Request for Comments On A Specification of
Trusted NFS (TNFS) Protocol Extensions

1. Status Of This Memo

This draft document specifies extensions to RFC 1094 [1] which support network file access in a multilevel secure (MLS) network environment[1]. This draft was approved by the Trusted Systems Interoperability Group (TSIG), whose charter is to promote multi-vendor trusted system interoperability.

2. Abstract

Additional functionality has been developed for UNIX† systems to address the TCSEC [2] requirements for trusted systems. New requirements are driving efforts to develop interoperable, networked solutions for trusted UNIX environments. A specific approach for addressing TCSEC MLS requirements is identified in the CMW requirements document [3]. Developing support for network interoperability among MLS classified systems is a primary goal of the trusted UNIX community.

Sun Microsystem's Network File System V2 protocol is an industry (de facto) standard network file access mechanism, and represents one of the key components of system interoperability in the current UNIX networking market. This draft document describes extensions to the NFS V2 protocol which support network file access in a MLS network environment. It will be submitted to the RFC editor as a protocol specification. Distribution of this draft document is unlimited. Please send comments to the author at the address identified in section 6 below.

3. MLS Security Extensions

MLS security functionality includes discretionary access control (DAC), subject and object security labeling, mandatory access control (MAC), authentication, auditing, and documentation. Exchanging information between MLS systems requires communicating additional security information along with the actual data.

The primary goal of this specification is to describe extensions to the NFS V2 protocol which support network file access between MLS systems with a minimal impact on the existing NFS V2 environment[2]. It is also intended that this MLS environment will permit unmodified NFS clients and servers to continue to be fully supported.

The general approach used in extending the NFS V2 protocol is to transport additional user context in the form of an extended NFS UNIX style credential between a Trusted NFS (TNFS) client and server, and to map that context into the appropriate server security policies which address file access. In addition, security file attributes are returned with each NFS (TNFS) procedure call. Otherwise, the NFS V2 protocol remains essentially unchanged.

Two companion documents [4][5] complete the set of documentation describing the TNFS environment.

3.1. The Extended User Context

The Sun RPC protocol [6][7] includes two authentication parameters in a request message:

an authentication credential - used to identify or present a client subject's credentials to a server along with a given request for access or information, and

an authentication verifier - used to validate the subject's credentials,

---

[1] Multilevel Secure systems include, for example, support for B1 and CMW security policies.

† UNIX is a registered trademark of A. T. & T.

[2] Revisions to the NFS V2 protocol have been specified and presented for comment to the NFS community; this document addresses extensions to the V2 protocol only.

and an authentication verifier in the RPC response message.

An NFS server uses the client subject's credentials to perform appropriate access checks prior to servicing the request. The verifier parameter in the RPC request message is used to authenticate the client subject's credentials[3].

Several styles of authentication are currently defined for NFS[4], and an NFS server may elect to support multiple authentication styles concurrently. A new RPC authentication style, AUTH_MLS, is defined for use in the TNFS environment. The definition of the AUTH_MLS credential combines the information in the AUTH_UNIX credential with extensions for the additional security attributes:

- audit id - immutable subject (user) identifier, not affected by modifications to either the real or effective user or group identifiers,
- sensitivity label - used with a MAC policy; a subject generally has a static, top-level clearance, but is permitted to execute processes at a sensitivity level different from (i.e. lower than) his/her actual clearance,
- information label - also used with a MAC policy; dynamically adjusted based upon the information content associated with the subject (or object),
- integrity label - used with commercial, multi-party security policy (eg. Clark-Wilson [8], Biba [9]),
- privilege mask - used to identify privileges (eg. chown, chmod) or "rights" granted to a given subject, generally to override an existing security policy, and
- national caveat label - used with multi-national security policy [10]

The additional security attributes will actually be represented within the AUTH_MLS credential by fixed size *tokens*, which can support multiple translation schemes through the use of an appropriate translation mechanism [5]. For instance, mechanisms such as M.I.T. Project Athena's Hesiod/BIND or Sun Microsystem's NIS[5] lookup service could be used to support the translation of tokens and security attribute information.

There are several advantages to the use of a token translation model. One major advantage is that the actual security attribute information may be defined within the translation service, while the attribute representation may be defined by a small, fixed sized token within the relatively small amount of unallocated space in the credential structure. A second advantage of a translation model is that it may accommodate multiple security policies and translations. Finally, a token translation model permits security policies to be developed independently from the translation mechanism. Tokens are transferred within the AUTH_MLS credential as opaque objects which are given context by the security policy mechanisms implemented by the TNFS clients and servers.

Note that although tokens are defined as opaque objects, tokens which represent the same security attribute and which reside within the same translation scheme may be compared for equality. This characteristic permits tokens representing a specific security attribute to be referenced in comparisons without requiring the tokens to be translated.

3.2. Network Provided Security Attribute Labels

Support for the transfer of MAC sensitivity labels for the Internet Protocol Suite has been addressed by the CIPSO [11], and RIPSO [12] documents. The security information defined within the AUTH_MLS credential, however, provides for the transfer of security attributes required to support MLS access policies without requiring the underlying network layer to provide security attribute information. Transferring security attributes within the RPC layer also provides for the support of a policy where data may be transferred with a security classification which is different from the security classification of the network

---

[3] Authentication of client and server identities is not currently addressed in this specification, but will be addressed in a future revision.

[4] Styles currently defined are AUTH_NONE, AUTH_UNIX, AUTH_SHORT, and AUTH_DES.

[5] Network Information Service, known previously as the Yellow Pages Service

layer. For instance, file data with a given security classification might first be encrypted and then transferred through a network with a lower security classification. If security attributes are provided by both the RPC layer and the underlying network layer, then the security attribute information provided within the AUTH_MLS credential shall be applied to the file data transferred within the RPC message.

3.3. Discretionary Access Control

A Discretionary Access Control (DAC) policy provides for the restriction of subject access to objects based on the identity of subjects and/or the groups for which they are members. Most secure systems address DAC requirements through the use of access control lists. Associated with each file is a list which identifies the set of user and group combinations authorized to access the file, along with the access privileges associated with each combination.

The information contained in the AUTH_MLS credential of a TNFS client request includes user and group identification sufficient to permit the server to apply appropriate DAC policies in controlling access to its shared, local file objects. For example, the subject represented by the user and/or group identifiers contained in the client request may be checked against the access control list information associated with the referenced file on the server. Access control list information is not required to be transmitted from the client to the server in support of a server based access control policy. Client based support for access control of server based file objects is discussed below in the section which describes the extended attribute cache.

3.4. Mandatory Access Control

A Mandatory Access Control (MAC) policy provides for the restriction of subject access to objects based on the sensitivity of the information contained in the objects. MAC policies thus include assigning levels of trust or clearance to system users (subjects), and levels of sensitivity to system objects, and then ensuring that only users with sufficient clearance can access the classified information.

3.4.1. Sensitivity Labels

When MAC policies are enabled, each system subject and object is created with a sensitivity label, and the system MAC policies compare the labels when determining access.

The AUTH_MLS credential contains the sensitivity label information associated with the TNFS client subject (application) making the access request. This information is sufficient to permit the MAC policy checking mechanism on the server to determine whether to permit access to the requested object or information.

3.4.2. Information Labels

Information labels represent the actual sensitivity of a given subject or object, and permit the additional identification of control markings for a given piece of information. The information label is dynamically adjusted on both subjects and objects to the highest sensitivity level reflected by a subject/object pair: if a subject issues a write request to an object, the information label of the object will be adjusted (if necessary) to the level defined by the information label of the subject; if a subject issues a read request to an object, the information label of the subject will be adjusted to the level defined by the information label of the object. Note that information labels are adjusted upwards as a result of these actions; information labels are never automatically adjusted to a lower level.

The AUTH_MLS credential in the RPC request message contains the current information label associated with a TNFS client application (subject), and permits a remote file's object information label to be adjusted (if necessary) as a result of a client generated *write* operation. The TNFS reply message includes a field for the information label associated with an accessed file object, permitting the subject's information label to be adjusted (if necessary) as a result of a client generated *read* operation.

These extensions are sufficient to support the MAC information label policies with respect to network file access.

3.5. MAC and DAC Coverage with TNFS

In an MLS environment, both DAC and MAC access control policies are applied in determining access to a given object. In a network environment of MLS systems participating in TNFS file access, the AUTH_MLS credential permits a TNFS server to apply both DAC and MAC policies in consideration of a request from a remote NFS client subject. Thus, MLS based network file access using the NFS V2 protocol can be supported through the use of the AUTH_MLS credential. Listing or modifying the DAC and MAC security attributes of a server's file from a client, however, requires additional protocol extensions. Identifying additional security access restrictions when a request is made to open a remote file is also considered to be a requirement. Extensions designed to satisfy these requirements are addressed by the TNFS protocol, and are described in the next subsections.

3.5.1. Remote Access to Extended File Attributes

The TCSEC notion of *appropriate privilege* is an integral part of the MLS environment. It is expected that a subject with appropriate privilege will want to gain access to the additional file attribute information for the purposes of modification and/or viewing of that information. Subject privileges are defined within the AUTH_MLS credential. Note, however, that the privileges associated with a given subject on a given client system may not be extended to the subject on a given server. Although most subjects will likely retain their privileges on the server, a client administrator, for example, may not be granted administrative privileges on the server.

The DAC and MAC security attribute information includes MAC and information labels, and access control list information (ACLs). Supporting remote access to this information is more difficult to address in the network environment, since:

- it requires transmitting additional file security attribute information (or its representation) "over the wire", and
- additional file attribute information cannot be accommodated in the existing NFS V2 protocol file attribute data structures; additional support setting and getting the extended security attributes is required

Thus, extensions to the NFS V2 protocol procedures have been defined to support access to the extended attributes of served files. The complete set of NFS protocol procedures and security extensions are referred to in this document as the TNFS protocol.

3.5.2. File Open Enhancement

Using the NFS V2 protocol, a client request to *open (2)* a remote file on the server may be translated by the client into a GETATTR procedure call for the current directory[6], followed by a LOOKUP procedure call for the file to be opened. If valid responses from these procedure calls are returned, the client's NFS file attribute cache is updated, and an open file descriptor may be returned to the requesting application.

Since the NFS V2 protocol does not transmit an actual open request to the server, however, an MLS server will not be able to apply the appropriate DAC and MAC policy at the time of the open request, and the application may find that it has successfully opened the file, but that it cannot access the file due to stronger access control policies being applied by the server in response to specific client access requests.

An *access* protocol procedure would permit the client to determine whether access to the file would be supported by the server, based on the application's open request type and the associated extended security attribute information. An additional TNFS protocol procedure has been defined to address this issue.

3.5.3. TNFS Protocol Extensions

Extensions to the NFS V2 protocol are defined in this section of the specification. These extensions are designed to support remote access to the security file attribute extensions, and to support the file open enhancement.

[6] Depends on the presence of valid attributes in the lookup cache (DNLC).

3.5.3.1. Data Structure Definitions

The definitions which support the MLS extensions are described in this section. Since the definitions for the TNFS protocol are an extension of the original NFS V2 protocol, this specification will include all of the extended data structure definitions, and a few of the original definitions for clarity. Note that the arguments and results are defined using the RPC language.

The following RPC constants are used to identify the TNFS extensions which support MLS security policies. The TNFS program will be registered as a separate service with the RPC port mapping service, but will share the same UDP [13] port number with the original NFS V2 service. Registration as a different service distinguishes the TNFS service from the original NFS V2 service. The use of a different version number distinguishes each request/response message.

```
PROGRAM 390086  /* TNFS Program Number */
VERSION      1  /* TNFS Version 1 */
PORT      2049  /* Original NFS Port */
```

The **stat** type is returned from every procedure call. A value of NFS_OK indicates the call completed successfully. Other values indicate that an error occurred during the servicing of the request. Note: this structure is unchanged from the NFS V2 Protocol Specification. It is (partially) reproduced here for clarity.

stat

```
enum stat {
        NFS_OK = 0,
        NFSERR_PERM = 1,
        NFSERR_NOENT = 2,
        . . .
        [other NFS errors as defined in the V2 protocol specification]
};
```

The credential parameter is included in each RPC request message, and is used to supply the client subject's credentials to the server. The AUTH_MLS credential will be used with the TNFS procedure calls and is defined as follows:

```
#define AUTH_MLS 200000     /* decimal */

#define MLS_TOKEN_SIZE 4    /* 4 octets or 32 bits */

typedef opaque t_token[MLS_TOKEN_SIZE]; /* tokens are opaque */

struct authmls_cred {
        u_long  auc_stamp;         /* arbitrary ID */
        char    auc_machname<255>; /* machine name */
        u_long  auc_uid;           /* effective uid */
        u_long  auc_gid;           /* effective gid */
        u_long  auc_len;           /* len of groups list */
        u_long  auc_gids<24>;      /* groups */
        u_long  auc_aid;           /* audit id */
        t_token auc_privs;         /* privileges token */
        t_token auc_sens;          /* sensitivity token */
        t_token auc_info;          /* information token */
        t_token auc_integ;         /* integrity token */
        t_token auc_ncs;           /* national caveat set token */
};
```

Note that if a given security attribute is not being exchanged, then the corresponding credential token values shall be set to all zeros. A given security policy may require that only a subset of the security attributes provided for in this specification be exchanged. For example, a C2 network

security policy requires the support of privileges, and might also require support for Access Control Lists (ACLs). In that case, the sensitivity, information, integrity, and national caveat token values shall be set to all zeros in the exchange messages.

The **fattr** structure defines the complete set of file attributes of a file. The extended **fattr** structure combines the NFS V2 **fattr** structure with additional fields for a file's security attributes. The security attributes are represented by tokens.

```
struct fattr {
        ftype   type;       /* file type */
        u_long  mode;       /* encoded access mode */
        u_long  nlink;      /* number of hard links */
        u_long  uid;        /* file's owner id */
        u_long  gid;        /* file's group id */
        u_long  size;       /* file size in bytes */
        u_long  blocksize; /* number bytes/block */
        u_long  rdev;       /* device number of the file */
        u_long  blocks;     /* current number of blocks */
        u_long  fsid;       /* file system id */
        u_long  fileid;     /* unique file identifier */
        timeval atime;      /* time of file's last access */
        timeval mtime;      /* time last modified (written) */
        timeval ctime;      /* time of last attribute change */
        t_token privs;      /* privileges token */
        t_token sens;       /* sensitivity token */
        t_token info;       /* information token */
        t_token integ;      /* integrity token */
        t_token ncs;        /* national caveat set token */
        t_token acl;        /* access control list token */
};
```

Note that if a given security attribute is not being exchanged, then the corresponding file attribute token values shall be set to all zeros.

Note also that the value of information token, **info**, in the *fattr* structure of the response message shall be non-zero if:

(1) the server supports an information label security policy, and

(2) the subject's (requester's) information label requires adjustment as a result of the support of that policy

Otherwise, the information token field shall be set to all zeros.

The **sattr** structure defines the file attributes which can be set from the client. The extended **sattr** structure combines the NFS V2 **sattr** structure with additional fields for the security attributes, which are represented by tokens. A token value of all zeros indicates that the token field is to be ignored.

```
struct sattr {
        u_long  mode;     /* encoded access mode */
        u_long  uid;      /* file's owner id */
        u_long  gid;      /* file's group id */
        u_long  size;     /* file size in bytes */
        timeval atime;    /* last access time */
        timeval mtime;    /* last data modify time */
        t_token privs;    /* privileges token */
        t_token sens;     /* sensitivity token */
        t_token info;     /* information token */
```

```
        t_token integ;   /* integrity token */
        t_token ncs;     /* national caveat set token */
        t_token acl;     /* access control list token */
};
```

The **sattrargs** structure is used by the SETATTR procedure. It contains the extended **sattr** structure definition.

```
struct sattrargs {
    fhandle file;
    sattr attributes;
};
```

The **attrstat** structure defines a common procedure result containing the status of the procedure call. It is returned with the results of GETATTR, SETATTR, and WRITE procedure calls. If the call was successful, **attrstat** contains the results for the specific procedure called, and the complete set of file attributes for the file on which the procedure was executed.

```
union attrstat switch (stat status) {
        case NFS_OK:
            fattr attributes;
        default:
            void;
};
```

The **diropres** structure defines the results of a directory procedure call. If the call was successful, **diropres** contains a new file handle **file** and the complete set of associated file **attributes**.

```
union diropres switch (stat status) {
        case NFS_OK:
            struct {
                fhandle file;
                fattr attributes;
            } diropok;
        default:
            void;
};
```

The **readlinkres** structure defines the results of a READLINK procedure call. If the call was successful, **readlinkres** contains the **data** in the symbolic link of the file identified by the file handle argument, and the complete set of associated file **attributes**. File attributes are returned with the READLINK procedure call to support the information label adjustment policy.

```
union readlinkres switch (stat status) {
        case NFS_OK:
            struct {
                path data;
                fattr attributes;
            } readlinkok;
        default:
            void;
};
```

The **readdirres** structure defines the results of a READDIR procedure call. If the call was successful, **readdirres** returns a variable number of directory **entries**, with a total size of up to the amount specified in the argument **count** of the **readdirargs** structure. Each entry contains a unique file identifier, and an

opaque "pointer" to the next entry in the directory. The **eof** flag has a value of TRUE if there are no more directory entries. File attributes are returned with the READDIR procedure call to support the information label adjustment policy.

```
union readdirres switch (stat status) {
      case NFS_OK:
         struct {
            entry *entries;
            bool eof;
            fattr attributes;
         } readdirok;
      default:
         void;
};
```

3.5.3.2. TNFS Protocol Procedure Definitions

The TNFS Protocol Definition integrates the use of:

- o the extended **fattr** and **sattr** structures,
- o an AUTH_MLS authentication style RPC credential,
- o a new TNFS protocol version number to differentiate between NFS V2 and the security extended TNFS protocol, and
- o a new protocol procedure, ACCESS, to support the file open enhancement described earlier

Other than these changes, however, the syntax and semantics of TNFS remain the same as in the original NFS V2 specification.

3.5.3.2.1. Access Procedure

The following descriptions are used to define the new ACCESS procedure.

Definitions used to identify the access request type:

```
#define READ    0x001
#define WRITE   0x002
#define EXEC    0x004
#define SEARCH  0x008
#define APPEND  0x010
```

Arguments for the remote access procedure:

accessargs

```
struct accessargs {
      fhandle  file;
      u_long   flag;
 };
```

Response from the remote access procedure:

accessres

```
union accessres switch ( stat status ) {
   case NFS_OK:
      struct {
         bool_t status;   /* access status: TRUE or FALSE  */
         fattr  attributes;  /* standard file attributes */
      }  accessok;
```

```
    default:
        void;

};
```

Procedure definition for checking remote access permission:

accessres
NFSPROC_ACCESS(accessargs) = 18

Description:

Determine if access as described by **flag** will be permitted on the remote served object **file** by the requester. Flag values are bit encoded as defined previously. READ access means that the data in **file** can be read, WRITE access means that the data in **file** can be modified (written), EXEC access means that **file** can be accessed and executed (local execution of a remote file), SEARCH access means that the directory **file** can be used as the argument to a LOOKUP operation, and APPEND means that the file size can be extended. If **status** is NFS_OK:

accessok.status will be set to **TRUE** if the access request would be allowed, and set to **FALSE** otherwise, and

attributes will contain the complete set of file attributes

Otherwise:

the NFSERR error number returned identifies the error condition

Implementation:

The ACCESS procedure provides a means for checking file access permission prior to issuing a subsequent set of file operations. For example, a TNFS client may issue an access procedure as a result of an application's file *open (2)* request to determine if subsequent file *reads* and/or *writes* by the application would be denied by the server as a result of the server's extended file access security policies. Note that the information returned by the server in response to an ACCESS procedure call is not static; subsequent file administrative procedures may result in the modification of the file's security attributes.

3.5.3.2.2. TNFS Service Routines

The TNFS protocol definition is defined below as a set of procedures, arguments, and results. All modified data structure definitions are included in this specification. Most NFS V2 protocol data definitions remain unchanged, and are documented in the NFS V2 protocol specification. The complete set of TNFS protocol procedures are defined below. The ACCESS procedure is new, but the other procedures are the same as those defined in the NFS V2 specification. The GETATTR, SETATTR, LOOKUP, READLINK, READ, WRITE, CREATE, MKDIR, READDIR, and ACCESS procedures for the TNFS protocol, however, include the extended file attribute structure *fattr* in the response message.

```
program TNFS_PROGRAM {
    version TNFS_VERSION {
        void        NFSPROC_NULL (void) = 0;
        attrstat    NFSPROC_GETATTR (fhandle) = 1;
        attrstat    NFSPROC_SETATTR (sattrargs) = 2;
        diropres    NFSPROC_LOOKUP (diropargs) = 4;
        readlinkres NFSPROC_READLINK (fhandle) = 5;
        readres     NFSPROC_READ (readargs) = 6;
        attrstat    NFSPROC_WRITE (writeargs) = 8;
        diropres    NFSPROC_CREATE (createargs) = 9;
        stat        NFSPROC_REMOVE (diropargs) = 10;
        stat        NFSPROC_RENAME (renameargs) = 11;
        stat        NFSPROC_LINK (linkargs) = 12;
        stat        NFSPROC_SYMLINK (symlinkargs) = 13;
```

```
        diropres    NFSPROC_MKDIR (createargs) = 14;
        stat        NFSPROC_RMDIR (diropargs) = 15;
        readdirres  NFSPROC_READDIR (readdirargs) = 16;
        statfsres   NFSPROC_STATFS (fhandle) = 17;
        accessres   NFSPROC_ACCESS (accessargs) = 18;
    } = 1;     /* Trusted NFS Version 1  */
} = 390086;    /* Trusted NFS Program Number */
```

3.5.4. Using TNFS

With the TNFS protocol procedures described above, listing and modifying remote extended file attributes is now supported. The definition of a new application programming interface (API) to support the display of a file's security attributes will permit either a new list command (e.g. lsacl, lsmac) or a modification to the existing *ls (2)* command to display the security attribute information associated with a remote file. Likewise, the definition of a new API for setting a file's security attributes will permit new change security attribute commands to be developed (e.g. chacl, chmac).

The file open enhancement discussed previously may now be supported. The open API will be translated into a GETATTR operation for the current directory, a LOOKUP operation for the file to be opened, and an ACCESS operation which returns a boolean value indicating whether the access requested would be permitted, along with the complete set of the file's attributes. Thus, the TNFS client can determine whether the application requesting to *open* the remote file will be able to access it based on the open request type and the application's security credentials. As described earlier, a server may choose to associate a set of privileges with the remote subject which are different from the privilege set associated with the subject on the client system. The ACCESS procedure call returns the server's assessment of the subject's access capabilities.

The information label adjustment policy is also supported, since the AUTH_MLS credential contains the subject's information label, and the TNFS reply message contains an extended file attribute structure which includes the file object's information label. Note that the subject's information label may require adjustment as a result of reading a remote file (READ), reading a remote directory (READDIR), or reading a remote symbolic link (READLINK). A remote file's (object) information label may be adjusted as a result of SETATTR, WRITE, CREATE, RENAME, LINK, SYMLINK, and MKDIR TNFS procedure calls.

3.5.5. The Extended Attribute Cache

NFS caching strategies are implementation specific, and are not part of the NFS protocol. Caching is also not required to support TNFS interoperability. This specification will therefore not include specific details on the issue of attribute caching. However, since the caching mechanisms are included in the NFS reference source code releases, and since attribute caching is critical for achieving NFS performance goals, several suggestions are included in this section.

In most NFS client implementations, remote file attributes are cached on the client, improving performance and reducing network traffic. The attribute cache is updated frequently, as most NFS procedures return file attributes along with other requested information.

A client side cache for the extended security file attributes should also be considered for similar reasons. Since all of the file's security attributes are returned with each TNFS file access request, an extended security attribute cache can now be maintained on the client.

Extending the attribute validation procedure to include validating the security file attributes permits the complete set of file attributes to be checked and refreshed if they are no longer valid. If the file's cached attributes are not valid, a **GETATTR** procedure call can be made. The TNFS reply to this procedure now includes the complete set of file attribute information, permitting all of the file's cached attributes to be refreshed. Cached attribute entries shall be aged and eventually flushed unless refreshed.

Note again that an attribute caching policy is not part of the protocol, and is an implementation technique used to improve performance. During the window of time that the cache entry is valid, the client system applies the MLS access control policies on behalf of the server. It is recommended that if an

implementation supports the use of client side attribute caching, it shall also support a mechanism for disabling the attribute cache. Additional implementation details are provided in [4].

3.5.6. TNFS Access Control Policy

The access control policy recommended by this proposal may be stated as follows:

- a client system *always* applies the access control policy to a local request for access to a local resource,
- a client system *may (temporarily)* apply the access control policy to a local request for access to a remote resource; this policy applies to the case of client side attribute caching
- a server system *always* applies the access control policy to a local request for access to a local resource,
- a server system *always* applies the access control policy to a remote access to a local resource

This TNFS access control policy ensures that no access will be made without the application of appropriate access control.

3.5.7. TNFS Auditing Policy

The auditing policy recommended by this proposal may be stated as follows. When the security auditing function is enabled:

- a client system *always* audits a local request for access to a local resource,
- a client system *may* audit a local request for access to a remote resource,
- a server system *always* audits a local request for access to a local resource,
- a server system *may* audit a remote request for access to a local resource
- **an implementation shall support:**
  - the option for auditing requests for local access to remote resources on the client, and
  - the option for auditing remote requests for access to local resources on the server

    Note: This option may require the auditing of the specific TNFS protocol procedure calls, since the protocol procedures are not translated into actual "system calls" in many server implementations.

This TNFS auditing policy ensures that both clients and servers have the ability to audit all access activity within their domain. In a given network environment, it may be desirable to optionally disable auditing of remote access on either the client or server to avoid duplication.

3.5.8. Support for NFS V2 Clients and Servers

The MLS environment described in this document assumes that most file access will take place between MLS modified clients and servers. It is still useful, however, to define the mechanism by which TNFS systems can continue to interoperate with NFS V2 systems through the use of an appropriate policy.

One such policy involves the use of a filter or gateway placed between the modified and unmodified systems. This gateway would insert or delete the appropriate security attribute information on behalf of the unmodified systems.

This specification assumes the existence of a local database on each MLS system which identifies:

- the hosts which that system will communicate with,
- the security attributes which it expects to use in the exchange of any data with a given host, and
- the translation scheme which will be used in translating tokens between this client/server pair.

This information is needed by all network applications, and is not limited to NFS file access. The use of such a database permits a given system to apply some intelligence in dealing with unmodified clients and

servers, and permits an additional verification (in terms of the expected security attributes) for MLS modified clients and servers. Since TNFS is registered as a different service with the RPC port mapping service, the mapping service may be queried to determine if the TNFS service is supported by a given server.

Based on the information obtained from this database and/or the RPC port mapping service, a TNFS client would not send any security extended NFS procedure calls to a server which did not support the service. A TNFS client should also refrain from sending extraneous security attribute information to a TNFS server that does not support an equivalent set of security attributes.

4. Conclusion

This document describes the set of extensions which support network file access in a network environment consisting of MLS systems using the proposed TNFS protocol extensions. Unmodified NFS clients and servers are supported using the de facto NFS V2 protocol.

With the previously defined extensions, the MLS network file access requirements are met. The extended structure definitions support the DAC and MAC attributes required for modifying or displaying the security attribute information. The enhanced file open operation and the information label adjustment policies are also supported.

Thus, a small set of extensions to the NFS V2 environment permits MLS access control policies to be supported. Agreement on these changes will permit the current base of NFS clients and servers to be accommodated in the secure environment with no changes, and for TNFS modified systems to interoperate using MLS policies.

This specification places no dependencies on the underlying network layer, but does acknowledge that security labeling information is provided by at least some network implementations.

5. Acknowledgements

I would like to acknowledge the members of the TSIG NFS Subcommittee, who were instrumental in evolving the MLS extended NFS Protocol Specification from the original proposal. Many comments were also made during the review of the later drafts which greatly improved the specification's readability. Contributing members included Morgan Clark, Jeff Edelheit, Fran Fadden, Tricia Jordan, Will Lees, Scott Norton, Mike Shipley, Carl Smith, Dave Summers, and Charlie Watt.

The specification was also reviewed by numerous persons outside of the subcommittee. I would like to acknowledge many of these persons as well, as a number of their comments are also reflected in the final version.

6. Author's Address

Fred Glover
Digital Equipment Corporation
110 Spit Brook Road ZK03-3/U14
Nashua, New Hampshire 03062-2698

Phone: 603-881-0388

EMail: fglover@decvax.dec.com

7. References

[1] Sun Microsystems, Inc., "Sun Network Filesystem Protocol Specification", Internet RFC 1094

[2] United States Department of Defense Intelligence Agency, "Trusted Computer Systems Evaluation Criteria"

[3] United States Department of Defense Intelligence Agency, "Compartmented Mode Workstation Requirements"

[4] Trusted Systems Interoperability Group, "The MLS NFS Implementor's Guide", TSIG Document

[5] Trusted Systems Interoperability Group, "The MLS Token Translation Specification", TSIG Document

[6] Sun Microsystems, Inc., "Sun Remote Procedure Call Specification", Internet RFC 1057

[7] Sun Microsystems, Inc., "Sun External Data Representation Specification", Internet RFC 1014

[8] Clark, D. D. and David R. Wilson, "A Comparison of Commercial and Military Computer Security Policies", Proceedings of the 1987 IEEE Symposium on Security and Privacy, Computer Society Press of the IEEE, Washington, DC.

[9] Biba, K. J., "Integrity Considerations for Secure Computer Systems", TR-76-372, HQ Electronic Systems Division, Hanscomb AFB, MA, April 1977

[10] UK Ministry of Defense, CHOTS contract

[11] Trusted Systems Interoperability Group, "Commerical IP Security Option", TSIG Document

[12] "The Revised IP Security Option", Internet RFCs 1038, 1108

[13] Postel, J., "User Datagram Protocol", Internet RFC 768