NetApp™

Go further, faster

NFSv4 Multi-Domain Access

Andy Adamson
andros@netapp.com
ABFAB WG, IETF 80 March 2011

Table of Contents

- Motivation
- NFSv4 Authentication Identity
- NFSv4 Authorization Identity
- Multi-Domain Aware File System
- Obtaining Authorization Information

Motivation

- The NFSv4 distributed file system can join multiple administrative domains
- Each with a separate name resolution service
- Each with separate security services
- Share a common multi-domain namespace
- How do we enable user access and ACLs on files across this multi-domain namespace?

NFSv4 Authentication Identity

- NFSv4 supports multiple authentication methods
- Each initiator principal (user or client machine) authenticates to an acceptor principal (always a server)
- Authentication is in the RPCSEC_GSS layer
  - Initiator identities are referred to as 'client principals'
  - Per security mechanism principal representation
- NFSv4 does NOT prescribe how client principals are represented in file systems

NFSv4 Authorization Identity

- File access decisions on the server are based on authenticated client principals' authorization information and the authorization meta-data of the file system.
- Authorization information or *context* consists of
  - User ID
  - Primary Group ID
  - Group ID list of groups that the user is a member of
  - Other privileges, granted authorizations, etc
  - name@domain form for all IDs

NFSv4 Authorization Meta Data

- Occurs at the NFSv4 protocol layer, above the RPCSEC_GSS layer (GETATTR, SETATTR)
  - owner, owner_group, acl, dacl and sacl attributes.
- ACLs are a list of ACEs
  - Each ACE has a 'who' field
- On-the-wire NFSv4 represents users and groups in name@domain form
  - name is the user or group name
  - domain is a DNS domain name.

Additional NFSv4 Restrictions

- The name portion of name@domain MUST be unique within the specified DNS domain
- Every local representation of a user or group MUST have a canonical name@domain
  - It MUST be possible to return the canonical name@domain for any identity stored on disk
- Due to ID collisions, AUTH_SYS can only be used in a namespace which shares a uidNumber and gidNumber translation service
  - Addressed in RPCSEC_GSS version 3

Multi-Domain Aware File System

- Multi-domain support starts at the fileserver where local ID forms need to be able to represent global identities from both local and remote domains.
- Most file systems use a numerical form for on-disk identity representation because user and group names change
  - Searching for and updating file authorization meta-data not trivial
  - Global namespace makes this worse

Local Representation of Global Identity

- Storing <name@domain-name> on disk
  - Not used due to user/group rename issue
- Storing <user-ID@domain-name> on disk
  - POSIX uidNumber, gidNumber is an exmple
    - Unique only within local domain
    - Construct global ID by adding domain name
- Storing <user-ID@domain-ID> on disk
  - Windows Security Identifier (SID) is an example

Multi-Domain Aware File System

- As noted, NFSv4 uses a string form for on-the-wire identity representation
  - Kerberos principal@REALM
  - NFSv4 name@domain
  - These names are *global*
- NFSv4 servers need perform two kinds of mapping
  - Between authentication identity and the authorization context
  - Between the on-the-wire and on-disk authorization identity

Resolving Cross-Domain Authorization Information

- In the cross-domain case, the authoritative source for client principal authorization context is a directory service in the client principal domain
- There are several ways the NFSv4 server can obtain (in a secure manner) the cross-domain authoritative authorization information for a client principal

GSS-API Authorization Payload

- A mechanism specific GSS-API authorization payload containing credential authorization data
  - Kerberos Windows PAC
  - New Kerberos Principal Authorization Data (PAD)
    - draft-sorce-krbwg-general-pac-01
- The KDC gathers the authorization information from the client principal directory service

Directory Service Queries

- An NFS server local domain directory query when there is a security agreement between the two cross-domain directory services plus regular update data feeds so that the NFS server local domain directory service is authoritative for the client principal domain
- A direct query from the NFS server to the client principal authoritative directory service
  - Multi-domain draft specifies LDAP requirements

Use of GSS-API Name Attributes

- The authorization data information SHOULD be obtained via the GSS-API name attribute interface
  - I-D.ietf-kitten-gssapi-naming-exts
- Could be a single attribute for the credential authorization data
- Could be discrete GSS-API name attributes corresponding to the authorization data elements
- Details for these attributes are TBD.