ACE Names and UID/GID/SIDs

Mapping NFSv4 ACE Names to Internal Identifiers

Or

How to Deal with Users and Groups From Muitple Domains on POSIX and Multi-Protocol File Servers and Clients

Nicolas.Williams@sun.com

TOC

- Introduction
- On-Demand Automatic Mapping
- The Concept
- Mapping Login Names to ACE Name Lists
- The Protocol

Introduction

- NFSv4 ACE Names
- POSIX UIDs and GIDs
- Windows SIDs
- Multi-Protocol File Servers

NFSv4 ACE Names

- NFSv4 ACL Entries (ACEs) use string representation of usernames and groupnames with mandatory domain qualifiers (e.g., janedoe@foobar.com, sales@foo.com)
- NFSv4 does not require that file servers support the use of usernames and groupnames from multiple domains on the same file server or shares, nor does it require that file servers don't
- NFSv4's domain-qualified ACE names are needed to properly support Windows domain model

POSIX UIDs and GIDs

- POSIX defines 32-bit ints as internal identifiers for users (UIDs), and, separately, groups (GIDs); POSIX systems normally store user/group references in file ownership/ACLs as UIDs/GIDs
- Status quo: UID and GID namespaces span a single user and group domain, so POSIX file servers have a hard time supporting the use of ACE names from multiple domains
- Metadirectories are not a scalable solution for intranets with lots of M&A activity

Windows SIDs

- Windows uses a structured type for internal forms of user/group identifiers: a domain identifier (OID-like) plus a “relative” identifier (RID) - a 32-bit int – to ID a user or group in that domain
  - RID reuse is not allowed
- This is called a Security IDentifier (SID)
- Domain members generally can lookup SIDs for users and groups
- Supports ACE names from multiple domains
  - Domains must be trusted by file server's domain

Multi-protocol File Servers

- File servers that supports CIFS and either NFSv2/3 or NFSv4 with POSIX clients, and file servers that support NFSv4 with Windows and POSIX clients and/or NFSv2/3 need to support ACE names from multiple domains
- But that means that POSIX systems need to support ACE names from multiple domains
- How?
  - See rest of this slide show

On-Demand, Automatic Mapping

- Well, just map those ACE names to UIDs and GIDs
- How?
  - Establish new mappings on-demand and auto-matically as needed
  - Do this per-POSIX system (for pure NFSv4 worlds) or per-domain, with the MAP-PER_PROG ONC/RPC protocol
- Draft-williams-nfsv4-ace-mapping describes this

The Concept

- First time you see an ACE name, talk to mapping service and get a mapping and cache it [forever]
  - First time the mapping service sees an ACE name [from a “trusted” domain] it assigns an unused UID or GID
- Next time you have the mapping in your cache and you can map the UID/GID or SID to the ACE name or vice-versa w/o further net I/O
  - And you can map SIDs to UIDs/GIDs and vice-versa
- UIDs/GIDs/SIDs (RIDs) MUST NOT be reused
- ACE names MUST NOT be reused
  - Because of caching of mappings

The Concept (cont.)

- UIDs, GIDs, SIDs are called, in the draft, Internal Security IDentifiers (ISIDs)
- ACE name renaming/aliasing is possible, provided that ACE name canonicalization is available
  - MAPPER_PROG supports ACE name canon
- ACE name and/or UID/GID/SID reuse is actually possible if time is set aside to cleanup existing references, but reuse of these SHOULD be avoided in case there are hidden dangling refs

The Concept (cont.)

- A “mapping domain” is a mapping service and DB and all of its clients
  - Mapping domains have names
    - Generally the same names as the default [DNS, LDAP] domain names of their clients
  - Clients cannot change mapping domains
  - All clients of a given mapping domain will see consistent UID/GID assignments to user/group names
- A mapping service and domain are like a partial meta-directory on auto-pilot and w/o the need for synchronization with other directories
  - Very low maintenance

Mapping Login Names to ACE Name Lists

- File servers also need to be able to map GSS-API principal names (akin to login names) to the sets of user and group ACE names associated with those principals
  - Generally hosts need to map login names to their UIDs/SIDs and the GIDs/SIDs of the groups they are members of
  - NFSv4 servers generally deal in GSS-API princ names
- MAPPER_PROG provides a facility to do this

Login Mapping Optimization for Kerberos V

- A KerberosV authorization data type is defined to optimize use of this facility (spec'ed in ASN.1)
  - AD-NFSv4-ACE-NAME-LIST authz-data contains a list of user ACE names and a list of group ACE names
  - AD-NFSv4-ACE-NAME-LIST MUST be sent inside a AD-KDC-ISSUED authz-data container (and, optionally, that inside a AD-IF-RELEVANT)
- By putting this data in Kerberos tickets servers are spared the trouble of separately mapping Kerberos principals to their ACE name lists
  - KDCs need only make the lookup once, when they issue TGTs

Principals, ACE Names, ACLs and ACEs

GSS Principal's ACE Names

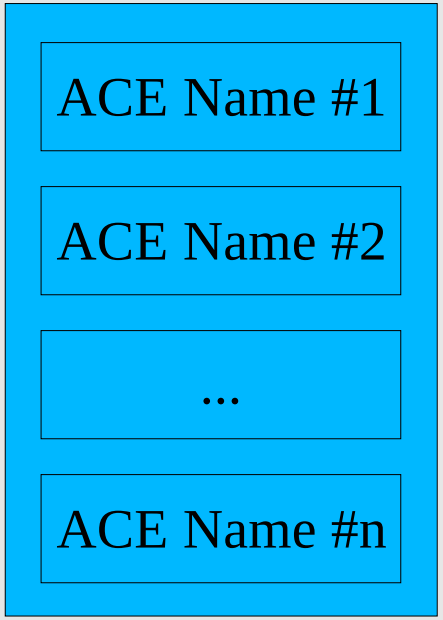

Access Control List (ACL)

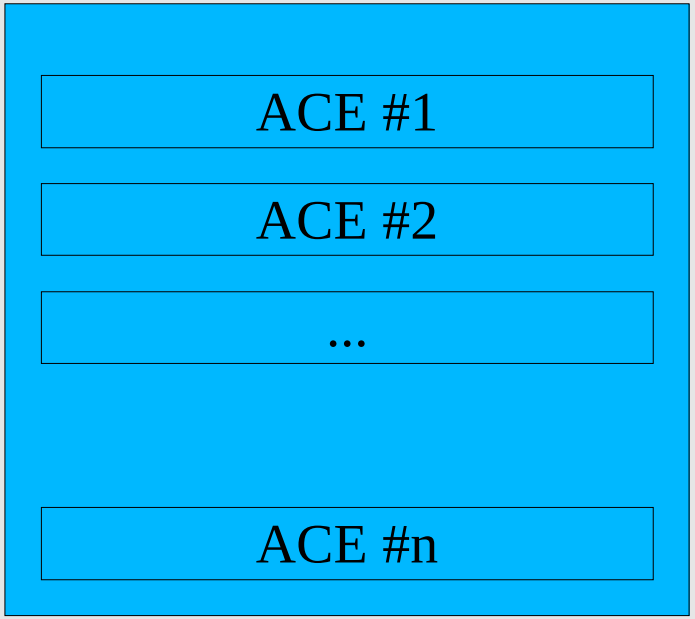

Principals, ACE Names, ACLs and ACEs

- **Principal names** map to lists or sets of ACE **Names**
- ACE **Names** map to internal sec. IDs (ISIDs)
  - Though servers could avoid use of ISIDs by storing ACE Names on disk
- Servers authorize file access through ACL eval.
  - ACL evaluation is simple: for each ACE in the ACL, check if the principal has that ACE Name – if yes, eval ACE, if not, next
- Principal names are for authentication
  - Their ACE **Name** lists are for authorization

MAPPER_PROG

- An ONC/RPC protocol
- Six procedures:
  - MAP_NULL
  - MAP_SECINFO (RPCSEC_GSS sec. triple nego.)
  - MAP_ACE_NAME (map ACE name → ISID)
  - MAP_ISID (map ISID → ACE name)
  - MAP_LOGIN_NAME (login/GSS name → ACE names)
  - MAP_GET_RETIREMENTS (get ACE name and ISID retirement/future reuse warnings)

MAPPER_PROG and Directories

- MAPPER can work with any kind of directory
- For best results MAPPER needs to be able to canonicalize user/group names
  - LDAP schemas need a new class for mapping ACE names to their canonical names
  - NIS doesn't need any changes (but who wants NIS?)
- MAPPER needs to be able to lookup Persons by SID in ActiveDirectory
- Mapping domains should be seeded with pre-existing mappings in their corresponding dirs
  - After that ISIDs in dirs become irrelevant

Remaining Work

- Specify LDAP schema for ACE name canonicalization/aliasing in Windows and POSIX
- Add pseudo-code to for MAPPER servers
- Add normative and informative references
- Fix typos and thinkos
- Make sure ONC/RPC and ASN.1 syntax compile

Security Considerations

- Mapping ISIDs to ACE names is a priviledge op for file servers; NFS clients should not need it
  - Because it allows user and group enumeration
- Diskless clients should not need to map login names / GSS princ names to ACE name lists
- ACE name and ISID reuse is bad
- RPCSEC_GSS MUST be used for mapping requests (mutual auth, integrity or privacy protected)

Questions

- Q/A