Security for Data-Intensive Computing

SOS-10

March, 2006

Ron A. Oldfield
Sandia National Laboratories

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Some security issues

- Security is not just a technology issue
  - There is a general lack of confidence/trust in existing mechanisms
  - Karen Haines mentioned that a lack of confidence in security is the reason they do not share datasets.
- True security is hard... it must be comprehensive
  - Richard Feynman foiled LANL security by walking through a hole in the fence.
- Security takes all the fun out of computing
  - With few exceptions, security is not part of the original designof most I/O systems

Security for HPC at Sandia

Sandia takes an “isolationist” approach

- Red/Black switching physically separates the system
- Separate networks for different types of computing
  - Open, restricted, classified
- Storage in “vaults”
- Sandia relies on firewalls and standard mechanisms for authentication (e.g., Kerberos)

- Security is the responsibility of the user
  - There was no “File System” option on the Lockheed security exam.
  - Consequences extend all the way to jail time.

Are there differences in the security requirements for computation vs. data-at-rest?

- Yes… well it depends…
- Most data accessed by HPC applications is meant to be “transient” (at least with the PFS)
  - Data is staged onto system for fast input.
  - Data is moved off the system for archival storage or analysis.
- Different mechanisms may provide security on different systems (policies may not change)

Are there opportunities to leverage commercial developments?

- **Yes… I think.**
- **As legislators/institutions define and mandate security policy, we will become more reliant on commercial mechanisms to guarantee compliance.**
- **HPC I/O systems (commercial or otherwise) need to easily integrate these mechanisms.**

Lightweight File System Architecture

Authentication Server
- Use standard mechanisms (e.g., GSS-API) to create/verify *credentials*
- Credentials are transferable
- Distributed at app launch

Authorization Server
- Creates/verifies *capabilities*
  - Coarse grained access controls (*containers*)
- Capabilities are tranferable
- "Immediate" revocation

Storage Servers
- Object interface (blobs of bytes)
- Enforce access-control policy
  - Cache capabilities

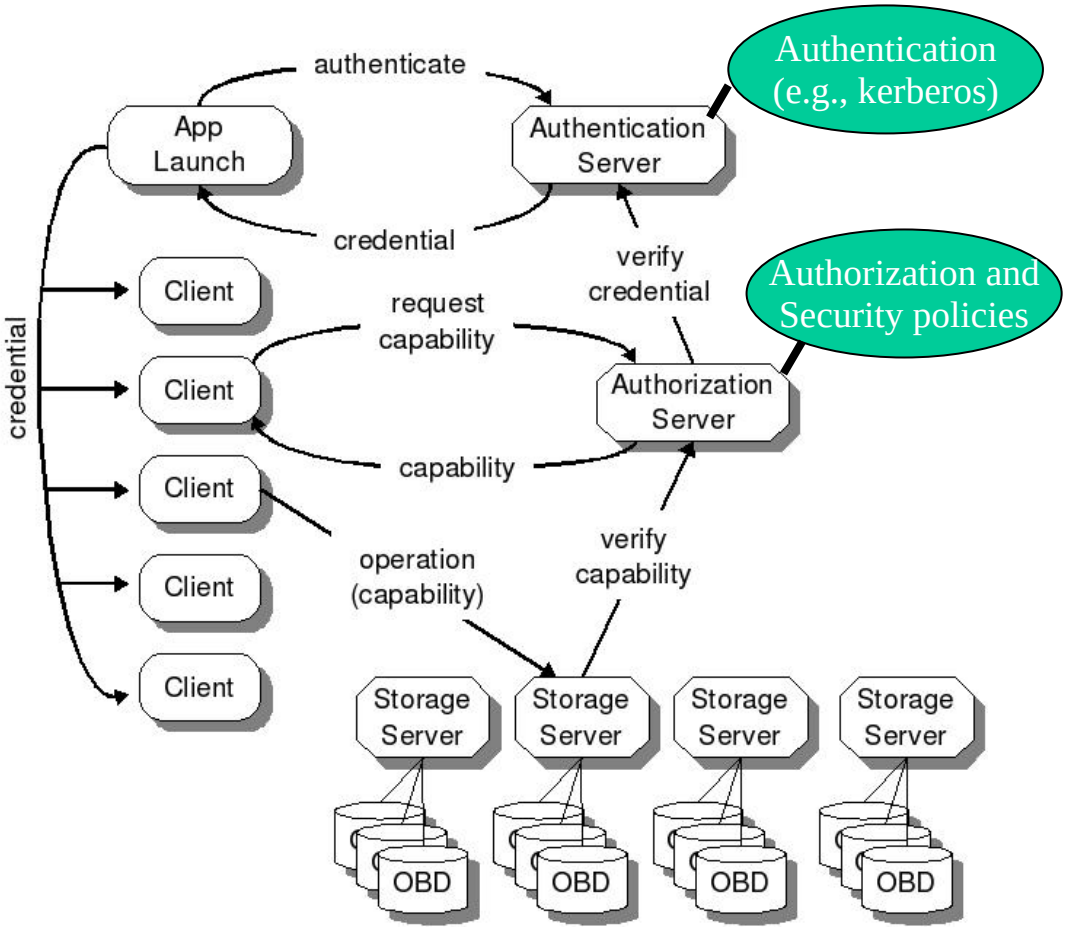

Sandia National Laboratories

HPC security needs that IT industry does not address: what are the gaps?

Security without the performance hit

- Need security mechanisms that are scalable
  - NASD approach works for authorization
  - GSS-API is not scalable (requires security context)

- Level of security required is dependent on classification of the data. Global policy enforcement (typical among security systems) hinders apps that don't require the security.
  - Permissions structure needs to express classification
  - I/O systems need to be more flexible.

HPC security needs that IT industry does not address: what are the gaps?

Security without isolation

- Pathforward goals for file systems
  - Global access
  - Integrated infrastructure for WAN access
  - Security
  - …
- In Rob's panel, security was not mentioned as a challenge for global file systems!