



SANS Institute

Information Security Reading Room

An Overview of Gnutella

Brenda Batkins

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

An Overview of Gnutella

Brenda L. Batkins

July 27, 2001

Version 1.2e

Introduction

A recent review of software loaded on a personal computer revealed a hitherto unknown (at least, by me) software program called "Toadnode." Nestling serenely among the standard Microsoft programs was an unknown entity. What the heck is Toadnode; what does it do?

My first step was to go to my favorite internet search engine – Altavista at <http://www.altavista.com>. By my keying in Toadnode, the search engine returned the home address as <http://www.toadnode.com>. When I reached the home page, I learned that Toadnode was a peer to peer file-sharing program, compatible with Gnutella. Since I was unfamiliar with Gnutella, I decided to base my certification research document on the Gnutella protocol. This document addresses the following issues: the origins of Gnutella, what it is and how it works, some Gnutella-compatible software, and some possible security concerns.

Where Did Gnutella Come From?

Research into the protocol led me to an article written by Catherine Bacon entitled "Download Now: Gnutella" which stated:

Gnutella, named for the GNU open-source movement and Nutella, the yummy hazelnut-and-chocolate spread, is a real-time, peer-based, file-sharing client that lets you search for and download files from other Gnutella users. Unlike Napster, Gnutella does not run on a server. It's not "based" anywhere...

...Gnutella originally was conceived, written, and released by Justin Frankel and Tom Pepper of Nullsoft, the company that makes Winamp, in March 2000. When AOL bought Nullsoft, it pulled the plug on Gnutella, realizing the potential for controversy. But programmers already had gotten their hands on the software and reverse engineered it, and soon versions of Gnutella began popping up for different operating systems./1

What Exactly is a Gnutella-type Model Concept?

The Gnutella protocol uses a somewhat different concept than the typical internet client server model. According to LimeWare's glossary of peer to peer terms, Gnutella uses a "servent" concept described as "A combination of a server and a client. In the old centralized file-sharing model, there were distributors of information, called servers, and requestors of information, called clients. In the decentralized gnutella model, each computer on the network is both a client and a server and is thus called a "servent." /2 In essence, a computer on a Gnutella network can both listen and respond when another computer talks.

How Gnutella Works

To utilize the Gnutella protocol and get connected to a Gnutella network, a user first must download and install a Gnutella compatible program. According to Toadnode's FAQ, Gnutella-compatible software works as follows:

The protocol defines the way that applications, such as Toadnode, communicate over the Internet. It is helpful to think of the P2P network as a conversation between computers. Some computers are "talking" while others are "listening". To coordinate this conversation packets are tagged with special descriptors so that each computer receiving the packet knows how to react. The current Gnutella protocol (version 0.4) defines five descriptors: Ping, Pong, Query, QueryHit and Push.

Step 1: Determining who is on the network

A "Ping" packet is used to announce your presence on the network. When another computer hears your Ping it will respond with a "Pong" packet. It will also forward your Ping packet to other computers to which it is connected and, in response, they too will send back Pong packets. Each Ping and Pong packet contains a Globally Unique Identifier (GUID). A Pong packet also contains an IP address, port number, and information about how much data is being shared by the computer that sent the Pong. Pong packets are not necessarily returned directly to the point of origin, instead they are sent from computer to computer via the same route as the initial Ping. After sending a Ping to one computer you will start receiving many Pong responses via that one computer. Now that the Pong packets have told you who your active peers are, you can start making searches.

Step 2: Searching

Gnutella is a protocol for distributed search. Gnutella "Query" packets allow you to search by asking other computers if they are sharing specific content (and have an acceptably fast network connection). A Query packet might ask, "Do you have any content that matches the string 'Homer'?" This question is sent to all the computers that sent you Pong packets. Each of these computers does two things. First, each computer checks to see if it has any content that matches the search string. In this case it looks to see if there are any files in a specified directory marked "sharable to the outside world" that have the letters "Homer" in its complete file path. Second, each computer sends your Query packet on to all the computers to which it is connected. These computers check their directories and send your Query packet to all their connected computers. This process continues until you run out of computers to ask or until the Query packet gets too old and times out. This last detail is important because without a pre-defined Time To Live (TTL) the Query packet could get bounced around for a very long time, potentially forever. Most servers, including Toadnode, allow you to adjust the TTL. GUIDs in each packet are used to make sure that the same message does not get passed to the same computer again and again, creating a loop.

Step 3: Downloading

By the time you are ready to download, the question you asked in your Query packet has been distributed to a huge number of computers. Each computer has checked its shared information and determined if it is sharing anything that matches "Homer". Let us say

that three computers that received your Query packet have a match for “Homer”. The last two packet descriptors, called “QueryHit” and “Push” are responsible for content delivery. Each of the three computers will send you a QueryHit packet via the same delivery route, computer-to-computer, that the Query packet originally traveled. The QueryHit packet contains the IP address and GUID of the computer that has the data as well as information about the file that matched your query.

When you receive a QueryHit packet your server software will display the name of the file for you and give you the option to download. File transfers use the HTTP protocol’s GET method directly between your computer and the computer that has the file you want. Normally, your computer will initiate the HTTP connection to the computer that has the file. Occasionally, due to a firewall, you will be unable to initiate a connection directly to the computer that has the file you want. In these cases the “Push” packet is used. The Push packet allows a message to be delivered to the computer that has the file you would like to download via the route that the QueryHit packet originally traveled, except in reverse. The Push packet tells this computer that you would like to download a file but cannot manage to initiate an HTTP connection. This computer then becomes the initiator, attempting to connect directly to you, which often is possible because the firewall between the machines is only limiting connections initiated from outside the firewall. /3

If a picture is worth a thousand words, then you might want to visit Toadnode’s flash animation of how Gnutella works at <http://www.toadnode.com/flashpage.asp>. /4

What Are Some Gnutella-like Programs?

To find listings of Gnutella-like software programs, I went to <http://www.gnutelliums.com> /5 and discovered the following brief descriptions:

Windows clients:

BearShare (<http://www.bearshare.com>) (July 23, 2001)

“BearShare is an exciting new Windows file sharing program from Free Peers, Inc. that lets you, your friends, and everyone in the world share files! Built on Gnutella technology, BearShare provides a simple, easy to use interface combined with a powerful connection and search engine that puts thousands of different files in easy reach!” /6

Gnotella (<http://www.gnotella.com>) (July 23, 2001)

“Gnotella is clone of Gnutella, a distributed real time search and file sharing program. Gnotella is for the Win32 environment, and offers extra benefits such as multiple searches, improved filtering/spam protection, bandwidth monitoring, enhanced statistics, upload throttling, and skinning, as well as more.” /6

Gnucleus (<http://gnucleus.sourceforge.net/>) (July 23, 2001)

“An open Gnutella client for an open network. Made for windows utilizing MFC (works in WINE). Constantly evolving, easy enough for the first time user and advanced enough to satisfy the experts.” /6

LimeWire (<http://www.limewire.com>) (July 23, 2001)

“LimeWire is a multi-platform Gnutella client with nice features like auto-connect, browse host, multiple search, upload throttling, connection quality control, library management and sophisticated filtering. It is built for the both the novice and power user.” /6

Phex (<http://www.konrad-haenel.de/phex/>) (July 23, 2001)

“Phex is entirely based on William W. Wong’s [Furi](#). As Furi has not been updated for over one year I decided to continue it’s development. But in case Wong is currently working on a new version of Furi i decided to rename my branch of the client to Phex.

FURI is a Gnutella protocol-compatible Java program that can participate in the Gnutella network. It is a full version program with a easy to use GUI interface that can perform most of the tasks of a Gnutella servant.” /6

Toadnode (<http://www.toadnode.com>) (July 23, 2001)

Toadnode described itself as “an extensible platform for peer-to-peer (P2P) networks. Its core functionality revolves around the ability to find, retrieve and distribute data between users across multiple networks. Toadnode pairs this ability to search, with an application layer to accommodate plug-ins that fully exploits and leverages the data that is distributed.

Toadnode is a [FREE](#) application that is designed to work with computers running most versions of Microsoft Windows.” /4b

Linux/Unix clients:

Gnutellium (<http://newtella.com/linux>) (July 23, 2001)

“Gnewtellium is the Linux/Unix port of [Newtella](#).

Newtella is the new way to share music over the internet. It combines a focus on music, like Napster, with a decentralized network of users, and is based on the gnutella protocol. The software is designed to retrieve and exchange only MP3 files. As such, it prevents the unrestricted duplication of viruses and self-executing trojan horses. It also prevents illicit uses (such as child pornography) of the gnutella network.” /7

Gnut (http://www.gnutelliums.com/linux_unix/gnut/) (July 23, 2001)

“Gnut is a command-line client which implements the gnutella protocol. It supports all features available in the original Nullsoft client, as well as many others. Bandwidth limiting, sorting of results, regular expression searching, are among the list. It will compile and run on a wide range of POSIX compliant (and not so compliant) systems including: SunOS, Linux, FreeBSD, HP-UX, and Win32.” /7

Hagelslag (<http://tiefighter.et.tudelft.nl/hagelslag>) (July 23, 2001)

“Hagelslag is an implementation of Gnutella. The main goals for this implementation are flexibility, stability and performance. The development of Hagelslag was primarily aimed at i386 machines running Linux, as of version 0.8, FreeBSD is supported as well.” /7

Phex (<http://www.konrad-haenel.de/phex/>) (July 23, 2001)

“Phex is entirely based on William W. Wong’s [Furi](#). As Furi has not been updated for over one year I decided to continue it’s development. But in case Wong is currently working on a new version of Furi i decided to rename my branch of the client to Phex.

FURI is a Gnutella protocol-compatible Java program that can participate in the Gnutella network. It is a full version program with a easy to use GUI interface that can perform most of the tasks of a Gnutella servant.” /7

Qtella (<http://www.qtella.net/>) (July 23, 2001)

“Qtella is a new Gnutella client for Linux written in C++ using the Qt libraries. It should be no problem to use Qtella on any platforms where Qt with thread support (library qt-mt must exists) is installed.” /7

Macintosh clients:

Mactella (<http://www.excc.com/>) (July 23, 2001)

“Mactella is the Mac version of Gnutella, an open-source file-sharing network that allows you to exchange an assortment of file formats with other users. It can operate on any port and has no centralized server. This program is capable of transferring any type of file that users put online, including ZIP, MPEG, ASF, MOV, QT, HQX, EXE, JAR, and SIT.” /8

Macintosh clients also included Phex and LimeWire. These two programs were also listed as Java Gnutella clients with LimeWire touting itself as a Solaris Gnutella client as well.

As a footnote, during the course of my research, I found that at <http://www.abctella.com> you can do file searches over the internet for Gnutella files.

Some Security Concerns

As you can tell from the information presented above, it appears that the Gnutella protocol is here to stay. Gnutella-compatible software is prolific and allows file searching and sharing of virtually any file you can imagine. These files can range from music to movies to games, to spreadsheets, word documents, pornography, executables, etc. Imagine, if you will, an employee downloading a file sharing program, and through either malicious or inadvertent intent, making proprietary company information available.

Because it is possible to get information through a firewall via the Gnutella protocol (see <http://gnutella.wego.com/go/wego.pages.page?groupId=116705&view=page&pageId=200448&f>

[olderId=118398&panelId=119597&action=view](#)) /9, corporations might not be as safe as they would like to believe and might want to consider other security countermeasures as well.

Concerning security issues, I found an interesting slide show entitled “Security Aspects of Napster and Gnutella” by Steven M. Bellovin at <http://www.research.att.com/~smb/talks/NapsterGnutella/sld001.htm>. /9 This document lists common functions of Napster and Gnutella, differences between the two and some possible security and privacy ramifications.

Furthermore, in a recent article at http://www.auditnet.org/articles/have_you_been_napstered.htm /10, Rob Harmer gives some very compelling arguments concerning the risks of file sharing, especially in light of a company’s information assets and resources. From illegal and unauthorized copies of software, to viruses, to files that might be considered unacceptable for the work environment (e.g., adult content), the Gnutella protocol will allow files to be shared. What are some steps that can be taken to mitigate the risk and to protect the company?

1. Develop a policy defining acceptable levels of use by employees of information resources, assets, and the consequences of failure to adhere to the policy.
2. Conduct periodic reviews of software loaded on employee’s personal computers.
3. Conduct periodic security awareness programs to enlighten the employees about emerging risks and sound security practices.
4. Maintain and review security logs of activity at the firewall and at the server level as well.
5. Ensure anti-virus software is current.
6. Block all ports that are not necessary.

References:

- 1/ Brown, Catherine, “Download Now: Gnutella”, June 23, 2000
<http://www.newmedia.com/nm-ie.asp?articleID=1145> (July 23, 2001)
- 2/ <http://www.limewire.com/index.jsp/glossary#servent> (July 23, 2001)
- 3/ <http://www.toadnode.com/FAQs.asp#gnutellaproto> (July 23, 2001)

- 4/ <http://www.toadnode.com/flashpage.asp> (July 23, 2001)
- 4b/ <http://www.toadnode.com> (July 23, 2001)
- 5/ <http://www.gnutelliums.com> (July 23, 2001)
- 6/ <http://www.gnutelliums.com/windows/> (July 23, 2001)
- 7/ http://www.gnutelliums.com/linux_unix/ (July 23, 2001)
- 8/ <http://www.gnutelliums.com/macintosh/> (July 23, 2001)
- 9/
<http://gnutella.wego.com/go/wego.pages.page?groupId=116705&view=page&pageId=200448&folderId=118398&panelId=119597&action=view> (July 23, 2001)
- 10/ Bellovin, Steven M. "Security Aspects of Napster and Gnutella"
<http://www.research.att.com/~smb/talks/NapsterGnutella/sld001.htm> (July 23, 2001)
- 11/ Harmer, Rob "Have You Been Napstered" (March 4, 2001)
http://www.auditnet.org/articles/have_you_been_napstered.htm

Other URL's Cited:

<http://www.altavista.com>

<http://www.bearshare.com/>

<http://www.gnotella.com/>

<http://gnucleus.sourceforge.net/>

<http://www.limewire.com/>

<http://www.konrad-haenel.de/phex/>

<http://newtella.com/linux/>

http://www.gnutelliums.com/linux_unix/gnut/

<http://tiefighter.et.tudelft.nl/hagelslag/>

<http://www.qtella.net/>

<http://www.abctella.com>