

UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION

Petitioners

v.

NONEND INVENTIONS, N.V.

Patent Owner

CASES: To Be Assigned

Patent Nos. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

I, Ziad Elkurjie, declare:

1. I have been asked by Petitioners to provide this Declaration to be used in the Petition for *Inter Partes* Review of U.S. Patent Nos. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2 (“the Patents”).

2. I have been compensated at the rate of \$250 per hour for my time to prepare, review, and sign this Declaration. My compensation is not contingent upon the outcome of any *inter partes* review, district court litigation, or any other proceeding related to the Patents.

I. BACKGROUND AND EXPERIENCE

3. I am an technology and product development executive with over 15 years of experience with fast growth business and systems operations. I have experience in business-to-consumer (B2C), e-commerce, and business systems architecture, development, and operations.

4. I earned a Bachelor of Science Degree in Computer Science from University of Toronto.

5. My public LinkedIn profile summarizing my employment and education experience is attached to this Declaration as Attachment A.

II. CLIP2 DISTRIBUTED SEARCH SERVICES

6. Between 1999 and November 2001, I was director of engineering at Clip2.com, Inc. (“Clip2”).

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

7. Clip2 operated a web site at <http://www.clip2.com> to enable users to publish Web guides on the topics that interest them. The web site was public and available for any user to participate on any topic. Visitors to the web site could perform keyword searching of the guide database, and guides were given numerical ratings based on their usage rates.

8. In April 2000, PC Magazine honored Clip2 it as one of the Top 100 sites on the internet.

9. In 2000, Clip2 recognized that Gnutella and related technologies opened the way for a revolutionary new generation of distributed open search.

10. Gnutella is a protocol for distributed search. Gnutella supports a traditional client/centralized server search. Gnutella is a peer-to-peer, decentralized model, where, every client is a server, and vice versa.

11. To address this technical innovation, Clip2 created a Distributed Search Services site at <http://dss.clip2.com>, providing technical data to the developer and user communities interested in distributed open search systems. The data reported by Clip2 on the Gnutella network is derived from network measurement, monitoring, and analysis technology developed by Clip2.

12. To facilitate its research in the Gnutella network, Clip2 also operated a public web site at <http://gnutellahosts.com/> which provided research articles published by Clip2 about Gnutella, a description of the Gnutella protocol, a

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

list of hosts to allow a user to connect to the Gnutella network, links to other web sites focused on Gnutella, and other information.

13. As director of engineering at Clip2, I lead efforts to write software to analyze and map the performance of the Gnutella network. The results of this analysis and mapping was published by Clip2 at <http://dss.clip2.com> and <http://gnutellahosts.com>. I have personal knowledge of the websites in 2000 because the research work I was involved with was posted to those websites. I visited and assisted with the development of those web sites in 2000.

III. GnutellaHosts WEB SITE AND DOCUMENTS

14. Petitioners' have provided me with several web pages from the <http://gnutellahosts.com>, archived on the Internet Archive Wayback Machine at various points between August 2000 and February 2001.

15. Attachment B is a true and correct copy of the "About Us: Distributed Search Services by Clip2" page on the on the <http://gnutellahosts.com> web site, http://gnutellahosts.com/dss_about.html, cached by the Internet Archive on August 18, 2000. I recognize and recall the format and content of the web page in Attachment B as a true and correct copy of the About Us: Distributed Search Services by Clip2" page of the GnutellaHosts website <http://gnutellahosts.com> as it appeared in August 2000. This web page described Clip2 and the Distributed Search Services project's function to provide reports

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

and information about Gnutella. Attachment B is still available on the Internet
at the following URL:

[https://web.archive.org/web/20000818053627/http://gnutellahosts.com/
dss_about.html](https://web.archive.org/web/20000818053627/http://gnutellahosts.com/dss_about.html)

16. Attachment C is a true and correct copy of the main page on the on the
<http://gnutellahosts.com> web site, <http://gnutellahosts.com>, cached by the
Internet Archive on August 16, 2000. I recognize and recall the format and
content of the web page in Attachment C as a true and correct copy of the
main page of the GnutellaHosts website <http://gnutellahosts.com> as it
appeared in August 2000. The main page contained a list of Gnutella Hosts to
allow a user to connect to the Gnutella network, Gnutella network statistics,
Gnutella resources including the Gnutella Protocol Specification v0.4
document, links to other web sites with Gnutella information, and other
information. As I explained before, <http://gnutellahosts.com> was provided by
Clip2 Distributed Search Services (“Clip2 DSS”). Attachment C is still
available on the Internet at the following URL:

<https://web.archive.org/web/20000816001015/http://gnutellahosts.com/>

17. Attachment D is a true and correct copy of the contact page on the on
the <http://gnutellahosts.com> web site, <http://gnutellahosts.com/contact.htm>,
cached by the Internet Archive on December 11, 2000. I recognize and recall

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

the format and content of the web page in Attachment D as a true and correct copy of contact page on the on the <http://gnutellahosts.com> web site as it appeared in December 2000 (some of the images on the clickable links to other portions of the web site are missing). The contact page further confirms my independent recollection that the Gnutella Hosts web site was operated by Clip2 Distributed Search Services. Attachment D is still available on the Internet at the following URL: <https://web.archive.org/web/20001211215800/http://gnutellahosts.com/contact.html>

18. Attachment E is a true and correct copy of the articles page on the on the <http://gnutellahosts.com> web site, <http://gnutellahosts.com/articles.htm>, cached by the Internet Archive on December 11, 2000. I recognize and recall the format and content of the web page in Attachment E as a true and correct copy of articles page on the on the <http://gnutellahosts.com> web site as it appeared in December 2000 (some of the images on the clickable links to other portions of the web site are missing). The articles page contains a list of Gnutella articles, including “Bandwidth Barriers to Gnutella Network Scalability” and “Gnutella: To the Bandwidth Barrier and Beyond” and the Gnutella Protocol Specification v0.4 that were public and available for download or viewing on or before December 2000. These articles and the

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

specification were authored by Clip2 as part of the Distributed Search Services project. Attachment E is still available on the Internet at the following URL:

<https://web.archive.org/web/20001211205900/http://gnutellahosts.com/articles.html>

19. Attachment F is a true and correct copy of the Gnutella Protocol Specification v0.4 on the on the <http://gnutellahosts.com> web site, <http://gnutellahosts.com/GnutellaProtocol04.pdf>, cached by the Internet Archive on August 18, 2000. I recognize and recall the format and content of the protocol specification in Attachment F as a true and correct copy of the protocol specification available for download on the <http://gnutellahosts.com> web site in mid-August 2000.

20. The metadata for Attachment F is shown below:

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

Document Properties

Description Security Fonts Initial View Custom Advanced

Description

File: GnutellaProtocol04.pdf

Title: GnutellaProtocol0.PDF

Author: achasin

Subject:

Keywords:

Created: 8/10/2000 4:20:00 PM

Modified:

Application: GnutellaProtocol0.4 - Microsoft Word

Additional Metadata...

Advanced

PDF Producer: Acrobat PDFWriter 4.0 for Windows NT

PDF Version: 1.2 (Acrobat 3.x)

Location: C:\

File Size: 31.56 KB (32,317 Bytes)

Page Size: 8.50 x 11.00 in

Number of Pages: 8

Tagged PDF: No

Fast Web View: No

Help OK Cancel

21. The metadata for Attachment F indicates that the GnutellaProtocol04.pdf document was created on August 10, 2000, which is the same as my independent recollection of when Clip2 created the Gnutella Protocol Specification v0.4.

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

22. The metadata for Attachment F indicates that the author of the GnutellaProtocol04.pdf document was “achasin.” I recognize this name as Andrew Chasin, a software engineer I supervised at Clip2, and the metadata is consistent with my independent recollection that Andrew Chasin was part of the engineering team who worked on the Gnutella Protocol Specification v0.4 document.

23. Attachment F is a description of version 0.4 of the Gnutella protocol, which indicates how Gnutella clients/servers (called “servents” because each every client is a server, and every server is a client) communicate with each other on the Gnutella network. The protocol was authored by the engineering team at Clip2 as part of the Distributed Search Services project. Attachment F is still available on the Internet at the following URL:

<https://web.archive.org/web/20000816001015/http://gnutellahosts.com/GnutellaProtocol04.pdf>

24. Attachment G is a true and correct copy of an Internet Archive page showing all of the times a copy of the GnutellaProtocol04.pdf, available at <http://gnutellahosts.com/GnutellaProtocol04.pdf>, was cached by the Internet Archive. Attachment G shows that the first time the Gnutella Protocol Specification v0.4, GnutellaProtocol04.pdf, was cached by the Internet Archive from <http://gnutellahosts.com> was August 18, 2000. This is

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

consistent with my independent recollection that the Gnutella Protocol Specification v0.4, GnutellaProtocol04.pdf was publicly posted and available on <http://gnutellahosts.com> at least as early as August 18, 2000. Attachment G is available on the Internet at the following URL: https://web.archive.org/web/20000315000000*/http://gnutellahosts.com/GnutellaProtocol04.pdf

25. After Clip2 created the Gnutella Protocol Specification v0.4, it was posted shortly thereafter on the <http://dss.clip2.com> and <http://gnutellahosts.com> and was accessible to the public. My independent recollection is confirmed by Attachment C, which shows a link to the document on the August 18, 2000 cached webpage. The Gnutella Protocol Specification v0.4 was accessed and downloaded by the public after it was posted, and became a widely cited document to explain the Gnutella protocol.

26. Attachment H is a true and correct copy of the an article authored by engineers at Clip2, titled “Bandwidth Barriers to Gnutella Network Scalability” posted on the on the <http://gnutellahosts.com> web site, http://www.gnutellahosts.com/dss_barrier.html, cached by the Internet Archive on January 16, 2001 (some of the images in the article and for the clickable links on the web page are missing). The article in Attachment H is dated September 8, 2000, and it was posted to shortly thereafter on the <http://dss.clip2.com> and later to <http://gnutellahosts.com>, and accessed by the public within a few days

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

after it was posted. I recognize and recall the format and content of the web page in Attachment H as a true and correct copy of article “Bandwidth Barriers to Gnutella Network Scalability” posted on the on the <http://gnutellahosts.com> web site as it appeared in September 2000 (some of the images on the clickable links to other portions of the web site are missing). Attachment H is available on the Internet at the following URL: https://web.archive.org/web/20010116190600/http://www.gnutellahosts.com/dss_barrier.html

27. My recollection about posting and public access of the “Bandwidth Barriers to Gnutella Network Scalability” article in Attachment H is confirmed by September 12, 2000 public commentary on this article on the slashdot.org web site, as shown in Attachment K, which also links to a copy of the article posted at http://dss.clip2.com/dss_barrier.html. Attachment K is available on the Internet at the following URL: <http://slashdot.org/story/00/09/12/1217200/bandwidth-barriers-to-gnutella-network-scalability>

28. Attachment I is a true and correct copy of the a web page article authored by engineers at Clip2, titled “Gnutella Protocol Extensions” posted on the on the <http://gnutellahosts.com> web site, http://gnutellahosts.com/protocol_extensions.html, cached by the Internet Archive on February 2, 2001 (some of the images for the clickable links on the web page are missing). I

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

recognize and recall the format and content of the web page in Attachment I as a true and correct copy of the article text as it was posted by February 2001 (the pictures accompanying the web page are missing). Attachment I is still available on the Internet at the following URL:

https://web.archive.org/web/20010202155400/http://gnutellahosts.com/protocol_extensions.html

29. Attachment J is a true and correct copy of the an article authored by engineers at Clip2, titled “Gnutella: To the Bandwidth Barrier and Beyond” posted on the on the <http://gnutellahosts.com> web site, <http://www.gnutellahosts.com/gnutella.html>, cached by the Internet Archive on December 13, 2000 (some of the images in the article and for the clickable links on the web page are missing). The article in Attachment J is dated November 6, 2000, and it was posted to shortly thereafter on the <http://dss.clip2.com> and to <http://gnutellahosts.com>, and accessed by the public within a few days after it was posted. I recognize and recall the format and content of the web page in Attachment J as a true and correct copy of article “Gnutella: To the Bandwidth Barrier and Beyond” posted on the on the <http://gnutellahosts.com> web site as it appeared in November 2000 (the pictures accompanying the web page are missing). Attachment H is available on the Internet at the following URL:

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

<https://web.archive.org/web/20001213064300/http://www.gnutellahosts.com/gnutella.html>

30. The web site <http://gnutellahosts.com>, including web site pages in Attachments B-F and H-J, were accessible to the public and were well known websites in 2000 and 2001 for information about Gnutella technology and information.

31. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

32. I declare under penalty of perjury that the foregoing is true and correct.

Executed on 26th day of September 2015 in Santa Barbara, California.



Ziad Elkurjie

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

- Attachment A: Ziad Elkurjie LinkedIn Profile, *available at* <https://www.linkedin.com/in/zkurjie> (last accessed Sept. 19, 2015)
- Attachment B: Internet Archive Wayback Machine August 18, 2000 capture of “About Us: Distributed Search Services by Clip2” web page of the Gnutella Hosts website, http://gnutellahosts.com/dss_about.html, *available at*: https://web.archive.org/web/20000818053627/http://gnutellahosts.com/dss_about.html (last accessed Sept. 19, 2015)
- Attachment C: Internet Archive Wayback Machine August 16, 2000 capture of the main web page of the Gnutella Hosts website, <http://gnutellahosts.com/>, *available at*: <https://web.archive.org/web/20000816001015/http://gnutellahosts.com/> (last accessed Sept. 19, 2015)
- Attachment D: Internet Archive Wayback Machine December 11, 2000 capture of the contact web page of the Gnutella Hosts website, <http://gnutellahosts.com/contact.html>, *available at*: <https://web.archive.org/web/20001211215800/http://gnutellahosts.com/contact.html> (last accessed Sept. 19, 2015)
- Attachment E: Internet Archive Wayback Machine December 11, 2000 capture of articles web page of the Gnutella Hosts website, <http://>

DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

gnutellahosts.com/articles.html, available at: [https://](https://web.archive.org/web/20001211205900/http://gnutellahosts.com/articles.html)

[web.archive.org/web/20001211205900/http://](http://web.archive.org/web/20001211205900/http://gnutellahosts.com/articles.html)

gnutellahosts.com/articles.html (last accessed Sept. 19, 2015)

Attachment F: Internet Archive Wayback Machine August 18, 2000 capture of
The Gnutella Protocol Specification v0.4 document posted on
the Gnutella Hosts website, [http://gnutellahosts.com/
GnutellaProtocol04.pdf](http://gnutellahosts.com/GnutellaProtocol04.pdf), available at: [https://web.archive.org/
web/20000816001015/http://gnutellahosts.com/
GnutellaProtocol04.pdf](https://web.archive.org/web/20000816001015/http://gnutellahosts.com/GnutellaProtocol04.pdf) (last accessed Sept. 19, 2015)

Attachment G: Internet Archive Wayback Machine, Captures of [http://](http://gnutellahosts.com/GnutellaProtocol04.pdf)
gnutellahosts.com/GnutellaProtocol04.pdf between August 18,
2000 and May 25, 2010, available at: [https://web.archive.org/
web/20000315000000*/http://gnutellahosts.com/
GnutellaProtocol04.pdf](https://web.archive.org/web/20000315000000*/http://gnutellahosts.com/GnutellaProtocol04.pdf) (last accessed Sept. 19, 2015)

Attachment H: Internet Archive Wayback Machine January 16, 2001 capture of
the “Bandwidth Barriers to Gnutella Network Scalability”
November 6, 2000 article posted on the Gnutella Hosts website,
http://www.gnutellahosts.com/dss_barrier.html, available at:
[https://web.archive.org/web/20010116190600/http://](https://web.archive.org/web/20010116190600/http://www.gnutellahosts.com/dss_barrier.html)


DECLARATION OF ZIAD ELKURJIE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NOS. 7,590,752 B2, 7,779,138 B2, 8,099,513 B2

www.gnutellahosts.com/dss_barrier.html (last accessed Sept. 19, 2015)

Attachment I: Internet Archive Wayback Machine February 2, 2001 capture of “Gnutella Protocol Extensions” web page of the Gnutella Hosts website, http://gnutellahosts.com/protocol_extensions.html, available at: https://web.archive.org/web/20010202155400/http://gnutellahosts.com/protocol_extensions.html (last accessed Sept. 19, 2015)

Attachment J: Internet Archive Wayback Machine December 13, 2000 capture of the “Gnutella: To the Bandwidth Barrier and Beyond” September 6, 2000 article posted on the Gnutella Hosts website, <http://www.gnutellahosts.com/gnutella.html>, available at: <https://web.archive.org/web/20001213064300/http://www.gnutellahosts.com/gnutella.html> (last accessed Sept. 19, 2015)

Attachment K: Bandwidth Barriers To Gnutella Network Scalability Slashdot, <http://slashdot.org/story/00/09/12/1217200/bandwidthbarrierstognutellanetworkscalability>, September 12, 2000 (last accessed Sept. 19, 2015)



Excite
January 1997 – December 1998 (2 years) | San Francisco Bay Area

E-commerce Databases Start-ups Mobile Applications

Software Development Product Development Scalability Online Marketing

Product Management Mobile Devices High Availability Agile Methodologies

Open Source System Architecture Web Applications [See 12+ >](#)

University of Toronto - New College
B.Sc., Computer Science
1989 – 1993



English	French
Arabic	

- See who you know in common
- Get introduced
- Contact **Ziad** directly

Not the Ziad Elkurjie you're looking for? [View more](#)

LinkedIn member directory: [a](#) [b](#) [c](#) [d](#) [e](#) [f](#) [g](#) [h](#) [i](#) [j](#) [k](#) [l](#) [m](#) [n](#) [o](#) [p](#) [q](#) [r](#) [s](#) [t](#) [u](#) [v](#) [w](#) [x](#) [y](#) [z](#) [more](#) | Browse members [by country](#)

[LinkedIn](#) © 2015 |
 [User Agreement](#) |
 [Privacy Policy](#) |
 [Community Guidelines](#) |
 [Cookie Policy](#) |
 [Copyright Policy](#) |
 [Guest Controls](#)

INTERNET ARCHIVE
WayBackMachine

[1 captures](#)
18 Aug 00 - 18 Aug 00

JUL **AUG** S
18
1999 2000 2

Alpha Release

DISTRIBUTED SEARCH SERVICES
BY **Clip2**

About Us

Clip2 is intensely passionate about enabling *open search*.

Our eponymous service at <http://www.clip2.com/> enables users to publish Web guides on the topics that interest them. Open for any user to participate on any topic, the site supports keyword searching of the guide database and qualifies guides with numerical ratings based on their usage rates. In April 2000, PC Magazine honored it as one of the Top 100 sites on the Web.

Recently, Gnutella and related technologies have opened the way for a revolutionary new generation of *distributed* open search. In the case of Gnutella, a user's search query is distributed by being passed from host to host on a network of machines that communicate using the Gnutella protocol. Search results return to the user from hosts that choose to respond. The public Gnutella network is open in the sense that any host may connect to it and respond to queries in any way that it sees fit. For a wide variety of applications, distributed search systems promise to connect information providers and consumers more efficiently than centralized systems.

Clip2 is advancing this revolution with its Distributed Search Services site at <http://dss.clip2.com/>, providing technical data to the developer and user communities interested in distributed open search systems. The site's current focus on Gnutella is reflective of the popularity of this protocol. The data we report on the Gnutella network is derived from proprietary network measurement, monitoring, and analysis technology developed by Clip2.

Clip2 is based in the Stanford Research Park just south of the Stanford campus in Palo Alto, California. Our investors include nearby VC firms [Greylock](#) and [Weiss, Peck & Greer Venture Partners](#) and prominent Silicon Valley angels. We are actively building our [team](#).

[Home](#) | [About Us](#) | [Join Us](#) | [Contact Us](#)

© Copyright Clip2.com, Inc. 2000. All rights reserved.

INTERNET ARCHIVE
WayBackMachine

93 captures
16 Aug 00 - 1 Aug 15

http://gnutellahosts.com/

Go

JUL AUG 16 1999 2000 20

DISTRIBUTED SEARCH SERVICES
BY Clip2

Alpha Release

GNUTELLA OPTIMIZED HOST LIST

Need a host? Clip2 DSS brings you the best host list available. To use it, simply add a connection to **gnutellahosts.com:6346** in your Gnutella client. Our server will fill your host catcher with the list.

What makes this list so good? The hosts in it are connected to the Gnutella network in such a way that they can reach more files than other hosts. We refresh the list *continuously* and last confirmed all hosts were responding as of 8/15/00 17:10:07 PDT.

IP ADDRESS	FILES REACHABLE	GB REACHABLE	% OF NETWORK CONTENT
141.84.102.25:6346	294,468	18,748	100.0%
195.24.30.203:6346	294,468	18,748	100.0%
134.102.206.55:6346	294,468	18,748	100.0%
200.248.248.200:6000	294,468	18,748	100.0%
147.52.113.13:6346	294,468	18,748	100.0%
165.230.77.83:6346	294,468	18,748	100.0%
18.244.1.196:6346	294,468	18,748	100.0%
163.176.52.37:6346	294,468	18,748	100.0%
207.235.120.176:5634	294,468	18,748	100.0%
148.241.61.11:5635	294,468	18,748	100.0%

GNUTELLA NETWORK STATISTICS

Stats as determined by our crawler:

Last Crawl 8/15/00 16:51:46 PDT

Hosts 1,968

Files* 294,468

Gigabytes* 18,748

Diameter 9

* As reported by network hosts

GNUTELLA SEARCH

TellaFind.com uses the Clip2 DSS host list to connect to the Gnutella network.

Search the Gnutella network via TellaFind:

Search

GNUTELLA LINKS

News gnutellanews.com

Downloads gnutelliums.com

File sharing portal zeropaid.com

De facto home gnutella.wego.com

Newsgroup alt.gnutella

GNUTELLA RESOURCES

- [Gnutella Protocol Specification v0.4 \(PDF\)](#)
- [Example Gnutella network map](#)
- [The "Napster Flood" of July 26-28, 2000](#)
- [What were Gnutella users looking for in June and August, 2000?](#)

[Home](#) | [About Us](#) | [Join Us](#) | [Contact Us](#)


© Copyright Clip2.com, Inc. 2000. All rights reserved.


INTERNET ARCHIVE
WayBackMachine

[3 captures](#)
11 Dec 00 - 2 Feb 01


<http://gnutellahosts.com/contact.html>

NOV **DEC** FEE
11
1999 **2000** 200


 Home



General Feedback dss-feedback@clip2.com
Protocol Comments dss-protocols@clip2.com
Reflector Support reflector-support@clip2.com
Resumes dss-jobs@clip2.com
Press Inquiries dss-press@clip2.com
All Other Matters corporate@clip2.com



Clip2.com, Inc.
3145 Porter Drive
Building C
Palo Alto, CA 94304



main line: 650.493.2920
main fax: 650.493.2926

© Copyright Clip2.com, Inc. 2000. All rights reserved.

INTERNET ARCHIVE
WayBackMachine

[14 captures](#)
11 Dec 00 - 18 Feb 05

http://gnutellahosts.com/articles.html


Go



NOV
11
1999


DEC
2000


FE
2001


Home



























Gnutella: To the Bandwidth Barrier and Beyond

The history of Gnutella can be divided into two eras: "pre-barrier" and "post-barrier," where the terms refer to the traffic level on the public Gnutella network relative to dial-up modem bandwidth, and the dividing line falls in August 2000. In this article, Clip2 DSS publishes previously unreleased data on the state of the pre-barrier network and provides additional information illustrating the barrier transition on which we first reported on September 8. In addition, we illuminate conditions on the poorly understood post-barrier network and examine the locations of network hosts. Finally, we introduce an analogy to describe the current state of the network and indicate future scenarios for network development. [Read it](#)

Query Stream Anomaly

Recent samples of the Gnutella query stream contained a significant surplus of four-character queries of unknown meaning. The source and reason for these queries are unknown. These anomalous queries constituted a sizable fraction of all queries and therefore contributed materially to the network average query rate. [Read it](#)

Bandwidth Barriers to Gnutella Network Scalability

The scalability of a Gnutella network to accommodate more users performing more searches is limited by the lowest bandwidth links prevalent within the network. Usage of the public Gnutella network has grown to the point that a "Dial-Up Modem Barrier" has been hit, with the result that network usability has degraded considerably. [Read it](#)

Gnutella Protocol Specification v0.4

A comprehensive guide to the Gnutella Protocol (Requires the free [Adobe Acrobat PDF reader](#)). [Read it](#)

Gnutella Network Map

A graphic visualization of the Gnutella network on July 27, 2000. [Read it](#)



The "Napster Flood" of July 26-28, 2000

News of U.S. District Judge Marilyn Hall Patel's order enjoining Napster from "copying or assisting or enabling or contributing to the copy or duplication of copyrighted songs and musical compositions..." began to reach the public late in the afternoon (PDT) on Wednesday, July 26, 2000. Thus began an exodus of refugees from the popular music file sharing service to alternatives including Gnutella. [Read it](#)



Gnutella Query Distribution

What are users looking for on the Gnutella network? Clip2 DSS sampled thousands of queries from the Gnutella network query stream to find out. [Read it](#)



© Copyright Clip2.com, Inc. 2000. All rights reserved.

The Gnutella Protocol Specification v0.4

Clip2 Distributed Search Services
<http://dss.clip2.com>
dss-protocols@clip2.com

Gnutella (pronounced "newtella") 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.

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.

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).

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

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

A servent wishing to accept the connection request must respond with

GNUTELLA OK\n\n

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.

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.

Descriptor Header

	Message ID		Payload Descriptor		TTL	Hops	Payload Length	
Byte offset	0	15	16		17	18	19	22

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 <code>Payload_Length</code> bytes from the end of this header i.e. there are no gaps or pad bytes in the Gnutella data stream.</i>

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.

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.

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

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.

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.

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.

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.*

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.

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.*

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).*

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.

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.

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.*

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.


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:

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.
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.
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.
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.
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.
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

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.

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
\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	Foobar.mp3\x00\x00

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
\r\n
```

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
```

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.

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




http://gnutellahosts.com/GnutellaProtocol04.pdf

BROWSE HISTORY

<http://gnutellahosts.com/GnutellaProtocol04.pdf>

Saved **6 times** between August 18, 2000 and May 25, 2010.

PLEASE DONATE TODAY. Your generosity preserves knowledge for future generations. Thank you.



Note

This calendar view maps the number of times <http://gnutellahosts.com/GnutellaProtocol04.pdf> was crawled by the Wayback Machine, *not* how many times the site was actually updated. More info in the [FAQ](#).



The Wayback Machine is an initiative of the [Internet Archive](#), a 501(c)(3) non-profit, building a digital library of Internet sites and other cultural artifacts in digital form. Other [projects](#) include [Open Library](#) & [archive-it.org](#).

Your use of the Wayback Machine is subject to the Internet Archive's [Terms of Use](#).

INTERNET ARCHIVE
WayBackMachine

3 captures
16 Jan 01 - 24 Apr 01

http://www.gnutellahosts.com/dss_barrier.html

Go

DEC
JAN
16
2000
2001



















Home

products

articles

events

developers



Bandwidth Barriers to Gnutella Network Scalability

September 8, 2000

Summary

The scalability of a Gnutella network to accommodate more users performing more searches is limited by the lowest bandwidth links prevalent within the network. Usage of the public Gnutella network has grown to the point that a "Dial-Up Modem Barrier" has been hit, with the result that network usability has degraded considerably. .


Queries and Average Query Rates

Typically, Gnutella users either run downloaded "servent" software that contains both client and server components, or they use one of several Web sites that provide Gnutella client capabilities without requiring any software installation. In either case, the client interface provides a means for a user to conduct a search of the network of Gnutella hosts reachable within a specified number of network hops. A Gnutella host that receives a query passes it on to all the other hosts to which the host is connected, and so on, until the query has traveled its allotted number of hops. Even if a particular host responds to a query, it still passes the query along so that others may respond. In this relaying of queries, hosts treat each other equally regardless of physical characteristics such as network connection speed or CPU power.

Clip2 DSS visited over a million hosts during a 34-day period in July and August, 2000 and counted the number of queries broadcast by each over a fixed time interval. This number included both queries originating at each host and queries being passed along by each host on behalf of others. We obtained an average query rate for each host, and then we averaged over all hosts discovered per day to obtain an average query rate for the Gnutella network for the day. The average query rate is a significant metric because it is a global measure of usage of the Gnutella network. Below, we plot the daily network average query rate:

Microsoft Exhibit 1024

1/4



The data show three trends:

- (1) The query rate grew by approximately 100% during the "Napster Flood" of July 26-28, the period during which Napster was threatened with imminent shutdown.
- (2) Over a period of days following Napster's reprieve, the query rate declined to pre-Flood levels.
- (3) The query rate hit a low on August 6 from which it has rebounded to a level comparable to that seen during the Napster Flood. The rebound included a sharp increase between August 13 and August 14.

Correlations and Scalability Barrier

We note that these trends qualitatively correlate with reports on the usability of the Gnutella network. In particular, during the Napster Flood and in the period since approximately August 14, numerous complaints were posted on Gnutella user forums. Users reported responses to their searches were fewer in number and slower to arrive than in the past. They also found their server software reporting smaller network sizes than they were accustomed to seeing.

Quantitatively, we have observed a correlation between the average query rate and the responsiveness of the network to Gnutella pings. As the average query rate increases, the fraction of hosts returning Gnutella pings in response to Gnutella pings decreases.

Based on these correlations, it appears the Gnutella network has a scalability barrier at an average query rate of approximately 10 queries per second. This barrier was hit during the Napster Flood and in the period from August 14 to the present.

Explanation of the Scalability Barrier We propose that the source of the barrier is a sufficiently prevalent and well-distributed class of hosts connected via dial-up modems at speeds below 56 kilobits per second. Below, we outline an approximate analysis supporting the plausibility of this hypothesis. Our approach is to show that even by a conservative estimation, a typical dial-up Gnutella host does not have sufficient bandwidth to effectively participate as a peer on the network when the average query rate exceeds approximately 10 queries per second.

We propose that the source of the barrier is a sufficiently prevalent and well-distributed class of hosts connected via dial-up modems at speeds below 56 kilobits per second. Below, we outline an approximate analysis

supporting the plausibility of this hypothesis. Our approach is to show that even by a conservative estimation, a typical dial-up Gnutella host does not have sufficient bandwidth to effectively participate as a peer on the network when the average query rate exceeds approximately 10 queries per second.

In a separate study, we found the most common number of connections per host to be three. A typical Gnutella query message is approximately 30 bytes long including header information, and these messages are encapsulated in messages of lower-level protocols. One TCP/IP wrapper alone requires approximately 40 bytes. To make a lower-bound estimate, we will assume one TCP/IP wrapper per Gnutella message and ignore byte contributions from other protocols that may be in play, such as PPP. Altogether, then, the bandwidth required to communicate 70-byte (560-bit) messages between 3 hosts at a rate of 10 per second is $560 \times 3 \times 10 = 16,800$ bits per second.

Gnutella pings are another sort of message broadcast on the network in the same way as queries. Gnutella hosts issue pings when they connect to other hosts in order to harvest Gnutella pongs that contain host information. Analyzing frequency of message types on the network, we have found pings to be approximately twice as common as queries, with a similar message size.

Gnutella traffic includes three other sorts of messages. Gnutella pongs are broadcast in response to pings, and query hits are routed in response to queries. We noted above both qualitative and quantitative results that query hits and pongs become less common as the average query rate increases. In addition, the older yet popular Nullsoft Gnutella v0.56 servent has been observed to broadcast yet another sort of Gnutella message ("push requests") that many other servent programs route. In sum, we estimate these other message types as a class occur at a bit rate no less than that consumed by queries, and this lower bound estimate will suffice for our purposes here.

In sum, these sources add up to at least $4 \times 16,800$ bits per second, or 67,200 bits per second, exceeding the bandwidth available to a dial-up modem. As the query rate rises into the range of 10 per second, it is clear a host connected to the network via dial-up modem will begin to fall behind in its processing of network traffic. The result would be exactly what has been observed: slow responses or no responses at all to queries and pings.

We reiterate that this is only an approximate analysis designed to demonstrate the plausibility of the hypothesis that dial-up modems are responsible for the observed scalability barrier. The plausibility is particularly strong given that we have made conservative, lower-bound estimates of several relevant quantities and still obtained a bit rate in excess of that typically achievable with a dial-up modem. To drive our point home, we note that the plausibility holds so long as the bit rate required at approximately 10 queries per second is merely in the regime of dial-up modems. We could ignore the 2X factor due to pings, gauge the total bit rate at 33,600 bits per second, and still have a reasonable argument indicting dial-up modems as the network bottleneck.

Finally, we point out that a barrier exists for each prevalent class of host connection. Were dial-up modem users to disappear from the network, the next barrier would be encountered at the next-highest-speed common class of connection, probably low-speed DSL in the low 100s of kilobits per

second range. This would accommodate growth of the average query rate by a factor of a few to 10 over the current average before the same problem would be encountered again.

[about](#) | [jobs](#) | [contact](#) | [terms of service](#)

© Copyright Clip2.com, Inc. 2001/2000. All rights reserved.

INTERNET ARCHIVE

WayBackMachine

2 captures

2 Feb 01 - 9 Mar 01

http://gnutellahosts.com/protocol_extensions.html

Go

JAN

FEB

2

2000

2001

201

Home

products

articles

events

developers



Gnutella Protocol Extensions

Clip2's [Gnutella protocol specification document](#) codifies common understanding of what constitutes version 0.4 of the Gnutella Protocol. Some developers have introduced elaborations relative to this specification, and we document these extensions below.

- [BearShare 1.3.0 and later](#)
- [Gnotella 0.73](#)

BearShare 1.3.0 and later

BearShare versions 1.3.0 (released December 27, 2000) and later place extra data in QHDs between the last double-nul-terminated filename and the servent GUID. A QueryHits descriptor received from a BearShare servent will have the following payload structure:

NumOfHits	Port	IP	Speed	ResultSet	BearShare Trailer	ServentID
-----------	------	----	-------	-----------	-------------------	-----------

and the BearShare Trailer is composed in this format:

VendorCode	OpenDataSize	OpenData	PrivateData
------------	--------------	----------	-------------

where:

- *VendorCode* is 4 bytes long depicting a vendor code in 4 case-insensitive characters. Example, "BEAR".
- *OpenDataSize* is 1 byte long containing the size of open data following this byte.
- *OpenData* Contains open data, currently 1 byte of 8 flags, one flag per bit. Currently, BearShare sets the first bit to indicate whether the servent is behind a firewall.
- *PrivateData* Contains undocumented BearShare specific data. The length of this field can be determined as follows:
QueryHits Descriptor Payload Size - (OpenDataSize + 4 + 1).

One way for developers to handle the extension is to

(1) Be aware that an incoming QHD may or may not contain additional data after the result set and before the servent GUID. No complete specification exists for the number of bytes that may be present, or their content. Use the payload-length field and count bytes as they are read from the stream to determine whether the extension bytes are present.

(2) If they are, read them from the stream, leaving 16 bytes for the servent GUID.

Microsoft

Exhibit 1024

1/2

(3) Process the QHD as usual.

Gnotella 0.73

Gnotella version 0.73 (released July 30, 2000) and possibly earlier versions place extra data in QueryHit descriptors (QHDs). According to Gnutella 0.4, each element of the result set in a QHD is supposed to be terminated by a double-nul. Gnotella may place extra data between the two nuls. The technical specification of this additional data is unknown, although it has been stated to represent the bitrate of MP3 encoding of the preceding file referenced in the result set. If the file referenced is not an MP3 file, it is not known what, if any, extra data is inserted.

Some servents may process Gnotella's extended QHDs without adverse consequences by simply disregarding data between nuls. Others may, upon not finding an element of the result terminated as expected, improperly read the descriptor. Once misread, it may be subsequently mishandled, and the connection that delivered it may be disconnected.

For information or feedback, send a note to dss-feedback@clip2.com.

[about](#) | [jobs](#) | [contact](#) | [terms of service](#)

Copyright © 2001 Clip2.com, Inc. All rights reserved.



NOV 13 DEC 14

[18 captures](#)
13 Dec 00 - 30 Aug 05

Gnutella: To the Bandwidth Barrier and Beyond

November 6, 2000

Overview

In this report, Clip2 DSS describes the evolution and present condition of the Gnutella peer-to-peer file-sharing network based on substantial data gathered over a five-month period. Significantly, we find the network has neither smoothly scaled nor catastrophically collapsed since average traffic grew to regularly exceed dial-up modem bandwidth in August 2000. Instead, the network persists in a fragmented state comprised of numerous continuously evolving responsive segments, the largest of which typically contains hundreds of hosts. We estimate at present that unique Gnutella users per day number no less than 10,000 and may range as high as 30,000. We suggest that further technical innovation and wide adoption of this innovation are necessary for the Gnutella network to scale beyond its present state.

The fraction of hosts providing files for download has varied between 15% and 40%, rebounding in October following a substantial slide in September. The majority of hosts are online for minutes to hours, with a minority component (measured at 30% in early August) connected on timescales of a day or longer. Users continue to query the network primarily for audio, video, image, and program files, and automated indexing schemes regularly constitute a few percent of query traffic.

In a large-scale survey, one out of three network hosts was found outside the US-centric top-level domains, with Germany and Japan dominating the non-US population. Internet Service Provider second-level domains accounted for most of the top sources of hosts on the COM and NET domains, and hosts on the @Home and Road Runner broadband services were especially prevalent. Gnutella hosts were found at over 550 institutions in the EDU domain, with MIT and Virginia Tech leading a broadly distributed pack.

Contents by Section

- Introduction
- Gnutella Genesis
- Pre-Barrier Gnutella
- Gnutella Hits the Wall
- Gnutella Beyond the Barrier
- Global Gnutella
- Gnutella Tomorrow

Contents by Frequently Asked Question

- How "concentrated" around one another are hosts on the Gnutella network?
- How long do hosts stay online?

- [What happened when Gnutella passed the bandwidth barrier? \(animation\)](#)
- [Do low servent host counts mean Gnutella lost a lot of users?](#)
- [Did spam push Gnutella over the dial-up bandwidth barrier?](#)
- [How many Gnutella users are there today?](#)
- [How big are the largest usable network segments?](#)
- [What can be done by and for dial-up users?](#)
- [How many users provide files for download?](#)
- [What are users looking for?](#)
- [How much of an international phenomenon is Gnutella?](#)
- [Aside from the US, in what countries is Gnutella most used?](#)
- [In the US, how do COM, NET, EDU, etc. compare?](#)
- [What ISPs within the COM domain have the most Gnutella users?](#)
- [What ISPs within the NET domain have the most Gnutella users?](#)
- [What universities within the EDU domain have the most Gnutella users?](#)
- [What is next for Gnutella?](#)

Introduction

The history of Gnutella can be divided into two eras: "pre-barrier" and "post-barrier," where the terms refer to the traffic level on the public Gnutella network relative to dial-up modem bandwidth, and the dividing line falls in August 2000. In this article, Clip2 DSS publishes previously unreleased data on the state of the pre-barrier network and provides additional information illustrating the barrier transition on which we first [reported](#) on September 8. In addition, we illuminate conditions on the poorly understood post-barrier network and examine the locations of network hosts. Finally, we introduce an analogy to describe the current state of the network and indicate future scenarios for network development.

Gnutella Genesis

On March 14, 2000 at 8:31 AM PST, a [message](#) posted on [Slashdot](#) spread the news that AOL's [Nullsoft](#) division had released an "open source Napster clone" named "Gnutella." On March 15 at 1:25 PM PST, Wired News [reported](#) that Nullsoft's distribution of "a file-sharing software tool which could be even more potent than Napster" had ceased. Nonetheless, third parties redistributed copies of Gnutella downloaded from Nullsoft, and in extraordinarily short order a number of "Gnutella clones" appeared to augment the "official" Nullsoft version. Clones spoke the same language - the Gnutella protocol - as the Nullsoft application, and could therefore communicate with each other and the Nullsoft version. As users began to run these programs and connect them to one another via the Internet, a network comprised of a mixture of Gnutella-speaking applications formed. Because these applications acted simultaneously as both servers and clients in a fully decentralized system, the term "servent" was adopted to describe them. In addition, "Gnutella" came to have several meanings:

- Gnutella = the servent produced by Nullsoft. A program identified only as "Gnutella" (often "Gnutella v0.56") has seen wide distribution; one venue alone, [Download.com](#), reports approximately 250,000 downloads to date.
- Gnutella = the protocol, of which Version 0.4 is in widespread use. Clip2 DSS has published a [protocol specification document](#).
- Gnutella = the network. At any time there is a single major publicly accessible network of Gnutella-speaking applications ("the" network) and a number of smaller and fully (and often private) disconnected networks. Clip2 DSS tracks the major network, and we report related information on our [home page](#).
- Gnutella = any combination of the above, or the system as a whole.

In this article, we discuss the evolution of Gnutella, the network.

Pre-Barrier Gnutella

Clip2 DSS began conducting systematic studies of the Gnutella network in June 2000. Most Gnutella servents provide a "host count" feature in which they report to the user a dynamically updated count of the hosts they have identified on the network. During the period from June through mid-August, servents connected to the public Gnutella network for at least a matter of minutes typically reported host counts on the order of 1,000 to 4,000 hosts. Clip2 DSS's Gnutella network crawler regularly found 1,000 to 8,000 hosts online during the course of a sub-1-hour full-network traversal. Were these hosts connected in a tight concentration or in a loosely knit and far-flung configuration? How many connections did a typical host have? We were able to answer these questions using the network graphs generated by our crawler.

The "concentration" of the network is a particularly interesting question because Gnutella servents typically issue queries with a "TTL" of 7, meaning a user's query will travel up to 7 hosts away on the Gnutella network from the originating computer. There is always a (not necessarily unique) shortest path between any two hosts on the network, and the longest such path is known in mathematical graph theory as the "diameter" of the network. Clearly, if the diameter were greater than $(7*2+1)=15$, there would be hosts from which a query could be launched and not reach the entire network before expiring. We saw such high-diameter networks in early July. Below, we show data for a crawl of 1,959 hosts on July 7. Here, we plot on a semi-log scale the number of pairs of hosts that had a given shortest-path distance (or "separation") between them. The diameter of the network discovered by this crawl was 22, indicating some regions were not in communication with others (assuming messages used the common TTL value). We also note that most pairs of hosts were separated by 7 hops.



In the latter days of the pre-barrier period, the network diameter fell to smaller values, typically 8 or 9. Below, we show network diameters for crawls made between July 25 and August 18. Note that while the period of larger diameters around July 28 corresponds to the "[Napster Flood](#)", a larger diameter network is not necessarily a direct result of more hosts connecting to the network. A network with few hosts can have a large diameter and vice versa; the diameter is purely a function of how the hosts are connected, not the number of hosts. The smaller diameters are significant in that they imply any host on the network could easily reach all other hosts using TTL=7.



Further examining host connectivity, we found a host was most likely to have a single connection, and hosts with higher numbers of connections were increasingly uncommon. In more technical terms, we saw roughly power-law degree distributions, where "power-law" means the number of hosts having a given degree varied as a power of the degree, and "degree" is shorthand for the total number of incoming and outgoing connections a given host has open. In addition to host separations, the degree distribution is another means of quantitatively assessing network connectivity, and we present below the degree distribution based on a 1,813-host crawl made on July 7. One particularly interesting coincidence is in that other researchers (e.g., [Broder et al. 1999](#)) have found power-law degree distributions for graphs in which the nodes are Web pages and the connections are Web links. In the plot, we compare measured data with a least-squares best fit of slope (power-law index) = -2.3.



As the above data show, the structure of the Gnutella network is in a continuous state of flux. Hosts come and go as quickly as users open and close Gnutella applications, and connections between hosts may only last for seconds. We found that half of an initial host population persisted after five hours, and that approximately 30% of the initial host population was stable on the timescale of 24 hours. We determined this result by comparing host populations found in a succession of crawls to an initial reference crawl and by repeating this analysis for multiple reference crawls; the plot below illustrates our findings.



Gnutella Hits the Wall

As we first reported in "[Bandwidth Barriers to Gnutella Network Scalability](#)" (September 8), average Gnutella network traffic began to regularly exceed the throughput capacity of dial-up modems in August. Hosts connected to the Internet via dial-up modems ceased to be able to effectively participate as peers on the Gnutella network. These hosts essentially became dead-ends, resulting in a widespread fragmentation of the Gnutella network into effectively disconnected components comprised of hosts with higher-speed Internet connections. The animation below illustrates the evolution of the network as traffic passed the dial-up barrier.



Data gathered by Clip2 DSS clearly illustrates the effective loss of dial-up hosts from the responsive portion of the Gnutella network. In one long-term experiment, we regularly visited hosts and issued probe-like Gnutella "ping" messages of TTL=2 to discover their neighbors. As shown below, unique hosts sending Gnutella "pong" messages in response to these small-TTL pings declined substantially and permanently in mid-to-late August.



As noted earlier, most Gnutella servents display a host count based on the number of pongs received since the application began running. As the barrier was passed, numerous [postings](#) on public user forums noted servents were displaying lower-than-usual host counts, often only in the tens or hundreds. Many users interpreted the decrease in host counts as meaning the total number of hosts on the network had decreased. However, as the above data show, the change in network responsiveness was rather abrupt. Such an abrupt transition is much more plausibly explained by the reaching of a technical barrier than a mass change in user behavior. Had users departed the network, we would have expected to see a decrease in the usage rate of the Clip2 DSS host list service; instead, we saw an unabated rise in usage. In addition, even though responses in the ping experiment had dropped, this experiment continued to find ten to twenty times the number of hosts that would be reported in a servent's host counter over a similar time interval. In sum, we found no evidence supporting a user exodus and multiple indications that the host population remained sizable but fragmented. The total number of hosts online can be substantially larger than the number reported in a servent's host counter, since a servent only sees out as far as the boundary of the responsive region to which the servent is connected.

The preceding discussion begs the question of the source of the traffic that

caused the network to reach the dial-up modem bandwidth barrier. Was it the result of continued growth in the number of Gnutella users, or was it the result of the introduction of programmatic sources of traffic, such as machine-generated spam? The question is difficult to answer due to a lack of comprehensive data. Clip2 DSS reported on one anomalous form of traffic seen on the network in early September that may have existed for some time prior. Since that report, we have regularly observed other forms of apparently automated messages on the network, including repeated series such as {a.mp3, b.mp3, c.mp3, ...} that appear to be network-indexing attempts. While this is an unresolved question, it is moot in a sense, because with continued growth in the user base, user-generated network traffic would have eventually reached the barrier level of its own accord.

Gnutella Beyond the Barrier

What is the state of the post-barrier Gnutella network? Among other sources, we can find some clues in data from the Clip2 DSS-operated "gnutellahosts.com" host list service, which publishes IP addresses of live Gnutella hosts. Approximately 10% of users access this list by visiting the Clip2 DSS home page; the remainder retrieve addresses by connecting their servents to a special-purpose Gnutella server operated by Clip2 DSS at gnutellahosts.com, port 6346. After responding to the incoming Gnutella servent with multiple IP addresses, the gnutellahosts.com server disconnects, and the servent proceeds to connect directly to the hosts at the provided addresses. As noted in the previous section, we have not observed any decrease in the traffic to gnutellahosts.com due to Gnutella having hit the barrier. On the contrary, traffic has continued to grow in the post-barrier period. Below, we show the long-term (3-month) straight-line trend in gnutellahosts.com usage.



How many users are there on the post-barrier network? How are they connected? From the number of callers to gnutellahosts.com and our probing experiments, we found no evidence of a sudden population collapse as the barrier was passed. Instead, we found evidence that the population had fragmented into multiple dynamically changing responsive and unresponsive segments. The sum of all data sources leads us to estimate that the total number of daily users of Gnutella numbers between 10,000 and 30,000, where the lower bound is a much better approximation than the upper bound. Below, we show the numbers of hosts in the largest responsive network segments we were able to identify over a range of post-barrier dates.



Since the fragmentation occurred, it has been a matter of chance whether or not a user manages to find and remain connected to a responsive segment. Typically, the gnutellahosts.com host list service has provided addresses of hosts in the largest identifiable responsive segment, although this region is a moving target. In order to track it, Clip2 DSS has refined its crawling strategy in recent weeks and regularly crawls the network on the timescale of every 15 minutes.

On the post-barrier network, dial-up modem users cannot effectively participate as peers throughout the network. What can be done to alleviate this situation? One solution is to connect these users to high-speed proxies that handle network traffic on their behalf. This is an underlying concept of the Clip2 Reflector(TM), a special-purpose Gnutella server, and

the network architecture that results is illustrated below:



Reflectors are programmed to maintain connectivity to the most responsive segment of the network by calling gnutellahosts.com (by default). Dial-up users singly connect to Reflectors that in turn maintain multiple outgoing network connections on their behalf. A list of running public-access Reflectors can be found on the [Clip2 DSS home page](#).

How different is user behavior in the post-barrier period relative to the pre-barrier period? One measure of behavior is the fraction of hosts serving a non-zero number of files. In a 24-hour period in early August, during the late pre-barrier era, researchers at Xerox PARC found 30% of hosts made available one or more files for download ([Adar & Huberman 2000](#)). Using a different methodology, Clip2 DSS independently measured the "serving fraction" before, during, and since the period of the PARC study. Our results confirm theirs during the period of their study. In the final days of the pre-barrier era, we observed an increase in the serving fraction to a maximum in excess of 40%. However, as the network evolved into the post-barrier period, we saw a substantial decrease in the serving fraction, down to a low of less than 15%. Notably, since early October we have observed a general rise in the serving fraction back to near pre-barrier levels. Below, we plot the serving fraction over time.



What are users searching for in the post-barrier period? Clip2 DSS analyzed three query stream samples of varying sizes taken on three different dates. Applying a subjective analysis to categorize 2,000 queries heard on September 19, we found the following breakdown:



Notes on these categories: The "gibberish" category includes queries consisting of non-alphanumeric characters; among other sources, we have seen such queries generated by poorly programmed clients that do not properly read and forward Gnutella query messages. The "automated indexing" category includes queries that appeared to be programmatically generated with the intent of indexing network content, such as "a.mp3", "b.mp3", etc. "File extension only" queries are just that, containing extensions such as "mp3" or "mpg" (possibly with accompanying periods, asterisks, or both) but no other content. "Song and artist+song" counts queries either containing only a song title or a song title along with an artist name. "Artist" counts queries containing just an artist name.

By objectively analyzing only those queries containing a file extension anywhere in the query string, we are able to analyze much larger data sets. We found from 20% to 40% of queries in three separate stream samples of 2,000, 30,000, and 150,000 queries contained a popular file extension somewhere in the query string. Below, among the set of queries that contained an extension, we show the frequencies of specific types of extensions.



Note the programmatically generated queries of the form "a.mp3", "b.mp3", etc. were a recurring feature in every sample, and the samples were spread over a 21-day period. These queries likely originated from a single

source and amounted to a few percent of total network query traffic.

Global Gnutella

We complete our survey of the Gnutella network by examining the access points of Gnutella hosts. Where are Gnutella hosts? To arrive at an answer, we utilized 3.3 million non-unique IP addresses gathered continuously by Clip2 DSS between July 27 and November 3, spanning both Gnutella eras. Of these addresses, 1.3 million (39%) were resolvable to non-numeric hostnames, and our reports below are on this resolvable subset. The populations we report below should be interpreted in probabilistic terms; the data have not been de-duplicated, so that the relative populations represent relative probabilities of host discovery within a given domain.

We divided top-level domains into US-Centric (COM, NET, EDU, MIL, GOV, US) and Non-US-Centric (all others) categories, although we note a number of non-US-based organizations operate domains in the US-Centric set. Our first finding is that Gnutella is a truly international phenomenon, with one out of three hosts located on a non-US-centric domain.



Among Non-US-Centric domains, 95% of hosts were found in just 17 country-specific top-level domains. European domains comprised 67% and the Asia-Pacific domains comprised 20%



Among US-Centric hosts, COM, NET, and EDU domains predictably dominated, although non-zero numbers of hosts were found on each of ORG, US, GOV, and MIL domains (in ratios 19:8:2:1, respectively).



In the case of the COM, NET, and EDU domains, we dug deeper to examine populations at the level of second-level domain names. The popular second-level COM domains were primarily broadband Internet service providers. The ISP @Home accounted for half of all Gnutella hosts in the COM domain, and Road Runner trailed in second place at nearly one quarter of hosts. Gnutella hosts with resolvable host names in the COM domain were therefore strongly concentrated in a small number of second-level domains, with 87% of hosts residing in the top 25. Among the top 25 were six second-level domains either representing non-ISP companies or organizations whose nature we could not determine. In total, Gnutella hosts were found on 1750 unique second-level domains within the COM domain.



The popular second-level NET domains were exclusively Internet service providers, both broadband and dial-up. The distribution among NET domains was less concentrated than among COM domains, with the leader, Road Runner, claiming less than 10% of Gnutella hosts on the NET domain. Note the second-place second-level NET domain was a German company, illustrating that while the NET domain is US-centric, it is not US-exclusive. Over 2000 unique second-level domain names were represented in the Gnutella host population.



The distribution of hosts among EDU domains was even less concentrated than among NET domains. The Massachusetts Institute of Technology led the pack at 3.8% by a sizable margin. Virginia Tech made a strong showing at slightly less than 3%, and the distribution declined smoothly from 3rd-place University of Southern California through the remainder of the list. In all, Gnutella hosts were discovered on over 550 second-level domains within the EDU domain.



In summary, major conclusions are that (1) the Gnutella network is an international phenomenon led by the US, Germany, and Japan; (2) substantial populations of hosts in COM and NET domains are on ISP second-level domains; and (3) hosts are widely distributed among EDU domains.

Gnutella Tomorrow

The Gnutella network is analogous to a continuous global rave: an informal, decentralized, unregulated gathering without a permanent location. Network hosts, like rave attendees, come and go unpredictably, connecting and disconnecting as fast as ravers switch dance partners. In the pre-barrier era, the Gnutella rave could accommodate all comers. The effect of the network hitting the dial-up modem barrier is analogous to a rave venue reaching capacity, with many would-be revelers being crammed shoulder-to-shoulder and left unable to dance, and still more spilling out the doors. Small regions within the crowd, corresponding to responsive segments of the Gnutella network, remain sufficiently open to enable movement. These pockets form and vanish, grow and shrink, and merge and split on a variety of timescales. In the post-barrier era, the Gnutella rave remains more popular than the typical raver, pressed in among the crowd, might realize. The crowding results in the potential of the gathering being released isolated bursts rather than in a continuous widespread discharge. While Gnutella has not scaled, we call attention to the fact that it has also not collapsed. Like many simple decentralized systems, it has remained remarkably robust in the face of technical adversity. In the present post-barrier period, Gnutella exists in an intermediate state between scaling and collapsing.

In the opinion of Clip2 DSS, this situation is probable to persist until either (1) the user population collapses and traffic falls to pre-barrier levels or (2) an "organizing principle" for connectivity that enables scaling takes root. In the former case, the improved performance that would result could potentially drive alienated users to return to the network, driving resurgence in traffic, another barrier crossing, and repetition of the entire cycle. In the latter case, examples of organizing connectivity principles include (1) dial-up users regularly connecting to broadband [Reflectors](#) and (2) widespread adoption of servents with sophisticated and consistently implemented connection management rules. However, to be widely and rapidly successful, any organizing principal must require no user action or change in behavior that is not immediately and powerfully rewarded, and it must not involve a change to the protocol that breaks the considerable installed application base. In the post-barrier era, there have been various initiatives to create new and smaller networks - spin-off raves - enabled by user adjustment of the "Gnutella handshake" mechanism in servents that support this feature. However, because the underlying technology is no different, if traffic on these networks were to grow sufficiently large, they would be subject to bandwidth barriers as well. These attempts to re-create pre-barrier conditions do so at the cost of sacrificing the relatively large user base on the main network and do not directly address the problem. To move beyond its present state, Gnutella awaits widely adopted technical

innovation.

© Copyright Clip2.com, Inc. 2000. All rights reserved.

Slashdot

[Stories](#)
[Firehose >](#)
[All Topics](#)
[Devices Popular](#)
[Build Polls](#)
[Entertainment Video](#)
[Technology Jobs](#)
[Open Source Deals](#)

[Search](#)

[Login](#) or [Sign up](#)

Follow [Slashdot blog](#) updates by [subscribing to our blog RSS feed](#)

Nickname:

Password: 6-20 characters long

☐ **Public Terminal**

[Forgot your password?](#)

Sign in with

Google

Facebook

Twitter

LinkedIn

[Close](#)

Build. Deploy. Manage. Find your
SAP HANA on the cloud.

JUST A VISIT AWAY

IBM Cloud marketplace



Slashdot Deals: 92% OFF to becoming a MSCE with the [Microsoft Certified Solutions Engineer: Server Infrastructure Bundle](#)

Bandwidth Barriers To Gnutella Network Scalability

Posted by [Hemos](#) on Tuesday September 12, 2000 @09:16AM from the all-must-have-high-speed dept.

[tinla](#) writes "[Clip2](#) have written a [mini-report](#) based upon sampling the Guntella Network during Aug and Sept. Includes the famous Napster Flood (when [Napster almost ceased](#)). Most telling is that Gnutella Network may have hit its scalability limit ALREADY thanks to all the modem connections searches must transverse."

For kicks, I've added a poll to the side - what do you use?



You may like to read:



[Hawking Warns Strong AI Could Threaten Humanity](#)

[The Brainteaser Elon Musk Asks New SpaceX Engineers](#)

[Is Advertising Morally Justifiable? The Importance of Protecting Our Attention](#)

[Debunking a Viral Internet Post About Breastfeeding Racism](#)

[Continued Cord Cutting Hits the Pay TV Business Hard](#)

This is the most recent story. Help us pick the next by [voting on submissions](#), or [submit your own](#).

REGISTER NOW

SAVE \$200

REGISTER HERE

Oct. 25-29, San Francisco
Discover All Things Java

ORACLE

This discussion has been archived. No new comments can be posted.

[Bandwidth Barriers To Gnutella Network Scalability](#) [More Login](#)

[Bandwidth Barriers To Gnutella Network Scalability](#)

Archived Discussion Load All Comments

Full Article Archived Log In Create an Account

Comments Filter:

- [Score:](#)
- [Insightful](#)
- [Informative](#)
- [Interesting](#)
- [Funny](#)

The Fine Print: The following comments are owned by whoever posted them. We are not responsible for them in any way.

0

[What about Freenet? \(Score:1\)](#)

[Moldyorgbunnet \(222311\)](#)

shouldn't freenet be on the list for this poll? Their spaces seem to me to overlap significantly.

0

[Re: What about Freenet? \(Score:1\)](#)

[Public Terminal Warns Strong AI Could Threaten Humanity](#)

574 comments [The Brainteaser Elon Musk Asks New SpaceX Engineers](#)

496 comments [Debunking a viral internet Post About Breastfeeding Racism](#)

351 comments [Continued Cord Cutting Hits the Pay TV Business Hard](#)

[Log In](#) [Forgot your password?](#)

[Close](#) [Close](#)

[next](#)

This is the most recent story. Help us pick the next by [voting on submissions](#), or [submit your own](#).

[previous](#)

0 comments



[Slashdot](#)

Archived Discussion

[Get more comments](#)

1 of 1 loaded

[Submit Story](#)

After the last of 16 mounting screws has been removed from an access cover, it will be discovered that the wrong access cover has been removed.

[FAQ](#)

[Story Archive](#)

[Hall of Fame](#)

[Advertising](#)

[Jobs](#)

[Terms](#)

[Privacy](#)

[Opt Out Choices](#)

[About](#)

[Feedback](#)

[Mobile View](#)

[Blog](#)

Trademarks property of their respective owners. Comments owned by the poster. Copyright © 2015 SlashdotMedia. All Rights Reserved. **DHI**

Slashdot is a [DHI service](#).

[Close](#)

[Slashdot](#)

Working...