

The Gnutella Protocol Specification v0.4

Clip2 Distributed Search Solutions

<http://dss.clip2.com>

dss-protocols@clip2.com

- ¶1 Gnutella¹ is a protocol for distributed search. Although the Gnutella protocol supports a traditional client/centralized server search paradigm, Gnutella's distinction is its peer-to-peer, decentralized model. In this model, every client is a server, and vice versa. These so-called Gnutella *servents* perform tasks normally associated with both clients and servers. They provide client-side interfaces through which users can issue queries and view search results, while at the same time they also accept queries from other servents, check for matches against their local data set, and respond with applicable results. Due to its distributed nature, a network of servents that implements the Gnutella protocol is highly fault-tolerant, as operation of the network will not be interrupted if a subset of servents goes offline.

¶2 Protocol Definition

The Gnutella protocol defines the way in which servents communicate over the network. It consists of a set of descriptors used for communicating data between servents and a set of rules governing the inter-servent exchange of descriptors. Currently, the following descriptors are defined:

| Descriptor | Description |
|------------|--|
| Ping | Used to actively discover hosts on the network. A servent receiving a Ping descriptor is expected to respond with one or more Pong descriptors. |
| Pong | The response to a Ping. Includes the address of a connected Gnutella servent and information regarding the amount of data it is making available to the network. |
| Query | The primary mechanism for searching the distributed network. A servent receiving a Query descriptor will respond with a QueryHit if a match is found against its local data set. |
| QueryHit | The response to a Query. This descriptor provides the recipient with enough information to acquire the data matching the corresponding Query. |
| Push | A mechanism that allows a firewalled servent to contribute file-based data to the network. |

- ¶3 A Gnutella servent connects itself to the network by establishing a connection with another servent currently on the network. The acquisition of another servent's address is not part of the protocol definition and will not be described here (Host cache services are currently the predominant way of automating the acquisition of Gnutella servent addresses).
- ¶4 Once the address of another servent on the network is obtained, a TCP/IP connection to the servent is created, and the following Gnutella connection request string (ASCII encoded) may be sent:

GNUTELLA CONNECT/<protocol version string>\n\n

¹ Typically pronounced "newtella" or, less commonly, "gnewtella" where the 'g' is not silent.

where <protocol version string> is defined to be the ASCII string "0.4" (or, equivalently, "\x30\x2e\x34") in this version of the specification.

- ¶5 A servent wishing to accept the connection request must respond with

GNUTELLA OK\n\n

- ¶6 Any other response indicates the servent's unwillingness to accept the connection. A servent may reject an incoming connection request for a variety of reasons - a servent's pool of incoming connection slots may be exhausted, or it may not support the same version of the protocol as the requesting servent, for example.
- ¶7 Once a servent has connected successfully to the network, it communicates with other servents by sending and receiving Gnutella protocol descriptors. Each descriptor is preceded by a Descriptor Header with the byte structure given below.

Note 1: All fields in the following structures are in network (i.e. big endian) byte order unless otherwise specified.

Note 2: All IP addresses in the following structures are in IPv4 format. For example, the IPv4 byte array

| | | | |
|--------|--------|--------|--------|
| 0xD0 | 0x11 | 0x32 | 0x04 |
| byte 0 | byte 1 | byte 2 | byte 3 |

represents the dotted address 208.17.50.4.

¶8 **Descriptor Header**

| | | | | | | | | | | |
|-------------|------------|----|--------------------|----|-----|----|------|----|----------------|----|
| | Message ID | | Payload Descriptor | | TTL | | Hops | | Payload Length | |
| Byte offset | 0 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 |

| | |
|---------------------------|--|
| Descriptor ID | <i>A 16-byte string uniquely identifying the descriptor on the network</i> |
| Payload Descriptor | <i>0x00 = Ping</i> <i>0x01 = Pong</i> <i>0x40 = Push</i> <i>0x80 = Query</i> <i>0x81 = QueryHit</i> |
| TTL | <i>Time To Live. The number of times the descriptor will be forwarded by Gnutella servents before it is removed from the network. Each servent will decrement the TTL before passing it on to another servent. When the TTL reaches 0, the descriptor will no longer be forwarded.</i> |
| Hops | <i>The number of times the descriptor has been forwarded. As a descriptor is passed from servent to servent, the TTL and Hops fields of the header must satisfy the following condition:</i> $TTL(0) = TTL(i) + Hops(i)$ <i>Where $TTL(i)$ and $Hops(i)$ are the value of the TTL and Hops fields of the header at the descriptor's i-th hop, for $i \geq 0$.</i> |
| Payload Length | <i>The length of the descriptor immediately following this header. The next descriptor header is located exactly Payload_Length bytes from the end of this header i.e. there are no gaps or pad bytes in the Gnutella data stream.</i> |

¶9 The TTL is the only mechanism for expiring descriptors on the network. Servents should carefully scrutinize the TTL field of received descriptors and lower them as necessary. Abuse of the TTL field will lead to an unnecessary amount of network traffic and poor network performance.

¶10 The Payload Length field is the only reliable way for a servent to find the beginning of the next descriptor in the input stream. The Gnutella protocol does not provide an "eye-catcher" string or any other descriptor synchronization method. Therefore, servents should rigorously validate the Payload Length field for each descriptor received (at least for fixed-length descriptors). If a servent becomes out of synch with its input stream, it should drop the connection associated with the stream since the upstream servent is either generating, or forwarding, invalid descriptors.

¶11 Immediately following the descriptor header, is a payload consisting of one of the following descriptors:

¶12 **Ping (0x00)**

Ping descriptors have no associated payload and are of zero length. A Ping is simply represented by a Descriptor Header whose Payload_Descriptor field is 0x00 and whose Payload_Length field is 0x00000000.

¶13 *A servent uses Ping descriptors to actively probe the network for other servents. A servent receiving a Ping descriptor may elect to respond with a Pong descriptor, which contains the address of an active Gnutella servent (possibly the one sending the Pong descriptor) and the amount of data it's sharing on the network.*

¶14 *This specification makes no recommendations as to the frequency at which a servent should send Ping descriptors, although servent implementers should make every attempt to minimize Ping traffic on the network.*

¶15 **Pong (0x01)**

| | Port | | IP Address | | | Number of Files Shared | | | Number of Kilobytes Shared | | |
|-------------|------|---|------------|---|---|------------------------|----|----|----------------------------|--|--|
| Byte offset | 0 | 1 | 2 | 5 | 6 | 9 | 10 | 13 | | | |

Port *The port number on which the responding host can accept incoming connections.*

IP Address *The IP address of the responding host.*

This field is in little endian format.

Number of Files Shared *The number of files that the servent with the given IP address and port is sharing on the network.*

Number of Kilobytes Shared *The number of kilobytes of data that the servent with the given IP address and port is sharing on the network.*

¶16 Pong descriptors are only sent in response to an incoming Ping descriptor. It is valid for more than one Pong descriptor to be sent in response to a single Ping descriptor. This enables host caches to send cached servent address information in response to a Ping request.

¶17 **Query (0x80)**

| | Minimum Speed | | Search criteria | |
|-------------|---------------|---|-----------------|-----|
| Byte offset | 0 | 1 | 2 | ... |

Minimum Speed *The minimum speed (in kB/second) of servents that should respond to this message. A servent receiving a Query descriptor with a Minimum Speed field of n kB/s should only respond with a QueryHit if it is able to communicate at a speed $\geq n$ kB/s*

Search Criteria *A nul (i.e. 0x00) terminated search string. The maximum length of this string is bounded by the Payload_Length field of the descriptor header.*

¶18 **QueryHit (0x81)**

| | Number of Hits | | Port | IP Address | Speed | Result Set | Servent Identifier | | | | |
|-------------|----------------|---|------|------------|-------|------------|--------------------|----|-----|---|--------|
| Byte offset | 0 | 1 | 2 | 3 | 6 | 7 | 10 | 11 | ... | n | n + 16 |

Number of Hits *The number of query hits in the result set (see below).*

Port *The port number on which the responding host can accept incoming connections.*

IP Address *The IP address of the responding host.*

Speed *The speed (in kB/second) of the responding host.*

Result Set *A set of responses to the corresponding Query. This set contains Number_of_Hits elements, each with the following structure:*

| | | | |
|-------------|-------------------------------|-----------|-----------|
| | File Index | File Size | File Name |
| Byte offset | 0 3 4 | 7 8 | ... |

- File Index** *A number, assigned by the responding host, which is used to uniquely identify the file matching the corresponding query.*
- File Size** *The size (in bytes) of the file whose index is File_Index*
- File Name** *The double-nul (i.e. 0x0000) terminated name of the file whose index is File_Index.*

The size of the result set is bounded by the size of the Payload_Length field in the Descriptor Header.

Servent Identifier *A 16-byte string uniquely identifying the responding servent on the network. This is typically some function of the servent's network address. The Servent Identifier is instrumental in the operation of the Push Descriptor (see below).*

- ¶19 QueryHit descriptors are only sent in response to an incoming Query descriptor. A servent should only reply to a Query with a QueryHit if it contains data that strictly meets the Query Search Criteria.
- ¶20 The Descriptor_Id field in the Descriptor Header of the QueryHit should contain the same value as that of the associated Query descriptor. This allows a servent to identify the QueryHit descriptors associated with Query descriptors it generated.

¶21 **Push (0x40)**

| | | | | |
|-------------|---------------------------------|------------|------------|------|
| | Servent Identifier | File Index | IP Address | Port |
| Byte offset | 0 15 16 | 19 20 | 23 24 | 25 |


- Servent Identifier** *The 16-byte string uniquely identifying the servent on the network who is being requested to push the file with index File_Index. The servent initiating the push request should set this field to the Servent_Identifier returned in the corresponding QueryHit descriptor. This allows the recipient of a push request to determine whether or not it is the target of that request.*
- File Index** *The index uniquely identifying the file to be pushed from the target servent. The servent initiating the push request should set this field to the value of one of the File_Index fields from the Result Set in the corresponding QueryHit descriptor.*
- IP Address** *The IP address of the host to which the file with File_Index should be pushed.*
- Port** *The port to which the file with index File_Index should be pushed.*

¶22 A servent may send a Push descriptor if it receives a QueryHit descriptor from a servent that doesn't support incoming connections. This might occur when the servent sending the QueryHit descriptor is behind a firewall. When a servent receives a Push descriptor, it may act upon the push request if and only if the Servent_Identifier field contains the value of its servent identifier. See the section below entitled "Firewalled Servents" for further details on the Push process.


¶23 **Descriptor Routing**

The peer-to-peer nature of the Gnutella network requires servents to route network traffic (queries, query replies, push requests, etc.) appropriately. A well-behaved Gnutella servent will route protocol descriptors according to the following rules:

- ¶24 1. Pong descriptors may only be sent along the same path that carried the incoming Ping descriptor. This ensures that only those servents that routed the Ping descriptor will see the Pong descriptor in response. A servent that receives a Pong descriptor with Descriptor ID = n , but has not seen a Ping descriptor with Descriptor ID = n should remove the Pong descriptor from the network.
- ¶25 2. QueryHit descriptors may only be sent along the same path that carried the incoming Query descriptor. This ensures that only those servents that routed the Query descriptor will see the QueryHit descriptor in response. A servent that receives a QueryHit descriptor with Descriptor ID = n , but has not seen a Query descriptor with Descriptor ID = n should remove the QueryHit descriptor from the network.
- ¶26 3. Push descriptors may only be sent along the same path that carried the incoming QueryHit descriptor. This ensures that only those servents that routed the QueryHit descriptor will see the Push descriptor. A servent that receives a Push descriptor with Descriptor ID = n , but has not seen a QueryHit descriptor with Descriptor ID = n should remove the Push descriptor from the network.
- ¶27 4. A servent will forward incoming Ping and Query descriptors to all of its directly connected servents, except the one that delivered the incoming Ping or Query.
- ¶28 5. A servent will decrement a descriptor header's TTL field, and increment its Hops field, before it forwards the descriptor to any directly connected servent. If, after decrementing the header's TTL field, the TTL field is found to be zero, the descriptor is not forwarded along any connection.
- ¶29 6. A servent receiving a descriptor with the same Payload Descriptor and Descriptor ID as one it has received before, should attempt to avoid forwarding the descriptor to any connected servent. Its intended recipients have already received such a descriptor, and sending it again merely wastes network bandwidth.



Example 1. Ping/Pong Routing



Example 2. Query/QueryHit/Push Routing

¶30 File Downloads

Once a servent receives a QueryHit descriptor, it may initiate the direct download of one of the files described by the descriptor's Result Set. Files are downloaded out-of-network i.e. a direct connection between the source and target servent is established in order to perform the data transfer. File data is never transferred over the Gnutella network.

- ¶31 The file download protocol is HTTP. The servent initiating the download sends a request string of the following form to the target server:

```
GET /get/<File Index>/<File Name>/ HTTP/1.0\r\n
Connection: Keep-Alive\r\n
Range: bytes=0-\r\n
User-Agent: Gnutella\r\n^
\r\n
```

where <File Index> and <File Name> are one of the File Index/File Name pairs from a QueryHit descriptor's Result Set. For example, if the Result Set from a QueryHit descriptor contained the entry

| | |
|-------------------|-------------------------------|
| File Index | 2468 |
| File Size | 4356789 |
| File Name | Foo bar .mp3\00\00 |

then a download request for the file described by this entry would be initiated as follows:

```
GET /get/2468/Foobar.mp3/ HTTP/1.0\r\n
Connection: Keep-Alive\r\n
Range: bytes=0-\r\n
```

User-Agent: Gnutella
\r\n

- ¶32 The server receiving this download request responds with HTTP 1.0 compliant headers such as

HTTP 200 OK\r\n
Server: Gnutella\r\n
Content-type: application/binary\r\n
Content-length: 4356789\r\n
\r\n

- ¶33 The file data then follows and should be read up to, and including, the number of bytes specified in the Content-length provided in the server's HTTP response.

- ¶34 The Gnutella protocol provides support for the HTTP Range parameter, so that interrupted downloads may be resumed at the point at which they terminated.

¶35 **Firewalled Servents**

It is not always possible to establish a direct connection to a Gnutella servent in an attempt to initiate a file download. The servent may, for example, be behind a firewall that does not permit incoming connections to its Gnutella port. If a direct connection cannot be established, the servent attempting the file download may request that the servent sharing the file "push" the file instead. A servent can request a file push by routing a Push request back to the servent that sent the QueryHit descriptor describing the target file. The servent that is the target of the Push request (identified by the *Servent Identifier* field of the Push descriptor) should, upon receipt of the Push descriptor, attempt to establish a new TCP/IP connection to the requesting servent (identified by the *IP Address* and *Port* fields of the Push descriptor). If this direct connection cannot be established, then it is likely that the servent that issued the Push request is itself behind a firewall. In this case, file transfer cannot take place.

- ¶36 If a direct connection can be established from the firewalled servent to the servent that initiated the Push request, the firewalled servent should immediately send the following:

GIV <File Index>:<Servent Identifier>/<File Name>\r\n

Where <File Index> and <Servent Identifier> are the values of the File Index and Servent Identifier fields respectively from the Push request received, and <File Name> is the name of the file in the local file table whose file index number is <File Index>. The servent receiving the GIV request header (i.e. the Push requester) should extract the <File Index> and <File Name> fields from the header and construct an HTTP GET request of the following form:

GET /get/<File Index>/<File Name>/ HTTP/1.0\r\n
Connection: Keep-Alive\r\n
Range: bytes=0-\r\n
User-Agent: Gnutella\r\n²
\r\n

- ¶37 The remainder of the file download process is identical to that described in the section entitled "File Downloads" above.

² The allowable values of the User-Agent string are defined by the HTTP standard. Servent developers cannot make any assumptions about the value here. The use of 'Gnutella' is for illustration purposes only.