Network Working Group R. Mesta
Internet-Draft Sun Microsystems, Inc.
Expires: April 15, 2006 Oct 12, 2005

A DNS RR for NFSv4 ID Domains
draft-ietf-nfsv4-dns-rr-00

Status of this Memo

Abstract

This document describes a new DNS Resource Record (RR) type that will be utilized by NFSv4 clients and servers to determine the domain string to utilize for on-the-wire user/group name attributes and ACL entry information. Discussion and suggestions for improvements requested.

Table of Contents

1. Introduction 3
2. NFS4ID Resource Record Definition 4
3. Example: Using the NFS4ID RR 5
  3.1. NFS4ID: RR Unavailable 5
  3.2. NFS4ID: RR Available 5
  3.3. NFS4ID: DNS Tree Traversal 6
4. IANA Considerations 7
5. Security Considerations 8
6. Acknowledgments 9
7. Normative References 9
8. Informative References 9
9. Author's Address 10
10. IPR Notices 11
11. Copyright Notice 11

1. Introduction

Version 4 of the Network File System (NFSv4) protocol specification [RFC3530] introduces a way for clients and servers to exchange file ownership and ACL entry information as string names qualified with a domain name, whereas earlier versions of the protocol used 32-bit integers for the same type of identifier meta data. Section 5.8 of [RFC3530] defines the generic format for string based identifiers to be "[user|group]@dns_domain".

The string identifier prescribed suggests that the domain to be used for the on-the-wire format be a DNS domain. However, the use of an NFSv4 client's and server's default DNS domain to qualify user/group names would be inappropriate on network configurations that utilize multiple DNS domains, but still use a common user/group name space throughout. This would lead to user/group name recognition failures across the network, at either client or server side, due to potentially mismatched domains. More succinctly, accessing NFSv4 managed files across multiple DNS domains can cause string identifiers to be mapped to "nobody", regardless of whether a common user/group name space is shared or not.

The challenge presented is then to have a mechanism for distributing a common domain configuration for use by NFSv4 implementations that only deal with domain-agnostic identifiers; more specifically, for NFSv4 clients and servers that are administratively controlled by distinct DNS domains.

A natural solution for this type of problem would be to have NFSv4 clients and servers query their configured DNS server for the specific "domain" to utilize for sending user/group and ACL attributes across DNS boundaries. Thus, in a properly configured deployment, having NFSv4 clients access NFSv4 servers on different DNS domains that still use a common user/group name space, would not lead to recognition failures due to the use of the same "domain" for NFSv4 user names, group names and ACL entry information.

A secondary benefit of using a DNS RR for the NFSv4 domain data store is that the resolver's searching mechanism can be leveraged to perform higher level domain traversal. This enables properly configured NFSv4 clients to perform searches on higher levels of the

DNS domain tree until either an NFS4ID RR is found or all possibilities have been exhausted.

This is the solution proposed by this memo.

2. NFS4ID Resource Record Definition

The general syntax for an NFS4ID resource record, whose type is expected to be IANA assigned as per [RFC2929], is:

```
        <owner>    <ttl>    <class>    NFS4ID    "dname string"
```

where:

- <owner>, <ttl> and <class> specify the zone, time-to-live and "IN" respectively, as defined in [RFC1034].

- The RDATA for this record is a string that will be used to specify the domain name to use in 'owner', 'owner_group' and ACL entry information, as defined by [RFC3530].

The proposed RR is meant for use solely by NFSv4; the use of the RDATA field to store additional class information will lead to the familiar sub-typing issues associated with the use of TXT RR's [RFC1464].

3. Example: Using the NFS4ID RR

As a real world example, assume that an enterprise has a top level domain of "example.com" and that it has multiple (perhaps geographically dispersed) DNS domains. For the sake of the current discussion, two domains is more than enough; "foo.example.com" and "bar.example.com". Assume further that NFSv4 has been deployed across these DNS domains and there are active NFSv4 mounts crossing the DNS domain boundary.

3.1. NFS4ID: RR Unavailable

Assuming that no NFS4ID RR's have been configured on either the "foo.example.com" nor "bar.example.com" name servers, then the NFSv4 clients and servers that have active cross-domain mounts should be sending user/group name attributes of the form "[user|group]@foo.example.com" or "[user|group]@bar.example.com".

If a user in client.foo.example.com wanted to access his/her files in server.bar.example.com, the user would find his/her files (seemingly) being owned by "nobody". The reason for this is that client.foo is trying to match server.bar's domain to its own, and since the domains are mismatched, that is, the DNS domain itself is being used for NFSv4 transactions, the client has no choice but to reject the user/group mapping.

3.2. NFS4ID: RR Available

The following configuration would be expected in order to make the NFS4ID RR available in both domains:

The "foo.example.com" domain zone file contains:

```
   $ORIGIN   foo.example.com.

   foo.example.com.        IN         NFS4ID       "example.com"
```

While the "bar.example.com" domain zone file contains:

```
   $ORIGIN   bar.example.com.

   bar.example.com.        IN         NFS4ID       "example.com"
```

Under this scenario, client.foo.example.com would access the user's data in server.bar.example.com; this time, however, the user and group name are of the form "[user|group@example.com" on-the-wire. The client will attempt to match the domain in the in-bound user/group attribute data and will match its own configured domain since

both client.foo and server.bar are utilizing the same domain for NFSv4 transactions.

3.3. NFS4ID: DNS Tree Traversal

Consider the case in which the top level domain zone file has the following NFS4ID entry:

```
   example.com.            IN         NFS4ID       "example.com"
```

As previously stated, the lower level DNS domains, "foo.example.com" and "bar.example.com", can each define their own NFS4ID RR's in order to override the NFS4ID record defined by the top level domain. To continue the example, assume that an NFS4ID record is only defined for domain "foo.example.com" and it is defined to be:

```
   foo.example.com.        IN         NFS4ID           "foo.foo"
```

Assuming the NFSv4 clients' /etc/resolv.conf 'search' parameter has been properly configured, an NFS4ID RR lookup in the "foo.example.com" domain will yield the string "foo.foo", whereas a lookup for the NFS4ID RR in the "bar.example.com" domain, will not yield any value and will propagate to the higher level domain as "example.com"; at this point, the string "example.com" will be returned for NFS4ID RR lookups in domain "bar.example.com".

4. IANA Considerations

IANA is requested to allocate RR type code TBD for NFS4ID from the standard RR type space.

5. Security Considerations

There are two main security considerations for this facility:

- o Denial of service attacks where clients and servers are made to disagree about their default NFSv4 domain and so ACL and file/

directory ownership manipulation can be made to fail.

- o Redirection attacks where a client is forced to use a different domain than it was otherwise intended to use while a multi-domain server can understand and distinguish between users (and groups) with the same names but in different domains. In this attack a user might be fooled into granting access to a file or directory to the wrong user or group. For example, a "chown joe somefile" command might be intended to reference "joe@one.domain" but the client may be made to use a different domain to qualify "joe", thus changing the ownership of 'somefile' to "jane@some.other.domain".

The latter is of particular concern as servers capable of operating in more than one domain are feasible and likely already exist.

The use of DNSSEC should foil both of these attacks, and thus, we recommend its use.

6. Acknowledgments

David Robinson, Spencer Shepler, Nico Williams, Bill Sommerfeld, and Olaf Kolkman.

7. Normative References

[RFC1034] Mockapetris, P., "Domain Names - Concepts And Facilities", RFC 1034, Nov 1987.

[RFC1464] Rosenbaum, R., "Using the Domain Name System To Store Arbitrary String Attributes", RFC 1464, May 1993.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2929] Eastlake, D., Brunner-Williams, E., and B. Manning, "Domain Name System (DNS) IANA Considerations", RFC 2929, Sep 2000.

[RFC3530] Shepler, S., Callaghan, B., Robinson, D., Thurlow, R., Beame, C., Eisler, M., and D. Noveck, "Network File System (NFS) version 4 Protocol", RFC 3530, April 2003.

8. Informative References

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", Oct 1998.

[RFC2535] Eastlake, D., "Domain Name System Security Extensions", March 1999.

9. Author's Address

Rick Mesta
Sun Microsystems, Inc.
5300 Riata Park Court
M/S: UAUS08-102
Austin, TX 78727
USA

Phone: +1 512-401-1076
Email: rick.mesta@sun.com

10. IPR Notices

11. Copyright Statement