

December 19, 1997

# A Common Internet File System (CIFS/1.0) Protocol

*Preliminary Draft*

## Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or made obsolete 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".

To learn the current status of any Internet-Draft, please check the "lid-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on [ftp.is.co.za](ftp://ftp.is.co.za) (Africa), [ftp.nordu.net](ftp://ftp.nordu.net) (Europe), [munnnari.oz.au](ftp://ftp.munnari.oz.au) (Pacific Rim), [ds.internic.net](ftp://ftp.ds.internic.net) (US East Coast), or [ftp.isi.edu](ftp://ftp.isi.edu) (US West Coast).

Distribution of this document is unlimited. Please send comments to the authors at [<cifs@microsoft.com>](mailto:cifs@microsoft.com). Discussion of CIFS is on the mailing list [<CIFS@DISCUSS.MICROSOFT.COM>](mailto:CIFS@DISCUSS.MICROSOFT.COM); subscribe by sending a message to [<ListServ@CIFS@DISCUSS.MICROSOFT.COM>](mailto:ListServ@CIFS@DISCUSS.MICROSOFT.COM) with a body of "subscribe CIFS you@your.domain". The mailing list archives are at [<http://discuss.microsoft.com/archives/cifs.html>](http://discuss.microsoft.com/archives/cifs.html). There is a CIFS home page at [<http://www.microsoft.com/intdev/cifs>](http://www.microsoft.com/intdev/cifs).

## Abstract

This document describes the CIFS file sharing protocol, version 1.0. Client systems use this protocol to request file and print services from server systems over a network. It is based on the Server Message Block protocol widely in use by personal computers and workstations running a wide variety of operating systems.

## Table Of Contents

1	Introduction .....	4
1.1	SUMMARY OF FEATURES .....	4
1.1.1	File access .....	5
1.1.2	File and record locking .....	5
1.1.3	Safe caching, read-ahead, and write-behind.....	5
1.1.4	File change notification .....	5
1.1.5	Protocol version negotiation.....	5
1.1.6	Extended attributes.....	5

1.1.7	<i>Distributed replicated virtual volumes</i> .....	5
1.1.8	<i>Server name resolution independence</i> .....	6
1.1.9	<i>Batched requests</i> .....	6
2	Protocol Operation Overview.....	6
2.1	SERVER NAME DETERMINATION.....	6
2.2	SERVER NAME RESOLUTION .....	7
2.3	SAMPLE MESSAGE FLOW .....	7
2.4	CIFS PROTOCOL DIALECT NEGOTIATION .....	7
2.5	MESSAGE TRANSPORT .....	8
2.5.1	<i>Connection Management</i> .....	8
2.6	OPPORTUNISTIC LOCKS .....	8
2.6.1	<i>Exclusive Oplocks</i> .....	9
2.6.2	<i>Batch Oplocks</i> .....	10
2.6.3	<i>Level II Oplocks</i> .....	10
2.7	SECURITY MODEL .....	11
2.8	AUTHENTICATION .....	11
2.9	DISTRIBUTED FILESYSTEM (DFS) SUPPORT .....	11
3	SMB Message Formats and Data Types.....	12
3.1	NOTATION .....	12
3.2	SMB HEADER.....	12
3.2.1	<i>Flags field</i> .....	13
3.2.2	<i>Flags2 Field</i> .....	14
3.2.3	<i>Tid Field</i> .....	14
3.2.4	<i>Pid Field</i> .....	15
3.2.5	<i>Uid Field</i> .....	15
3.2.6	<i>Mid Field</i> .....	15
3.2.7	<i>Status Field</i> .....	15
3.2.8	<i>Timeouts</i> .....	15
3.2.9	<i>Data Buffer (BUFFER) and String Formats</i> .....	16
3.3	FILE NAMES .....	16
3.4	WILDCARDS .....	17
3.5	DFS PATHNAMES .....	17
3.6	TIME AND DATE ENCODING.....	18
3.7	ACCESS MODE ENCODING.....	18
3.8	ACCESS MASK ENCODING.....	19
3.9	OPEN FUNCTION ENCODING.....	20
3.10	OPEN ACTION ENCODING .....	20
3.11	FILE ATTRIBUTE ENCODING .....	21
3.12	EXTENDED FILE ATTRIBUTE ENCODING .....	21
3.13	BATCHING REQUESTS ("ANDX" MESSAGES).....	22
3.14	"TRANSACTION" STYLE SUBPROTOCOLS.....	23
3.14.1	<i>SMB_COM_TRANSACTION2 Format</i> .....	24
3.14.2	<i>3.13.2 SMB_COM_NT_TRANSACTION Formats</i> .....	26
3.14.3	<i>Functional Description</i> .....	27
3.15	VALID SMB REQUESTS BY NEGOTIATED DIALECT.....	29
4	SMB Requests .....	30
4.1	SESSION REQUESTS .....	31
4.1.1	<i>NEGOTIATE: Negotiate Protocol</i> .....	31

4.1.2	<i>SESSION_SETUP_ANDX: Session Setup</i> .....	34
4.1.3	<i>LOGOFF_ANDX: User Logoff</i> .....	38
4.1.4	<i>TREE_CONNECT_ANDX: Tree Connect</i> .....	39
4.1.5	<i>TREE_DISCONNECT: Tree Disconnect</i> .....	41
4.1.6	<i>TRANS2_QUERY_FS_INFORMATION: Get File System Information</i> .....	42
4.1.7	<i>ECHO: Ping the Server</i> .....	45
4.1.8	<i>NT_CANCEL: Cancel request</i> .....	46
4.2	FILE REQUESTS .....	47
4.2.1	<i>NT_CREATE_ANDX: Create or Open File</i> .....	47
4.2.2	<i>NT_TRANSACT_CREATE: Create or Open File with EAs or SD</i> .....	50
4.2.3	<i>CREATE_TEMPORARY: Create Temporary File</i> .....	51
4.2.4	<i>READ_ANDX: Read Bytes</i> .....	51
4.2.5	<i>WRITE_ANDX: Write Bytes to file or resource</i> .....	53
4.2.6	<i>LOCKING_ANDX: Lock or Unlock Byte Ranges</i> .....	54
4.2.7	<i>SEEK: Seek in File</i> .....	56
4.2.8	<i>FLUSH: Flush File</i> .....	57
4.2.9	<i>CLOSE: Close File</i> .....	58
4.2.10	<i>DELETE: Delete File</i> .....	59
4.2.11	<i>RENAME: Rename File</i> .....	60
4.2.12	<i>MOVE: Rename File</i> .....	60
4.2.13	<i>COPY: Copy File</i> .....	62
4.2.14	<i>TRANS2_QUERY_PATH_INFORMATION: Get File Attributes given Path</i> .....	63
4.2.15	<i>TRANS2_QUERY_FILE_INFORMATION: Get File Attributes Given FID</i> .....	69
4.2.16	<i>TRANS2_SET_PATH_INFORMATION: Set File Attributes given Path</i> .....	69
4.2.17	<i>TRANS2_SET_FILE_INFORMATION: Set File Attributes Given FID</i> .....	71
4.3	DIRECTORY REQUESTS .....	72
4.3.1	<i>TRANS2_CREATE_DIRECTORY: Create Directory (with optional EAs)</i> .....	72
4.3.2	<i>DELETE_DIRECTORY: Delete Directory</i> .....	73
4.3.3	<i>CHECK_DIRECTORY: Check Directory</i> .....	73
4.3.4	<i>TRANS2_FIND_FIRST2: Search Directory using Wildcards</i> .....	74
4.3.5	<i>TRANS2_FIND_NEXT2: Resume Directory Search Using Wildcards</i> .....	79
4.3.6	<i>FIND_CLOSE2: Close Directory Search</i> .....	79
4.3.7	<i>NT_TRANSACT_NOTIFY_CHANGE: Request Change Notification</i> .....	80
4.4	DFS OPERATIONS .....	81
4.4.1	<i>TRANS2_GET_DFS_REFERRAL: Retrieve Distributed Filesystem Referral</i> .....	81
4.4.2	<i>TRANS2_REPORT_DFS_INCONSISTENCY: Inform a server about DFS Error</i> .....	83
4.5	MISCELLANEOUS OPERATIONS .....	84
4.5.1	<i>NT_TRANSACT_IOCTL</i> .....	84
4.5.2	<i>NT_TRANSACT_QUERY_SECURITY_DESC</i> .....	85
4.5.3	<i>NT_TRANSACT_SET_SECURITY_DESC</i> .....	85
5	SMB Symbolic Constants .....	85
5.1	SMB COMMAND CODES .....	85
5.2	SMB_COM_TRANSACTION2 SUBCOMMAND CODES .....	87
5.3	SMB_COM_NT_TRANSACTION SUBCOMMAND CODES .....	87
5.4	SMB PROTOCOL DIALECT CONSTANTS .....	88
6	Error Codes and Classes .....	88
7	Legal Notice .....	91

8	References .....	91
9	Authors' Addresses .....	92
10	Appendix A -- NETBIOS transport over TCP .....	92
10.1	CONNECTION ESTABLISHMENT.....	92
10.1.1	<i>Backwards compatability</i> .....	92
10.2	SERVER-SIDE CONNECTION PROCEDURES .....	92
11	Appendix B -- TCP transport .....	93
12	Appendix C -- Share Level Server Security .....	93

# 1 Introduction

This document describes the file sharing protocol for a proposed Common Internet File System (CIFS). CIFS is intended to provide an open cross-platform mechanism for client systems to request file services from server systems over a network. It is based on the standard Server Message Block (SMB) protocol widely in use by personal computers and workstations running a wide variety of operating systems. An earlier version of this protocol was documented as part of the X/OPEN (now Open Group) CAE series of standards [7]; this document updates the specification to include the latest shipping versions, and is published to allow the creation of implementations that inter-operate with those implementations.

The scope of this specification is limited to describing requests and responses for file services. Separate specifications exist for clients requesting services other than file services, e.g. print services.

Use of the Internet and the World Wide Web has been characterized by read-only access. Existing protocols such as FTP are good solutions for one-way file transfer. However, new read/write interfaces will become increasingly necessary as the Internet becomes more interactive and collaborative. Adoption of a common file sharing protocol having modern semantics such as shared files, byte-range locking, coherent caching, change notification, replicated storage, etc. would provide important benefits to the Internet community.

## 1.1 Summary of features

The protocol supports the following features:

- o File access
- o File and record locking
- o Safe caching, read-ahead, and write-behind
- o File change notification
- o Protocol version negotiation
- o Extended attributes
- o Distributed replicated virtual volumes

- o Server name resolution independence
- o Batched requests
- o Unicode file names

### 1.1.1 File access

The protocol supports the usual set of file operations: open, close, read, write, and seek.

### 1.1.2 File and record locking

The protocol supports file and record locking, as well as unlocked access to files. Applications that lock files can not be improperly interfered with by applications that do not; once a file or record is locked, non-locking applications are denied access to the file.

### 1.1.3 Safe caching, read-ahead, and write-behind

The protocol supports caching, read-ahead, and write-behind, even for unlocked files, as long as they are safe. All these optimizations are safe as long as only one client is accessing a file; read-caching and read-ahead are safe with many clients accessing a file as long as all are just reading. If many clients are writing a file simultaneously, then none are safe, and all file operations have to go to the server. The protocol notifies all clients accessing a file of changes in the number and access mode of clients accessing the file, so that they can use the most optimized safe access method.

### 1.1.4 File change notification

Applications can register with a server to be notified if and when file or directory contents are modified. They can use this to (for example) know when a display needs to be refreshed, without having to constantly poll the server.

### 1.1.5 Protocol version negotiation

There are several different versions and sub-versions of this protocol; a particular version is referred to as a *dialect*. When two machines first come into network contact they negotiate the dialect to be used. Different dialects can include both new messages as well as changes to the fields and semantics of existing messages in other dialects.

### 1.1.6 Extended attributes

In addition to many built-in file attributes, such as creation and modification times, non-file system attributes can be added by applications, such as the author's name, content description, *etc.*

### 1.1.7 Distributed replicated virtual volumes

The protocol supports file system subtrees which look like to clients as if they are on a single volume and server, but which actually span multiple volumes and servers. The files and directories of such a subtree can be physically moved to different servers, and their names do not have to change, isolating clients from changes in the server configuration. These subtrees can also

be transparently replicated for load sharing and fault tolerance. When a client requests a file, the protocol uses referrals to transparently direct a client to the server that stores it.

### 1.1.8 Server name resolution independence

The protocol allows clients to resolve server names using any name resolution mechanism. In particular, it allows using the DNS, permitting access to the file systems of other organizations over the Internet, or hierarchical organization of servers' names within an organization. Earlier versions of the protocol only supported a flat server name space.

### 1.1.9 Batched requests

The protocol supports the batching of multiple requests into a single message, in order to minimize round trip latencies, even when a later request depends on the results of an earlier one.

## 2 Protocol Operation Overview

In order to access a file on a server, a client has to:

- o Parse the full file name to determine the server name, and the relative name within that server.
- o Resolve the server name to a transport address (this may be cached)
- o Make a connection to the server (if no connection is already available)
- o Exchange CIFS messages (see below for an example)

This process may be repeated as many times as desired. Once the connection has been idle for a while, it may be torn down.

### 2.1 Server Name Determination

How the client determines the name of the server and the relative name within the server is outside of the scope of this specification. However, just for expository purposes, here are three examples.

In the URL "file://fs.megacorp.com/users/fred/stuff.txt", the client could take the part between the leading double slashes and the next slash as the server name and the remainder as the relative name -- in this example "fs.megacorp.com" and "/users/fred/stuff.txt", respectively.

In the path name "\\corpserver\public\policy.doc" the client could take the part between the leading double backslashes and the next slash as the server name, and the remainder as the relative name -- in this example, "corpserver" and "\public\policy.doc" respectively.

In the path name "x:\policy.doc" the client could use "x" as an index into a table that contains a server name and a file name prefix. If the contents of such a table for "x" were "corpserver" and "\public", then the server name and relative name would be the same as in the previous example.

## 2.2 Server Name Resolution

Like server name determination, how the client resolves the name to the transport address of the server is outside the scope of this specification. All that is required by CIFS is that a CIFS client **MUST** have some means to resolve the name of a CIFS server to a transport address, and that a CIFS server **MUST** register its name with a name resolution service known its clients.

Some examples of name resolution mechanisms include: using the Domain Name System (DNS) [1,2], and using NETBIOS name resolution (see RFC 1001 and RFC 1002 [3,4]). The server name might also be specified as the string form of an IPv4 address in the usual dotted decimal notation, *e.g.*, "157.33.135.101"; in this case, "resolution" consists of converting to the 32 bit IPv4 address.

Which method is used is configuration dependent; the default **SHOULD** be DNS to encourage interoperability over the Internet.

Note: The name resolution mechanism used may place constraints on the form of the server name; for example, in the case of NETBIOS, the server name must be 15 characters or less, and be upper case.

## 2.3 Sample Message Flow

The following illustrates a typical message exchange sequence for a client connecting to a user level server, opening a file, reading its data, closing the file, and disconnecting from the server. Note: using the CIFS request batching mechanism (called the "AndX" mechanism), the second to sixth messages in this sequence can be combined into one, so there are really only three round trips in the sequence, and the last one can be done asynchronously by the client.

Client Command =====	Server Response =====
SMB_COM_NEGOTIATE	Must be the first message sent by client to the server. Includes a list of SMB dialects supported by the client. Server response indicates which SMB dialect should be used.
SMB_COM_SESSION_SETUP_ANDX	Transmits the user's name and credentials to the server for verification. Successful server response has Uid field set in SMB header used for subsequent SMBs on behalf of this user.
SMB_COM_TREE_CONNECT_ANDX	Transmits the name of the disk share the client wants to access. Successful server response has Tid field set in SMB header used for subsequent SMBs referring to this resource.
SMB_COM_OPEN_ANDX	Transmits the name of the file, relative to Tid, the client wants to open. Successful server response includes a file id (Fid) the client should supply for subsequent operations on this file.
SMB_COM_READ	Client supplies Tid, Fid, file offset, and number of bytes to read. Successful server response includes the requested file data.
SMB_COM_CLOSE	Client closes the file represented by Tid and Fid. Server responds with success code.
SMB_COM_TREE_DISCONNECT	Client disconnects from resource represented by Tid.

## 2.4 CIFS Protocol Dialect Negotiation

The first message sent from an CIFS client to an CIFS server must be one whose *Command* field is SMB\_COM\_NEGOTIATE. The format of this client request includes an array of NULL terminated strings indicating the dialects of the CIFS protocol which the client supports. The server compares this list against the list of dialects the server supports and returns the index of the chosen dialect in the response message.

## 2.5 Message Transport

CIFS is transport independent. The CIFS protocol assumes:

- o a reliable connection oriented message-stream transport, and makes no higher level attempts to ensure sequenced delivery of messages between the client and server.
- o a well known endpoint for the CIFS service
- o some mechanism to detect failures of either the client or server node, and to deliver such an indication to the client or server software so they can clean up state. When a reliable transport connection from a client terminates, all work in progress by that client is terminated by the server and all resources open by that client on the server are closed.

It can run over any transport that meets these requirements. Some transports do not natively meet all the requirements, and a standard encapsulation of CIFS for that transport may need to be defined. Appendix A defines how to run CIFS over NETBIOS over TCP; Appendix B defines how to run CIFS over TCP.

### 2.5.1 Connection Management

Once a connection is established, the rules for reliable transport connection dissolution are:

- o If a server receives a transport establishment request from a client with which it is already conversing, the server may terminate all other transport connections to that client. This is to recover from the situation where the client was suddenly rebooted and was unable to cleanly terminate its resource sharing activities with the server.
- o A server may drop the transport connection to a client at any time if the client is generating malformed or illogical requests. However, wherever possible the server should first return an error code to the client indicating the cause of the abort.
- o If a server gets a hard error on the transport (such as a send failure) the transport connection to that client may be aborted.
- o A server may terminate the transport connection when the client has no open resources on the server, however, we recommend that the termination be performed only after some time has passed or if resources are scarce on the server. This will help performance in that the transport connection will not need to be reestablished if activity soon begins anew. Client software is expected to be able to automatically reconnect to the server if this happens.

## 2.6 Opportunistic Locks

Network performance can be increased if a client does not need to inform the server immediately about every change it makes to a file, or have to worry that other clients can make its information about the file out of date. For example, a client does not have to immediately write information into a file on the server if the client knows that no other process is accessing the data. Likewise, the client can buffer read-ahead data from the file if the client knows that no other process is writing the data.

The mechanism which allows clients to dynamically alter their buffering strategy in a consistent manner is known as "opportunistic locks", or *oplocks* for short. Versions of the CIFS file sharing protocol including and newer than the "LANMAN1.0" dialect support oplocks. (Note, however, that an implementation, even of these later dialects, can implement oplocks trivially by always refusing to grant them.)

There are three different types of oplocks:



- o A *Level II* oplock, when held, informs a client that there are multiple concurrent clients of a file, and none has yet modified it. It allows the client to perform reads and file attribute fetches using cached or read-ahead local information, but all other requests have to be sent to the server.
- o An *exclusive* oplock, when held, informs a client that it is the only one to have a file open. It allows the client to perform all file operations using cached or read-ahead local information until it closes the file, at which time the server has to be updated with any changes made to the state of the file (contents and attributes).
- o A *batch* oplock, when held, informs a client that it is the only one to have a file open. It allows the client to perform all file operations on cached or read-ahead local information (including opens and closes).

If a client holds no oplocks, all requests other than reads must be sent to the server. Reads may be performed using cached or read-ahead data as long as the byte range has been locked by the client; otherwise they too must be sent to the server.

When a client opens a file, it may request that the server grant it an exclusive or batch oplock on the file. The response from the server indicates the type of oplock granted to the client. If cached or read-ahead information was retained after the file was last closed, the client must verify that the last modified time is unchanged when the file is reopened before using the retained information.

The SMB\_COM\_LOCKING\_ANDX SMB is used to convey oplock break requests and acknowledgements (as well as lock and unlock requests).

## 2.6.1 Exclusive Oplocks

The exclusive oplock protocol is:

Client		<->	Server
A	B		
=====	=====	===	=====
Open ("foo")	Open("foo")	->	Open OK. Exclusive oplock granted.  read data  oplock break to A lock(s) response(s) write(s) response(s) open response to B
<locks, writes>		<-	
read (large)		->	
<reads from read-ahead >		<-	
lock(s)		->	
write(s)		<-	
close or oplock ack		->	
		<-	
		->	
		<-	

When client A opens the file, it can request an exclusive oplock. Provided no one else has the file open on the server, then the server MAY grant the oplock to client A.

If, at some point in the future, another client, such as client B, requests an open of the same file, or requests a path name based operation on the file, then the server MUST tell client A to relinquish its exclusive oplock. If client B's request will not modify the state of the file, the server MAY tell client A that its exclusive oplock has been replaced by a level II oplock.

When a client's exclusive oplock is broken, it must synchronize the server to the local state of the file (contents and attributes) and any locks it holds on the file, and then acknowledge the oplock break request. After the server receives the acknowledgement, it can process B's request.

## 2.6.2 Batch Oplocks

The batch oplock protocol is:

Client		<->	Server
A	B		
=====	=====	=====	=====
Open("foo")		->	Open OK. Batch oplock granted.
Read		<-	
<close>		->	
<open>		<-	
<seek>			read data
read		->	
<close>		<-	data
	Open("foo")	->	
		<-	Oplock break to A
Close		->	
		<-	Close OK to A
		<-	Open OK to B

When client A opens the file, it can request a batch oplock. Provided no one else has the file open on the server, then the server MAY grant the oplock to client A.

If, at some point in the future, another client, such as client B, requests *any* operation on the same file, then the server MUST tell client A to relinquish its batch oplock. If client B's request will not modify the state of the file (or rename it), the server MAY tell client A that its batch oplock has been replaced by a level II oplock.

If A has the file open at the time the oplock break request is received, its actions will be the same as if it had an exclusive oplock. If A does not have the file open at the time the oplock break request is received, it sends a close to the server. Once the file is actually closed at the server, client B's open request can be processed.

## 2.6.3 Level II Oplocks

The Level II oplock protocol is:

Client		<->	Server
A	B		
=====	=====	=====	=====
Open("foo")	Open("foo")	->	Open OK. Exclusive oplock granted.  data  Break to Level II oplock to A  lock(s) response(s)  Open OK. Oplock II oplock granted to B
Read		<-	
		->	
		<-	
		->	
lock(s)		<-	
		->	
oplock ack		<-	

When a client opens a file, it may request an exclusive or batch oplock. If the requested oplock cannot be granted, then the server MAY grant a Level II oplock if the file currently has an oplock on it. If there is currently an exclusive or batch oplock on the file, it must be broken and the break acknowledged before the open is processed. If there is currently a Level II oplock on the file, it does not need to be broken, and the open may be processed immediately.

If any client sends a request to modify the state of a file that has a Level II oplock, the server must ask all clients holding an oplock on the file to break it, but need not wait for an acknowledgement.

## 2.7 Security Model

Each server makes a set of resources available to clients on the network. A resource being shared may be a directory tree, printer, etc. So far as clients are concerned, the server has no storage or service dependencies on any other servers; a client considers the server to be the sole provider of the file (or other resource) being accessed.

The CIFS protocol requires server authentication of users before file accesses are allowed, and each server authenticates its own users. A client system must send authentication information to the server before the server will allow access to its resources.

A server requires the client to provide a user name and some proof of identity (often something cryptographically derived from a password) to gain access. The granularity of authorization is up to the server. For example, it may use the account name to check access control lists on individual files, or may have one access control list that applies to all files in the directory tree.

When a server validates the account name and password presented by the client, an identifier representing that authenticated instance of the user is returned to the client in the *Uid* field of the response SMB. This *Uid* must be included in all further requests made on behalf of the user from that client.

## 2.8 Authentication

The information on authentication that was in previous revisions of this document has been moved to a different specification.

## 2.9 Distributed Filesystem (DFS) Support

Protocol dialects of NT LM 0.1.2 and later support distributed filesystem operations. The distributed filesystem gives a way for this protocol to use a single consistent file naming scheme which may span a collection of different servers and shares. The

distributed filesystem model employed is a referral - based model. This protocol specifies the manner in which clients receive referrals.

The client can set a flag in the request SMB header indicating that the client wants the server to resolve this SMB's paths within the DFS known to the server. The server attempts to resolve the requested name to a file contained within the local directory tree indicated by the TID of the request and proceeds normally. If the request pathname resolves to a file on a different system, the server returns the following error:

`STATUS_DFS_PATH_NOT_COVERED` - the server does not support the part of the DFS namespace needed to resolved the pathname in the request. The client should request a referral from this server for further information.

A client asks for a referral with the `TRANS2_DFS_GET_REFERRAL` request containing the DFS pathname of interest. The response from the server indicates how the client should proceed.

The method by which the topological knowledge of the DFS is stored and maintained by the servers is not specified by this protocol.

### 3 SMB Message Formats and Data Types

Clients exchange messages with a server to access resources on that server. These messages are called Server Message Blocks (SMBs), and every SMB message has a common format.

This section describes the entire set of SMB commands and responses exchanged between CIFS clients and servers. It also details which SMBs are introduced into the protocol as higher dialect levels are negotiated.

#### 3.1 Notation

This specification makes use of "C"-like notation to describe the formats of messages. Unlike the "C" language, which allows for implementation flexibility in laying out structures, this specification adopts the following rules. Multi-byte values are always transmitted least significant byte first. All fields, except "bit-fields", are aligned on the nearest byte boundary (even if longer than a byte), and there is no implicit padding. Fields using the "bit field" notation are defined to be laid out within the structure with the first-named field occupying the lowest order bits, the next named field the next lowest order bits, and so on.

#### 3.2 SMB header

While each SMB command has specific encodings, there are some fields in the SMB header which have meaning to all SMBs. These fields and considerations are described in the following sections.

```
typedef unsigned char UCHAR;           // 8 unsigned bits
typedef unsigned short USHORT;         // 16 unsigned bits
typedef unsigned long ULONG;           // 32 unsigned bits

typedef struct {
    ULONG LowPart;
    LONG HighPart;
} LARGE_INTEGER;                      // 64 bits of data

typedef struct {
    UCHAR Protocol[4];                 // Contains 0xFF, 'SMB'
    UCHAR Command;                     // Command code
```

```

union {
    struct {
        UCHAR ErrorClass;           // Error class
        UCHAR Reserved;             // Reserved for future use
        USHORT Error;               // Error code
    } DosError;
    ULONG Status;                  // 32-bit error code
} Status;
UCHAR Flags;                     // Flags
USHORT Flags2;                   // More flags
union {
    USHORT Pad[6];                 // Ensure section is 12 bytes long
    struct {
        USHORT PidHigh;            // High part of PID
        UCHAR SecuritySignature[8]; // reserved for security
    } Extra;
};
USHORT Tid;                      // Tree identifier
USHORT Pid;                      // Caller's process id
USHORT Uid;                      // Unauthenticated user id
USHORT Mid;                      // multiplex id
UCHAR WordCount;                 // Count of parameter words
USHORT ParameterWords[ WordCount ]; // The parameter words
USHORT ByteCount;                // Count of bytes
UCHAR Buffer[ ByteCount ];        // The bytes
} SMB_HEADER;

```

All SMBs in this specification have identical format up to the *ParameterWords* fields. (Some obsolescent ones do not.) Different SMBs have a different number and interpretation of *ParameterWords* and *Buffer*. All reserved fields in the SMB header must be zero.

- o *Command* is the operation code that this SMB is requesting or responding to.

### 3.2.1 Flags field

This field contains 8 individual flags, numbered from least significant to most significant, which have the following meanings:

Bit ===	Meaning =====	Earliest Dialect =====
0	Reserved for obsolescent requests. (LOCK_AND_READ, WRITE_AND_CLOSE)	LANMAN1.0
1	Reserved (must be zero).	
2	Reserved (must be zero).	
3	When on, all pathnames in this SMB must be treated as case-less. When off, the pathnames are case sensitive.	LANMAN1.0
4	Reserved (clients must send as zero; servers must ignore).	
5	Reserved for obsolescent requests. (SMB_COM_OPEN, SMB_COM_CREATE and SMB_COM_CREATE_NEW)	LANMAN1.0
6	Reserved for obsolescent requests. (SMB_COM_OPEN, SMB_COM_CREATE and SMB_COM_CREATE_NEW)	LANMAN1.0
7	SMB_FLAGS_SERVER_TO_REDIR. When on, this SMB is being sent from the server in response to a client request. The Command field usually contains the same value in a protocol request from the client to the server as in the matching response from the server to the client. This bit unambiguously distinguishes the command request from the command response	PC NETWORK PROGRAM 1.0

### 3.2.2 Flags2 Field

This field contains six individual flags, numbered from least significant bit to most significant bit, which are defined below. Flags which not defined must be set to zero.

Bit ===	Meaning =====	Earliest Dialect =====
0	If set in a request, the server may return long components in path names in the response.	
1	If set, the client is aware of extended attributes.	
11	If set, the client is aware of Extended Security	NT LM 0.12
12	If set, any request pathnames in this SMB should be resolved in the Distributed File System.	NT LM 0.12
13	If set, indicates that a read will be permitted if the client does not have read permission but does have execute permission. This flag is only useful on a read request.	
14	If set, specifies that the returned error code is a 32 bit error code in Status.Status. Otherwise the Status.DosError.ErrorClass and Status.DosError.Error fields contain the DOS-style error information. When passing NT status codes is negotiated, this flag should be set for every SMB.	NT LM 0.12
15	If set, any fields of datatype STRING in this SMB message are encoded as UNICODE. Otherwise, they are in ASCII.	NT LM 0.12

### 3.2.3 Tid Field

*Tid* represents an instance of an authenticated connection to a server resource. The server returns *Tid* to the client when the client successfully connects to a resource, and the client uses *Tid* in subsequent requests referring to the resource.

In most SMB requests, *Tid* must contain a valid value. Exceptions include prior to getting a *Tid* established including SMB\_COM\_NEGOTIATE, SMB\_COM\_TREE\_CONNECT, SMB\_COM\_ECHO, and SMB\_COM\_SESSION\_SETUP\_ANDX. 0xFFFF should be used for *Tid* for these situations. The server is always responsible for enforcing use of a valid *Tid* where appropriate.

### 3.2.4 Pid Field

*Pid* is the caller's process id, and is generated by the client to uniquely identify a process within the client computer. Concurrency control is associated with *Pid* (and *PidHigh*) -- sharing modes and locks are arbitrated using the *Pid*. For example, if a file is successfully opened for exclusive access, subsequent opens from other clients or from the same client with a different *Pid* will be refused.

Clients inform servers of the creation of a new process by simply introducing a new *Pid* value into the dialogue for new processes. The client operating system must ensure that the appropriate close and cleanup SMBs will be sent when the last process referencing a file closes it. From the server's point of view, there is no concept of *Fids* "belonging to" processes. A *Fid* returned by the server to one process may be used by any other process using the same transport connection and *Tid*.

It is up to the client operating system to ensure only authorized client processes gain access to *Fids* (and *Tids*). On SMB\_COM\_TREE\_DISCONNECT (or when the client and server session is terminated) with a given *Tid*, the server will invalidate any files opened by any process on that client.

### 3.2.5 Uid Field

*Uid* is a user ID assigned by the server after a user authenticates to it, and that it will associate with that user until the client requests the association be broken. After authentication to the server, the client SHOULD make sure that the *Uid* is not used for a different user than the one that authenticated. (It is permitted that a single user have more than one *Uid*.) Requests that do authorization, such as open requests, will perform access checks using the identity associated with the *Uid*.

### 3.2.6 Mid Field

The multiplex ID (*Mid*) is used along with *Pid* to allow multiplexing the single client and server connection among the client's multiple processes, threads, and requests per thread. Clients may have many outstanding requests (up to the negotiated number) at one time. Servers MAY respond to requests in any order, but a response message MUST always contain the same *Mid* and *Pid* values as the corresponding request message. The client MUST NOT have multiple outstanding requests to a server with the same *Mid* and *Pid*.

### 3.2.7 Status Field

An SMB returns error information to the client in the *Status* field. Protocol dialects prior to NT LM 0.12 return status to the client using the combination of *Status.DosError.ErrorClass* and *Status.DosError.Error*. Beginning with NT LM 0.12 CIFS servers can return 32 bit error information to clients using *Status.Status* if the incoming client SMB has bit 14 set in the *Flags2* field of the SMB header. The contents of response parameters are not guaranteed in the case of an error return, and must be ignored. For write-behind activity, a subsequent write or close of the file may return the fact that a previous write failed. Normally write-behind failures are limited to hard disk errors and device out of space.

### 3.2.8 Timeouts

In general, SMBs are not expected to block at the server; they should return "immediately". But some SMB requests do indicate timeout periods for the completion of the request on the server. If a server implementation can not support timeouts, then an error can be returned just as if a timeout had occurred if the resource is not available immediately upon request.

### 3.2.9 Data Buffer (*BUFFER*) and String Formats

The data portion of SMBs typically contains the data to be read or written, file paths, or directory paths. The format of the data portion depends on the message. All fields in the data portion have the same format. In every case it consists of an identifier byte followed by the data.

Identifier =====	Description =====	Value =====
Data Block	See Below	1
Dialect	Null terminated String	2
Pathname	Null terminated String	3
ASCII	Null terminated String	4
Variable block	See Below	5

When the identifier indicates a data block or variable block then the format is a word indicating the length followed by the data.

In all dialects prior to NT LM 0.12, all strings are encoded in ASCII. If the agreed dialect is NT LM 0.12 or later, Unicode strings may be exchanged. Unicode strings include file names, resource names, and user names. This applies to null-terminated strings, length specified strings and the type-prefixed strings. In all cases where a string is passed in Unicode format, the Unicode string must be word-aligned with respect to the beginning of the SMB. Should the string not naturally fall on a two-byte boundary, a null byte of padding will be inserted, and the Unicode string will begin at the next address. In the description of the SMBs, items that may be encoded in Unicode or ASCII are labeled as STRING. If the encoding is ASCII, even if the negotiated string is Unicode, the quantity is labeled as UCHAR.

For type-prefixed Unicode strings, the padding byte is found after the type byte. The type byte is 4 (indicating SMB\_FORMAT\_ASCII) independent of whether the string is ASCII or Unicode. For strings whose start addresses are found using offsets within the fixed part of the SMB (as opposed to simply being found at the byte following the preceding field,) it is guaranteed that the offset will be properly aligned.

Strings that are never passed in Unicode are:

- o The protocol strings in the Negotiate SMB request.
- o The service name string in the Tree\_Connect\_AndX SMB.

When Unicode is negotiated, bit 15 should be set in the *Flags2* field of every SMB header.

Despite the flexible encoding scheme, no field of a data portion may be omitted or included out of order. In addition, neither an *WordCount* nor *ByteCount* of value 0 at the end of a message may be omitted.

## 3.3 File Names

File names in the CIFS protocol consist of components separated by a backslash ('\'). Early clients of the CIFS protocol required that the name components adhere to an 8.3 format name. These names consist of two parts: a basename of no more than 8 characters, and an extension of no more than 3 characters. The basename and extension are separated by a '.'. All characters are legal in the basename and extension *except* the space character (0x20) and:

" . / \ [ ] : + | < > = , \* ?



If the client has indicated long name support by setting *bit2* in the *Flags2* field of the SMB header, this indicates that the client is not bound by the 8.3 convention. Specifically this indicates that any SMB which returns file names to the client may return names which do not adhere to the 8.3 convention, and have a total length of up to 255 characters. This capability was introduced with the LM1.2X002 protocol dialect.

## 3.4 Wildcards

Some SMB requests allow wildcards to be given for the filename. The wildcard allows a number of files to be operated on as a unit without having to separately enumerate the files and individually operate on each one from the client.

If the client is using 8.3 names, each part of the name ( base (8) or extension (3) ) is treated separately. For long filenames the . in the name is significant even though there is no longer a restriction on the size of each of the components.

The ? character is a wild card for a single character. If a filename part commences with one or more "?"s then exactly that number of characters will be matched by the wildcards, e.g., "??x" equals "abx" but not "abcx" or "ax". When a filename part has trailing "?"s then it matches the specified number of characters or less, e.g., "x??" matches "xab", "xa" and "x", but not "xabc". If only "?"s are present in the filename part, then it is handled as for trailing "?"s

The \* character matches an entire part of the name, as does an empty specification for that part. A part consisting of \* means that the rest of the component should be filled with ? and the search should be performed with this wildcard character. For example, "\*.abc" or ".abc" match any file with an extension of "abc". ".\*", "\*" or "null" match all files in a directory.

If the negotiated dialect is "NT LM 0.12" or later, and the client requires MS-DOS wildcard matching semantics, UNICODE wildcards should be translated according to the following rules:

Translate the ? literal to >

Translate the . literal to " if it is followed by a ? or a \*

Translate the \* literal to < if it is followed by a .

The translation can be performed in-place.

## 3.5 DFS Pathnames

A DFS pathname adheres to the standard described in the FileNames section. A DFS enabled client accessing a DFS share should set the *Flags2* bit 12 in all name based SMB requests indicating to the server that the enclosed pathname should be resolved in the Distributed File System namespace. The pathname should always have the full file name, including the server name and share name. If the server can resolve the DFS name to a piece of local storage, the local storage will be accessed. If the server determines that the DFS name actually maps to a different server share, the access to the name will fail with the 32 bit status STATUS\_PATH\_NOT\_COVERED (0xC0000257), or DOS error ERRsrv/ERRbadpath.

On receiving this error, the DFS enabled client should ask the server for a *referral* (see TRANS2\_GET\_DFS\_REFERRAL). The referral request should contain the full file name.

The response to the request will contain a list of server and share names to try, and the part of the request file name that junctions to the list of server shares. If the ServerType field of the referral is set to 1 (SMB server), then the client should resubmit the request with the *original* file name to one of the server shares in the list, once again setting the *Flags2* bit 12 in the SMB. If the ServerType field is not 1, then the client should strip off the part of the file name that junctions to the server share before resubmitting the request to one of servers in the list.

A response to a referral request may elicit a response that does *not* have the StorageServers bit set. In that case, the client should resubmit the *referral request* to one of the servers in the list, until it finally obtains a referral response that has the StorageServers bit set, at which point the client can resubmit the request SMB to one of the listed server shares.

If, after getting a referral with the StorageServers bit set and resubmitting the request to one of the server shares in the list, the server fails the request with STATUS\_PATH\_NOT\_COVERED, it must be the case that there is an inconsistency between the view of the DFS namespace held by the server granting the referral and the server listed in that referral. In this case, the client may inform the server granting the referral of this inconsistency via the TRANS2\_REPORT\_DFS\_INCONSISTENCY SMB.

## 3.6 Time And Date Encoding

When SMB requests or responses encode time values, the following describes the various encodings used.

```
struct {  
    USHORT Day : 5;  
    USHORT Month : 4;  
    USHORT Year : 7;  
} SMB_DATE;
```

The Year field has a range of 0-119, which represents years 1980 - 2099. The Month is encoded as 1-12, and the day ranges from 1-31

```
struct {  
    USHORT TwoSeconds : 5;  
    USHORT Minutes : 6;  
    USHORT Hours : 5;  
} SMB_TIME;
```

Hours ranges from 0-23, Minutes range from 0-59, and TwoSeconds ranges from 0-29 representing two second increments within the minute.

```
typedef struct {  
    ULONG LowTime;  
    LONG HighTime;  
} TIME;
```

TIME indicates a signed 64-bit integer representing either an absolute time or a time interval. Times are specified in units of 100ns. A positive value expresses an absolute time, where the base time (the 64-bit integer with value 0) is the beginning of the year 1601 AD in the Gregorian calendar. A negative value expresses a time interval relative to some base time, usually the current time.

```
typedef unsigned long UTIME;
```

UTIME is the number of seconds since Jan 1, 1970, 00:00:00.0.

## 3.7 Access Mode Encoding

Various client requests and server responses, such as SMB\_COM\_OPEN, pass file access modes encoded into a USHORT. The encoding of these is as follows:

1111 11  
 5432 1098 7654 3210  
 rWrC rLLL rSSS rAAA

where:

W - Write through mode. No read ahead or write behind allowed on this file or device. When the response is returned, data is expected to be on the disk or device.

S - Sharing mode:

- 0 - Compatibility mode
- 1 - Deny read/write/execute (exclusive)
- 2 - Deny write
- 3 - Deny read/execute
- 4 - Deny none

A - Access mode

- 0 - Open for reading
- 1 - Open for writing
- 2 - Open for reading and writing
- 3 - Open for execute

rSSSrAAA = 11111111 (hex FF) indicates FCB open (???)

C - Cache mode

- 0 - Normal file
- 1 - Do not cache this file

L - Locality of reference

- 0 - Locality of reference is unknown
- 1 - Mainly sequential access
- 2 - Mainly random access
- 3 - Random access with some locality
- 4 to 7 - Currently undefined

## 3.8 Access Mask Encoding

The **ACCESS\_MASK** structure is one 32 bit value containing standard, specific, and generic rights. These rights are used in access-control entries (ACEs) and are the primary means of specifying the requested or granted access to an object.

The bits in this value are allocated as follows:

Bits	Meaning
0 - 15	Specific rights. Contains the access mask specific to the object type associated with the mask.
16 - 23	Standard rights. Contains the object's standard access rights and can be a combination of the following predefined flags:

Bit	Flag	Meaning
-----	------	---------

16	DELETE	Delete access
17	READ_CONTROL	Read access to the owner, group, and discretionary access-control list (ACL) of the security descriptor
18	WRITE_DAC	Write access to the discretionary access-control list (ACL)
19	WRITE_OWNER	Write access to owner
20	SYNCHRONIZE	Windows NT: Synchronize access

Bits	Meaning
24	Access system security (ACCESS_SYSTEM_SECURITY). This flag is not a typical access type. It is used to indicate access to a system ACL. This type of access requires the calling process to have a specific privilege.
25	Maximum allowed (MAXIMUM_ALLOWED)
26, 27	Reserved
28	Generic all (GENERIC_ALL)
29	Generic execute (GENERIC_EXECUTE)
30	Generic write (GENERIC_WRITE)
31	Generic read (GENERIC_READ)

### 3.9 Open Function Encoding

*OpenFunction* specifies the action to be taken depending on whether or not the file exists. This word has the following format:

bits:

```
1111 11
5432 1098 7654 3210
rrrr rrrr rrrC rrOO
```

where:

C - Create (action to be taken if file does not exist).

- 0 -- Fail.
- 1 -- Create file.

r - reserved (must be zero).

O - Open (action to be taken if file exists).

- 0 - Fail.
- 1 - Open file.
- 2 - Truncate file.

### 3.10 Open Action Encoding

*Action* in the response to an open or create request describes the action taken as a result of the request. It has the following format:

bits:

```
1111 11
5432 1098 7654 3210
Lrrr rrrr rrrr rrOO
```

where:

L - Lock (single user total file lock status).

0 -- file opened by another user (or mode not supported by server).

1 -- file is opened only by this user at the present time.

r - reserved (must be zero).

O - Open (action taken on Open).

1 - The file existed and was opened.

2 - The file did not exist but was created.

3 - The file existed and was truncated.

## 3.11 File Attribute Encoding

When SMB messages exchange file attribute information, it is encoded in 16 bits as:

Value =====	Description =====
0x01	Read only file
0x02	Hidden file
0x04	System file
0x08	Volume
0x10	Directory
0x20	Archive file
others	Reserved - must be 0

## 3.12 Extended File Attribute Encoding

The extended file attributes is a 32 bit value composed of attributes and flags.

Any combination of the following attributes is acceptable, except all other file attributes override FILE\_ATTR\_NORMAL:

Name =====	Value =====	Meaning =====
ATTR_ARCHIVE	0x020	The file has not been archived since it was last modified. Applications use this attribute to mark files for backup or removal.
ATTR_COMPRESSED	0x800	The file or directory is compressed. For a file, this means that all of the data in the file is compressed. For a directory, this means that compression is the default for newly created files and subdirectories.
ATTR_NORMAL	0x080	The file has no other attributes set. This attribute is valid only if used alone.
ATTR_HIDDEN	0x002	The file is hidden. It is not to be included in an ordinary directory listing.
ATTR_READONLY	0x001	The file is read only. Applications can read the file but cannot write to it or delete it.
ATTR_TEMPORARY	0x100	The file is temporary
ATTR_DIRECTORY	0x010	The file is a directory
ATTR_SYSTEM	0x004	The file is part of or is used exclusively by the operating system.

Any combination of the following flags is acceptable:

Name =====	Value =====	Meaning =====
WRITE_THROUGH	0x80000000	Instructs the operating system to write through any intermediate cache and go directly to the file. The operating system can still cache write operations, but cannot lazily flush them.
NO_BUFFERING	0x20000000	Requests the server to open the file with no intermediate buffering or caching; the server is not obliged to honor the request. An application must meet certain requirements when working with files opened with <code>FILE_FLAG_NO_BUFFERING</code> . File access must begin at offsets within the file that are integer multiples of the volume's sector size; and must be for numbers of bytes that are integer multiples of the volume's sector size. For example, if the sector size is 512 bytes, an application can request reads and writes of 512, 1024, or 2048 bytes, but not of 335, 981, or 7171 bytes.
RANDOM_ACCESS	0x10000000	Indicates that the application intends to access the file randomly. The server MAY use this flag to optimize file caching.
SEQUENTIAL_SCAN	0x08000000	Indicates that the file is to be accessed sequentially from beginning to end. Windows uses this flag to optimize file caching. If an application moves the file pointer for random access, optimum caching may not occur; however, correct operation is still guaranteed. Specifying this flag can increase performance for applications that read large files using sequential access. Performance gains can be even more noticeable for applications that read large files mostly sequentially, but occasionally skip over small ranges of bytes.
DELETE_ON_CLOSE	0x04000000	Requests that the server delete the file immediately after all of its handles have been closed.
BACKUP_SEMANTICS	0x02000000	Indicates that the file is being opened or created for a backup or restore operation. The server SHOULD allow the client to override normal file security checks, provided it has the necessary permission to do so.
POSIX_SEMANTICS	0x01000000	Indicates that the file is to be accessed according to POSIX rules. This includes allowing multiple files with names differing only in case, for file systems that support such naming. (Use care when using this option because files created with this flag may not be accessible by applications written for MS-DOS, Windows 3.x, or Windows NT.)

### 3.13 Batching Requests ("AndX" Messages)

LANMAN1.0 and later dialects of the CIFS protocol allow multiple SMB requests to be sent in one message to the server. Messages of this type are called AndX SMBs, and they obey the following rules:

- o The embedded command does not repeat the SMB header information. Rather the next SMB starts at the *WordCount* field.
- o All multiple (chained) requests must fit within the negotiated transmit size. For example, if `SMB_COM_TREE_CONNECT_ANDX` included `OPENandX` `SMB_COM_OPEN_ANDX` which included `SMB_COM_WRITE` were sent, they would all have to fit within the negotiated buffer size. This would limit the size of the write.
- o There is one message sent containing the chained requests and there is one response message to the chained requests. The server may NOT elect to send separate responses to each of the chained requests.

- o All chained responses must fit within the negotiated transmit size. This limits the maximum value on an embedded SMB\_COM\_READ for example. It is the client's responsibility to not request more bytes than will fit within the multiple response.
- o The server will implicitly use the result of the first command in the "X" command. For example the *Tid* obtained via SMB\_COM\_TREE\_CONNECT\_ANDX would be used in the embedded SMB\_COM\_OPEN\_ANDX and the *Fid* obtained in the SMB\_COM\_OPEN\_ANDX would be used in the embedded SMB\_COM\_READ.
- o Each chained request can only reference the same *Fid* and *Tid* as the other commands in the combined request. The chained requests can be thought of as performing a single (multi-part) operation on the same resource.
- o The first *Command* to encounter an error will stop all further processing of embedded commands. The server will not back out commands that succeeded. Thus if a chained request contained SMB\_COM\_OPEN\_ANDX and SMB\_COM\_READ and the server was able to open the file successfully but the read encountered an error, the file would remain open. This is exactly the same as if the requests had been sent separately.
- o If an error occurs while processing chained requests, the last response (of the chained responses in the buffer) will be the one which encountered the error. Other unprocessed chained requests will have been ignored when the server encountered the error and will not be represented in the chained response. Actually the last valid *AndXCommand* (if any) will represent the SMB on which the error occurred. If no valid *AndXCommand* is present, then the error occurred on the first request/response and *Command* contains the command which failed. In all cases the error information are returned in the SMB header at the start of the response buffer.
- o Each chained request and response contains the offset (from the start of the SMB header) to the next chained request/response (in the *AndXOffset* field in the various "and X" protocols defined later e.g. SMB\_COM\_OPEN\_ANDX). This allows building the requests unpacked. There may be space between the end of the previous request (as defined by *WordCount* and *ByteCount*) and the start of the next chained request. This simplifies the building of chained protocol requests. Note that because the client must know the size of the data being returned in order to post the correct number of receives (e.g. SMB\_COM\_TRANSACTION, SMB\_COM\_READ\_MPX), the data in each response SMB is expected to be truncated to the maximum number of 512 byte blocks (sectors) which will fit (starting at a 32 bit boundary) in the negotiated buffer size with the odd bytes remaining (if any) in the final buffer.

### 3.14 "Transaction" Style Subprotocols

The "transaction" style subprotocols are used for commands that potentially need to transfer a large amount of data (greater than 64K bytes).

### 3.14.1 SMB\_COM\_TRANSACTION2 Format

Primary Client Request =====	Description =====
Command UCHAR WordCount; USHORT TotalParameterCount; USHORT TotalDataCount; USHORT MaxParameterCount; USHORT MaxDataCount; UCHAR MaxSetupCount; UCHAR Reserved; USHORT Flags;  ULONG Timeout; USHORT Reserved2; USHORT ParameterCount; USHORT ParameterOffset; USHORT DataCount; USHORT DataOffset; UCHAR SetupCount; UCHAR Reserved3; USHORT Setup[SetupCount]; USHORT ByteCount; STRING Name[]; UCHAR Pad[]; UCHAR Parameters[ ParameterCount]; UCHAR Pad1[]; UCHAR Data[ DataCount ];	SMB_COM_TRANSACTION2 Count of parameter words; value = (14 + SetupCount) Total parameter bytes being sent Total data bytes being sent Max parameter bytes to return Max data bytes to return Max setup words to return  Additional information: bit 0 - also disconnect TID in <i>TID</i>  Parameter bytes sent this buffer Offset (from header start) to Parameters Data bytes sent this buffer Offset (from header start) to data Count of setup words Reserved (pad above to word) Setup words (# = SetupWordCount) Count of data bytes Must be NULL Pad to SHORT or LONG Parameter bytes (# = ParameterCount) Pad to SHORT or LONG Data bytes (# = DataCount)

Interim Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0



Secondary Client Request =====	Description =====
Command	SMB_COM_TRANSACTION_SECONDARY
UCHAR WordCount; USHORT TotalParameterCount; USHORT TotalDataCount; USHORT ParameterCount; USHORT ParameterOffset; USHORT ParameterDisplacement; USHORT DataCount; USHORT DataOffset; USHORT DataDisplacement; USHORT Fid;  USHORT ByteCount; UCHAR Pad[]; UCHAR Parameters[ParameterCount]; UCHAR Pad1[]; UCHAR Data[DataCount];	Count of parameter words = 8 Total parameter bytes being sent Total data bytes being sent Parameter bytes sent this buffer Offset (from header start) to Parameters Displacement of these Parameter bytes Data bytes sent this buffer Offset (from header start) to data Displacement of these data bytes <i>FID</i> for handle based requests, else 0xFFFF. This field is present only if this is an SMB_COM_TRANSACTION2 request. Count of data bytes Pad to SHORT or LONG Parameter bytes (# = ParameterCount) Pad to SHORT or LONG Data bytes (# = DataCount)

Server Response =====	Description =====
UCHAR WordCount; USHORT TotalParameterCount; USHORT TotalDataCount; USHORT Reserved; USHORT ParameterCount; USHORT ParameterOffset; USHORT ParameterDisplacement; USHORT DataCount; USHORT DataOffset; USHORT DataDisplacement; UCHAR SetupCount; UCHAR Reserved2; USHORT Setup[SetupWordCount]; USHORT ByteCount; UCHAR Pad[]; UCHAR Parameters[ParameterCount]; UCHAR Pad1[]; UCHAR Data[DataCount];	Count of data bytes; value = 10 + <i>SetupCount</i> Total parameter bytes being sent Total data bytes being sent  Parameter bytes sent this buffer Offset (from header start) to Parameters Displacement of these Parameter bytes Data bytes sent this buffer Offset (from header start) to data Displacement of these data bytes Count of setup words Reserved (pad above to word) Setup words (# = SetupWordCount) Count of data bytes Pad to SHORT or LONG Parameter bytes (# = ParameterCount) Pad to SHORT or LONG Data bytes (# = DataCount)

### 3.14.2 3.13.2 SMB\_COM\_NT\_TRANSACTION Formats

Primary Client Request =====	Description =====
UCHAR WordCount; UCHAR MaxSetupCount; USHORT Reserved; ULONG TotalParameterCount; ULONG TotalDataCount; ULONG MaxParameterCount; ULONG MaxDataCount; ULONG ParameterCount; ULONG ParameterOffset; ULONG DataCount; ULONG DataOffset; UCHAR SetupCount; USHORT Function; UCHAR Buffer[1]; USHORT Setup[SetupWordCount]; USHORT ByteCount; UCHAR Pad1[]; UCHAR Parameters[ParameterCount]; UCHAR Pad2[]; UCHAR Data[DataCount]; Data bytes	Count of parameter words; value = (19 + SetupCount) Max setup words to return  Total parameter bytes being sent Total data bytes being sent Max parameter bytes to return Max data bytes to return Parameter bytes sent this buffer Offset (from header start) to Parameters Data bytes sent this buffer Offset (from header start) to data Count of setup words The transaction function code  Setup words Count of data bytes Pad to LONG Parameter bytes Pad to LONG

Interim Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

Secondary Client Request =====	Description =====
UCHAR WordCount; UCHAR Reserved[3]; ULONG TotalParameterCount; ULONG TotalDataCount; ULONG ParameterCount; ULONG ParameterOffset; ULONG ParameterDisplacement;  ULONG DataCount; ULONG DataOffset; ULONG DataDisplacement;  UCHAR Reserved1; USHORT ByteCount; UCHAR Pad1[]; UCHAR Parameters[ParameterCount]; UCHAR Pad2[]; UCHAR Data[DataCount];	Count of parameter words = 18 MBZ Total parameter bytes being sent Total data bytes being sent Parameter bytes sent this buffer Offset (from header start) to Parameters Specifies the offset from the start of the overall parameter block to the parameter bytes that are contained in this message  Data bytes sent this buffer Offset (from header start) to data Specifies the offset from the start of the overall data block to the data bytes that are contained in this message.  Count of data bytes Pad to LONG Parameter bytes Pad to LONG Data bytes

Server Response =====	Description =====
UCHAR WordCount; UCHAR Reserved[3]; ULONG TotalParameterCount; ULONG TotalDataCount; ULONG ParameterCount; ULONG ParameterOffset; ULONG ParameterDisplacement;  ULONG DataCount; ULONG DataOffset; ULONG DataDisplacement;  UCHAR SetupCount; USHORT Setup[SetupWordCount]; USHORT ByteCount; UCHAR Pad1[]; UCHAR Parameters[ParameterCount]; UCHAR Pad2[]; UCHAR Data[DataCount];	Count of data bytes; value = 18 + SetupCount  Total parameter bytes being sent Total data bytes being sent Parameter bytes sent this buffer Offset (from header start) to Parameters Specifies the offset from the start of the overall parameter block to the parameter bytes that are contained in this message  Data bytes sent this buffer Offset (from header start) to data Specifies the offset from the start of the overall data block to the data bytes that are contained in this message.  Count of setup words Setup words Count of data bytes Pad to LONG Parameter bytes Pad to SHORT or LONG Data bytes

### 3.14.3 Functional Description

The transaction *Setup* information and/or *Parameters* define functions specific to a particular resource on a particular server. Therefore the functions supported are not defined by the transaction sub-protocol. The transaction protocol simply provides a means of delivering them and retrieving the results.

The number of bytes needed in order to perform the transaction request may be more than will fit in a single buffer.

At the time of the request, the client knows the number of parameter and data bytes expected to be sent and passes this information to the server via the primary request (*TotalParameterCount* and *TotalDataCount*). This may be reduced by lowering the total number of bytes expected (*TotalParameterCount* and *TotalDataCount*) in each (if any) secondary request.

When the amount of parameter bytes received (total of each *ParameterCount*) equals the total amount of parameter bytes expected (smallest *TotalParameterCount*) received, then the server has received all the parameter bytes.

Likewise, when the amount of data bytes received (total of each *DataCount*) equals the total amount of data bytes expected (smallest *TotalDataCount*) received, then the server has received all the data bytes.

The parameter bytes should normally be sent first followed by the data bytes. However, the server knows where each begins and ends in each buffer by the offset fields (*ParameterOffset* and *DataOffset*) and the length fields (*ParameterCount* and *DataCount*). The displacement of the bytes (relative to start of each) is also known (*ParameterDisplacement* and *DataDisplacement*). Thus the server is able to reassemble the parameter and data bytes should the individual requests be received out of sequence.

If all parameter bytes and data bytes fit into a single buffer, then no interim response is expected and no secondary request is sent.

The client knows the maximum amount of data bytes and parameter bytes which the server may return (from *MaxParameterCount* and *MaxDataCount* of the request). Thus the client initializes its bytes expected variables to these values. The server then informs the client of the actual amounts being returned via each message of the server response (*TotalParameterCount* and *TotalDataCount*). The server may reduce the expected bytes by lowering the total number of bytes expected (*TotalParameterCount* and/or *TotalDataCount*) in each (any) response.

When the amount of parameter bytes received (total of each *ParameterCount*) equals the total amount of parameter bytes expected (smallest *TotalParameterCount*) received, then the client has received all the parameter bytes.

Likewise, when the amount of data bytes received (total of each *DataCount*) equals the total amount of data bytes expected (smallest *TotalDataCount*) received, then the client has received all the data bytes.

The parameter bytes should normally be returned first followed by the data bytes. However, the client knows where each begins and ends in each buffer by the offset fields (*ParameterOffset* and *DataOffset*) and the length fields (*ParameterCount* and *DataCount*). The displacement of the bytes (relative to start of each) is also known (*ParameterDisplacement* and *DataDisplacement*). The client is able to reassemble the parameter and data bytes should the server responses be received out of sequence.

The flow for these transactions over a connection oriented transport is:

1. The client sends the primary client request identifying the total bytes (both parameters and data) which are expected to be sent and contains the set up words and as many of the parameter and data bytes as will fit in a negotiated size buffer. This request also identifies the maximum number of bytes (setup, parameters and data) the server is to return on the transaction completion. If all the bytes fit in the single buffer, skip to step 4.
2. The server responds with a single interim response meaning "OK, send the remainder of the bytes" or (if error response) terminate the transaction.
3. The client then sends another buffer full of bytes to the server. This step is repeated until all of the bytes are sent and received.
4. The Server sets up and performs the transaction with the information provided.
5. Upon completion of the transaction, the server sends back (up to) the number of parameter and data bytes requested (or as many as will fit in the negotiated buffer size). This step is repeated until all result bytes have been returned.

The flow for the transaction protocol when the request parameters and data do not all fit in a single buffer is:

Client =====	<-> =====	Server =====
Primary TRANSACTION request	->	Interim Server Response
Secondary TRANSACTION request 1	<-	
Secondary TRANSACTION request 2	->	
Secondary TRANSACTION request 2	->	
Secondary TRANSACTION request N	->	
	<-	TRANSACTION response 1
	<-	TRANSACTION response 2
	<-	TRANSACTION response m

The flow for the transaction protocol when the request parameters and data does all fit in a single buffer is:

Client =====	<-> =====	Server =====
Primary TRANSACTION request	->	TRANSACTION response 1 TRANSACTION response 2 TRANSACTION response m
	<-	
	<-	
	<-	

The primary transaction request through the final response make up the complete transaction exchange, thus the *Tid*, *Pid*, *Uid* and *Mid* must remain constant and can be used as appropriate by both the server and the client. Of course, other SMB requests may intervene as well.

There are (at least) three ways that actual server responses have been observed to differ from what might be expected. First, some servers will send Pad bytes to move the DataOffset to a 2- or 4-byte boundary even if there are no data bytes; the point here is that the ByteCount must be used instead of ParameterOffset plus ParameterCount to infer the actual message length. Second, some servers always return MaxParameterCount bytes even if the particular Transact2 has no parameter response. Finally, in case of an error, some servers send the "traditional WordCount==0/ByteCount==0" response while others generate a Transact response format.

### 3.15 Valid SMB Requests by Negotiated Dialect

CIFS clients and servers may exchange the following SMB messages if the "PC NETWORK PROGRAM 1.0" dialect is negotiated:

SMB_COM_CREATE_DIRECTORY	SMB_COM_DELETE_DIRECTORY
SMB_COM_OPEN	SMB_COM_CREATE
SMB_COM_CLOSE	SMB_COM_FLUSH
SMB_COM_DELETE	SMB_COM_RENAME
SMB_COM_QUERY_INFORMATION	SMB_COM_SET_INFORMATION
SMB_COM_READ	SMB_COM_WRITE
SMB_COM_LOCK_BYTE_RANGE	SMB_COM_UNLOCK_BYTE_RANGE
SMB_COM_CREATE_TEMPORARY	SMB_COM_CREATE_NEW
SMB_COM_CHECK_DIRECTORY	SMB_COM_PROCESS_EXIT
SMB_COM_SEEK	SMB_COM_TREE_CONNECT
SMB_COM_TREE_DISCONNECT	SMB_COM_NEGOTIATE
SMB_COM_QUERY_INFORMATION_DISK	SMB_COM_SEARCH
SMB_COM_OPEN_PRINT_FILE	SMB_COM_WRITE_PRINT_FILE
SMB_COM_CLOSE_PRINT_FILE	SMB_COM_GET_PRINT_QUEUE

If the "LANMAN 1.0" dialect is negotiated, all of the messages in the previous list must be supported. Clients negotiating LANMAN 1.0 and higher dialects will probably no longer send SMB\_COM\_PROCESS\_EXIT, and the response format for SMB\_COM\_NEGOTIATE is modified as well. New messages introduced with the LANMAN 1.0 dialect are:

SMB_COM_LOCK_AND_READ	SMB_COM_WRITE_AND_UNLOCK
SMB_COM_READ_RAW	SMB_COM_READ_MPX
SMB_COM_WRITE_MPX	SMB_COM_WRITE_RAW
SMB_COM_WRITE_COMPLETE	SMB_COM_WRITE_MPX_SECONDARY
SMB_COM_SET_INFORMATION2	SMB_COM_QUERY_INFORMATION2
SMB_COM_LOCKING_ANDX	SMB_COM_TRANSACTION
SMB_COM_TRANSACTION_SECONDARY	SMB_COM_IOCTL
SMB_COM_IOCTL_SECONDARY	SMB_COM_COPY
SMB_COM_MOVE	SMB_COM_ECHO
SMB_COM_WRITE_AND_CLOSE	SMB_COM_OPEN_ANDX
SMB_COM_READ_ANDX	SMB_COM_WRITE_ANDX
SMB_COM_SESSION_SETUP_ANDX	SMB_COM_TREE_CONNECT_ANDX
SMB_COM_FIND	SMB_COM_FIND_UNIQUE
SMB_COM_FIND_CLOSE	

The "LM1.2X002" dialect introduces these new SMBs:

SMB_COM_TRANSACTION2	SMB_COM_TRANSACTION2_SECONDARY
SMB_COM_FIND_CLOSE2	SMB_COM_LOGOFF_ANDX

"NT LM 0.12" dialect introduces:

SMB_COM_NT_TRANSACT	SMB_COM_NT_TRANSACT_SECONDARY
SMB_COM_NT_CREATE_ANDX	SMB_COM_NT_CANCEL
SMB_COM_NT_RENAME	

## 4 SMB Requests

This section lists the "best practice" SMB requests -- ones that would permit a client to exercise full CIFS functionality and optimum performance when interoperating with a server speaking the latest dialect as of this writing ("NT LM 0.12").

Note that, as of this writing, no existing client restricts itself to only these requests, so no useful server can be written that supports just them. The classification is provided so that future clients will be written to permit future servers to be simpler.

## 4.1 Session Requests

### 4.1.1 NEGOTIATE: Negotiate Protocol

Client Request =====	Description =====
UCHAR WordCount; USHORT ByteCount; struct { UCHAR BufferFormat; UCHAR DialectName[]; } Dialects[];	Count of parameter words = 0 Count of data bytes; min = 2  0x02 -- Dialect ASCII null-terminated string

The Client sends a list of dialects that it can communicate with. The response is a selection of one of those dialects (numbered 0 through n) or -1 (hex FFFF) indicating that none of the dialects were acceptable. The negotiate message is binding on the virtual circuit and must be sent. One and only one negotiate message may be sent, subsequent negotiate requests will be rejected with an error response and no action will be taken.

The protocol does not impose any particular structure to the dialect strings. Implementers of particular protocols may choose to include, for example, version numbers in the string.

If the server does not understand any of the dialect strings, or if `PC NETWORK PROGRAM 1.0` is the chosen dialect, the response format is

Server Response =====	Description =====
UCHAR WordCount; USHORT DialectIndex; USHORT ByteCount;	Count of parameter words = 1 Index of selected dialect Count of data bytes = 0

If the chosen dialect is greater than core up to and including LANMAN2.1, the protocol response format is

Server Response =====	Description =====
UCHAR WordCount; USHORT DialectIndex; USHORT SecurityMode;  USHORT MaxBufferSize; USHORT MaxMpxCount; USHORT MaxNumberVcs; USHORT RawMode;  ULONG SessionKey; SMB_TIME ServerTime; SMB_DATE ServerDate; USHORT ServerTimeZone; USHORT EncryptionKeyLength; USHORT Reserved; USHORT ByteCount UCHAR EncryptionKey[]; STRING PrimaryDomain[];	Count of parameter words = 13 Index of selected dialect Security mode: bit 0: 0 = share, 1 = user bit 1: 1 = use challenge/response authentication Max transmit buffer size (>= 1024) Max pending multiplexed requests Max VCs between client and server Raw modes supported: bit 0: 1 = Read Raw supported bit 1: 1 = Write Raw supported Unique token identifying this session Current time at server Current date at server Current time zone at server MBZ if this is not LM2.1 MBZ Count of data bytes The challenge encryption key The server's primary domain

*MaxBufferSize* is the size of the largest message which the client can legitimately send to the server

If *bit0* of the *Flags* field is set in the negotiate response, this indicates the server supports the SMB\_COM\_LOCK\_AND\_READ and SMB\_COM\_WRITE\_AND\_UNLOCK client requests.

If the *SecurityMode* field indicates the server is running in *user mode*, the client must send appropriate SMB\_COM\_SESSION\_SETUP\_ANDX requests before the server will allow the client to access resources. If the *SecurityMode* fields indicates the client should use challenge/response authentication, the client should use the authentication mechanism specified in section 2.10.

Clients using the "MICROSOFT NETWORKS 1.03" dialect use a different form of raw reads than documented here, and servers are better off setting *RawMode* in this response to 0 for such sessions.

If the negotiated dialect is "DOS LANMAN2.1" or "LANMAN2.1", then *PrimaryDomain* string should be included in this response.

If the negotiated dialect is NT LM 0.12, the response format is



Server Response =====	Description =====
UCHAR WordCount; USHORT DialectIndex; UCHAR SecurityMode;  USHORT MaxMpxCount; USHORT MaxNumberVcs; ULONG MaxBufferSize; ULONG MaxRawSize; ULONG SessionKey; ULONG Capabilities; ULONG SystemTimeLow; ULONG SystemTimeHigh; USHORT ServerTimeZone; UCHAR EncryptionKeyLength; USHORT ByteCount; UCHAR EncryptionKey[];  UCHAR OemDomainName[];  UCHAR GUID[16]  UCHAR SecurityBlob[]	Count of parameter words = 17 Index of selected dialect Security mode: bit 0: 0 = share, 1 = user bit 1: 1 = encrypt passwords bit 2: 1 = Security Signatures (SMB sequence numbers) enabled bit 3: 1 = Security Signatures (SMB sequence numbers) required Max pending multiplexed requests Max VCs between client and server Max transmit buffer size Maximum raw buffer size Unique token identifying this session Server capabilities System (UTC) time of the server (low). System (UTC) time of the server (high). Time zone of server (min from UTC) Length of encryption key. Count of data bytes The challenge encryption key; Present only for Non Extended Security i.e. CAP_EXTENDED_SECURITY is off in the Capabilities field The name of the domain (in OEM chars); Present only for Non Extended Security i.e. CAP_EXTENDED_SECURITY is off in the Capabilities field A globally unique identifier assigned to the server; present only when CAP_EXTENDED_SECURITY is on in the Capabilities field Opaque Security Blob associated with the security package; present only when CAP_EXTENDED_SECURITY is on in the Capabilities field

In addition to the definitions above, *MaxBufferSize* is the size of the largest message which the client can legitimately send to the server. If the client is using a connectionless protocol, *MaxBufferSize* must be set to the smaller of the server's internal buffer size and the amount of data which can be placed in a response packet.

*MaxRawSize* specifies the maximum message size the server can send or receive for SMB\_COM\_WRITE\_RAW or SMB\_COM\_READ\_RAW.

Connectionless clients must set *Sid* to 0 in the SMB request header.

*Capabilities* allows the server to tell the client what it supports. The bit definitions are:

Capability Name =====	Encoding =====	Meaning =====
CAP_RAW_MODE	0x0001	The server supports SMB_COM_READ_RAW and SMB_COM_WRITE_RAW
CAP_MPX_MODE	0x0002	The server supports SMB_COM_READ_MPX and SMB_COM_WRITE_MPX
CAP_UNICODE	0x0004	The server supports Unicode strings
CAP_LARGE_FILES	0x0008	The server supports large files with 64 bit offsets
CAP_NT_SMBS	0x0010	The server supports the SMBs particular to the NT LM 0.12 dialect
CAP_RPC_REMOTE_APIS	0x0020	The sever supports remote API requests via RPC
CAP_STATUS32	0x0040	The server can respond with 32 bit status codes in Status.Status
CAP_LEVEL_II_OPLOCKS	0x0080	The server supports level 2 oplocks
CAP_LOCK_AND_READ	0x0100	The server supports the SMB_COM_LOCK_AND_READ SMB
CAP_NT_FIND	0x0200	
CAP_DFS	0x1000	This server is DFS aware
CAP_BULK_TRANSFER	0x20000000	This server supports SMB_BULK_READ, SMB_BULK_WRITE
CAP_COMPRESSED_DATA	0x40000000	This server supports compressed data transfer (BULK_TRANSFER capability is required in order to support compressed data transfer)
CAP_EXTENDED_SECURITY	0x80000000	This server supports extended security validation

#### 4.1.1.1 Errors

SUCCESS/SUCCESS

ERRSRV/ERRerror

#### 4.1.2 SESSION\_SETUP\_ANDX: Session Setup

This SMB is used to further "Set up" the session normally just established via the negotiate protocol.

One primary function is to perform a "user logon" in the case where the server is in *user level* security mode. The *Uid* in the SMB header is set by the client to be the userid desired for the *AccountName* and validated by the *AccountPassword*.

If the negotiated protocol is prior to NT LM 0.12, the format of SMB\_COM\_SESSION\_SETUP\_ANDX is:

Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT MaxBufferSize; USHORT MaxMpxCount; USHORT VcNumber; ULONG SessionKey; USHORT PasswordLength; ULONG Reserved; USHORT ByteCount; UCHAR AccountPassword[]; STRING AccountName[]; STRING PrimaryDomain[]; STRING NativeOS[]; STRING NativeLanMan[];	Count of parameter words = 10 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Client maximum buffer size Actual maximum multiplexed pending requests 0 = first (only), nonzero=additional VC number Session key (valid iff VcNumber != 0) Account password size Must be 0 Count of data bytes; min = 0 Account Password Account Name Client's primary domain Client's native operating system Client's native LAN Manager type

and the response is:

Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Action;  USHORT ByteCount; STRING NativeOS[]; STRING NativeLanMan[]; STRING PrimaryDomain[];	Count of parameter words = 3 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Request mode: bit0 = logged in as GUEST Count of data bytes Server's native operating system Server's native LAN Manager type Server's primary domain

If the server is in "share level security mode", the account name and passwd should be ignored by the server.

If challenge/response authentication is not being used, *AccountPassword* should be a null terminated ASCII string with *PasswordLength* set to the string size including the null; the password will case insensitive. If challenge/response authentication is being used (see section 2.8), then *AccountPassword* will be the response to the server's challenge, and *PasswordLength* should be set to its length.

The server validates the name and password supplied and if valid, it registers the user identifier on this session as representing the specified *AccountName*. The *Uid* field in the SMB header will then be used to validate access on subsequent SMB requests. The SMB requests where permission checks are required are those which refer to a symbolically named resource such as SMB\_COM\_OPEN, SMB\_COM\_RENAME, SMB\_COM\_DELETE, etc.. The value of the *Uid* is relative to a specific client/server session so it is possible to have the same *Uid* value represent two different users on two different sessions at the server.

Multiple session setup commands may be sent to register additional users on this session. If the server receives an additional SMB\_COM\_SESSION\_SETUP\_ANDX, only the *Uid*, *AccountName* and *AccountPassword* fields need contain valid values (the server MUST ignore the other fields).

The client writes the name of its domain in *PrimaryDomain* if it knows what the domain name is. If the domain name is unknown, the client either encodes it as a NULL string, or as a question mark.

If *bit0* of *Action* is set, this informs the client that although the server did not recognize the *AccountName*, it logged the user in as a guest. This is optional behavior by the server, and in any case one would ordinarily expect guest privileges to limited.

Another function of the Session Set Up protocol is to inform the server of the maximum values which will be utilized by this client. Here *MaxBufferSize* is the maximum message size which the client can receive. Thus although the server may support 16k buffers (as returned in the SMB\_COM\_NEGOTIATE response), if the client only has 4k buffers, the value of *MaxBufferSize* here would be 4096. The minimum allowable value for *MaxBufferSize* is 1024. The SMB\_COM\_NEGOTIATE response includes the server buffer size supported. Thus this is the maximum SMB message size which the client can send to the server. This size may be larger than the size returned to the server from the client via the SMB\_COM\_SESSION\_SETUP\_ANDX protocol which is the maximum SMB message size which the server may send to the client. Thus if the server's buffer size were 4k and the client's buffer size were only 2K, the client could send up to 4k (standard) write requests but must only request up to 2k for (standard) read requests.

The *VcNumber* field specifies whether the client wants this to be the first VC or an additional VC.

The values for *MaxBufferSize*, *MaxMpxCount*, and *VcNumber* must be less than or equal to the maximum values supported by the server as returned in the SMB\_COM\_NEGOTIATE response.

If the server gets a SMB\_COM\_SESSION\_SETUP\_ANDX request with *VcNumber* of 0 and other VCs are still connected to that client, they will be aborted thus freeing any resources held by the server. This condition could occur if the client was rebooted and reconnected to the server before the transport level had informed the server of the previous VC termination.

If the negotiated SMB dialect is "NT LM 0.12" and the server supports ExtendedSecurity i.e. the CAP\_EXTENDED\_SECURITY flag is set in the Capabilities field of the Negotiate Response SMB, the Extended Security SessionSetup SMB format is:

Client Request =====	Description =====
UCHAR WordCount;	Count of parameter words = 12
UCHAR AndXCommand;	Secondary (X) command; 0xFF = none
UCHAR AndXReserved;	Reserved (must be 0)
USHORT AndXOffset;	Offset to next command WordCount
USHORT MaxBufferSize;	Client's maximum buffer size
USHORT MaxMpxCount;	Actual maximum multiplexed pending requests
USHORT VcNumber;	0 = first (only), nonzero=additional VC number
ULONG SessionKey;	Session key (valid iff VcNumber != 0)
USHORT SecurityBlobLength;	Length of opaque security blob
ULONG Reserved;	must be 0
ULONG Capabilities;	Client capabilities
USHORT ByteCount;	Count of data bytes; min = 0
UCHAR SecurityBlob[]	The opaque security blob
STRING NativeOS[];	Client's native operating system, Unicode
STRING NativeLanMan[];	Client's native LAN Manager type, Unicode

The response is:

Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Action;  USHORT SecurityBlobLength USHORT ByteCount; UCHAR SecurityBlob[]  STRING NativeOS[]; STRING NativeLanMan[]; STRING PrimaryDomain[];	Count of parameter words = 3 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Request mode: bit0 = logged in as GUEST length of Security Blob that follows in a later field Count of data bytes SecurityBlob of length specified in field SecurityBlobLength Server's native operating system Server's native LAN Manager type Server's primary domain

There may be multiple parts involved in the security blob exchange. In that case, the server may return an error STATUS\_MORE\_PROCESSING\_REQUIRED (a value of 0xC0000016) in the SMB status. The client can then repeat the SessionSetupAndX SMB with the rest of the security blob.

If the negotiated SMB dialect is "NT LM 0.12" or later and the server does not support Extended Security (i.e. the CAP\_EXTENDED\_SECURITY flag in the Capabilities field of the Negotiate Response SMB is not set), the format of the response SMB is unchanged, but the request is:

Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT MaxBufferSize; USHORT MaxMpxCount; USHORT VcNumber; ULONG SessionKey; USHORT CaseInsensitivePasswordLength; USHORT CaseSensitivePasswordLength; ULONG Reserved; ULONG Capabilities; USHORT ByteCount; UCHAR CaseInsensitivePassword[]; UCHAR CaseSensitivePassword[]; STRING AccountName[]; STRING PrimaryDomain[]; STRING NativeOS[]; STRING NativeLanMan[];	Count of parameter words = 13 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Client's maximum buffer size Actual maximum multiplexed pending requests 0 = first (only), nonzero=additional VC number Session key (valid iff VcNumber != 0) Account password size, ANSI Account password size, Unicode must be 0 Client capabilities Count of data bytes; min = 0 Account Password, ANSI Account Password, Unicode Account Name, Unicode Client's primary domain, Unicode Client's native operating system, Unicode Client's native LAN Manager type, Unicode

The client expresses its capabilities to the server encoded in the *Capabilities* field:

Capability Name =====	Encoding =====	Description =====
CAP_UNICODE	0x0004	The client can use UNICODE strings
CAP_LARGE_FILES	0x0008	The client can deal with files having 64 bit offsets
CAP_NT_SMBS	0x0010	The client understands the SMBs introduced with the NT LM 0.12 dialect. Implies CAP_NT_FIND.
CAP_NT_FIND	0x0200	
CAP_STATUS32	0x0040	The client can receive 32 bit errors encoded in <i>Status.Status</i>
CAP_LEVEL_II_OPLOCKS	0x0080	The client understands Level II oplocks

The entire message sent and received including the optional ANDX SMB must fit in the negotiated maximum transfer size. The following are the only valid SMB commands for *AndXCommand* for SMB\_COM\_SESSION\_SETUP\_ANDX

SMB_COM_TREE_CONNECT_ANDX	SMB_COM_OPEN
SMB_COM_OPEN_ANDX	SMB_COM_CREATE
SMB_COM_CREATE_NEW	SMB_COM_CREATE_DIRECTORY
SMB_COM_DELETE	SMB_COM_DELETE_DIRECTORY
SMB_COM_FIND	SMB_COM_FIND_UNIQUE
SMB_COM_COPY	SMB_COM_RENAME
SMB_COM_NT_RENAME	SMB_COM_CHECK_DIRECTORY
SMB_COM_QUERY_INFORMATION	SMB_COM_SET_INFORMATION
SMB_COM_NO_ANDX_COMMAND	SMB_COM_OPEN_PRINT_FILE
SMB_COM_GET_PRINT_QUEUE	SMB_COM_TRANSACTION

#### 4.1.2.1 Errors

ERRSRV/ERRerror	- no NEG_PROT issued
ERRSRV/ERRbadpw	- password not correct for given username
ERRSRV/ERRtoomanyuids	- maximum number of users per session exceeded
ERRSRV/ERRnosupport	- chaining of this request to the previous one is not supported

#### 4.1.3 LOGOFF\_ANDX: User Logoff

This SMB is the inverse of SMB\_COM\_SESSION\_SETUP\_ANDX.

Client Request =====	Description =====
UCHAR WordCount;	Count of parameter words = 2
UCHAR AndXCommand;	Secondary (X) command; 0xFF = none
UCHAR AndXReserved;	Reserved (must be 0)
USHORT AndXOffset;	Offset to next command WordCount
USHORT ByteCount;	Count of data bytes = 0

Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT ByteCount;	Count of parameter words = 2 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Count of data bytes = 0

The user represented by *Uid* in the SMB header is logged off. The server closes all files currently open by this user, and invalidates any outstanding requests with this *Uid*.

SMB\_COM\_SESSION\_SETUP\_ANDX is the only valid *AndXCommand*. for this SMB.

#### 4.1.3.1 Errors

ERRSRV/invid - TID was invalid

ERRSRV/baduid - UID was invalid

#### 4.1.4 TREE\_CONNECT\_ANDX: Tree Connect

Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Flags;  USHORT PasswordLength; USHORT ByteCount; UCHAR Password[]; STRING Path[]; STRING Service[];	Count of parameter words = 4 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Additional information bit 0 set = disconnect Tid Length of Password[] Count of data bytes; min = 3 Password Server name and share name Service name

The serving machine verifies the combination and returns an error code or an identifier. The full name is included in this request message and the identifier identifying the connection is returned in the *Tid* field of the SMB header. The *Tid* field in the client request is ignored. The meaning of this identifier (*Tid*) is server specific; the client must not associate any specific meaning to it.

If the negotiated dialect is LANMAN1.0 or later, then it is a protocol violation for the client to send this message prior to a successful SMB\_COM\_SESSION\_SETUP\_ANDX, and the server ignores *Password*.

If the negotiated dialect is prior to LANMAN1.0 and the client has not sent a successful SMB\_COM\_SESSION\_SETUP\_ANDX request when the tree connect arrives, a user level security mode server must nevertheless validate the client's credentials as discussed earlier in this document.

*Path* follows UNC style syntax, that is to say it is encoded as \\server\share and it indicates the name of the resource to which the client wishes to connect.

Because *Password* may be an authentication response, it is a variable length field with the length specified by *PasswordLength*. If authentication is not being used, *Password* should be a null terminated ASCII string with *PasswordLength* set to the string size including the terminating null.

The server can enforce whatever policy it desires to govern share access. Typically, if the server is paused, administrative privilege is required to connect to any share; if the server is not paused, administrative privilege is required only for administrative shares (C\$, etc.). Other such policies may include valid times of day, software usage license limits, number of simultaneous server users or share users, etc.

The Service component indicates the type of resource the client intends to access. Valid values are:

Service =====	Description =====	Earliest Dialect Allowed =====
A:	disk share	PC NETWORK PROGRAM 1.0
LPT1:	printer	PC NETWORK PROGRAM 1.0
IPC	named pipe	MICROSOFT NETWORKS 3.0
COMM	communications device	MICROSOFT NETWORKS 3.0
????	any type of device	MICROSOFT NETWORKS 3.0

If *bit0* of *Flags* is set, the tree connection to *Tid* in the SMB header should be disconnected. If this tree disconnect fails, the error should be ignored.

If the negotiated dialect is earlier than DOS LANMAN2 . 1, the response to this SMB is:

Server Response =====	Description =====
UCHAR WordCount;	Count of parameter words = 2
UCHAR AndXCommand;	Secondary (X) command; 0xFF = none
UCHAR AndXReserved;	Reserved (must be 0)
USHORT AndXOffset;	Offset to next command WordCount
USHORT ByteCount;	Count of data bytes; min = 3

If the negotiated is DOS LANMAN2 . 1 or later, the response to this SMB is:

Server Response =====	Description =====
UCHAR WordCount;	Count of parameter words = 3
UCHAR AndXCommand;	Secondary (X) command; 0xFF = none
UCHAR AndXReserved;	Reserved (must be 0)
USHORT AndXOffset;	Offset to next command WordCount
USHORT OptionalSupport;	Optional support bits
USHORT ByteCount;	Count of data bytes; min = 3
UCHAR Service[];	Service type connected to. Always ANSI.
STRING NativeFileSystem[];	Native file system for this tree

*NativeFileSystem* is the name of the filesystem; values to be expected include FAT, NTFS, etc.

*OptionalSupport* bits has the encoding:



Name =====	Encoding =====	Description =====
SMB_SUPPORT_SEARCH_BITS	0x0001	
SMB_SHARE_IS_IN_DFS	0x0002	

Some servers negotiate "DOS LANMAN2.1" dialect or later and still send the "downlevel" (i.e. wordcount==2) response. Valid AndX following commands are

SMB_COM_OPEN	SMB_COM_OPEN_ANDX	SMB_COM_CREATE
SMB_COM_CREATE_NEW	SMB_COM_CREATE_DIRECTORY	SMB_COM_DELETE
SMB_COM_DELETE_DIRECTORY	SMB_COM_FIND	SMB_COM_COPY
SMB_COM_FIND_UNIQUE	SMB_COM_RENAME	
SMB_COM_CHECK_DIRECTORY	SMB_COM_QUERY_INFORMATION	
SMB_COM_GET_PRINT_QUEUE	SMB_COM_OPEN_PRINT_FILE	
SMB_COM_TRANSACTION	SMB_COM_NO_ANDX_CMD	
SMB_COM_SET_INFORMATION	SMB_COM_NT_RENAME	

#### 4.1.4.1 Errors

ERRDOS/ERRnomem  
 ERRDOS/ERRbadpath  
 ERRDOS/ERRinvdevice  
 ERRSRV/ERRaccess  
 ERRSRV/ERRbadpw  
 ERRSRV/ERRinvnetname

#### 4.1.5 TREE\_DISCONNECT: Tree Disconnect

This message informs the server that the client no longer wishes to access the resource connected to with a prior SMB\_COM\_TREE\_CONNECT or SMB\_COM\_TREE\_CONNECT\_ANDX.

Client Request =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

The resource sharing connection identified by *Tid* in the SMB header is logically disconnected from the server. *Tid* is invalidated; it will not be recognized if used by the client for subsequent requests. All locks, open files, etc. created on behalf of *Tid* are released.

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

#### 4.1.5.1 Errors

ERRSRV/ERRinvnid  
ERRSRV/ERRbaduid

#### 4.1.6 TRANS2\_QUERY\_FS\_INFORMATION: Get File System Information

This transaction requests information about a filesystem on the server.

Client Request =====	Value =====
WordCount; TotalParameterCount; MaxSetupCount; SetupCount; Setup[0];	15 2 or 4 0 1 or 2 TRANS2_QUERY_FS_INFORMATION
Parameter Block Encoding =====	Description =====
USHORT Information Level;	Level of information requested

The filesystem is identified by *Tid* in the SMB header.

*MaxDataCount* in the transaction request must be large enough to accommodate the response.

The encoding of the response parameter block depends on the *InformationLevel* requested. Information levels whose values are greater than 0x102 are mapped to corresponding calls to *NtQueryVolumeInformationFile* calls by the server. The two levels below 0x102 are described below. The requested information is placed in the *Data* portion of the transaction response.

InformationLevel =====	Value =====
SMB_INFO_ALLOCATION	1
SMB_INFO_VOLUME	2
SMB_QUERY_FS_VOLUME_INFO	0x102
SMB_QUERY_FS_SIZE_INFO	0x103
SMB_QUERY_FS_DEVICE_INFO	0x104
SMB_QUERY_FS_ATTRIBUTE_INFO	0x105

The following sections describe the *InformationLevel* dependent encoding of the data part of the transaction response.

**4.1.6.1 SMB\_INFO\_ALLOCATION**

Data Block Encoding =====	Description =====
ULONG idFileSystem; ULONG cSectorUnit; ULONG cUnit; ULONG cUnitAvail; USHORT cbSector;	File system identifier. NT server always returns 0 Number of sectors per allocation unit Total number of allocation units Total number of available allocation units Number of bytes per sector

**4.1.6.2 SMB\_INFO\_VOLUME**

Data Block Encoding =====	Description =====
ULONG ulVsn; UCHAR cch; STRING Label;	Volume serial number Number of characters in Label The volume label

**4.1.6.3 SMB\_QUERY\_FS\_VOLUME\_INFO**

Data Block Encoding =====	Description =====
LARGE_INTEGER ULONG ULONG	Volume Creation Time Volume Serial Number Length of Volume Label in bytes
BYTE	Reserved
BYTE	Reserved
STRING Label;	The volume label

**4.1.6.4 SMB\_QUERY\_FS\_SIZE\_INFO**

Data Block Encoding =====	Description =====
LARGE_INTEGER LARGE_INTEGER ULONG	Total Number of Allocation units on the Volume Number of free Allocation units on the Volume Number of sectors in each Allocation unit
ULONG	Number of bytes in each sector

**4.1.6.5 SMB\_QUERY\_FS\_DEVICE\_INFO**

Data Block Encoding =====	Value =====
ULONG ULONG	DeviceType; Values as specified below Characteristics of the device; Values as specified below

For DeviceType, note that the values 0-32767 are reserved for the exclusive use of Microsoft Corporation. The following device types are currently defined:

FILE_DEVICE_BEEP	0x00000001
FILE_DEVICE_CD_ROM	0x00000002
FILE_DEVICE_CD_ROM_FILE_SYSTEM	0x00000003
FILE_DEVICE_CONTROLLER	0x00000004
FILE_DEVICE_DATALINK	0x00000005
FILE_DEVICE_DFS	0x00000006
FILE_DEVICE_DISK	0x00000007
FILE_DEVICE_DISK_FILE_SYSTEM	0x00000008
FILE_DEVICE_FILE_SYSTEM	0x00000009
FILE_DEVICE_INPORT_PORT	0x0000000a
FILE_DEVICE_KEYBOARD	0x0000000b
FILE_DEVICE_MAILSLOT	0x0000000c
FILE_DEVICE_MIDI_IN	0x0000000d
FILE_DEVICE_MIDI_OUT	0x0000000e
FILE_DEVICE_MOUSE	0x0000000f
FILE_DEVICE_MULTI_UNC_PROVIDER	0x00000010
FILE_DEVICE_NAMED_PIPE	0x00000011
FILE_DEVICE_NETWORK	0x00000012
FILE_DEVICE_NETWORK_BROWSER	0x00000013
FILE_DEVICE_NETWORK_FILE_SYSTEM	0x00000014
FILE_DEVICE_NULL	0x00000015
FILE_DEVICE_PARALLEL_PORT	0x00000016
FILE_DEVICE_PHYSICAL_NETCARD	0x00000017
FILE_DEVICE_PRINTER	0x00000018
FILE_DEVICE_SCANNER	0x00000019
FILE_DEVICE_SERIAL_MOUSE_PORT	0x0000001a
FILE_DEVICE_SERIAL_PORT	0x0000001b
FILE_DEVICE_SCREEN	0x0000001c
FILE_DEVICE_SOUND	0x0000001d
FILE_DEVICE_STREAMS	0x0000001e
FILE_DEVICE_TAPE	0x0000001f
FILE_DEVICE_TAPE_FILE_SYSTEM	0x00000020
FILE_DEVICE_TRANSPORT	0x00000021
FILE_DEVICE_UNKNOWN	0x00000022
FILE_DEVICE_VIDEO	0x00000023
FILE_DEVICE_VIRTUAL_DISK	0x00000024
FILE_DEVICE_WAVE_IN	0x00000025
FILE_DEVICE_WAVE_OUT	0x00000026
FILE_DEVICE_8042_PORT	0x00000027
FILE_DEVICE_NETWORK_REDIRECTOR	0x00000028
FILE_DEVICE_BATTERY	0x00000029
FILE_DEVICE_BUS_EXTENDER	0x0000002a
FILE_DEVICE_MODEM	0x0000002b
FILE_DEVICE_VDM	0x0000002c

Some of these device types are not currently accessible over the network and may never be accessible over the network. Some may change to be accessible over the network. The values for device types that may never be accessible over the network may be redefined to be just reserved at some date in the future.

*Characteristics* is the sum of any of the following:

FILE_REMOVABLE_MEDIA	0x00000001
FILE_READ_ONLY_DEVICE	0x00000002
FILE_FLOPPY_DISKETTE	0x00000004
FILE_WRITE_ONCE_MEDIA	0x00000008
FILE_REMOTE_DEVICE	0x00000010
FILE_DEVICE_IS_MOUNTED	0x00000020
FILE_VIRTUAL_VOLUME	0x00000040

#### 4.1.6.6 SMB\_QUERY\_FS\_ATTRIBUTE\_INFO

Data Block Encoding =====	Description =====
ULONG	File System Attributes; possible values described below
LONG	Maximum length of each file name component in number of bytes
ULONG	Length, in bytes, of the name of the file system
STRING	Name of the file system

Where *FileSystemAttributes* is the sum of any of the following:

FILE_CASE_SENSITIVE_SEARCH	0x00000001
FILE_CASE_PRESERVED_NAMES	0x00000002
FILE_PRERSISTENT_ACLS	0x00000004
FILE_FILE_COMPRESSION	0x00000008
FILE_VOLUME_QUOTAS	0x00000010
FILE_DEVICE_IS_MOUNTED	0x00000020
FILE_VOLUME_IS_COMPRESSED	0x00008000

#### 4.1.6.7 Errors

ERRSRV/invnid - TID was invalid  
 ERRSRV/baduid - UID was invalid  
 ERRHRD/ERRnotready - the file system has been removed  
 ERRHRD/ERRdata - disk I/O error  
 ERRSRV/ERRaccess - user does not have the right to perform this operation  
 ERRSRV/ERRinvdevice - resource identified by TID is not a file system

#### 4.1.7 ECHO: Ping the Server

This request is used to test the connection to the server, and to see if the server is still responding.

Client Request =====	Description =====
UCHAR WordCount; USHORT EchoCount; USHORT ByteCount; UCHAR Buffer[1];	Count of parameter words = 1 Number of times to echo data back Count of data bytes; min = 1 Data to echo

Server Response =====	Description =====
UCHAR WordCount; USHORT SequenceNumber; USHORT ByteCount; UCHAR Buffer[1];	Count of parameter words = 1 Sequence number of this echo Count of data bytes; min = 4 Echoed data

Each response echoes the data sent, though *ByteCount* may indicate no data. If *EchoCount* is zero, no response is sent.

*Tid* in the SMB header is ignored, so this request may be sent to the server even if there are no valid tree connections to the server.

The flow for the ECHO protocol is:

Client Request =====	<-> =====	Server Response =====
Echo Request (EchoCount == n)	-> <- <- <-	Echo Response 1 Echo Response 2 Echo Response n

#### 4.1.7.1 Errors

ERRSRV/ERRbaduid - UID was invalid  
 ERRSRV/ERRnoaccess - session has not been established  
 ERRSRV/ERRnosupport - ECHO function is not supported

#### 4.1.8 NT\_CANCEL: Cancel request

This SMB allows a client to cancel a request currently pending at the server.

Client Request =====	Description =====
UCHAR WordCount; USHORT ByteCount;	No words are sent (== 0) No bytes (==0)

The *Sid*, *Uid*, *Pid*, *Tid*, and *Mid* fields of the SMB are used to locate an pending server request from this session. If a pending request is found, it is "hurried along" which may result in success or failure of the original request. No other response is generated for this SMB.

## 4.2 File Requests

### 4.2.1 NT\_CREATE\_ANDX: Create or Open File

This command is used to create or open a file or a directory.

Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; UCHAR Reserved; USHORT NameLength; ULONG Flags;  ULONG RootDirectoryFid; ACCESS_MASK DesiredAccess; LARGE_INTEGER AllocationSize; ULONG ExtFileAttributes; ULONG ShareAccess; ULONG CreateDisposition; ULONG CreateOptions; ULONG ImpersonationLevel; UCHAR SecurityFlags;  USHORT ByteCount; STRING Name[];	Count of parameter words = 24 Secondary command; 0xFF = None Reserved (must be 0) Offset to next command WordCount Reserved (must be 0) Length of Name[] in bytes Create bit set: 0x02 - Request an oplock 0x04 - Request a batch oplock 0x08 - Target of open must be directory If non-zero, open is relative to this directory access desired Initial allocation size File attributes Type of share access Action to take if file exists or not Options to use if creating a file Security QOS information Security tracking mode flags: 0x1 - SECURITY_CONTEXT_TRACKING 0x2 - SECURITY_EFFECTIVE_ONLY Length of byte parameters File to open or create

The *DesiredAccess* parameter is specified in section 3.8 on Access Mask Encoding.

If no value is specified, it still allows an application to query attributes without actually accessing the file.

The *ExtFileAttributes* parameter specifies the file attributes and flags for the file. The parameter's value is the sum of allowed attributes and flags defined in section 3.12 on Extended File Attribute Encoding

The *ShareAccess* field Specifies how this file can be shared. This parameter must be some combination of the following values:

Name	Value	Meaning
	0	Prevents the file from being shared.
FILE_SHARE_READ	0x00000001	Other open operations can be performed on the file for read access.
FILE_SHARE_WRITE	0x00000002	Other open operations can be performed on the file for write access.
FILE_SHARE_DELETE	0x00000004	Other open operations can be performed on the file for delete access.

The *CreateDisposition* parameter can contain one of the following values:

CREATE_NEW	Creates a new file. The function fails if the specified file already exists.
CREATE_ALWAYS	Creates a new file. The function overwrites the file if it exists.
OPEN_EXISTING	Opens the file. The function fails if the file does not exist.
OPEN_ALWAYS	Opens the file, if it exists. If the file does not exist, act like CREATE_NEW.
TRUNCATE_EXISTING	Opens the file. Once opened, the file is truncated so that its size is zero bytes. The calling process must open the file with at least GENERIC_WRITE access. The function fails if the file does not exist.

The *ImpersonationLevel* parameter can contain one or more of the following values:

SECURITY_ANONYMOUS	Specifies to impersonate the client at the Anonymous impersonation level.
SECURITY_IDENTIFICATION	Specifies to impersonate the client at the Identification impersonation level.
SECURITY_IMPERSONATION	Specifies to impersonate the client at the Impersonation impersonation level.
SECURITY_DELEGATION	Specifies to impersonate the client at the Delegation impersonation level.

The *SecurityFlags* parameter can have either of the following two flags set:

SECURITY_CONTEXT_TRACKING	Specifies that the security tracking mode is dynamic. If this flag is not specified, Security Tracking Mode is static.
SECURITY_EFFECTIVE_ONLY	Specifies that only the enabled aspects of the client's security context are available to the server. If you do not specify this flag, all aspects of the client's security context are available. This flag allows the client to limit the groups and privileges that a server can use while impersonating the client.

The response is as follows:



Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; Secondary command; UCHAR AndXReserved; USHORT AndXOffset; UCHAR OplockLevel;  USHORT Fid; ULONG CreateAction; TIME CreationTime; TIME LastAccessTime; TIME LastWriteTime; TIME ChangeTime; ULONG ExtFileAttributes; LARGE_INTEGER AllocationSize; LARGE_INTEGER EndOfFile; USHORT FileType; USHORT DeviceState; BOOLEAN Directory; USHORT ByteCount;	Count of parameter words = 26 0xFF = None MBZ Offset to next command WordCount The oplock level granted 0 - No oplock granted 1 - Exclusive oplock granted 2 - Batch oplock granted 3 - Level II oplock granted The file ID The action taken The time the file was created The time the file was accessed The time the file was last written The time the file was last changed The file attributes The number of bytes allocated The end of file offset  state of IPC device (e.g. pipe) TRUE if this is a directory = 0

The following SMBs may follow SMB\_COM\_NT\_CREATE\_ANDX:

SMB\_COM\_READ            SMB\_COM\_READ\_ANDX  
 SMB\_COM\_IOCTL

## 4.2.2 NT\_TRANSACT\_CREATE: Create or Open File with EAs or SD

This command is used to create or open a file or a directory, when EAs or an SD must be applied to the file.

Request Parameter Block Encoding =====	Description =====
ULONG Flags; ULONG RootDirectoryFid; ACCESS_MASK DesiredAccess; LARGE_INTEGER AllocationSize; ULONG ExtFileAttributes; ULONG ShareAccess; ULONG CreateDisposition; ULONG CreateOptions; ULONG SecurityDescriptorLength; ULONG EaLength; ULONG NameLength; ULONG ImpersonationLevel; UCHAR SecurityFlags; STRING Name[NameLength];	Creation flags (see below) Optional directory for relative open Desired access The initial allocation size in bytes, if file created The extended file attributes The share access Action to take if file exists or not Options for creating a new file Length of SD in bytes Length of EA in bytes Length of name in characters Security QOS information Security QOS information The name of the file (not NULL terminated)
Data Block Encoding =====	Description =====
UCHAR SecurityDescriptor[ SecurityDescriptorLength]; UCHAR ExtendedAttributes[EaLength];	

Creation Flag Name =====	Value =====	Description =====
NT_CREATE_REQUEST_OPLOCK	0x02	Level I oplock requested
NT_CREATE_REQUEST_OPBATCH	0x04	Batch oplock requested
NT_CREATE_OPEN_TARGET_DIR	0x08	Target for open is a directory

Output Parameter Block Encoding =====	Description =====
UCHAR OplockLevel; UCHAR Reserved; USHORT Fid; ULONG CreateAction; ULONG EaErrorOffset; TIME CreationTime; TIME LastAccessTime; TIME LastWriteTime; TIME ChangeTime; ULONG ExtFileAttributes; LARGE_INTEGER AllocationSize; LARGE_INTEGER EndOfFile; USHORT FileType; USHORT DeviceState; BOOLEAN Directory;	The oplock level granted  The file ID The action taken Offset of the EA error The time the file was created The time the file was accessed The time the file was last written The time the file was last changed The file attributes The number of bytes allocated The end of file offset  state of IPC device (e.g. pipe) TRUE if this is a directory

See the description of NT\_CREATE\_ANDX for the definition of the parameters.

### 4.2.3 CREATE\_TEMPORARY: Create Temporary File

The server creates a data file in *Directory* relative to *Tid* in the SMB header and assigns a unique name to it.

Client Request =====	Server Response =====
UCHAR WordCount; USHORT reserved; UTIME CreationTime; USHORT ByteCount; UCHAR BufferFormat; STRING DirectoryName[];	Count of parameter words = 3 Ignored by the server New file's creation time stamp Count of data bytes; min = 2 0x04 Directory name

Server Response =====	Description =====
UCHAR WordCount; USHORT Fid; USHORT ByteCount; UCHAR BufferFormat; STRING Filename[];	Count of parameter words = 1 File handle Count of data bytes; min = 2 0x04 File name

*Fid* is the returned handle for future file access. *Filename* is the name of the file which was created within the requested *Directory*. It is opened in compatibility mode with read/write access for the client.

Support of *CreationTime* by the server is optional.

### 4.2.4 READ\_ANDX: Read Bytes

Large File Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Fid; ULONG Offset; USHORT MaxCount; USHORT MinCount; ULONG Reserved; USHORT Remaining; ULONG OffsetHigh; USHORT ByteCount;	Count of parameter words = 10 or 12 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount File handle Offset in file to begin read Max number of bytes to return Reserved Must be 0 Reserved Upper 32 bits of offset (only if WordCount is 12) Count of data bytes = 0

Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Remaining; USHORT DataCompactionMode; USHORT Reserved; USHORT DataLength; USHORT DataOffset; USHORT Reserved[5]; USHORT ByteCount; UCHAR Pad[]; UCHAR Data[ DataLength];	Count of parameter words = 12 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Reserved -- must be -1  Reserved (must be 0) Number of data bytes (min = 0) Offset (from header start) to data Reserved (must be 0) Count of data bytes  Data from resource

If the negotiated dialect is NT LM 0.12 or later, the client may use the 12 parameter word version of the request. This version allows specification of 64 bit file offsets.

If CAP\_LARGE\_READX was indicated by the server in the negotiate protocol response, the request's *MaxCount* field may exceed the negotiated buffer size if *Fid* refers to a disk file. The server may arbitrarily elect to return fewer than *MaxCount* bytes in response.

The following SMBs may follow SMB\_COM\_READ\_ANDX:  
 SMB\_COM\_CLOSE

#### 4.2.4.1 Errors

ERRDOS/ERRnoaccess  
 ERRDOS/ERRbadfid  
 ERRDOS/ERRlock  
 ERRDOS/ERRbadaccess  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid

## 4.2.5 WRITE\_ANDX: Write Bytes to file or resource

Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Fid; ULONG Offset; ULONG Reserved; USHORT WriteMode;  USHORT Remaining; USHORT Reserved; USHORT DataLength; USHORT DataOffset; ULONG OffsetHigh; USHORT ByteCount; UCHAR Pad[]; UCHAR Data[DataLength];	Count of parameter words = 12 or 14 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount File handle Offset in file to begin write Must be 0 Write mode bits: 0 - write through Bytes remaining to satisfy request  Number of data bytes in buffer (>=0) Offset to data bytes Upper 32 bits of offset (only present if WordCount = 14) Count of data bytes Pad to SHORT or LONG Data to write

Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Count; USHORT Remaining; ULONG Reserved; USHORT ByteCount;	Count of parameter words = 6 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Number of bytes written Reserved  Count of data bytes = 0

A *ByteCount* of 0 does not truncate the file. Rather a zero length write merely transfers zero bytes of information to the file. A request such as SMB\_COM\_WRITE must be used to truncate the file.

If *WriteMode* has bit0 set in the request and *Fid* refers to a disk file, the response is not sent from the server until the data is on stable storage.

If the negotiated dialect is NT LM 0.12 or later, the 14 word format of this SMB may be used to access portions of files requiring offsets expressed as 64 bits. Otherwise, the *OffsetHigh* field must be omitted from the request.

The following are the valid *AndXCommand* values for this SMB:

SMB_COM_READ	SMB_COM_READ_ANDX
SMB_COM_LOCK_AND_READ	SMB_COM_WRITE_ANDX
SMB_COM_CLOSE	

#### 4.2.5.1 Errors

ERRDOS/ERRnoaccess  
 ERRDOS/ERRbadfid  
 ERRDOS/ERRlock  
 ERRDOS/ERRbadaccess  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid

#### 4.2.6 LOCKING\_ANDX: Lock or Unlock Byte Ranges

SMB\_COM\_LOCKING\_ANDX allows both locking and/or unlocking of file range(s).

Client Request =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT Fid; UCHAR LockType; UCHAR OplockLevel; ULONG Timeout; USHORT NumberOfUnlocks; USHORT NumberOfLocks; USHORT ByteCount; LOCKING_ANDX_RANGE Unlocks[]; LOCKING_ANDX_RANGE Locks[];	Count of parameter words = 8 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount File handle See LockType table below The new oplock level Milliseconds to wait for unlock Num. unlock range structs following Num. lock range structs following Count of data bytes Unlock ranges Lock ranges

LockType Flag Name =====	Value =====	Description =====
LOCKING_ANDX_SHARED_LOCK	0x01	Read-only lock
LOCKING_ANDX_OPLOCK_RELEASE	0x02	Oplock break notification
LOCKING_ANDX_CHANGE_LOCKTYPE	0x04	Change lock type
LOCKING_ANDX_CANCEL_LOCK	0x08	Cancel outstanding request
LOCKING_ANDX_LARGE_FILES	0x10	Large file locking format

LOCKING_ANDX_RANGE Format =====	
USHORT Pid; ULONG Offset; ULONG Length;	PID of process "owning" lock Offset to bytes to [un]lock Number of bytes to [un]lock

Large File LOCKING_ANDX_RANGE Format =====	
USHORT Pid; USHORT Pad; ULONG OffsetHigh; ULONG OffsetLow; ULONG LengthHigh; ULONG LengthLow;	PID of process "owning" lock Pad to DWORD align (mbz) Offset to bytes to [un]lock (high) Offset to bytes to [un]lock (low) Number of bytes to [un]lock (high) Number of bytes to [un]lock (low)

Server Response =====	Description =====
UCHAR WordCount; UCHAR AndXCommand; UCHAR AndXReserved; USHORT AndXOffset; USHORT ByteCount;	Count of parameter words = 2 Secondary (X) command; 0xFF = none Reserved (must be 0) Offset to next command WordCount Count of data bytes = 0

Locking is a simple mechanism for excluding other processes read/write access to regions of a file. The locked regions can be anywhere in the logical file. Locking beyond end-of-file is permitted. Any process using the *Fid* specified in this request's *Fid* has access to the locked bytes, other processes will be denied the locking of the same bytes.

The proper method for using locks is not to rely on being denied read or write access on any of the read/write protocols but rather to attempt the locking protocol and proceed with the read/write only if the locks succeeded.

Locking a range of bytes will fail if any subranges or overlapping ranges are locked. In other words, if any of the specified bytes are already locked, the lock will fail.

If *NumberOfUnlocks* is non-zero, the *Unlocks* vector contains *NumberOfUnlocks* elements. Each element requests that a lock at *Offset of Length* be released. If *NumberOfLocks* is nonzero, the *Locks* vector contains *NumberOfLocks* elements. Each element requests the acquisition of a lock at *Offset of Length*.

*Timeout* is the maximum amount of time to wait for the byte range(s) specified to become unlocked. A timeout value of 0 indicates that the server should fail immediately if any lock range specified is locked. A timeout value of -1 indicates that the server should wait as long as it takes for each byte range specified to become unlocked so that it may be again locked by this protocol. Any other value of *smb\_timeout* specifies the maximum number of milliseconds to wait for all lock range(s) specified to become available.

If any of the lock ranges timeout because of the area to be locked is already locked (or the lock fails), the other ranges in the protocol request which were successfully locked as a result of this protocol will be unlocked (either all requested ranges will be locked when this protocol returns to the client or none).

If *LockType* has the `LOCKING_ANDX_SHARED_LOCK` flag set, the lock is specified as a shared lock. Locks for both read and write (where `LOCKING_ANDX_SHARED_LOCK` is clear) should be prohibited, but other shared locks should be permitted. If shared locks can not be supported by a server, the server should map the lock to a lock for both read and write. Closing a file with locks still in force causes the locks to be released in no defined order.

If *LockType* has the `LOCKING_ANDX_LARGE_FILES` flag set and if the negotiated protocol is NT LM 0.12 or later, then the *Locks* and *Unlocks* vectors are in the Large File `LOCKING_ANDX_RANGE` format. This allows specification of 64 bit offsets for very large files.

If the one and only member of the *Locks* vector has the `LOCKING_ANDX_CANCEL_LOCK` flag set in the *LockType* field, the client is requesting the server to cancel a previously requested, but not yet responded to, lock.

If `LockType` has the `LOCKING_ANDX_CHANGE_LOCKTYPE` flag set, the client is requesting that the server atomically change the lock type from a shared lock to an exclusive lock or vice versa. If the server can not do this in an atomic fashion, the server must reject this request. NT and W95 servers do not support this capability.

Oplocks are described in the "Opportunistic Locks" section elsewhere in this document. A client requests an oplock by setting the appropriate bit in the `SMB_COM_OPEN_ANDX` request when the file is being opened in a mode which is not exclusive. The server responds by setting the appropriate bit in the response SMB indicating whether or not the oplock was granted. By granting the oplock, the server tells the client the file is currently only being used by this one client process at the current time. The client can therefore safely do read ahead and write behind as well as local caching of file locks knowing that the file will not be accessed/changed in any way by another process while the oplock is in effect. The client will be notified when any other process attempts to open or modify the oplocked file.

When another user attempts to open or otherwise modify the file which a client has oplocked, the server delays the second attempt and notifies the client via an `SMB_LOCKING_ANDX` SMB asynchronously sent from the server to the client. This message has the `LOCKING_ANDX_OPLOCK_RELEASE` flag set indicating to the client that the oplock is being broken. *OplockLevel* indicates the type of oplock the client now owns. If *OplockLevel* is 0, the client possesses no oplocks on the file at all, if *OplockLevel* is 1 the client possesses a Level II oplock. The client is expected to flush any dirty buffers to the server, submit any file locks and respond to the server with either an `SMB_LOCKING_ANDX` SMB having the `LOCKING_ANDX_OPLOCK_RELEASE` flag set, or with a file close if the file is no longer in use by the client. If the client sends an `SMB_LOCKING_ANDX` SMB with the `LOCKING_ANDX_OPLOCK_RELEASE` flag set and *NumberOfLocks* is zero, the server does not send a response. Since a close being sent to the server and break oplock notification from the server could cross on the wire, if the client gets an oplock notification on a file which it does not have open, that notification should be ignored.

Due to timing, the client could get an "oplock broken" notification in a user's data buffer as a result of this notification crossing on the wire with a `SMB_COM_READ_RAW` request. The client must detect this (use length of msg, "FFSMB", MID of -1 and *Command* of `SMB_COM_LOCKING_ANDX`) and honor the "oplock broken" notification as usual. The server must also note on receipt of an `SMB_COM_READ_RAW` request that there is an outstanding (unanswered) "oplock broken" notification to the client and return a zero length response denoting failure of the read raw request. The client should (after responding to the "oplock broken" notification), use a standard read protocol to redo the read request. This allows a file to actually contain data matching an "oplock broken" notification and still be read correctly.

The entire message sent and received including the optional second protocol must fit in the negotiated maximum transfer size. The following are the only valid SMB commands for *AndXCommand* for `SMB_COM_LOCKING_ANDX`:

<code>SMB_COM_READ</code>	<code>SMB_COM_READ_ANDX</code>
<code>SMB_COM_WRITE</code>	<code>SMB_COM_WRITE_ANDX</code>
<code>SMB_COM_FLUSH</code>	

#### 4.2.6.1 Errors

ERRDOS/ERRbadfile  
ERRDOS/ERRbadfid  
ERRDOS/ERRlock  
ERRDOS/ERRinvdevice  
ERRSRV/ERRinvid  
ERRSRV/ERRbaduid

#### 4.2.7 SEEK: Seek in File

The seek message is sent to set the current file pointer for *Fid*.



Client Request =====	Description =====
UCHAR WordCount; USHORT Fid; USHORT Mode;  LONG Offset; USHORT ByteCount;	Count of parameter words = 4 File handle Seek mode: 0 = from start of file 1 = from current position 2 = from end of file Relative offset Count of data bytes = 0

The starting point of the seek is set by *Mode*:

- 0 seek from start of file
- 1 seek from current file pointer
- 2 seek from end of file

The "current position" reflects the offset plus data length specified in the previous read, write or seek request, and the pointer set by this command will be replaced by the offset specified in the next read, write or seek command.

Server Response =====	Description =====
UCHAR WordCount; ULONG Offset; USHORT ByteCount;	Count of parameter words = 2 Offset from start of file Count of data bytes = 0

The response returns the new file pointer in *Offset* which is expressed as the offset from the start of the file, and may be beyond the current end of file. An attempt to seek to before the start of file sets the current file pointer to start of the file.

This request should generally only be issued by clients wishing to find the size of a file, since all read and write requests include the read or write file position as part of the SMB. This request is inappropriate for very large files, as the offsets specified are only 32 bits. A seek which results in an Offset which can not be expressed in 32 bits returns the least significant.

#### 4.2.7.1 Errors

ERRDOS/ERRbadfid  
 ERRDOS/ERRnoaccess  
 ERRSRV/ERRinvdevice  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid

#### 4.2.8 FLUSH: Flush File

The flush SMB is sent to ensure all data and allocation information for the corresponding file has been written to stable storage. When the *Fid* has a value -1 (hex FFFF) the server performs a flush for all file handles associated with the client and *Pid*. The response is not sent until the writes are complete.

Client Request =====	Description =====
UCHAR WordCount; USHORT Fid; USHORT ByteCount;	Count of parameter words = 1 File handle Count of data bytes = 0

This client request is probably expensive to perform at the server, since the server's operating system is generally scheduling disk writes in a way which is optimal for the system's read and write activity integrated over the entire population of clients. This message from a client "interferes" with the server's ability to optimally schedule the disk activity; clients are discouraged from overuse of this SMB request.

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

#### 4.2.8.1 Errors

ERRDOS/ERRbadfid  
ERRSRV/ERRinvid  
ERRSRV/ERRbaduid

#### 4.2.9 CLOSE: Close File

The close message is sent to invalidate a file handle for the requesting process. All locks or other resources held by the requesting process on the file should be released by the server. The requesting process can no longer use *Fid* for further file access requests.

Client Request =====	Description =====
UCHAR WordCount; USHORT Fid; UTIME LastWriteTime USHORT ByteCount;	Count of parameter words = 3 File handle Time of last write Count of data bytes = 0

If *LastWriteTime* is 0, the server should allow its local operating system to set the file's times. Otherwise, the server should set the time to the values requested. Failure to set the times, even if requested by the client in the request message, should not result in an error response from the server.

If *Fid* refers to a print spool file, the file should be spooled to the printer at this time.

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

#### 4.2.9.1 Errors

ERRDOS/ERRbadfid  
 ERRSRV/ERRinvdevice  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid

### 4.2.10 DELETE: Delete File

The delete file message is sent to delete a data file. The appropriate *Tid* and additional pathname are passed. Read only files may not be deleted, the read-only attribute must be reset prior to file deletion.

Client Request =====	Description =====
UCHAR WordCount; USHORT SearchAttributes; USHORT ByteCount; UCHAR BufferFormat; STRING FileName[];	Count of parameter words = 1  Count of data bytes; min = 2 0x04 File name

Multiple files may be deleted in response to a single request as SMB\_COM\_DELETE supports wildcards

*SearchAttributes* indicates the attributes that the target file(s) must have. If the attribute is zero then only normal files are deleted. If the system file or hidden attributes are specified then the delete is inclusive -both the specified type(s) of files and normal files are deleted. Attributes are described in the "Attribute Encoding" section of this document.

If *bit0* of the *Flags2* field of the SMB header is set, a pattern is passed in, and the file has a long name, then the passed pattern must match the long file name for the delete to succeed. If *bit0* is clear, a pattern is passed in, and the file has a long name, then the passed pattern must match the file's short name for the deletion to succeed.

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

#### 4.2.10.1 Errors

ERRDOS/ERRbadpath  
 ERRDOS/ERRbadfile  
 ERRDOS/ERRnoaccess  
 ERRHRD/ERRnowrite  
 ERRSRV/ERRaccess  
 ERRSRV/ERRinvdevice  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid

## 4.2.11 RENAME: Rename File

The rename file message is sent to change the name of a file.

Client Request =====	Description =====
UCHAR WordCount; USHORT SearchAttributes; USHORT ByteCount; UCHAR BufferFormat1; STRING OldFileName[]; UCHAR BufferFormat2; STRING NewFileName[];	Count of parameter words = 1 Target file attributes Count of data bytes; min = 4 0x04 Old file name 0x04 New file name

Files *OldFileName* must exist and *NewFileName* must not. Both pathnames must be relative to the *Tid* specified in the request. Open files may be renamed.

Multiple files may be renamed in response to a single request as Rename File supports wildcards in the file name (last component of the pathname).

*SearchAttributes* indicates the attributes that the target file(s) must have. If *SearchAttributes* is zero then only normal files are renamed. If the system file or hidden attributes are specified then the rename is inclusive -both the specified type(s) of files and normal files are renamed. The encoding of *SearchAttributes* is described in section 3.11 - File Attribute Encoding.

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

### 4.2.11.1 Errors

ERRDOS/ERRbadpath  
 ERRDOS/ERRbadfile  
 ERRDOS/ERRnoaccess  
 ERRDOS/ERRdiffdevice  
 ERRHRD/ERRnowrite  
 ERRSRV/ERRaccess  
 ERRSRV/ERRinvdevice  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid

## 4.2.12 MOVE: Rename File

The source file is copied to the destination and the source is subsequently deleted.

Client Request =====	Description =====
UCHAR WordCount; USHORT Tid2; USHORT OpenFunction; USHORT Flags;  USHORT ByteCount; UCHAR Format1; STRING OldFileName[]; UCHAR FormatNew; STRING NewFileName[];	Count of parameter words = 3 Second (target) file id what to do if target file exists Flags to control move operations: 0 - target must be a file 1 - target must be a directory 2 - reserved (must be 0) 3 - reserved (must be 0) 4 - verify all writes Count of data bytes; min = 2 0x04 Old file name 0x04 New file name

*OldFileName* is copied to *NewFileName*, then *OldFileName* is deleted. Both *OldFileName* and *NewFileName* must refer to paths on the same server. *NewFileName* can refer to either a file or a directory. All file components except the last must exist; directories will not be created.

*NewFileName* can be required to be a file or a directory by the Flags field.

The *Tid* in the header is associated with the source while *Tid2* is associated with the destination. These fields may contain the same or differing valid values. *Tid2* can be set to -1 indicating that this is to be the same *Tid* as in the SMB header. This allows use of the move protocol with SMB\_TREE\_CONNECT\_ANDX.

Server Response =====	Description =====
UCHAR WordCount; USHORT Count; USHORT ByteCount; UCHAR ErrorFileFormat; STRING ErrorFileName[];	Count of parameter words = 1 Number of files moved Count of data bytes; min = 0 0x04 (only if error) Pathname of file where error occurred

The source path must refer to an existing file or files. Wildcards are permitted. Source files specified by wildcards are processed until an error is encountered. If an error is encountered, the expanded name of the file is returned in *ErrorFileName*. Wildcards are not permitted in *NewFileName*.

*OpenFunction* controls what should happen if the destination file exists. If  $(OpenFunction \& 0x30) == 0$ , the operation should fail if the destination exists. If  $(OpenFunction \& 0x30) == 0x20$ , the destination file should be overwritten.

#### 4.2.12.1 Errors

ERRDOS/ERRfilexists  
 ERRDOS/ERRbadfile  
 ERRDOS/ERRnoaccess  
 ERRDOS/ERRnofiles  
 ERRDOS/ERRbadshare  
 ERRHRD/ERRnowrite

ERRSRV/ERRnoaccess  
 ERRSRV/ERRinvdevice  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid  
 ERRSRV/ERRnosupport  
 ERRSRV/ERRaccess

#### 4.2.13 COPY: Copy File

Client Request =====	Description =====
UCHAR WordCount; USHORT Tid2; USHORT OpenFunction; USHORT Flags;  USHORT ByteCount; UCHAR SourceFileNameFormat; STRING SourceFileName; UCHAR TargetFileNameFormat; STRING TargetFileName;	Count of parameter words = 3 Second (target) path TID What to do if target file exists Flags to control copy operation: bit 0 - target must be a file bit 1 - target must be a dir. bit 2 - copy target mode: 0 = binary, 1 = ASCII bit 3 - copy source mode: 0 = binary, 1 = ASCII bit 4 - verify all writes bit 5 - tree copy Count of data bytes; min = 2 0x04 Pathname of source file 0x04 Pathname of target file

The file at *SourceName* is copied to *TargetFileName*, both of which must refer to paths on the same server.

The *Tid* in the header is associated with the source while *Tid2* is associated with the destination. These fields may contain the same or differing valid values. *Tid2* can be set to -1 indicating that this is to be the same *Tid* as in the SMB header. This allows use of the move protocol with SMB\_TREE\_CONNECT\_ANDX.

Server Response =====	Description =====
UCHAR WordCount; USHORT Count; USHORT ByteCount; UCHAR ErrorFileFormat; STRING ErrorFileName;	Count of parameter words = 1 Number of files copied Count of data bytes; min = 0 0x04 (only if error)

The source path must refer to an existing file or files. Wildcards are permitted. Source files specified by wildcards are processed until an error is encountered. If an error is encountered, the expanded name of the file is returned in *ErrorFileName*. Wildcards are not permitted in *TargetFileName*. *TargetFileName* can refer to either a file or a directory.

The destination can be required to be a file or a directory by the bits in *Flags*. If neither *bit0* nor *bit1* are set, the destination may be either a file or a directory. *Flags* also controls the copy mode. In a binary copy for the source, the copy stops the first time an

EOF (control-Z) is encountered. In a binary copy for the target, the server must make sure that there is exactly one EOF in the target file and that it is the last character of the file.

If the destination is a file and the source contains wildcards, the destination file will either be truncated or appended to at the start of the operation depending on bits in *OpenFunction* (see section 3.7). Subsequent files will then be appended to the file.

If the negotiated dialect is `LM1.2X002` or later, *bit5* of *Flags* is used to specify a tree copy on the remote server. When this option is selected the destination must not be an existing file and the source mode must be binary. A request with *bit5* set and either *bit0* or *bit3* set is therefore an error. When the tree copy mode is selected, the *Count* field in the server response is undefined.

#### 4.2.13.1 Errors

ERRDOS/ERRfilexists  
 ERRDOS/ERRshare  
 ERRDOS/ERRnofids  
 ERRDOS/ERRbadfile  
 ERRDOS/ERRnoaccess  
 ERRDOS/ERRnofiles  
 ERRDOS/ERRbadshare  
 ERRSRV/ERRnoaccess  
 ERRSRV/ERRinvdevice  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid  
 ERRSRV/ERRaccess

#### 4.2.14 TRANS2\_QUERY\_PATH\_INFORMATION: Get File Attributes given Path

This request is used to get information about a specific file or subdirectory.

Client Request =====	Value =====
WordCount MaxSetupCount SetupCount Setup[0]	15 0 1 TRANS2_QUERY_PATH_INFORMATION
Parameter Block Encoding =====	Description =====
USHORT InformationLevel; ULONG Reserved; STRING FileName;	Level of information requested Must be zero File or directory name

The following InformationLevels may be requested:

Information Level	Value
=====	=====
SMB_INFO_STANDARD	1
SMB_INFO_QUERY_EA_SIZE	2
SMB_INFO_QUERY_EAS_FROM_LIST	3
SMB_INFO_QUERY_ALL_EAS	4
SMB_INFO_IS_NAME_VALID	6
SMB_QUERY_FILE_BASIC_INFO	0x101
SMB_QUERY_FILE_STANDARD_INFO	0x102
SMB_QUERY_FILE_EA_INFO	0x103
SMB_QUERY_FILE_NAME_INFO	0x104
SMB_QUERY_FILE_ALL_INFO	0x107
SMB_QUERY_FILE_ALT_NAME_INFO	0x108
SMB_QUERY_FILE_STREAM_INFO	0x109
SMB_QUERY_FILE_COMPRESSION_INFO	0x10B

The requested information is placed in the Data portion of the transaction response. For the information levels greater than 0x100, the transaction response has 1 parameter word which should be ignored by the client.

The following sections describe the *InformationLevel* dependent encoding of the data part of the transaction response.



**4.2.14.1 SMB\_INFO\_STANDARD & SMB\_INFO\_QUERY\_EA\_SIZE**

Data Block Encoding =====	Description =====
SMB_DATE CreationDate; SMB_TIME CreationTime; SMB_DATE LastAccessDate; SMB_TIME LastAccessTime; SMB_DATE LastWriteDate; SMB_TIME LastWriteTime; ULONG DataSize; ULONG AllocationSize; USHORT Attributes; ULONG EaSize;	Date when file was created Time when file was created Date of last file access Time of last file access Date of last write to the file Time of last write to the file File Size Size of filesystem allocation unit File Attributes Size of file's EA information (SMB_INFO_QUERY_EA_SIZE)

**4.2.14.2 SMB\_INFO\_QUERY\_EAS\_FROM\_LIST & SMB\_INFO\_QUERY\_ALL\_EAS**

Response Field =====	Value =====
MaxDataCount	Length of EAList found (minimum value is 4)
Parameter Block Encoding =====	Description =====
USHORT EaErrorOffset	Offset into EAList of EA error
Data Block Encoding =====	Description =====
ULONG ListLength; UCHAR EaList[]	Length of the remaining data The extended attributes list

**4.2.14.3 SMB\_INFO\_IS\_NAME\_VALID**

This requests checks to see if the name of the file contained in the request's *Data* field has a valid path syntax. No parameters or data are returned on this information request. An error is returned if the syntax of the name is incorrect. *Success* indicates the server accepts the path syntax, but it does not ensure the file or directory actually exists.

**4.2.14.4 SMB\_QUERY\_FILE\_BASIC\_INFO**

Data Block Encoding =====	Description =====
LARGE_INTEGER CreationTime; LARGE_INTEGER LastAccessTime; LARGE_INTEGER LastWriteTime; LARGE_INTEGER ChangeTime USHORT Attributes;	Time when file was created Time of last file access Time of last write to the file Time when file was last changed File Attributes

**4.2.14.5 SMB\_QUERY\_FILE\_STANDARD\_INFO**

Data Block Encoding =====	Description =====
LARGE_INTEGER AllocationSize LARGE_INTEGER EndOfFile; ULONG NumberOfLinks BOOLEAN DeletePending BOOLEAN Directory	Allocated size of the file in number of bytes Offset to the first free byte in the file Number of hard links to the file Indicates whether the file is marked for deletion Indicates whether the file is a directory

**4.2.14.6 SMB\_QUERY\_FILE\_EA\_INFO**

Data Block Encoding =====	Description =====
ULONG EASize	Size of the file's extended attributes in number of bytes

**4.2.14.7 SMB\_QUERY\_FILE\_NAME\_INFO**

Data Block Encoding =====	Description =====
ULONG FileNameLength STRING FileName	Length of the file name in number of bytes Name of the file

**4.2.14.8 SMB\_QUERY\_FILE\_ALL\_INFO**

Data Block Encoding =====	Description =====
LARGE_INTEGER CreationTime;	Time when file was created
LARGE_INTEGER LastAccessTime;	Time of last file access
LARGE_INTEGER LastWriteTime;	Time of last write to the file
LARGE_INTEGER ChangeTime	Time when file was last changed
USHORT Attributes;	File Attributes
LARGE_INTEGER AllocationSize	Allocated size of the file in number of bytes
LARGE_INTEGER EndOfFile;	Offset to the first free byte in the file
ULONG NumberOfLinks	Number of hard links to the file
BOOLEAN DeletePending	Indicates whether the file is marked for deletion
BOOLEAN Directory	Indicates whether the file is a directory
LARGE_INTEGER Index Number	A file system unique identifier
ULONG EASize	Size of the file's extended attributes in number of bytes
ULONG AccessFlags	Access that a caller has to the file; Possible values and meanings are specified below
LARGE_INTEGER Index Number	A file system unique identifier
LARGE_INTEGER CurrentByteOffset	Current byte offset within the file
ULONG Mode	Current Open mode of the file handle to the file; possible values and meanings are detailed below
ULONG AlignmentRequirement	Buffer Alignment required by device; possible values detailed below
ULONG FileNameLength	Length of the file name in number of bytes
STRING FileName	Name of the file

The AccessFlags specifies the access permissions a caller has to the file and can have any suitable combination of the following values:

Value	Meaning
FILE_READ_DATA	0x00000001 Data can be read from the file
FILE_WRITE_DATA	0x00000002 Data can be written to the file
FILE_APPEND_DATA	0x00000004 Data can be appended to the file
FILE_READ_EA	0x00000008 Extended attributes associated with the file can be read
FILE_WRITE_EA	0x00000010 Extended attributes associated with the file can be written
FILE_EXECUTE	0x00000020 Data can be read into memory from the file using system paging I/O
FILE_READ_ATTRIBUTES	0x00000080 Attributes associated with the file can be read
FILE_WRITE_ATTRIBUTES	0x00000100 Attributes associated with the file can be written
DELETE	0x00010000 The file can be deleted
READ_CONTROL	0x00020000 The access control list and ownership associated with the file can be read
WRITE_DAC	0x00040000 The access control list and ownership associated with the file can be written.
WRITE_OWNER	0x00080000 Ownership information associated with the file can be written
SYNCHRONIZE	0x00100000 The file handle can wait on to synchronize with the completion of an input/output request

The Mode field specifies the mode in which the file is currently opened. The possible values may be a suitable and logical combination of the following:

Value		Meaning
FILE_WRITE_THROUGH	0x00000002	File is opened in mode where data is written to file before the driver completes a write request
FILE_SEQUENTIAL_ONLY	0x00000004	All access to the file is sequential
FILE_SYNCHRONOUS_IO_ALERT	0x00000010	All operations on the file are performed synchronously
FILE_SYNCHRONOUS_IO_NONALERT	0x00000020	All operations on the file are to be performed synchronously. Waits in the system to synchronize I/O queuing and completion are not subject to alerts.

The AlignmentRequirement field specifies buffer alignment required by the device and can have any one of the following values:

Value		Meaning
FILE_BYTE_ALIGNMENT	0x00000000	The buffer needs to be aligned on a byte boundary
FILE_WORD_ALIGNMENT	0x00000001	The buffer needs to be aligned on a word boundary
FILE_LONG_ALIGNMENT	0x00000003	The buffer needs to be aligned on a 4 byte boundary
FILE_QUAD_ALIGNMENT	0x00000007	The buffer needs to be aligned on an 8 byte boundary
FILE_OCTA_ALIGNMENT	0x0000000f	The buffer needs to be aligned on a 16 byte boundary
FILE_32_BYTE_ALIGNMENT	0x0000001f	The buffer needs to be aligned on a 32 byte boundary
FILE_64_BYTE_ALIGNMENT	0x0000003f	The buffer needs to be aligned on a 64 byte boundary
FILE_128_BYTE_ALIGNMENT	0x0000007f	The buffer needs to be aligned on a 128 byte boundary
FILE_256_BYTE_ALIGNMENT	0x000000ff	The buffer needs to be aligned on a 256 byte boundary
FILE_512_BYTE_ALIGNMENT	0x000001ff	The buffer needs to be aligned on a 512 byte boundary

#### 4.2.14.9 SMB\_QUERY\_FILE\_ALT\_NAME\_INFO

Data Block Encoding	Description
=====	=====
ULONG FileNameLength	Length of the file name in number of bytes
STRING FileName	Name of the file

**4.2.14.10 SMB\_QUERY\_FILE\_STREAM\_INFO**

Data Block Encoding =====	Description =====
ULONG NextEntryOffset ULONG StreamNameLength LARGE_INTEGER StreamSize LARGE_INTEGER StreamAllocationSize STRING FileName	Offset to the next entry (in bytes) Length of the stream name in number of bytes Size of the stream in number of bytes Allocated size of the stream in number of bytes Name of the stream

**4.2.14.11 SMB\_QUERY\_FILE\_COMPRESSION\_INFO**

Data Block Encoding =====	Description =====
LARGE_INTEGER CompressedFileSize USHORT CompressionFormat  UCHAR CompressionUnitShift UCHAR ChunkShift UCHAR ClusterShift  UCHAR Reserved[3]	Size of the compressed file in number of bytes A constant signifying the compression algorithm used. Possible values are: 0 - There is no compression 2- Compression Format is LZNT  stored in log2 format. $1 \ll \text{ChunkShift} = \text{ChunkSizeInBytes}$ indicates how much space must be saved to successfully compress a compression unit

**4.2.15 TRANS2\_QUERY\_FILE\_INFORMATION: Get File Attributes Given FID**

This request is used to get information about a specific file or subdirectory given a handle to it.

Client Request =====	Value =====
WordCount MaxSetupCount SetupCount Setup[0]	15 0 1 TRANS2_QUERY_FILE_INFORMATION
Parameter Block Encoding =====	Description =====
USHORT Fid; USHORT InformationLevel;	Handle of file for request Level of information requested

The available information levels, as well as the format of the response are identical to TRANS2\_QUERY\_PATH\_INFORMATION.

**4.2.16 TRANS2\_SET\_PATH\_INFORMATION: Set File Attributes given Path**

This request is used to set information about a specific file or subdirectory.

Client Request =====	Value =====
WordCount MaxSetupCount SetupCount Setup[0]	15 0 1 TRANS2_SET_PATH_INFORMATION
Parameter Block Encoding =====	Description =====
USHORT InformationLevel; ULONG Reserved; STRING FileName;	Level of information to set Must be zero File or directory name

The following *Information Levels* may be set:

Information Level =====	Value =====
SMB_INFO_STANDARD	1
SMB_INFO_QUERY_EA_SIZE	2
SMB_INFO_QUERY_ALL_EAS	4

The response formats are:

#### 4.2.16.1 SMB\_INFO\_STANDARD & SMB\_INFO\_QUERY\_EA\_SIZE

Parameter Block Encoding =====	Description =====
USHORT Reserved	0
Data Block Encoding =====	Description =====
SMB_DATE CreationDate; SMB_TIME CreationTime; SMB_DATE LastAccessDate; SMB_TIME LastAccessTime; SMB_DATE LastWriteDate; SMB_TIME LastWriteTime; ULONG DataSize; ULONG AllocationSize; USHORT Attributes; ULONG EaSize;	Date when file was created Time when file was created Date of last file access Time of last file access Date of last write to the file Time of last write to the file File Size Size of filesystem allocation unit File Attributes Size of file's EA information (SMB_INFO_QUERY_EA_SIZE)

#### 4.2.16.2 SMB\_INFO\_QUERY\_ALL\_EAS

Response Field =====	Value =====
MaxDataCount	Length of FEAList found (minimum value is 4)
Parameter Block Encoding =====	Description =====
USHORT EaErrorOffset	Offset into EAList of EA error
Data Block Encoding =====	Description =====
ULONG ListLength; UCHAR EaList[]	Length of the remaining data The extended attributes list

#### 4.2.17 TRANS2\_SET\_FILE\_INFORMATION: Set File Attributes Given FID

This request is used to set information about a specific file or subdirectory given a handle to the file or subdirectory.

Client Request =====	Value =====
WordCount MaxSetupCount SetupCount Setup[0]	15 0 1 TRANS2_SET_FILE_INFORMATION
Parameter Block Encoding =====	Description =====
USHORT Fid; USHORT InformationLevel; USHORT Reserved;	Handle of file for request Level of information requested Ignored by the server

The following *InformationLevels* may be set:

Information Level =====	Value =====
SMB_INFO_STANDARD	1
SMB_INFO_QUERY_EA_SIZE	2
SMB_SET_FILE_BASIC_INFO	0x101
SMB_SET_FILE_DISPOSITION_INFO	0x102
SMB_SET_FILE_ALLOCATION_INFO	0x103
SMB_SET_FILE_END_OF_FILE_INFO	0x104

The two levels below 0x101 are as described in the NT\_SET\_PATH\_INFORMATION transaction. The requested information is placed in the Data portion of the transaction response. For the information levels greater than 0x100, the transaction response has 1 parameter word which should be ignored by the client.

#### 4.2.17.1 SMB\_FILE\_DISPOSITION\_INFO

Response Field =====	Value =====
BOOLEAN FileIsDeleted	A boolean which is TRUE if the file is marked for deletion

#### 4.2.17.2 SMB\_FILE\_ALLOCATION\_INFO

Response Field =====	Value =====
LARGE_INTEGER	File Allocation size in number of bytes

#### 4.2.17.3 SMB\_FILE\_END\_OF\_FILE\_INFO

Response Field =====	Value =====
LARGE_INTEGER	The total number of bytes that need to be traversed from the beginning of the file in order to locate the end of the file

### 4.3 Directory Requests

#### 4.3.1 TRANS2\_CREATE\_DIRECTORY: Create Directory (with optional EAs)

This requests the server to create a directory relative to *Tid* in the SMB header, optionally assigning extended attributes to it.

Client Request =====	Value =====
WordCount	15
MaxSetupCount	0
SetupCount	1
Setup[0]	TRANS2_CREATE_DIRECTORY
Parameter Block Encoding =====	Description =====
ULONG Reserved;	Reserved--must be zero
STRING Name[];	Directory name to create
UCHAR Data[];	Optional FEAList for the new directory



Response Parameter Block =====	Description =====
USHORT EaErrorOffset	Offset into FEAList of first error which occurred while setting EAs

### 4.3.2 DELETE\_DIRECTORY: Delete Directory

The delete directory message is sent to delete an empty directory. The appropriate *Tid* and additional pathname are passed. The directory must be empty for it to be deleted.

Client Request =====	Description =====
UCHAR WordCount; USHORT ByteCount; UCHAR BufferFormat; STRING DirectoryName[];	Count of parameter words = 0 Count of data bytes; min = 2 0x04 Directory name

The directory to be deleted cannot be the root of the share specified by *Tid*.

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

### 4.3.3 CHECK\_DIRECTORY: Check Directory

This SMB is used to verify that a path exists and is a directory. No error is returned if the given path exists and the client has read access to it. Client machines which maintain a concept of a "working directory" will find this useful to verify the validity of a "change working directory" command. Note that the servers do NOT have a concept of working directory for a particular client. The client must always supply full pathnames relative to the *Tid* in the SMB header.

Client Request =====	Description =====
UCHAR WordCount; USHORT ByteCount; UCHAR BufferFormat; STRING DirectoryPath[];	Count of parameter words = 0 Count of data bytes; min = 2 0x04 Directory path

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

DOS clients, in particular, depend on the SMB\_ERR\_BAD\_PATH return code if the directory is not found.

#### 4.3.3.1 Errors

ERRDOS/ERRbadfile  
ERRDOS/ERRbadpath

ERRDOS/ERRnoaccess  
 ERRHRD/ERRdata  
 ERRSRV/ERRinvid  
 ERRSRV/ERRbaduid  
 ERRSRV/ERRaccess

#### 4.3.4 TRANS2\_FIND\_FIRST2: Search Directory using Wildcards

Client Request =====	Value =====
WordCount TotalDataCount SetupCount Setup[0]	15 Total size of extended attribute list 1 TRANS2_FIND_FIRST2
Parameter Block Encoding =====	Description =====
USHORT SearchAttributes; USHORT SearchCount; USHORT Flags;  USHORT InformationLevel; ULONG SearchStorageType; STRING FileName; UCHAR Data[ TotalDataCount ]	Maximum number of entries to return Additional information: Bit 0 - close search after this request Bit 1 - close search if end of search reached Bit 2 - return resume keys for each entry found Bit 3 - continue search from previous ending place Bit 4 - find with backup intent See below  Pattern for the search FEAList if InformationLevel is QUERY_EAS_FROM_LIST

Response Parameter Block =====	Description =====
USHORT Sid; USHORT SearchCount; USHORT EndOfSearch; USHORT EaErrorOffset; USHORT LastNameOffset;  UCHAR Data[ TotalDataCount ]	Search handle Number of entries returned Was last entry returned? Offset into EA list if EA error Offset into data to file name of last entry, if server needs it to resume search; else 0 Level dependent info about the matches found in the search

This request allows the client to search for the file(s) which match the file specification. The search can be continued if necessary with TRANS2\_FIND\_NEXT2. There are numerous levels of information which may be obtained for the returned files, the desired level is specified in the *InformationLevel* field of the request.

InformationLevel Name =====	Value =====
SMB_INFO_STANDARD	1
SMB_INFO_QUERY_EA_SIZE	2
SMB_INFO_QUERY_EAS_FROM_LIST	3
SMB_FIND_FILE_DIRECTORY_INFO	0x101
SMB_FIND_FILE_FULL_DIRECTORY_INFO	0x102
SMB_FIND_FILE_NAMES_INFO	0x103
SMB_FIND_FILE_BOTH_DIRECTORY_INFO	0x104

The following sections detail the data returned for each InformationLevel. The requested information is placed in the *Data* portion of the transaction response. Note: a client which does not support long names can only request SMB\_INFO\_STANDARD.

A four-byte resume key precedes each data item (described below) if bit 2 in the Flags field is set, i.e. if the request indicates the server should return resume keys.

#### 4.3.4.1 SMB\_INFO\_STANDARD

Response Field =====	Description =====
SMB_DATE CreationDate; SMB_TIME CreationTime; SMB_DATE LastAccessDate; SMB_TIME LastAccessTime; SMB_DATE LastWriteDate; SMB_TIME LastWriteTime; ULONG DataSize; ULONG AllocationSize; USHORT Attributes; UCHAR FileNameLength; STRING FileName;	Date when file was created Time when file was created Date of last file access Time of last file access Date of last write to the file Time of last write to the file File Size Size of filesystem allocation unit File Attributes Length of filename in bytes Name of found file

#### 4.3.4.2 SMB\_INFO\_QUERY\_EA\_SIZE

Response Field =====	Description =====
SMB_DATE CreationDate; SMB_TIME CreationTime; SMB_DATE LastAccessDate; SMB_TIME LastAccessTime; SMB_DATE LastWriteDate; SMB_TIME LastWriteTime; ULONG DataSize; ULONG AllocationSize; USHORT Attributes; ULONG EaSize; UCHAR FileNameLength; STRING FileName;	Date when file was created Time when file was created Date of last file access Time of last file access Date of last write to the file Time of last write to the file File Size Size of filesystem allocation unit File Attributes Size of file's EA information Length of filename in bytes Name of found file

#### 4.3.4.3 SMB\_INFO\_QUERY\_EAS\_FROM\_LIST

This request returns the same information as `SMB_INFO_QUERY_EA_SIZE`, but only for files which have an EA list which match the EA information in the *Data* part of the request.

**4.3.4.4 SMB\_FIND\_FILE\_DIRECTORY\_INFO**

Response Field =====	Description =====
ULONG NextEntryOffset; ULONG FileIndex; LARGE_INTEGER CreationTime; LARGE_INTEGER LastAccessTime; LARGE_INTEGER LastWriteTime; LARGE_INTEGER ChangeTime; LARGE_INTEGER EndOfFile; LARGE_INTEGER AllocationSize; ULONG ExtFileAttributes; ULONG FileNameLength; STRING FileName;	Offset from this structure to beginning of next one  file creation time last access time last write time last attribute change time file size size of filesystem allocation information Extended file attributes (see section 3.12) Length of filename in bytes Name of the file

**4.3.4.5 SMB\_FIND\_FILE\_FULL\_DIRECTORY\_INFO**

Response Field =====	Description =====
ULONG NextEntryOffset; ULONG FileIndex; LARGE_INTEGER CreationTime; LARGE_INTEGER LastAccessTime; LARGE_INTEGER LastWriteTime; LARGE_INTEGER ChangeTime; LARGE_INTEGER EndOfFile; LARGE_INTEGER AllocationSize; ULONG ExtFileAttributes; ULONG FileNameLength; ULONG EaSize; STRING FileName;	Offset from this structure to beginning of next one  file creation time last access time last write time last attribute change time file size size of filesystem allocation information Extended file attributes (see section 3.12) Length of filename in bytes Size of file's extended attributes Name of the file

**4.3.4.6 SMB\_FIND\_FILE\_BOTH\_DIRECTORY\_INFO**

Response Field =====	Description =====
ULONG NextEntryOffset; ULONG FileIndex; LARGE_INTEGER CreationTime; LARGE_INTEGER LastAccessTime; LARGE_INTEGER LastWriteTime; LARGE_INTEGER ChangeTime; LARGE_INTEGER EndOfFile; LARGE_INTEGER AllocationSize; ULONG ExtFileAttributes; ULONG FileNameLength; ULONG EaSize; UCHAR ShortNameLength; UCHAR Reserved WCHAR ShortName[12]; STRING FileName;	Offset from this structure to beginning of next one  file creation time last access time last write time last attribute change time file size size of filesystem allocation information Extended file attributes (see section 3.12) Length of FileName in bytes Size of file's extended attributes Length of file's short name in bytes  File's 8.3 conformant name in Unicode Files full length name

**4.3.4.7 SMB\_FIND\_FILE\_NAMES\_INFO**

Response Field =====	Description =====
ULONG NextEntryOffset; ULONG FileIndex; ULONG FileNameLength; STRING FileName;	Offset from this structure to beginning of next one  Length of FileName in bytes Files full length name

### 4.3.5 TRANS2\_FIND\_NEXT2: Resume Directory Search Using Wildcards

This request resumes a search which was begun with a previous TRANS2\_FIND\_FIRST2 request.

Client Request =====	Value =====
WordCount SetupCount Setup[0]	15 1 TRANS2_FIND_NEXT2
Parameter Block Encoding =====	Description =====
USHORT Sid; USHORT SearchCount; USHORT InformationLevel; ULONG ResumeKey; USHORT Flags;  STRING FileName;	Search handle Maximum number of entries to return Levels described in TRANS2_FIND_FIRST2 request Value returned by previous find2 call Additional information: bit set- 0 - close search after this request 1 - close search if end of search reached 2 - return resume keys for each entry found 3 - resume/continue from previous ending place 4 - find with backup intent Resume file name

*Sid* is the value returned by a previous successful TRANS2\_FIND\_FIRST2 call. If *Bit3* of *Flags* is set, then *FileName* may be the NULL string, since the search is continued from the previous TRANS2\_FIND request. Otherwise, *FileName* must not be more than 256 characters long.

Response Field =====	Description =====
USHORT SearchCount; USHORT EndOfSearch; USHORT EaErrorOffset; USHORT LastNameOffset;  UCHAR Data[TotalDataCount]	Number of entries returned Was last entry returned? Offset into EA list if EA error Offset into data to file name of last entry, if server needs it to resume search; else 0 Level dependent info about the matches found in the search

### 4.3.6 FIND\_CLOSE2: Close Directory Search

This SMB closes a search started by the TRANS2\_FIND\_FIRST2 transaction request.

Client Request =====	Description =====
UCHAR WordCount; USHORT Sid; USHORT ByteCount;	Count of parameter words = 1 Find handle Count of data bytes = 0

Server Response =====	Description =====
UCHAR WordCount; USHORT ByteCount;	Count of parameter words = 0 Count of data bytes = 0

#### 4.3.6.1 Errors

ERRDOS/ERRbadfid  
ERRSRV/ERRinvid  
ERRSRV/ERRaccess

### 4.3.7 NT\_TRANSACT\_NOTIFY\_CHANGE: Request Change Notification

Client Setup Words =====	Description =====
ULONG CompletionFilter; USHORT Fid; BOOLEAN WatchTree; UCHAR Reserved;	Specifies operation to monitor Fid of directory to monitor TRUE = watch all subdirectories too MBZ

This command notifies the client when the directory specified by *Fid* is modified. It also returns the name(s) of the file(s) that changed. The command completes once the directory has been modified based on the supplied *CompletionFilter*. The command is a "single shot" and therefore needs to be reissued to watch for more directory changes.

A directory file must be opened before this command may be used. Once the directory is open, this command may be used to begin watching files and subdirectories in the specified directory for changes. The first time the command is issued, the *MaxParameterCount* field in the transact header determines the size of the buffer that will be used at the server to buffer directory change information between issuances of the notify change commands.

When a change that is in the *CompletionFilter* is made to the directory, the command completes. The names of the files that have changed since the last time the command was issued are returned to the client. The *ParameterCount* field of the response indicates the number of bytes that are being returned. If too many files have changed since the last time the command was issued, then zero bytes are returned and an alternate status code is returned in the *Status* field of the response.

The *CompletionFilter* is a mask created as the sum of any of the following flags:

FILE_NOTIFY_CHANGE_FILE_NAME	0x00000001
FILE_NOTIFY_CHANGE_DIR_NAME	0x00000002
FILE_NOTIFY_CHANGE_NAME	0x00000003
FILE_NOTIFY_CHANGE_ATTRIBUTES	0x00000004
FILE_NOTIFY_CHANGE_SIZE	0x00000008
FILE_NOTIFY_CHANGE_LAST_WRITE	0x00000010
FILE_NOTIFY_CHANGE_LAST_ACCESS	0x00000020
FILE_NOTIFY_CHANGE_CREATION	0x00000040
FILE_NOTIFY_CHANGE_EA	0x00000080
FILE_NOTIFY_CHANGE_SECURITY	0x00000100
FILE_NOTIFY_CHANGE_STREAM_NAME	0x00000200
FILE_NOTIFY_CHANGE_STREAM_SIZE	0x00000400
FILE_NOTIFY_CHANGE_STREAM_WRITE	0x00000800



Server Response =====	Description =====
ParameterCount Parameters[ ParameterCount ]	# of bytes of change data FILE_NOTIFY_INFORMATION structures

The response contains FILE\_NOTIFY\_INFORMATION structures, as defined below. The NextEntryOffset field of the structure specifies the offset, in bytes, from the start of the current entry to the next entry in the list. If this is the last entry in the list, this field is zero. Each entry in the list must be longword aligned, so NextEntryOffset must be a multiple of four.

```
typedef struct {
    ULONG NextEntryOffset;
    ULONG Action;
    ULONG FileNameLength;
    WCHAR FileName[1];
} FILE_NOTIFY_INFORMATION;
```

Where *Action* describes what happened to the file named *FileName*:

FILE_ACTION_ADDED	0x00000001
FILE_ACTION_REMOVED	0x00000002
FILE_ACTION_MODIFIED	0x00000003
FILE_ACTION_RENAMED_OLD_NAME	0x00000004
FILE_ACTION_RENAMED_NEW_NAME	0x00000005
FILE_ACTION_ADDED_STREAM	0x00000006
FILE_ACTION_REMOVED_STREAM	0x00000007
FILE_ACTION_MODIFIED_STREAM	0x00000008

## 4.4 DFS Operations

### 4.4.1 TRANS2\_GET\_DFS\_REFERRAL: Retrieve Distributed Filesystem Referral

The client sends this request to ask the server to convert *RequestFilename* into an alternate name for this file. This request can be sent to the server if the server response to the NEGOTIATE SMB included the CAP\_DFS capability. The TID of the request must be IPC\$. *Bit15 of Flags2* in the SMB header must be set, indicating this is a UNICODE request.

Client Request =====	Description =====
WordCount TotalDataCount  SetupCount Setup[0]	15 0  1 TRANS2_GET_DFS_REFERRAL
Parameter Block Encoding =====	Description =====
USHORT MaxReferralLevel WCHAR RequestFileName;	Latest referral version number understood DFS name of file for which referral is sought

Response Data Block =====	Description =====
USHORT PathConsumed; USHORT NumberOfReferrals; USHORT Flags;  REFERRAL_LIST Referrals[]	Number of <i>RequestFilename</i> bytes client Number of referrals contained in this response bit0 - The servers in <i>Referrals</i> are capable of fielding TRANS2_GET_DFS_REFERRAL. bit1 - The servers in <i>Referrals</i> should hold the storage for the requested file. Set of referrals for this file
UNICODESTRING Strings	Used to hold the strings pointed to by Version 2 Referrals in <i>REFERRALS</i> .

The server response is a list of *Referrals* which inform the client where it should resubmit the request to obtain access to the file. *PathConsumed* in the response indicates to the client how many characters of *RequestFilename* have been consumed by the server. When the client chooses one of the referrals to use for file access, the client may need to strip the leading *PathConsumed* characters from the front of *RequestFilename* before submitting the name to the target server. Whether or not the pathname should be trimmed is indicated by the individual referral as detailed below.

*Flags* indicates how this referral should be treated. If *bit0* is clear, any entity in the *Referrals* list holds the storage for *RequestFilename*. If *bit0* is set, any entity in the *Referrals* list has further referral information for *RequestFilename* – a TRANS2\_GET\_DFS\_REFERRAL request should be sent to an entity in the *Referrals* list for further resolution.

The format of an individual referral contains version and length information allowing the client to skip referrals it does not understand. *MaxReferralLevel* indicates to the server the latest version of referral which the client can digest. Since each referral has a uniform element, *MaxReferralLevel* is advisory only. Each element in *Referrals* has this envelope:

REFERRAL_LIST element =====	
USHORT VersionNumber	Version of this referral element
USHORT ReferralSize	Size of this referral element

The following referral element versions are defined:

Version 1 Referral Element Format =====	
USHORT ServerType	Type of <i>Node</i> handling referral: 0 - Don't know 1 - SMB Server 2 - Netware Server 3 - Domain
USHORT ReferralFlags	Flags which describe this referral: 01 - Strip off <i>PathConsumed</i> characters before submitting <i>RequestFileName</i> to <i>Node</i>
UNICODESTRING Node	Name of entity to visit next

Version 2 Referral Element Format =====	
USHORT ServerType	Type of <i>Node</i> handling referral: 0 - Don't know 1 - SMB Server 2 - Netware Server 3 - Domain
USHORT ReferralFlags	Flags which describe this referral: 01 - Strip off <i>PathConsumed</i> characters before submitting <i>RequestFileName</i> to <i>Node</i>
ULONG Proximity	A hint describing the proximity of this server to the client. 0 indicates the closest, higher numbers indicate increasingly "distant" servers. The number is only relevant within the context of the servers listed in <i>this</i> particular SMB.
ULONG TimeToLive	Number of seconds for which the client can cache this referral.
USHORT DfsPathOffset	Offset, in bytes from the beginning of this referral, of the DFS Path that matched <i>PathConsumed</i> bytes of the <i>RequestFileName</i> .
USHORT DfsAlternatePathOffset	Offset, in bytes from the beginning of this referral, of an alternate name (8.3 format) of the DFS Path that matched <i>PathConsumed</i> bytes of the <i>RequestFileName</i> .
USHORT NetworkAddressOffset	Offset, in bytes from the beginning of this referral, of the entity to visit next.

The CIFS protocol imposes no referral selection policy.

#### 4.4.2 TRANS2\_REPORT\_DFS\_INCONSISTENCY: Inform a server about DFS Error

As part of the Distributed Name Resolution algorithm, a DFS client may discover a knowledge inconsistency between the referral server (i.e., the server that handed out a referral), and the storage server (i.e., the server to which the client was redirected to by the referral server). When such an inconsistency is discovered, the DFS client optionally sends this SMB to the referral server, allowing the referral server to take corrective action.

Client Request =====	Description =====
WordCount MaxParameterCount SetupCount Setup[0]	15 0 1 TRANS2_REPORT_DFS_INCONSISTENCY
Parameter Block Encoding =====	Description =====
UNICODESTRING RequestFileName;	DFS Name of file for which referral was sought

The data part of this request contains the referral element (Version 1 format only) believed to be in error. These are encoded as described in the `TRANS2_GET_DFS_REFERRAL` response. If the server returns success, the client can resubmit the `TRANS2_GET_DFS_REFERRAL` request to this server to get a new referral. It is not mandatory for the DFS knowledge to be automatically repaired – the client must be prepared to receive further errant referrals and must not wind up looping between this request and the `TRANS2_GET_DFS_REFERRAL` request.

*Bit15* of *Flags2* in the SMB header must be set, indicating this is a UNICODE request.

## 4.5 Miscellaneous Operations

### 4.5.1 NT\_TRANSACT\_IOCTL

This command allows device and file system control functions to be transferred transparently from client to server.

Setup Words Encoding =====	Description =====
ULONG FunctionCode; USHORT Fid; BOOLEAN IsFsctl;  UCHAR IsFlags;	NT device or file system control code Handle for io or fs control. Unless <i>BIT0</i> of <i>ISFLAGS</i> is set. Indicates whether the command is a device control (FALSE) or a file system control (TRUE). <i>BIT0</i> - command is to be applied to share root handle. Share must be a DFS share.
Data Block Encoding =====	Description =====
Data[ TotalDataCount ]	Passed to the Fsctl or Ioctl

Server Response =====	Description =====
SetupCount Setup[0] DataCount Data[ DataCount ]	1 Length of information returned by io or fs control Length of information returned by io or fs control The results of the io or fs control

## 4.5.2 NT\_TRANSACT\_QUERY\_SECURITY\_DESC

This command allows the client to retrieve the security descriptor on a file.

Client Parameter Block =====	Description =====
USHORT Fid; USHORT Reserved; ULONG SecurityInformation;	FID of target MBZ Fields of descriptor to set

NtQuerySecurityObject() is called, requesting *SecurityInformation*. The result of the call is returned to the client in the *Data* part of the transaction response.

## 4.5.3 NT\_TRANSACT\_SET\_SECURITY\_DESC

This command allows the client to change the security descriptor on a file.

Client Parameter Block Encoding =====	Description =====
USHORT Fid; USHORT Reserved; ULONG SecurityInformation;	FID of target MBZ Fields of SD that to set
Data Block Encoding =====	Description =====
Data[TotalDataCount]	Security Descriptor information

*Data* is passed directly to NtSetSecurityObject(), with *SecurityInformation* describing which information to set. The transaction response contains no parameters or data.

# 5 SMB Symbolic Constants

## 5.1 SMB Command Codes

The following values have been assigned for the SMB Commands.

SMB_COM_CREATE_DIRECTORY	0x00
SMB_COM_DELETE_DIRECTORY	0x01
SMB_COM_OPEN	0x02
SMB_COM_CREATE	0x03
SMB_COM_CLOSE	0x04
SMB_COM_FLUSH	0x05
SMB_COM_DELETE	0x06
SMB_COM_RENAME	0x07
SMB_COM_QUERY_INFORMATION	0x08
SMB_COM_SET_INFORMATION	0x09
SMB_COM_READ	0x0A
SMB_COM_WRITE	0x0B
SMB_COM_LOCK_BYTE_RANGE	0x0C
SMB_COM_UNLOCK_BYTE_RANGE	0x0D
SMB_COM_CREATE_TEMPORARY	0x0E
SMB_COM_CREATE_NEW	0x0F
SMB_COM_CHECK_DIRECTORY	0x10
SMB_COM_PROCESS_EXIT	0x11
SMB_COM_SEEK	0x12
SMB_COM_LOCK_AND_READ	0x13
SMB_COM_WRITE_AND_UNLOCK	0x14
SMB_COM_READ_RAW	0x1A
SMB_COM_READ_MPX	0x1B
SMB_COM_READ_MPX_SECONDARY	0x1C
SMB_COM_WRITE_RAW	0x1D
SMB_COM_WRITE_MPX	0x1E
SMB_COM_WRITE_COMPLETE	0x20
SMB_COM_SET_INFORMATION2	0x22
SMB_COM_QUERY_INFORMATION2	0x23
SMB_COM_LOCKING_ANDX	0x24
SMB_COM_TRANSACTION	0x25
SMB_COM_TRANSACTION_SECONDARY	0x26
SMB_COM_IOCTL	0x27
SMB_COM_IOCTL_SECONDARY	0x28
SMB_COM_COPY	0x29
SMB_COM_MOVE	0x2A
SMB_COM_ECHO	0x2B
SMB_COM_WRITE_AND_CLOSE	0x2C
SMB_COM_OPEN_ANDX	0x2D
SMB_COM_READ_ANDX	0x2E
SMB_COM_WRITE_ANDX	0x2F
SMB_COM_CLOSE_AND_TREE_DISC	0x31
SMB_COM_TRANSACTION2	0x32
SMB_COM_TRANSACTION2_SECONDARY	0x33
SMB_COM_FIND_CLOSE2	0x34
SMB_COM_FIND_NOTIFY_CLOSE	0x35
SMB_COM_TREE_CONNECT	0x70
SMB_COM_TREE_DISCONNECT	0x71
SMB_COM_NEGOTIATE	0x72
SMB_COM_SESSION_SETUP_ANDX	0x73
SMB_COM_LOGOFF_ANDX	0x74
SMB_COM_TREE_CONNECT_ANDX	0x75
SMB_COM_QUERY_INFORMATION_DISK	0x80
SMB_COM_SEARCH	0x81

SMB_COM_FIND	0x82
SMB_COM_FIND_UNIQUE	0x83
SMB_COM_NT_TRANSACT	0xA0
SMB_COM_NT_TRANSACT_SECONDARY	0xA1
SMB_COM_NT_CREATE_ANDX	0xA2
SMB_COM_NT_CANCEL	0xA4
SMB_COM_OPEN_PRINT_FILE	0xC0
SMB_COM_WRITE_PRINT_FILE	0xC1
SMB_COM_CLOSE_PRINT_FILE	0xC2
SMB_COM_GET_PRINT_QUEUE	0xC3
SMB_COM_READ_BULK	0xD8
SMB_COM_WRITE_BULK	0xD9
SMB_COM_WRITE_BULK_DATA	0xDA

## 5.2 SMB\_COM\_TRANSACTION2 Subcommand codes

The subcommand code for SMB\_COM\_TRANSACTION2 request is placed in Setup[0]. The parameters associated with any particular request are placed in the *Parameters* vector of the request. The defined subcommand codes are:

Setup[0] Transaction2 Subcommand Code =====	Value =====	Description =====
TRANS2_OPEN2	0x00	Create file with extended attributes
TRANS2_FIND_FIRST2	0x01	Begin search for files
TRANS2_FIND_NEXT2	0x02	Resume search for files
TRANS2_QUERY_FS_INFORMATION	0x03	Get file system information
	0x04	Reserved
TRANS2_QUERY_PATH_INFORMATION	0x05	Get information about a named file or directory
TRANS2_SET_PATH_INFORMATION	0x06	Set information about a named file or directory
TRANS2_QUERY_FILE_INFORMATION	0x07	Get information about a handle
TRANS2_SET_FILE_INFORMATION	0x08	Set information by handle
TRANS2_FSCTL	0x09	Not implemented by NT server
TRANS2_IOCTL2	0x0A	Not implemented by NT server
TRANS2_FIND_NOTIFY_FIRST	0x0B	Not implemented by NT server
TRANS2_FIND_NOTIFY_NEXT	0x0C	Not implemented by NT server
TRANS2_CREATE_DIRECTORY	0x0D	Create directory with extended attributes
TRANS2_SESSION_SETUP	0x0E	Session setup with extended security information
TRANS2_GET_DFS_REFERRAL	0x10	Get a DFS referral
TRANS2_REPORT_DFS_INCONSISTENCY	0x11	Report a DFS knowledge inconsistency

## 5.3 SMB\_COM\_NT\_TRANSACTION Subcommand Codes

For these transactions, *Function* in the primary client request indicates the operation to be performed. It may assume one of the following values:

SubCommand Code =====	Value =====	Description =====
NT_TRANSACT_CREATE	1	File open/create
NT_TRANSACT_IOCTL	2	Device IOCTL
NT_TRANSACT_SET_SECURITY_DESC	3	Set security descriptor
NT_TRANSACT_NOTIFY_CHANGE	4	Start directory watch
NT_TRANSACT_RENAME	5	Reserved (Handle-based rename)
NT_TRANSACT_QUERY_SECURITY_DESC	6	Retrieve security descriptor info

## 5.4 SMB Protocol Dialect Constants

This is the list of CIFS protocol dialects, ordered from least functional (earliest) version to most functional (most recent) version:

Dialect Name =====	Comment =====
PC NETWORK PROGRAM 1.0	The original MSNET SMB protocol (otherwise known as the "core protocol")
PCLAN1.0	Some versions of the original MSNET defined this as an alternate to the core protocol name
MICROSOFT NETWORKS 1.03	This is used for the MS-NET 1.03 product. It defines Lock&Read, Write&Unlock, and a special version of raw read and raw write.
MICROSOFT NETWORKS 3.0	This is the DOS LANMAN 1.0 specific protocol. It is equivalent to the LANMAN 1.0 protocol, except the server is required to map errors from the OS/2 error to an appropriate DOS error.
LANMAN1.0	This is the first version of the full LANMAN 1.0 protocol
LM1.2X002	This is the first version of the full LANMAN 2.0 protocol
DOS LM1.2X002	This is the DOS equivalent of the LM1.2X002 protocol. It is identical to the LM1.2X002 protocol, but the server will perform error mapping to appropriate DOS errors.
DOS LANMAN2.1	DOS LANMAN2.1
LANMAN2.1	OS/2 LANMAN2.1
Windows for Workgroups 3.1a	Windows for Workgroups Version 1.0
NT LM 0.12	The SMB protocol designed for NT networking. This has special SMBs which duplicate the NT semantics.

CIFS servers select the most recent version of the protocol known to both client and server. Any CIFS server which supports dialects newer than the original core dialect must support all the messages and semantics of the dialects between the core dialect and the newer one. This is to say that a server which supports the NT LM 0.12 dialect must also support all of the messages of the previous 10 dialects. It is the client's responsibility to ensure it only sends SMBs which are appropriate to the dialect negotiated. Clients must be prepared to receive an SMB response from an earlier protocol dialect -- even if the client used the most recent form of the request.

## 6 Error Codes and Classes

This section lists all of the valid values for *Status.DosError.ErrorClass*, and most of the error codes for *Status.DosError.Error*.

The following error classes may be returned by the server to the client.



Class =====	Code =====	Comment =====
SUCCESS	0	The request was successful.
ERRDOS	0x01	Error is from the core DOS operating system set.
ERRSRV	0x02	Error is generated by the server network file manager.
ERRHRD	0x03	Error is an hardware error.
ERRCMD	0xFF	Command was not in the "SMB" format.

The following error codes may be generated with the SUCCESS error class.

Class =====	Code =====	Comment =====
SUCCESS	0	The request was successful.

The following error codes may be generated with the ERRDOS error class.

Error =====	Code =====	Description =====
ERRbadfunc	1	Invalid function. The server did not recognize or could not perform a system call generated by the server, e.g., set the DIRECTORY attribute on a data file, invalid seek mode.
ERRbadfile	2	File not found. The last component of a file's pathname could not be found.
ERRbadpath	3	Directory invalid. A directory component in a pathname could not be found.
ERRnofids	4	Too many open files. The server has no file handles available.
ERRnoaccess	5	Access denied, the client's context does not permit the requested function. This includes the following conditions: invalid rename command write to Fid open for read only read on Fid open for write only attempt to delete a non-empty directory
ERRbadfid	6	Invalid file handle. The file handle specified was not recognized by the server.
ERRbadmcb	7	Memory control blocks destroyed.
ERRnomem	8	Insufficient server memory to perform the requested function.
ERRbadmem	9	Invalid memory block address.
ERRbadenv	10	Invalid environment.
ERRbadformat	11	Invalid format.
ERRbadaccess	12	Invalid open mode.
ERRbaddata	13	Invalid data (generated only by IOCTL calls within the server).
ERRbaddrive	15	Invalid drive specified.
ERRremcd	16	A Delete Directory request attempted to remove the server's current directory.
ERRdiffdevice	17	Not same device (e.g., a cross volume rename was attempted)
ERRnofiles	18	A File Search command can find no more files matching the specified criteria.
ERRbadshare	32	The sharing mode specified for an Open conflicts with existing FIDs on the file.
ERRlock	33	A Lock request conflicted with an existing lock or specified an invalid mode, or an Unlock requested attempted to remove a lock held by another process.
ERRfilexists	80	The file named in the request already exists.

The following error codes may be generated with the ERRSRV error class.

Error =====	Code =====	Description =====
ERRerror	1	Non-specific error code. It is returned under the following conditions: <ul style="list-style-type: none"> <li>• resource other than disk space exhausted (e.g. TIDs)</li> <li>• first SMB command was not negotiate</li> <li>• multiple negotiates attempted</li> <li>• internal server error</li> </ul>
ERRbadpw	2	Bad password - name/password pair in a Tree Connect or Session Setup are invalid.
ERRaccess	4	The client does not have the necessary access rights within the specified context for the requested function.
ERRinvnid	5	The Tid specified in a command was invalid.
ERRinvnetname	6	Invalid network name in tree connect.
ERRinvdevice	7	Invalid device - printer request made to non-printer connection or non-printer request made to printer connection.
ERRqfull	49	Print queue full (files) -- returned by open print file.
ERRqtoobig	50	Print queue full -- no space.
ERRqeof	51	EOF on print queue dump.
ERRinvpfid	52	Invalid print file FID.
ERRsmbcmd	64	The server did not recognize the command received.
ERRsrverror	65	The server encountered an internal error, e.g., system file unavailable.
ERRfilespecs	67	The Fid and pathname parameters contained an invalid combination of values.
ERRbadpermits	69	The access permissions specified for a file or directory are not a valid combination. The server cannot set the requested attribute.
ERRsetattrmode	71	The attribute mode in the Set File Attribute request is invalid.
ERRpaused	81	Server is paused. (reserved for messaging)
ERRmsgoff	82	Not receiving messages. (reserved for messaging).
ERRnoroom	83	No room to buffer message. (reserved for messaging).
ERRrmuns	87	Too many remote user names. (reserved for messaging).
ERRtimeout	88	Operation timed out.
ERRnoresource	89	No resources currently available for request.
ERRtoomanyuids	90	Too many Uids active on this session.
ERRbaduid	91	The Uid is not known as a valid user identifier on this session.
ERRusempx	250	Temporarily unable to support Raw, use MPX mode.
ERRusestd	251	Temporarily unable to support Raw, use standard read/write.
ERRcontmpx	252	Continue in MPX mode.
ERRnosupport	65535	Function not supported.

The following error codes may be generated with the ERRHRD error class.

Error =====	Code =====	Description =====
ERRnowrite	19	Attempt to write on write-protected media
ERRbadunit	20	Unknown unit.
ERRnotready	21	Drive not ready.
ERRbadcmd	22	Unknown command.
ERRdata	23	Data error (CRC).
ERRbadreq	24	Bad request structure length.
ERRseek	25	Seek error.
ERRbadmedia	26	Unknown media type.
ERRbadsector	27	Sector not found.
ERRnopaper	28	Printer out of paper.
ERRwrite	29	Write fault.
ERRread	30	Read fault.
ERRgeneral	31	General failure.
ERRbadshare	32	A open conflicts with an existing open.
ERRlock	33	A Lock request conflicted with an existing lock or specified an invalid mode, or an Unlock requested attempted to remove a lock held by another process.
ERRwrongdisk	34	The wrong disk was found in a drive.
ERRFCBUnavail	35	No FCBs are available to process request.
ERRsharebufexc	36	A sharing buffer has been exceeded.

## 7 Legal Notice

Microsoft does not know of any third-party rights that are violated by this contribution. Microsoft makes no other representations regarding this contribution.

## 8 References

- [1] P. Mockapetris, "Domain Names - Concepts And Facilities", RFC 1034, November 1987
- [2] P. Mockapetris, "Domain Names - Implementation And Specification", RFC 1035, November 1987
- [3] Karl Auerbach, "Protocol Standard For A Netbios Service On A Tcp/Udp Transport: Concepts And Methods", RFC 1001, March 1987
- [4] Karl Auerbach, "Protocol Standard For A Netbios Service On A Tcp/Udp Transport: Detailed Specifications", RFC 1002, March 1987
- [5] US National Bureau of Standards, "Data Encryption Standard", Federal Information Processing Standard (FIPS) Publication 46-1, January 1988
- [6] Rivest, R. - MIT and RSA Data Security, Inc., "The MD4 Message Digest Algorithm", RFC 1320, April 1992

[7] X/Open Company Ltd., "X/Open CAE Specification - Protocols for X/Open PC Interworking: SMB, Version 2", X/Open Document Number: CAE 209, September 1992.

## 9 Authors' Addresses

Paul Leach  
Dilip Naik  
Microsoft  
1 Microsoft Way  
Redmond, WA 98052  
paulle@microsoft.com  
dilipn@microsoft.com

## 10 Appendix A -- NETBIOS transport over TCP

When operating CIFS over the NETBIOS transport over TCP, connections are established and messages transferred as specified in RFC 1001 and RFC 1002.

Message transport is done using NETBIOS session service (see section 5.3 of RFC 1001 and section 4.3 of RFC 1002).

### 10.1 Connection Establishment

After the server name has been resolved to an IP address, then a connection to the server needs to be established if one has not already been set up. Connection establishment is done using the NETBIOS session service, which requires the client to provide a "calling name" and a "called name". The calling name is not significant in CIFS, except that an identical name from the same transport address is assumed to represent the same client; the called name is always "\*SMBSERVER". Connection establishment results in a "Session Request" packet to port 139 (see section 4.3.2 of RFC 1002).

#### 10.1.1 Backwards compatability

If a CIFS client wishes to inter-operate with older SMB servers, then if the server rejects the session request, it can retry with a new called name. The choice of the new called name depends on the name resolution mechanism used. If DNS was used, the called name should be constructed from the first component of the server's DNS name, truncated to 15 characters if necessary, and then padded to 16 characters with blank (20 hex) characters. If NETBIOS was used, then the called name is just the NETBIOS name. If these fail, then a NETBIOS "Adapter Status" request may be made to obtain the server's NETBIOS name, and the connection establishment retried with that as the called name.

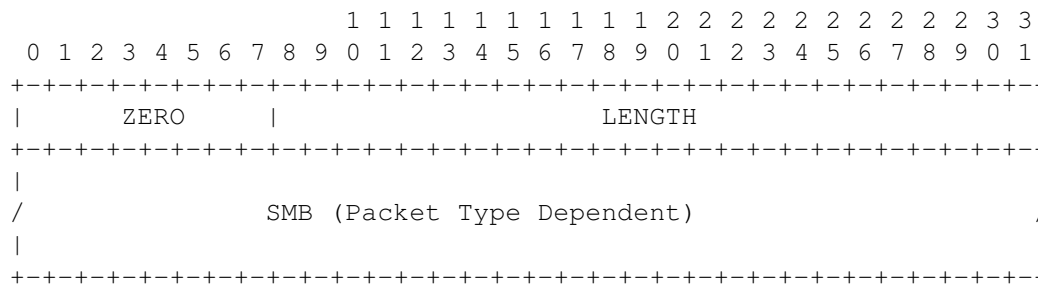
### 10.2 Server-side Connection Procedures

A CIFS server running over NETBIOS MUST accept any session request specifying a called name of "\*SMBSERVER".

In addition, if it wishes to support older SMB clients, it MAY have one or more NETBIOS names and accept session request specifying them as the called name.

## 11 Appendix B -- TCP transport

When operating CIFS over TCP, connections are established to TCP port TBD, and each message is framed as follows:



Each CIFS request starts with a 4 byte field encoded as above: a byte of zero, followed by three bytes of length; after that follows the body of the request.

## 12 Appendix C -- Share Level Server Security

Each server makes a set of resources available to clients on the network. A resource being shared may be a directory tree, named pipe, printer, etc. So far as clients are concerned, the server has no storage or service dependencies on any other servers; a client considers the server to be the sole provider of the file (or other resource) being accessed.

The CIFS protocol requires server authentication of users before file accesses are allowed, and each server authenticates its own users. A client system must send authentication information to the server before the server will allow access to its resources.

The CIFS protocol used to define two methods that can be selected by the server for security: *share level* and *user level*. User level security is the only non-obsolescent method, and is what was described in section 2.8.

A *share level* server makes some directory on a disk device (or other resource) available. An optional password may be required to gain access. Thus any user on the network who knows the name of the server, the name of the resource and the password has access to the resource. Share level security servers may use different passwords for the same shared resource with different passwords allowing different levels of access.

Share level only clients do not send SESSION\_SETUP\_ANDX requests. Instead, they send TREE\_CONNECT\_ANDX requests that include a password, or which use challenge/response authentication to prove that they know a password.

When a *user level* server validates the account name and password presented by the client, an identifier representing that authenticated instance of the user is returned to the client in the *Uid* field of the response SMB. In contrast, a *share level* server returns no useful information in the *Uid* field.

If the server is executing in share level security mode, *Tid* is the only thing used to allow access to the shared resource. Thus if the user is able to perform a successful connection to the server specifying the appropriate netname and passwd (if any) the resource may be accessed according to the access rights associated with the shared resource (same for all who gained access this way).

The user level security model was added after the original dialect of the CIFS protocol was issued, and subsequently some clients may not be capable of sending account name and passwords to the server. A server in user level security mode communicating with one of these clients MAY allow a client to connect to resources even if the client has not sent account name information:

1. If the client's computer name is identical to an account-name known on the server, and if the password supplied or authenticated via challenge/response to connect to the shared resource matches that account's password, an implicit "user logon" will be performed using those values. If the above fails, the server may fail the request or assign a default account name of its choice.

2. The value of *Uid* in subsequent requests by the client will be ignored and all access will be validated assuming the account name selected above.