

SMB File Sharing URI Scheme

Intellectual Property Rights Disclosure Acknowledgment

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Status of this Memo

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 obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Discussions regarding this document and the SMB URI scheme should take place on the uri@w3.org mailing list.

This Internet-Draft will expire on January 8th, 2007.

Abstract

The Server Message Block (SMB) protocol is one of the most widely used network file system protocols in existence. This document describes a scheme for an SMB Uniform Resource Indicator (SMB URI) which can be used to identify SMB workgroups, servers, server shares, directories, files, inter-process communications ports (named pipes), and devices--the objects in the SMB network file system space.

Hertel

Expires January 8, 2007

[Page 1]

Table of Contents

| | | |
|-----------------------------|--|--------------------|
| 1. | Introduction. | 3 |
| 1.1. | Purpose. | 3 |
| 2. | SMB URI Scheme Overview | 4 |
| 3. | NetBIOS, NBT, and NBT Transport Semantics | 4 |
| 3.1. | The NetBIOS Name Service | 5 |
| 3.2. | The NetBIOS Datagram Service | 5 |
| 3.3. | The NetBIOS Session Service. | 5 |
| 3.4. | Accommodating NBT in the SMB URI | 5 |
| 4. | NetBIOS-Based Workgroups. | 6 |
| 5. | SMB URI Definition. | 9 |
| 6. | SMB URI Syntax Elements | 10 |
| 6.1. | scheme | 10 |
| 6.2. | smb-service. | 10 |
| 6.3. | auth-domain. | 11 |
| 6.4. | smb-srv-name | 11 |
| 6.5. | port | 12 |
| 6.6. | scope-id | 12 |
| 6.7. | nbt-context. | 13 |
| 7. | The Relationship Between the SMB URI and the UNC Format | 15 |
| 8. | Authentication. | 16 |
| 9. | Security Considerations | 16 |
| 10. | Character Encoding Issues. | 16 |
| 11. | Acknowledgments. | 17 |
| 12. | References | 18 |
| 13. | Author's Address | 18 |
| 14. | Copyright Notice | 19 |
| 15. | Disclaimer of Validity | 19 |
| Appendix A. | Working with NetBIOS Names (Implementation Notes). | 20 |
| A.1. | NetBIOS Names. | 20 |
| A.2. | SMB Sessions via NBT | 20 |
| A.3. | Resolving DNS names and IP addresses to SMB server names | 21 |
| A.4. | Determining the Namespace of the smb-srv-name. | 22 |
| A.5. | Workgroup vs. SMB Server Names | 24 |

Hertel

Expires January 8, 2007

[Page 2]

1. Introduction

The Server Message Block protocol (SMB) was created in the 1980's by Dr. Barry Feigenbaum at IBM Corporation. It was later extended by various contributors at 3Com, IBM, Intel, and Microsoft. During the mid 1990's SMB was renamed CIFS (Common Internet File System). Both names are in general use today. In common parlance, however, the term "SMB" is typically used when referring to the core filesharing protocol itself, while "CIFS" includes the set of sub-protocols and extensions used to create the complete filesharing suite.

SMB was originally carried via a proprietary network transport, the interface to which was called NetBIOS (Network Basic Input Output System). Two Internet RFCs ([\[RFC1001\]](#), [\[RFC1002\]](#)) were published which describe a mechanism for implementing the NetBIOS API on top of TCP and UDP, thus allowing SMB to be transported over IP inter-networks. Those RFCs are now known collectively as Internet Standard #19, and the transport protocol that they describe is commonly called NBT (or, sometimes, NetBT) for "NetBIOS over TCP/IP."

In addition to transport via NBT, newer implementations of SMB typically support SMB over TCP/IP without the intervening NetBIOS emulation layer. This is known as SMB over "native TCP" or "naked transport".

Several attempts have been made to document and even standardize the CIFS suite ([\[XOPENSMB\]](#), [\[ONET\]](#), [\[SNIACIFS\]](#), [\[IMPCIFS\]](#)), yet the further development of CIFS remains under the control of Microsoft. Despite its proprietary nature, the workings of SMB are sufficiently well known that SMB file sharing has been successfully implemented by several third-party commercial vendors and in Open Source. SMB server and client software is available for a wide variety of operating system platforms. The very large number of systems which support this form of file sharing make an SMB URI scheme both practical and desirable.

1.1. Purpose

This document does not attempt to detail the proper implementation of SMB/CIFS or the workings of the NBT transport layer. The goal is to present the syntax of the SMB URI and to describe how that syntax is mapped to the semantics of the CIFS suite. A limited set of implementation guidelines is provided in the appendices.

Hertel

Expires January 8, 2007

[Page 3]

2. SMB URI Scheme Overview

An SMB URI is identified by one of two scheme names: "smb" or "cifs". There are existing implementations that support one or the other or both of these names, so new and updated implementations must, unfortunately, provide support for both. As of this writing, the "smb" name appears to be the more popular of the two and is, therefore, the preferred form to use when transcribing SMB URIs.

The SMB URI conforms to the general syntax of URI described in [\[RFC3986\]](#). The scheme is comparable to other URI schemes that are in common use, and should be familiar to anyone who uses common user agents such as web browsers. For example, consider the following URIs:

```
ftp://server/dirpath/file.name  
smb://server/dirpath/file.name
```

Both strings identify the same file, differing only in the protocol used to access the file.

As with other URI schemes, servers (hosts) may be identified by DNS name, IPv4 address, or IPv6 address. The SMB URI scheme also supports the use of NetBIOS names to identify SMB services. NetBIOS names are resolved via the NBT Name Service, as described in [\[RFC1001\]](#) and detailed in [\[RFC1002\]](#).

3. NetBIOS, NBT, and NBT Transport Semantics

NBT was created so that applications and services that made use of the NetBIOS API could communicate over IP internetworks. The NBT transport maps the semantics of the NetBIOS API onto TCP and UDP, hiding the fact that the underlying transport mechanism has changed.

In order to make this work, [RFC1001/1002](#) define three key services:

- The NetBIOS Name Service
- The NetBIOS Datagram Service
- The NetBIOS Session Service

These three services collaborate to create a "virtual NetBIOS LAN" on top of IP networks.

The use of NBT as a transport for SMB is widespread, so the syntax and semantics of the SMB URI scheme have been adapted to accommodate NetBIOS naming conventions and NBT context requirements.

Hertel

Expires January 8, 2007

[Page 4]

3.1. The NetBIOS Name Service

The NetBIOS API uses names (strings of 16 octets) to identify communications endpoints. These names are the addresses used in "NetBIOS LANs". The NetBIOS Name Service is responsible for transparently mapping these names to the IP addresses at which they are registered, so that NetBIOS applications and services can find one another.

The set of nodes that are connected to the same virtual NetBIOS LAN is known as the "scope" of the LAN. NBT scopes are always given a name, or Scope Identifier (Scope ID), though this normally goes unnoticed because the default Scope ID is the empty string (""). Multiple scopes, each having a different Scope ID, may intersect across the same IP internetwork without conflict.

The SMB URI scheme provides a mechanism for specifying the NBT Scope ID in the URI string.

3.2. The NetBIOS Datagram Service

The SMB URI scheme does not make direct use of the Datagram Service, so there is no special syntax to support the handling of NetBIOS datagram messages.

3.3. The NetBIOS Session Service

The Session Service provides a one-to-one mapping of NetBIOS Sessions to TCP sessions. NetBIOS sessions over TCP are initiated by sending an NBT SESSION REQUEST message. An NBT POSITIVE SESSION RESPONSE indicates that the session has been established successfully. (See [\[RFC1001\]](#), [section 16.2](#).)

The SMB URI scheme supports the exchange of SMB messages over either NBT sessions or "naked" TCP sessions. Other than the SESSION REQUEST/POSITIVE SESSION RESPONSE, there is no difference between the two transport types.

3.4 Accommodating NBT in the SMB URI

The SMB URI scheme supports both NBT and native TCP transport of SMB messages. The syntax of the scheme provides for the inclusion of NBT context information so that NetBIOS names can be properly resolved and NetBIOS sessions established.

Hertel

Expires January 8, 2007

[Page 5]

The scheme provides the ability to specify the following NBT context information:

- Broadcast Address for NBT broadcast name resolution (B mode)
- NBNS Address for point-to-point name resolution (P mode)
- Name resolution mode selection (B, P, M, or H mode)
- NBT Scope ID
- NetBIOS CALLED name (destination address)
- NetBIOS CALLING name (source address)

Most SMB implementations use configuration files or DHCP to establish an initial NBT context. The starting context is referred to as the "base context" in the remainder of this document. NBT context information given in absolute URIs is applied against the base context to create the "current context". NBT context information given in relative URIs is applied against the current context to modify the current context.

The current NBT context is always used to interpret the meaning of a given URI string. A relative URI containing updated NBT context information will cause the resulting URI to be re-evaluated.

4. NetBIOS-Based Workgroups

The "workgroup" system is a part of the SMB/CIFS protocol suite. Workgroups are built on top of NetBIOS, and are identified by their NetBIOS names. The workgroup system allows SMB file servers to be organized into named groups, with the goal of making it easier to locate resources by categorizing them.

Within each workgroup, a list of member servers is maintained. In addition to the server list, each workgroup maintains a list of all other known workgroups. The combined list is known as the "browse list". A copy of the browse list may be obtained by sending a specific query via SMB.

The SMB URI scheme views the SMB file sharing environment hierarchically. Conceptually, the hierarchy is arranged as follows:

- List of workgroups (from the browse list)
 - + List of servers within a given workgroup (from the browse list)
 - + List of shares (shared objects) offered by a server
 - + Directories, files, etc. within a share

Hertel

Expires January 8, 2007

[Page 6]

That hierarchy is mapped to the SMB URI scheme as follows:

| URI Format | Indicates |
|---------------------------|------------------------------------|
| ===== | ===== |
| smb:// | List of known workgroups |
| smb://smb-wrkgrp/ | + SMB servers within the workgroup |
| smb://smb-server/ | + Shares offered by an SMB server |
| smb://smb-server/abs-path | + Directories, files, etc. |

As shown above, the SMB URI provides syntax that indicates requests for subsets of the browse list. In particular, the form:

```
smb://
```

represents a request for the list of all known workgroups, while the form:

```
smb://smb-wrkgrp/
```

represents a request for the list of servers that are members of a particular workgroup.

A problem arises, however, because the syntax used for requesting the list of servers in a workgroup is indistinguishable from that of a request for the list of shares (shared objects, such as exported filesystems) offered by an SMB file server. Thus, the two requests must be differentiated semantically. Consider the following example:

```
smb://corgi/
```

If the name "corgi" is a NetBIOS name and it resolves (via the NBT Name Service) to a workgroup name, then a user agent would return a list of servers in the CORGI workgroup. If, however, the name resolves to a server name, then the user agent would return a list of shares offered by the SMB server named CORGI.

It is rare, but possible (in a misconfigured NBT network), that a NetBIOS name will represent both a workgroup and an SMB file server. In this situation, SMB file services take precedence. Some user agents may be capable of returning both the list of servers in the workgroup and the list of shares provided by the SMB file server, allowing the user to determine which is the desired result. This is the recommended approach.

There is a special case to be considered when using relative references to move between a workgroup reference and a reference to a server in the workgroup. Consider a workgroup named "corgis" and a server named "cue" that is a member of that workgroup.

Hertel

Expires January 8, 2007

[Page 7]

Presented with the URI string

```
smb://corgis/
```

a user agent may return a list of servers that are members of the CORGIS workgroup, including node CUE, and allow the user to select one of those SMB servers.

The relative reference used to transition from

```
smb://corgis/
```

to

```
smb://cue/
```

would be "../cue".

When moving upward in the hierarchy, one might expect:

```
"smb://cue/" + ".." ==> "smb://"
```

In this example, however, the user agent is aware that node "cue" is a member of the "corgis" workgroup, so:

```
"smb://cue/" + ".." ==> "smb://corgis/"
```

The NBT workgroup membership of an SMB server may be determined either by sending a Node Status Request query to the server (see [\[RFC1001\], section 15.1.4](#)) or by maintaining a local cache of workgroup information, or both. Obviously, the choice is implementation dependent. If the server's workgroup membership is not available via either of these methods, then it is acceptable to move directly to the top of the hierarchy (smb://).

Hertel

Expires January 8, 2007

[Page 8]

5. SMB URI Definition

The following grammar defines the syntax of the SMB URI. It is based upon the grammar given in [Appendix A of \[RFC3986\]](#). Refer to that RFC for token definitions missing from the grammar below.

```

smb-URI      = ( smb-absURI / smb-relURI )
smb-absURI   = scheme "://" smb-service [ "?" [ nbt-context ] ]
smb-relURI   = ( path-absolute / path-rootless )
               [ "?" [ nbt-context ] ]

scheme        = "smb" / "cifs"
smb-service   = ( smb-wrkgrp / smb-net-path )

smb-wrkgrp    = [ smb-userinfo "@" ] [ smb-srv-name ]
               [ ":" port ] [ "/" ]
smb-net-path  = smb-server [ path-absolute ]
smb-server    = [ smb-userinfo "@" ] smb-srv-name [ ":" port ]

smb-srv-name  = nbt-name / host
nbt-name      = netbiosname [ "." scope-id ]
netbiosname   = netbiosnamec 1*14( netbiosnamec / "*" )
netbiosnamec  = ( ALPHA / DIGIT / pct-encoded
                 / "-" / "_" / "~"
                 / "!" / "$" / "&" / "'" / "(" / ")"
                 / "+" / "," / ";" / "="
               )

scope-id      = [ scope-label *( "." scope-label ) ]
scope-label   = 1*63( scope-char )
scope-char    = ALPHA / DIGIT / "-" / "_" / "~"
               / sub-delims / pct-encoded

smb-userinfo  = [ auth-domain ";" ] userinfo-nosem
auth-domain   = smb-srv-name
userinfo-nosem = *( unreserved / pct-encoded
                   / "!" / "$" / "&" / "'" / "(" / ")"
                   / "*" / "+" / "," / "=" / ";"
                 )

nbt-context   = nbt-param *("; " nbt-param )
nbt-param     = ( "BROADCAST=" IPv4address [ ":" port ]
                 / "CALLED=" netbiosname
                 / "CALLING=" netbiosname
                 / ( "NBNS=" / "WINS=" ) host [ ":" port ]
                 / "NODETYPE=" ( "B" / "P" / "M" / "H" )
                 / ( "SCOPEID=" / "SCOPE=" ) scope-id
               )

```

Hertel

Expires January 8, 2007

[Page 9]

6. SMB URI Syntax Elements

The SMB URI scheme is more or less comparable to other URI schemes used for remote filesystem access. It differs primarily in its support for the NBT transport and NBT workgroups. This section provides further explanation and description of those syntax elements that are most likely to require clarification.

6.1. scheme

As described in [section 2](#), above, an SMB URI is identified by one of two scheme names: "smb" or "cifs". These are equivalent, but the former is preferred.

6.2. smb-service

The SMB URI can be used to access workgroup information or SMB file server services. There are minor differences in SMB URI syntax depending up on which of these service types is being accessed.

It is possible, for instance, to request workgroup information without specifying a destination server name. In particular, the URI:

```
smb://
```

represents a request for the list of locally available NBT workgroups.

In some situations the workgroup list may not be available to unauthenticated users, so the SMB URI scheme allows inclusion of smb-userinfo information without the need to specify an smb-srv-name (a workgroup name). Thus, the following is permitted:

```
smb://user@/
```

In the above, the username "user" is being supplied. (The user agent should prompt for a password to prevent the password from being exposed.)

As with the smb-userinfo field, an SMB URI may include a port reference without an smb-srv-name, as in the following example:

```
smb://:4220/
```

(This is an example of an attempt to retrieve an NBT workgroup list via SMB using destination TCP port 4220.)

Hertel

Expires January 8, 2007

[Page 10]

Another difference between workgroup and SMB file server references is that workgroup references can not be followed by a path. The browse list does not offer shares, directories, or files so an SMB URI string such as the following cannot represent a workgroup query:

```
smb://corgis/puppies/
```

6.3. auth-domain

The auth-domain field is passed to the underlying SMB layer for interpretation. It is used to specify the SMB authentication authority (typically a "Domain Controller"). The interpretation of this field is specific to the workings of the SMB protocol and should be handled by the underlying SMB implementation.

6.4. smb-srv-name

The SMB URI supports the use of NetBIOS names and Scope IDs to identify SMB servers and services. When included as part of an SMB URI, the syntax of the NetBIOS name is a superset of the syntax of a DNS domain name label. For example:

```
smb://jcifs/
```

Syntactically, the string "jcifs" in the smb-srv-name field of the above string may be seen as either a DNS host name (unqualified), or as a NetBIOS name. The underlying SMB implementation must determine the namespace of the name. (This is a common problem in SMB implementations and is typically solved by first attempting to resolve the name as a NetBIOS name and then, if that fails, as a DNS host name.)

Likewise, given:

```
smb://jcifs.samba.org/
```

the string "jcifs.samba.org" may be interpreted either as a qualified DNS name, or as a NetBIOS name with appended Scope ID.

A NetBIOS name is simply a string of octets with a maximum length of 15 octets. (The actual maximum length of the NetBIOS name is 16-octets, but the 16th is reserved.) In practice, the only restriction on the syntax of a NetBIOS name is that it may not begin with an ASCII asterisk character (0x2A). Octet values that are permitted by NetBIOS name syntax but excluded by the SMB URI syntax must be escaped. Note, in particular, that the dot character (0x2E) must be escaped if used in a NetBIOS name.

The resolution of NetBIOS names to IP addresses is described in

[[RFC1001](#)] and [[RFC1002](#)].

Hertel

Expires January 8, 2007

[Page 11]

6.5. port

[RFC1001](#)/1002 includes a mechanism for retargeting Session Service connections to alternate ports (see [\[RFC1001\], section 16.1.1.](#)) which means that non-standard ports may be used for SMB over NBT transport. There may be other valid reasons for providing SMB services on on-standard ports.

The URI port field may be used to specify an alternate service port for SMB over either NBT or native TCP transport. (The transport type must be detected by the underlying SMB implementation.)

6.6. scope-id

The correct syntax of an NBT Scope Identifier is described in [section 9 of \[RFC1001\]](#) as

"...a character string meeting the requirements of the domain name system for domain names."

The intent was that the Scope ID should match the "preferred syntax" for domain names, as given in Appendix 1 of [\[RFC883\]](#) (which has since been superceded). Specifically:

"The labels must follow the rules for ARPANET host names. They must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and hyphen. There are also some restrictions on the length. Labels must be 63 characters or less."

Several implementations of NBT now exist that disregard the character set restrictions described above. The SMB URI is, therefore, a bit more flexible with regard to the octet values that may be used. In fact, the grammar given above allows the use of pct-encoded values, so any octet value may be specified. Note, however, that implementations MUST discard zero values ("%00") in the Scope ID because many implementations interpret the Scope ID as being a nul-terminated string.

Although the syntax allows Scope IDs that do not match the original preferred syntax, the use of nonconforming Scope IDs is generally considered unwise because the interpretation of nonconforming Scope IDs is likely to be inconsistent.

The requirement that the Scope ID consist of dot-delimited labels of one to sixty-three octets is enforced by the on-the-wire encoding used by NBT.

Hertel

Expires January 8, 2007

[Page 12]

The SMB URI scheme provides two mechanisms for specifying an NBT Scope ID. The first, as shown in the grammar above, is to append the Scope ID to the NetBIOS name as part of the smb-srv-name field, using a dot (".") as a delimiter. This mechanism is included to support existing implementations.

The other mechanism is to specify the Scope ID as part of the nbt-context. The two examples given below are equivalent:

```
smb://netbios.scope.id/  
smb://netbios/?SCOPE=scope.id
```

The latter format is less ambiguous and, therefore, preferred. User agents that rewrite URI strings for display purposes should rewrite SMB URI strings that contain a Scope ID to conform to the nbt-context format.

Note also that the scope-id syntax specifically permits a Scope ID that is the empty string (""). The empty string is a valid Scope ID and is, in fact, the default on all known implementations.

6.7. nbt-context

NBT context information is appended to the tail end of an SMB URI string in the form of a URI query string. Context information is specified using key/value pairs. Multiple context elements may be specified by separating the key/value pairs with semi-colons.

The nbt-context may be used to provide information about the NBT transport layer and related support servers. Information provided in the nbt-context overrides the current NBT context maintained by the user agent. The nbt-context is interpreted locally by the user agent.

The nbt-context is made up of zero or more nbt-params fields, which are specified as key/value pairs. For example:

```
smb://jcifs/?CALLED=VIRTSESV;NBNS=172.24.19.18
```

In the above example, the CALLED parameter is assigned a value of "VIRTSESV", and the NBNS parameter is assigned a value of "172.24.19.18".

Hertel

Expires January 8, 2007

[Page 13]

The following keywords are defined:

BROADCAST: The IPv4 broadcast address to which to send NBT broadcast name queries. This may, for example, be used on multi-homed hosts to specify a target subnet.

The value assigned to the BROADCAST keyword may optionally include a port number (delimited by a colon). The standard port for NBT name resolution is UDP/137. It is rare that a different port will be used for broadcast name resolution.

CALLED: Specifies the NetBIOS name of the SMB server (the NetBIOS destination address.) A CALLED name is required by the NBT Session Request message (see [\[RFC1002\], Section 4.3.2](#)).

If NBT transport is used and the CALLED name is not specified within the URI string, the underlying SMB implementation must deduce the CALLED name from available information. (See [Appendix A](#), below.)

CALLING: Specifies the NetBIOS name of the client (the NetBIOS source address.) This value is only used with NBT transport. It is required by the NBT Session Request message (see [\[RFC1002\], Section 4.3.2](#)).

If NBT transport is used and the CALLING name is not specified in the current NBT context, the underlying SMB implementation must generate a suitable name. (Typically, this will be based on the system's host name.)

NBNS: Specifies the NetBIOS Name Server (NBNS) to be used for point-to-point NBT Name Resolution. The NBNS may be specified using a DNS name or an IP address. See [\[RFC1001\]](#) for information on the NBNS.

The value assigned to this parameter may, optionally, include a port number (delimited by a colon). The standard port for NBT name resolution is UDP/137. The use of a non-standard port for point-to-point NBT name resolution is rare (but less so than it is for broadcast name resolution).

NODETYPE: One of B, P, M, H, or the empty string. These represent the different mechanisms by which a NetBIOS name may be resolved to an IP address on an NBT network. The first three types are defined in [RFC](#)

[1001](#)/1002. H mode is the inverse of M mode (in H mode

Hertel

Expires January 8, 2007

[Page 14]

the NBNS is queried before a broadcast query is sent). An empty NODETYPE indicates that NBT name resolution should not be attempted (use DNS name resolution instead). Some examples:

```
smb://smedley/?NBNS=172.24.19.18;NODETYPE=H
smb://corgis/?NODETYPE=B
smb://jcifs.samba.org/?NODETYPE=;CALLED=SMBSESV
```

SCOPE: Specifies the NBT Scope Identifier. Use of the SCOPE keyword is preferred over inclusion of the Scope ID in the nbt-name field. User agents must support both mechanisms.

The default Scope ID is the empty string. This can be specified in the SMB URI by assigning an empty value to the SCOPE keyword. For example:

```
smb://bran/SCOPE=
smb://marika/SCOPE=;NODETYPE=B
```

SCOPEID: A synonym for SCOPE.

WINS: A synonym for NBNS.

Although all of the keywords and values are shown in upper case, case is not significant.

The client implementation should provide a means for setting the base context. The nbt-context is used to override default values or to supply values missing from the local configuration. Most of all, the nbt-context makes it possible for an SMB URI string to maintain a consistent interpretation as it travels from one NBT scope to another.

7. The Relationship Between the SMB URI and the UNC Format

Some operating systems support a format known as Universal Naming Convention (UNC). UNC is a means for identifying network resources. SMB is one of the protocols supported by UNC.

In general, a UNC string specifying a resource available via SMB protocol can be converted into an SMB URI string by simply adding the "smb:" (or "cifs:") prefix and reversing the direction of all of the separating slashes. For example:

| UNC form | URI form |
|-----------------------|---------------------------|
| ----- | ----- |
| \\corgis\docs\ | smb://corgis/docs/ |
| \\corgis\docs\jolyon\ | smb://corgis/docs/jolyon/ |

\\corgis\docs\jolyon\rabbit.txt smb://corgis/docs/jolyon/rabbit.txt

Hertel

Expires January 8, 2007

[Page 15]

8. Authentication

SMB authentication can be divided into the following categories:

- o None
- o Share-based
- o User-based
- o Authentication-Server-based (NT Domain, AD Domain, or Kerberos)

The authentication mechanism to be used is negotiated during client/server session setup. Client applications, therefore, are aware of the server's authentication requirements and may prompt for appropriate input (username, password, authentication domain). By prompting for authentication information, an application ensures that such information is entered by the user in a controlled manner, and that security measures (if any), such as password encryption or password hash generation, are applied by the SMB protocol handler before the data are transmitted.

Some authentication values may also be provided within the SMB URI string. In particular, the following fields may optionally be included in the URI:

- auth-domain - The authentication domain (single-signon database server) to use for authorization
- userinfo - User account identifier (username)

9. Security Considerations

All of the implementations of the SMB URI that are known to exist at the time of this writing support inclusion of a password in the URI string. This is generally considered to be dangerous, since it may encourage exposure of the plaintext password. User agents should provide a mechanism for prompting the user to enter passwords, separate from the URI itself, in a reasonably secure fashion.

All other general security considerations relate to the workings of the SMB/CIFS protocol suite and are beyond the scope of this document.

10. Character Encoding Issues

The only restriction that [RFC1001/1002](#) places on the octet values that may be used in a NetBIOS name is that the name may not begin with an asterisk ('*', ASCII value 0x2A). No other values are excluded by those RFCs. For historical reasons, however, some implementations disallow the use of a nul byte (0x00) within a NetBIOS name. NetBIOS names are interpreted as fixed-length

strings of octets, so common mutli-byte character sets may cause

problems with older implementations.

Octet values less than 128 (0x80) in a NetBIOS name are typically interpreted as US-ASCII characters. The interpretation of octet values above 127 is dependent upon host configuration; there is no protocol mechanism to specify which codepage or character set is in use. In general, URI escape sequences should be used to represent characters with octet values above 127.

NetBIOS names, share names, and the directory paths and filenames offered by an SMB server may all contain characters from outside the 7-bit US-ASCII character set. Applications **MUST** support the use of the URI escape sequence as described in [[RFC3986](#)] to accommodate octet values that represent non-US-ASCII characters.

The SMB protocol has evolved over time to include support for various character encoding schemes. A complete discussion of SMB and NBT character encoding issues is way beyond the scope of this document.

11. Acknowledgments

The creation of this document would not have been possible without the help and guidance of

Michael B. Allen
David Farmer
Roy T. Fielding
Steven French
Larry Masinter
Richard Sharpe

and the aggregate knowledge and wisdom of

The jCIFS Team
The Samba Team
The Samba-TNG Team
The SNIA CIFS Working Group
The samba-technical and jCIFS mailing list participants
The IETF URI-review and W3C URI mailing list participants

Funding for the RFC Editor function is currently provided by the Internet Society.

Hertel

Expires January 8, 2007

[Page 17]

12. References

- [RFC883] Mockapetris, P., "DOMAIN NAMES - IMPLEMENTATION and SPECIFICATION", [RFC 883](#), November 1983.
- [RFC1001] Karl Auerbach, et. al., "Protocol Standard For a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods", [RFC 1001](#), March 1987.
- [RFC1002] Karl Auerbach, et. al., "Protocol Standard For a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications", [RFC 1002](#), March 1987.
- [RFC3986] Berners-Lee, T., Fielding, R. and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", [RFC 3986](#), January 2005.
- [XOPENSMB] "Protocols for X/Open PC Interworking: SMB, Version 2", ISBN 1-872630-45-6, The Open Group, October 1992.
- [ONET] Microsoft Corporation, Intel Corporation, "Microsoft Networks/OpenNET Filesharing Protocol", Document Version 2, Intel Part No. 138446, November 7, 1988.
- [SNIACIFS] Storage Network Industry Association CIFS Documentation Work Group, "Common Internet File System (CIFS) Technical Reference", Version: CIFS-TR 1.0, March 1, 2002.
- [IMPCIFS] Hertel, Christopher R., "Implementing CIFS -- the Common Internet File System", ISBN 0-13-047116-X, Prentice Hall PTR, August 2003
(<http://ubiqx.org/cifs/>)

13. Author's Address

Christopher R. Hertel
Alvarri, Inc.
15 South Fifth Street
Suite 1010
Minneapolis MN 55402

E'mail: crh@samba.org
crh@ubiqx.mn.org

Hertel

Expires January 8, 2007

[Page 18]

14. Copyright Notice

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

15. Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Appendix A. Working with NetBIOS Names (Implementation Notes)

The information presented in this section is intended as a guide for implementors.

Name resolution, particularly with the inclusion of support for [RFC1001](#)/1002 NBT naming, may result in ambiguous meaning for some SMB URI strings. This problem is reduced if correct NBT context information is included in URI strings, and can be eliminated if all implementations follow the same basic sequence when resolving server names to addresses.

A.1. NetBIOS Names

NetBIOS names are addresses. They represent communication end-points within a NetBIOS LAN. [\[RFC1001\]](#) and [\[RFC1002\]](#) provide a mechanism for creating virtual NetBIOS LANs over TCP and UDP transport. The core of that mechanism is the NetBIOS Name Service, which provides for mapping between NetBIOS names and the IP addresses at which they are registered. A given host system may register several NetBIOS names, each representing an application or service that can communicate with other applications or services through the NetBIOS API.

A.2. SMB Sessions via NBT

SMB sessions are established and messages transferred via the NetBIOS session service (see [\[RFC1001\]](#), [section 5.3](#) and [\[RFC1002\]](#) [section 4.3](#)). The system that originates the connection is the "calling" node, and the target node is the "called" node. In order to establish an SMB session, a TCP connection must be established between the calling and called nodes.

Before a NetBIOS session can be established, the calling node must discover the IP address of the called node. This is done using the NetBIOS Name Service (see [\[RFC1001\]](#) [section 5.2](#) and [\[RFC1002\]](#) [section 4.2](#)). NetBIOS names are always 16 octets, padded with spaces (0x20) if necessary, as specified in the RFCs. By convention, however, the 16th octet is reserved for use as a service type indicator. This field is known as the "suffix".

The suffix byte is NEVER specified in an SMB URI string, but is appended by the client implementation.

A.3. Resolving DNS names and IP addresses to SMB server names

The NetBIOS Session Service requires that the client provide the NetBIOS names of both the calling and called nodes in the NBT SESSION REQUEST. When connecting to an SMB server, the calling name (source address) is typically the default NetBIOS name of the client, space padded as described, with a suffix byte value of 0x00. The called name (destination address) is the NetBIOS name of the server with a suffix byte value of 0x20.

Applications which support the SMB URI must include support for the use of DNS names or IP addresses in addition to NetBIOS names when initiating SMB connections via NetBIOS over TCP/IP (NBT) transport. This functionality is an extension to the NetBIOS over TCP/IP behavior specified in [RFC 1001](#) and [RFC 1002](#), and is not part of that standard. It is, however, a common extension and must be supported for compatibility reasons, and to provide access to SMB shares in situations in which the NetBIOS name space cannot be guaranteed to be consistent.

As stated above, the Session Request packet requires a called and a calling name, both of which are NetBIOS names. In order to create an NBT Session Request packet, the DNS name or IP address of the server must be reverse-mapped to the server's NetBIOS name. Mechanisms for doing so include:

- Issuing a NetBIOS Adapter Status Query

A NetBIOS Adapter Status Query is sent to the target IP address. (See [\[RFC1001\] section 15.1.4](#) and [\[RFC1002\] sections 4.2.17 & 4.2.18.](#)) If a response is received, and if the target node is running an SMB server service, then the response will include a NetBIOS name with a suffix byte value of 0x20. This NetBIOS name may be used as the called name in a Session Request packet.

It is possible that multiple entries will have a suffix byte of 0x20. If this is the case each name may be tried in turn, or one of the other methods must be used to discover the name of the SMB server service.

- Generic Server Name

This method is not supported by all SMB server implementations.

Some SMB servers will accept the generic SMB server name "*SMBSERVER". A client can simply use the name "*SMBSERVER" as the called name in a Session Request packet. As with all SMB server NetBIOS names, the "*SMBSERVER" name must be space padded and terminated with a suffix byte value of 0x20.

Hertel

Expires January 8, 2007

[Page 21]

The "*SMBSERVER" name begins with an asterisk character, so it is an illegal NetBIOS name (see [\[RFC1001\], section 5.2](#)). Thus, it is never registered with the NetBIOS Name Service and will not be returned in a NetBIOS Adapter Status Response.

If the target does not support the "*SMBSERVER" generic name, or if it is not running SMB services, it will return a CALLED NAME NOT PRESENT error.

Some SMB servers are capable of providing multiple SMB file services, each under a different NetBIOS name. In order to support the generic server name, these servers must designate one service as a default that will answer to "*SMBSERVER".

- Parsing the DNS Name or IP address (guessing)

This is the least reliable method for discovering an SMB server name.

Systems which support STD 19 transport will often use the same base host name within the DNS and NetBIOS name spaces. Thus, the first label of the DNS name is a good guess at the NetBIOS name of the target. If the input is an IP address rather than a DNS name, the a reverse lookup against the DNS may be performed to determine the DNS name.

The first label of the DNS name consists of the initial portion of the DNS name string up to but not including the first dot character ('.'). If the label is greater than 15 bytes in length, it cannot be a NetBIOS name. The label must be space padded to a total of 15 bytes, with a suffix value 0x20 (space) added. This forms a valid NetBIOS name that may be used as a called name in a Session Request packet.

If the target returns a CALLED NAME NOT PRESENT error, then the DNS name guess is incorrect.

Any of the above may be tried in any order.

A.4. Determining the Namespace of the smb-srv-name

NetBIOS names, DNS names, and IP addresses can not be easily distinguished syntactically. Unfortunately, given the nature and history of the SMB/CIFS suite, the appropriate mechanism for distinguishing between these server specifier types is the trial-and-error method.

Hertel

Expires January 8, 2007

[Page 22]

Implementations should begin with the assumption that the specifier is a NetBIOS name. The following process is used to test this assumption:

If the NODETYPE (which is part of the NBT context) is the empty string, then no NetBIOS name resolution mechanism has been selected and the name cannot be resolved as a NetBIOS name. Exit.

If the name string contains dot characters ('.', ASCII 0x2E), then separate the name into NetBIOS name and Scope ID at the first dot. Otherwise use the entire string as the NetBIOS name, and assume an Scope ID of "" unless the Scope ID is specified in the nbt-context.

REPEAT

If the resulting NetBIOS name is greater than 15 octets in length, then the name is not a NetBIOS name. Exit.

Issue STD 19 Name Queries using the NetBIOS name and Scope ID. Suffix values of 0x20, 0x1B, and/or 0x1D should be used. (See section A.5., below.)

If a Positive Name Query Response is received, then the name is a NetBIOS name. Exit, indicating success and returning the NetBIOS name and scope ID as parsed.

END

If the server specifier is not a NetBIOS name, then it is either a DNS name or an IP address. These are semantically equivalent.

Hertel

Expires January 8, 2007

[Page 23]

A.5. Workgroup vs. SMB Server Names

If the URI string is of the form

```
smb://smb-srv-name/
```

then the smb-srv-name may represent either an SMB server name or a workgroup name. The name MUST NOT be interpreted as a workgroup name if:

- There is path information following the trailing slash.

Workgroups do not make shares or directories available.

- The server field is entered as a DNS name or an IP address.

A workgroup name is a NetBIOS group name. Workgroups, conceptually, represent a group of servers rather than an individual server, and the browse list may be retrieved from one or more browse servers.

In these cases, the smb_srv_name is interpreted as a reference to an SMB server only. Workgroups may only be accessed via their NetBIOS names.

When testing the name using the algorithm presented in section A.4, a NetBIOS name suffix value of 0x20 is used to find an SMB server, and a suffix value of 0x1D or 0x1B is used to find a workgroup browse server.

A system operating in B mode will use the 0x1D suffix to search for a Local Master Browser operating on the same subnet. A system operating in P mode must use the 0x1B suffix to query the NBNS for the Domain Master Browser. An M mode system will first send a broadcast query for the 0x1D name and, if that fails, query the NBNS for the 0x1B name. H mode behavior is the opposite of M mode.

Hertel

Expires January 8, 2007

[Page 24]