Network File System Version 4 C. Lever
Internet-Draft Oracle
Intended status: Standards Track May 2, 2016
Expires: November 3, 2016

Federated Filesystem Security Addendum
draft-cel-nfsv4-federated-fs-security-addendum-05

Abstract

This document addresses critical security-related items that are missing from existing FedFS Proposed Standards.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 3, 2016.

Copyright Notice

Table of Contents

1. Introduction

Requirements for federated filesystems are described in [RFC5716]. Specification of the protocol used by administrators to configure fileservers and construct namespaces is provided in [RFC7533]. Specification of the protocol allowing fileservers to store namespace information is provided in [RFC7532].

These documents are now immutable. However, some security-related concerns have arisen that should be addressed immediately rather than waiting for another version of these protocols to be ratified.

1.1. Problem Statement: GSSAPI service name for ADMIN

After IESG review, the Security Considerations chapter of [RFC7533] now specifically requires that implementations of this protocol support GSSAPI security mechanisms.

ADMIN protocol clients must use a service principal to establish a GSS context shared with an ADMIN server. To construct the service principal, clients need to know a priori the protocol's GSSAPI service name. The form of that service name is described in section 4.1 of [RFC2743].

Also according to the final paragraph of section 4.1, requesting an addition to the "GSSAPI/Kerberos/SASL Service Names" registry

requires a specification. Because [RFC7533] cannot be changed, a new specification must be provided.

1.2. Problem Statement: GSSAPI service name for NSDB

[RFC7532] specifies that NSDB services must be based on the LDAP protocol [RFC4511]. [RFC7532] and [RFC7533] already specify a mechanism to protect NSDB connections using x.509 [RFC4513].

In some cases, it is inconvenient for domain administrators to provide x.509 certificates for NSDBs. One reason might be that administrators have no access to a public trusted Certificate Authority. If a Kerberos TGT service is available locally, for example, that could be a more logical choice than x.509 for managing NSDB server identity.

The RPC [RFC5531] and LDAP protocols have GSSAPI in common. The present document clarifies the use of existing SASL GSSAPI mechanisms when deployed with NSDBs. It does not address how the ADMIN protocol can specify SASL GSSAPI in NSDB connection parameters.

1.3. Problem Statement: Compromised NSDBs

The FedFS ADMIN RPC protocol provides a mechanism for provisioning NSDBs on remote fileservers. The operations it provides are FEDFS_SET_NSDB_PARAMS, FEDFS_GET_NSDB_PARAMS, and FEDFS_GET_LIMITED_NSDB_PARAMS.

FEDFS_SET_NSDB_PARAMS specifies the name of an NSDB and the security mode to use when connecting to this NSDB. The fileserver connects to an NSDB in order to resolve a FedFS junction. The ADMIN protocol specification further says:

> On success, this operation returns FEDFS_OK. When the operation returns, the new connection parameters SHOULD be used for all subsequent LDAP connections to the given NSDB. Existing connections MAY be terminated and re-established using the new connection parameters. The connection parameters SHOULD be durable across fileserver reboots.

There are two security modes defined in the protocol specification: FEDFS_SEC_NONE, which does not authenticate the LDAP server; and FEDFS_SEC_TLS, which uses START_TLS (RFC 4513) to authenticate the LDAP server.

When FEDFS_SEC_TLS is specified with the FEDFS_SET_NSDB_PARAMS operation, an x.509v3 certificate chain is also provided to the fileserver. The fileserver uses the provided certificate to

authenticate subsequent connections to this NSDB. The FEDFS_SET_NSDB_PARAMS operation can change the connection security used by a fileserver to connect to a particular NSDB from NONE to TLS or TLS to NONE.

Over time, domain administrators add NSDB connection parameters to each of their fileservers to enable FedFS junction resolution. The specified NSDB may be the domain's own, or it might be an NSDB in a foreign domain.

Many junctions on multiple fileservers can be created that use a particular NSDB. There is no way to find such junctions without an exhaustive search. Since filesystem namespace topology can evolve arbitrarily over time, a recorded pathname of any junction is almost guaranteed to become stale.

Now suppose we have two FedFS domains: example.net and university.edu. Suppose university.edu fileservers have a number of junctions that refer to locations maintained by example.net, and thus university.net's fileservers are configured to resolve junctions on example.net's NSDB.

One day Mallory compromises example.net's NSDB, but the domain administrator there is on a long vacation. The administrator at university.net discovers the compromise immediately, but has no control over the foreign NSDB and cannot create a fresh x.509 certificate or verify that the contents of the NSDB are unmolested. The only choice is to find and remove every junction in the university.edu domain that contains the compromised NSDB.

If university.edu is using a good implementation of FedFS, the administrative tools it provides might allow an administrator to simply visit each of its fileservers and mark the example.net NSDB as compromised. Any junction resolution that attempts to use that NSDB would fail, but all junctions remain in place. When example.net's administrator gets back from holiday and cleans up the mess, the university.edu administrator can then update each of her fileservers with fresh connection parameters for that NSDB.

However, none of this can be done remotely using the FedFS ADMIN protocol. It does not have a mechanism for removing NSDB connection parameters or for fencing a compromised NSDB.

2. GSSAPI Service Name for the FedFS ADMIN protocol

Section 6 of [RFC7533] requires a FedFS ADMIN server to support the RPCSEC_GSS framework [RFC2203]. The present document specifies the

GSSAPI service name, as described in Section 4.1 of [RFC2743], to be used for the FedFS ADMIN protocol.

Regardless of what security mechanism under RPCSEC_GSS is in use, a FedFS ADMIN server MUST identify itself in GSSAPI via a GSS_C_NT_HOSTBASED_SERVICE name type. GSS_C_NT_HOSTBASED_SERVICE names are of the form:

```
service@hostname
```

For the ADMIN protocol, the "service" element is

```
fedfs-admin
```

Implementations of security mechanisms will convert fedfs-admin@hostname to various different forms. For Kerberos V5, the following form is RECOMMENDED:

```
fedfs-admin/hostname
```

This service name SHOULD NOT be used to authenticate other GSSAPI services.

3. GSSAPI Service Name for the FedFS NSDB protocol

Section 5.2.1.1 of [RFC4513] specifies the GSS service name for LDAP. LDAP servers acting as NSDBs MUST use this service name, which is of the form:

```
service@hostname
```

When accessing an NSDB service, the "service" element is

```
ldap
```

Implementations of security mechanisms will convert ldap@hostname to various different forms. For Kerberos V5, the following form is RECOMMENDED:

```
ldap/hostname
```

FedFS-enabled file servers act as NSDB clients when resolving FedFS junctions. In order to access NSDBs via SASL GSSAPI, such clients would first authenticate to a KDC. To avoid a requirement for human interaction (say, to enter a Kerberos password), such clients should utilize a key stored in a keytab. Clients MAY use nfs/hostname, but MUST NOT use fedfs-admin/hostname.

3.1. Cross-realm considerations

Note that the target NSDB's REALM is not specified above. When authenticating a GSSAPI service, NSDB clients typically have a service name (in this case "ldap") and the fully qualified domain name of the NSDB server. The underlying LDAP client library will either:

1. Find the server's REALM based on local configuration, or

2. Request a referral from the local KDC if the NSDB server's FQDN is not registered in the default REALM.

Therefore, a pre-existing trust relationship must exist between the REALM of a FedFS-enabled file server and the REALMs containing foreign NSDBs containing junctions that file server wants to resolve. In this instance, an x.509 certificate may be a preferrable approach.

4. Fencing Compromised NSDBs

An NSDB is considered "foreign" relative to a particular FedFS domain if that domain's administrator has no administrative access to that NSDB.

When a FedFS domain administrator is faced with a foreign NSDB that is compromised or otherwise unusable, and in the absense of an implementation-provided mechanism for fencing an NSDB, the administrator can fence that NSDB using the following technique.

1. The administrator locally generates a new certificate for the compromised foreign NSDB. The certificate can be self-signed, or signed by the administrator's local certificate authority.

2. The administrator distributes this certificate to all of her domain's fileservers using the FedFS ADMIN protocol or some other secure means. The connection security for the foreign NSDB is set to FEDFS_SEC_TLS on each of the local domain's fileservers.

3. The administrator requests fresh certificate material from the administrator of the foreign NSDB.

4. When the threat has passed and the foreign NSDB is safe to use again, the administrator can distribute the new valid certificate material to her domain's fileservers.

No change to the ADMIN protocol as specified in [RFC7533] is required to fence a compromised NSDB. Step 2 guarantees that, on fileservers in the administrator's local FedFS domain, resolving a junction that

references the compromised foreign NSDB will fail until updated certificate material is provided.

5. Security Considerations

When deploying FedFS, the use of security mechanisms that maintain the confidentiality of all network communications is recommended. This includes the use of any pseudoflavor that supports the rpc_gss_svc_privacy service for the FedFS ADMIN protocol, and the use of TLS message encryption for the NSDB protocol.

When creating x.509 certificates for authenticating NSDBs, implementations should utilize keys that are as large as practical, especially if certificate lifetimes are long.

Operational security is further enhanced by ensuring that all hardware entropy sources are verified for cryptographic use. This recommendation applies to the creation of x.509 certificate material, random-variant UUIDs, and handshake keys used to secure transports, for example.

Information stored in fedfsDescr and fedfsAnnotation attributes are readable by any unauthenticated user of an NSDB, and therefore should contain no sensitive information.

6. IANA Considerations

In accordance with Section 4.1 of [RFC2743], the service name "fedfs-admin" will be registered in the GSSAPI Service Name registry at http://www.iana.org/assignments/gssapi-service-names/ gssapi-service-names.xml

The new entry should reference the present document as the specification.

7. Acknowledgements

The author of this document gratefully acknowledges the contributions of Simo Sorce, Nico Williams, Robert Thurlow, Spencer Shepler, Tom Haynes, and David Noveck.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2203] Eisler, M., Chiu, A., and L. Ling, "RPCSEC_GSS Protocol Specification", RFC 2203, DOI 10.17487/RFC2203, September 1997, <http://www.rfc-editor.org/info/rfc2203>.

[RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, DOI 10.17487/RFC2743, January 2000, <http://www.rfc-editor.org/info/rfc2743>.

[RFC4511] Sermersheim, J., Ed., "Lightweight Directory Access Protocol (LDAP): The Protocol", RFC 4511, DOI 10.17487/RFC4511, June 2006, <http://www.rfc-editor.org/info/rfc4511>.

[RFC4513] Harrison, R., Ed., "Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms", RFC 4513, DOI 10.17487/RFC4513, June 2006, <http://www.rfc-editor.org/info/rfc4513>.

[RFC5531] Thurlow, R., "RPC: Remote Procedure Call Protocol Specification Version 2", RFC 5531, DOI 10.17487/RFC5531, May 2009, <http://www.rfc-editor.org/info/rfc5531>.

[RFC7532] Lentini, J., Tewari, R., and C. Lever, Ed., "Namespace Database (NSDB) Protocol for Federated File Systems", RFC 7532, DOI 10.17487/RFC7532, March 2015, <http://www.rfc-editor.org/info/rfc7532>.

[RFC7533] Lentini, J., Tewari, R., and C. Lever, Ed., "Administration Protocol for Federated File Systems", RFC 7533, DOI 10.17487/RFC7533, March 2015, <http://www.rfc-editor.org/info/rfc7533>.

8.2. Informative References

[RFC5716] Lentini, J., Everhart, C., Ellard, D., Tewari, R., and M. Naik, "Requirements for Federated File Systems", RFC 5716, DOI 10.17487/RFC5716, January 2010, <http://www.rfc-editor.org/info/rfc5716>.

Author's Address

Charles Lever
Oracle Corporation
1015 Granger Avenue
Ann Arbor, MI 48104
USA

Phone: +1 734 274 2396
Email: chuck.lever@oracle.com