Network Working Group
Internet-Draft
Expires: August 5, 2005

R. Mesta
Sun Microsystems, Inc.
Feb 01, 2005

NFSv4 Cross-Domain Considerations
draft-mesta-nfsv4-domain-01

Status of this Memo

Abstract

The purpose of this document is to elicit discussion on configuration schemes for determining the domain name to be used by NFSv4 implementations that do not natively support users and groups from multiple domains.

This document also describes a method by which NFSv4 clients and servers can discover a domain name value appropriate for qualifying NFSv4 user and group names, by leveraging DNS TXT resource records.

Table of Contents

1. Introduction 3
2. Proposed Configuration Scheme 4
2.1 The _nfsv4idmapdomain DNS TXT Resource Record 4
2.2 DNS Tree Lookup Traversal 4
2.3 IETF DNS Community Considerations 5
3. Motivation 6
3.1 multi-DNS Domain Environments with Configured TXT RR's . . 6
3.2 multi-DNS Domain Environments w/o Configured TXT RR's . . 6
4. Security Considerations 7
5. Acknowledgments 8
6. Normative References 8
7. Informative References 8
8. Author's Address 8
9. IPR Notices 9
10. Copyright Notice 9

1. Introduction

Version 4 of the Network File System (NFSv4) protocol [RFC3530] introduces a way for clients and servers to exchange file ownership and ACL entry information as string names qualified with a DNS 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 format for string based identifiers which are intended to be the most flexible representation of file ownership between the different translation implementations. Further, [RFC3530] suggests that the 'domain' portion of the string identifier should be a DNS domain name:

> "The "dns_domain" portion of the owner string is meant to be a DNS domain name. For example, user@ietf.org. Servers should accept as valid a set of users for at least one domain. A server may treat other domains as having no valid translations. A more general service is provided when a server is capable of accepting users for multiple domains, or for all domains, subject to security constraints."

Some NFSv4 implementations do not support the notion of domain qualified user and group identifiers. These implementations are still required by [RFC3530] to qualify user and group names in NFSv4 protocol data. Additionally, these implementations can use the 'domain' qualifier to discover user/group name space boundaries.

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 use multiple DNS domains and sub-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.

This presents the problem of how to distribute the configuration of a domain name for use by NFSv4 implementations, which only deal with domain-agnostic identifiers, for qualifying user and group names.

This document describes one such configuration information distribution method using DNS TXT resource records.

2. Proposed Configuration Scheme

In order to mitigate NFSv4 deployment and promote the highest level of interoperability between NFSv4 implementations while a general-purpose method for mapping multi-domain user/groups to security identifiers is achieved, we propose a DNS TXT [RFC1464] resource record (RR) be adopted as convention between implementors. DNS RR's make the most sense since most customers manage their naming infrastructure via DNS.

2.1 The _nfsv4idmapdomain DNS TXT Resource Record

As stated in [RFC1464], the general syntax for a TXT resource record is:

```
<owner> <class> <ttl> <TXT> <"attribute name=attribute value">
```

Thus, following the syntax above, we propose the specific TXT record <owner> name of '_nfsv4idmapdomain' in order to to minimize the probability of TXT record name collision and to follow established practices when DNS TXT records are used. Kerberos utilizes the '_kerberos' DNS TXT RR <owner> name when performing realm-to-name mapping [KERB5]. Similarly, XFN also utilized DNS TXT records to hold subordinate naming system information [XFNDOC].

Thus, the general form of DNS TXT resource record syntax for NFSv4 domain configuration is prescribed:

```
_nfsv4idmapdomain.soa_domain.  IN   TXT   "domain.name"
```

where "domain.name" will be configured to the desired domain name to be used and/or exchanged in 'owner' and 'owner_group' attribute strings.

2.2 DNS Tree Lookup Traversal

From careful examination of the proposed DNS TXT RR, it can be readily seen that the proposed <owner> field inherits the SOA record's domain. This simple, but powerful side-effect, of having a DNS TXT record as the configuration scheme, allows deployments with multiple DNS domains to override any setting from a parent DNS domain.

For example, assume a customer configuration has a top level DNS domain of "foo.bar" and a corresponding DNS TXT RR has been defined as:

```
_nfsv4idmapdomain   IN   TXT   "foo.bar"
```

Assume further that there are two lower level domains; "ding.foo.bar" and "dong.foo.bar". These lower level DNS domains can in turn each define their own DNS TXT RR's in order to override the TXT record defined by the top level DNS domain. To continue the example, assume that a DNS TXT record is only defined for domain "ding.foo.bar" and it is defined to be:

```
_nfsv4idmapdomain   IN   TXT   "ding.ding"
```

Thus, assuming the 'search' parameter on the client's /etc/resolv.conf file has been properly configured, a DNS TXT RR lookup for "_nfsv4idmapdomain.ding.foo.bar" will yield the string "ding.ding" whereas a lookup for the "_nfsv4idmapdomain.dong.foo.bar" DNS TXT RR will not yield any value and will propagate to the higher level domain as "_nfsv4idmapdomain.foo.bar"; At this point, the string "foo.bar" will be returned for lookups in domain "dong.foo.bar".

2.3 IETF DNS Community Considerations

Discussion within the NFSv4 working group has supported the position that the use of a DNS RR is a reasonable way to distribute a common NFSv4 domain across a NFSv4 deployment. However, there is widespread agreement that overloading a DNS TXT RR is not the proper way to distribute the NFSv4 domain due to the well known sub-typing problem.

The DNS sub-typing problem limits a DNS client to query for all RR's, of the same type, that are available. The client then has to sift thru each reply to look for the desired sub-type. This incurs unnecessary overhead and is not the preferred method of extending DNS.

To alleviate this problem and still utilize DNS as a distribution mechanism for the NFSv4 domain, a new draft will be introduced to propose an NFSv4 application specific RR to the IETF DNS working group.

3. Motivation

As of the date of this memo, there is currently no known general-purpose solution for mapping multi-domain user/groups to security identifiers that can be leveraged. It is also expected that the majority of NFSv4 customer configurations are likely to leverage DNS for name resolution.

As such, the current Solaris NFSv4 implementation leverages the use of the aforementioned DNS TXT RR to configure an arbitrary string that will be used as the NFSv4 id mapping domain. Solaris uses this DNS TXT RR to mitigate NFSv4 deployment at the enterprise IT level and it is expected to be used by system administrators to configure the NFSv4 mapping domain to utilize when client(s) and server(s) exchange 'owner' and 'owner_group' attribute data.

3.1 multi-DNS Domain Environments with Configured TXT RR's

NFSv4 deployments within multi-DNS domain environments can leverage the use of the proposed DNS TXT RR to obtain an NFSv4 domain, that is unified across the different DNS domains, to use for the 'owner' and 'owner_group' attribute strings. The above will hold true whether the multi-DNS domain deployments share a common user/group administrative domain or not.

3.2 multi-DNS Domain Environments w/o Configured TXT RR's

NFSv4 deployments within multi-DNS domain environments in which the DNS TXT RR has not been set up will most likely utilize the DNS domain itself for the 'domain' portion of the attribute strings.

Client(s) and server(s) that interoperate within the same DNS domain boundary will properly map attribute strings to the local system's representation since a common NFSv4 domain is shared. However, client(s) and server(s) that interoperate across DNS domain boundaries, will more than likely map the attribute strings to "nobody" due to mismatched NFSv4 domains.

4. Security Considerations

While this memo raises no security issues, the use of DNSSEC [RFC2535] is recommended.

5. Acknowledgments

David Robinson, Spencer Shepler and Nicolas Williams for their insight and content contributions.

6. Normative References

[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.

[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.

7. Informative References

[KERB5] Garman, J., "Kerberos: The Definitive Guide, pp. 79", Aug 2003.

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

[XFNDOC] Solaris 2.5 Product Documentation, "DNS Text Record Format for XFN References", Nov 1995.

8. 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

9. IPR Notices

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

10. Copyright Notice