
  - Started several years ago
    - Based upon DSF work (started in 1996)
  - A working prototype has been developed
    - Implements our security model
      - All operations are secured by a credential (provided by a third party security Admin)
    - Algorithms and implementations to ensure recoverability of internal metadata
    - Pure software solution
- **Object Store technology promises to address several significant SAN issues**
  - Most significantly security
- **There is no other solution for SAN security if we truly wish to leverage the promise of SANs**