Peer2Peer - Networks

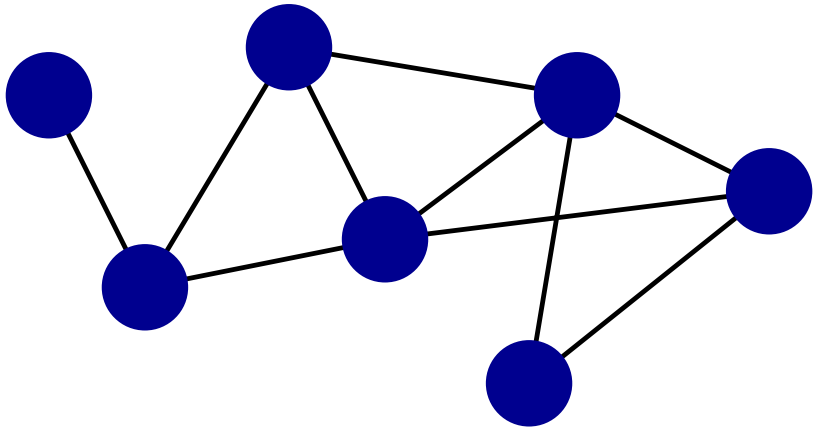

The New Internet

Teleseminar SS 2002

Dirk Zugenmaier

dirk@zugenmaier.de

Inhalt I

- Grundidee von Peer2Peer-Netzen
- Servergestütztes Peer2Peer
- Peer2Peer in Reinform
- Gnutella-Protokoll
- Kritik am Gnutella-Netzwerk
- Verbesserungen
- Quellen

Grundidee von Peer2Peer

- Verteilte Daten zugänglich machen
- Verteilte Netzlast
- Mutliple Datenhaltung

Servergestütztes Peer2Peer

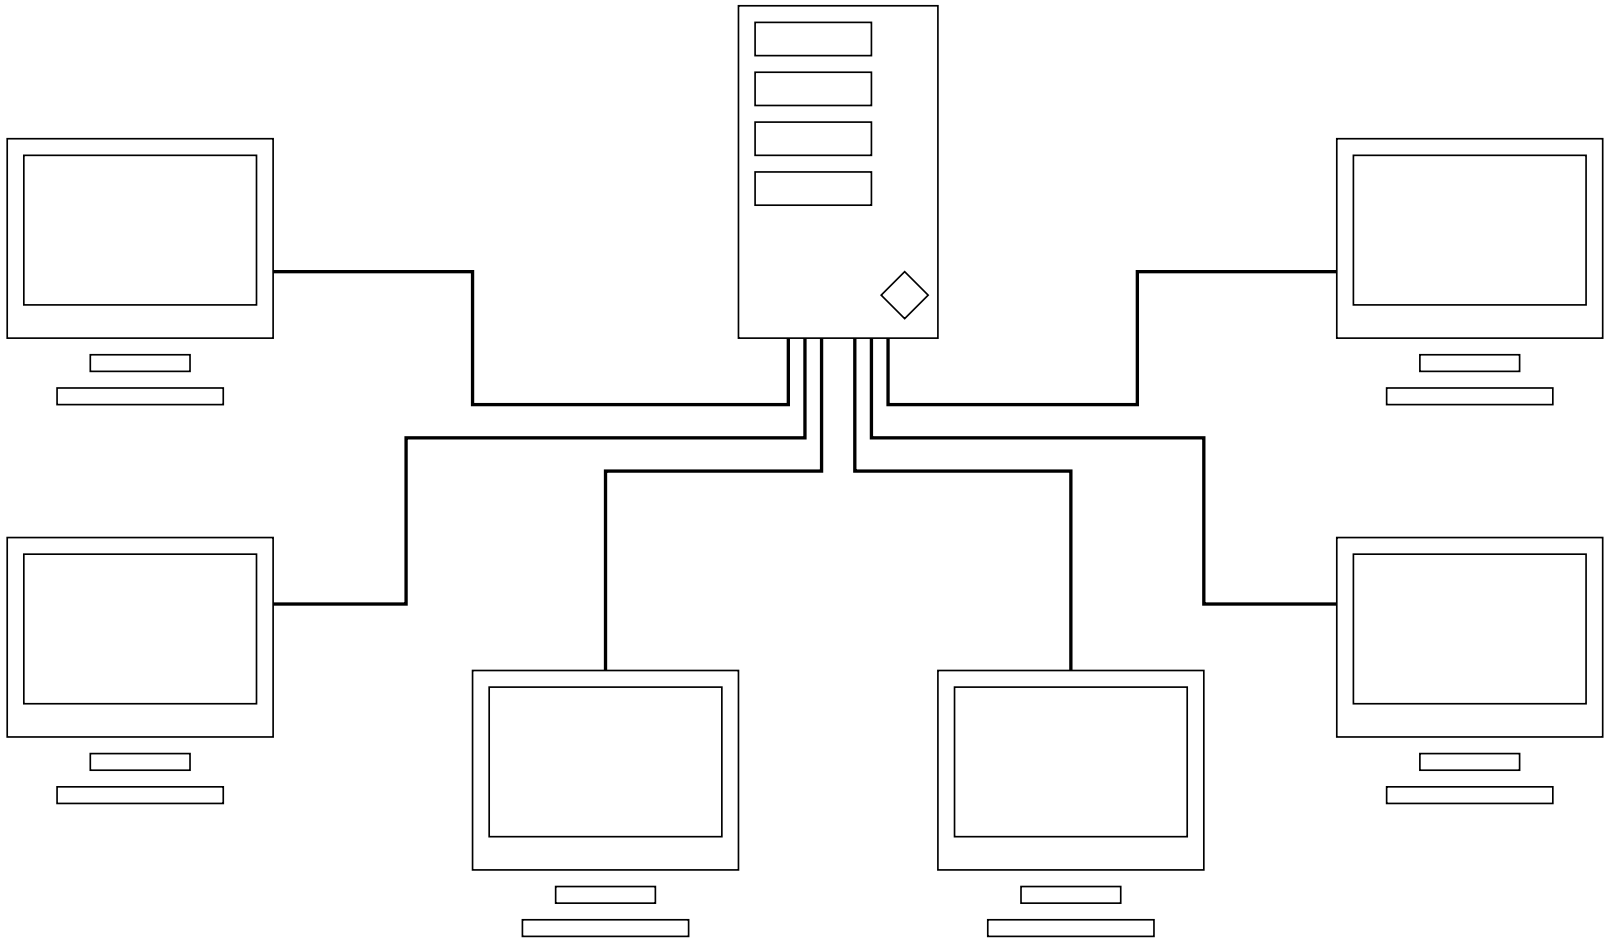

Topologie

Servergestütztes Peer2Peer

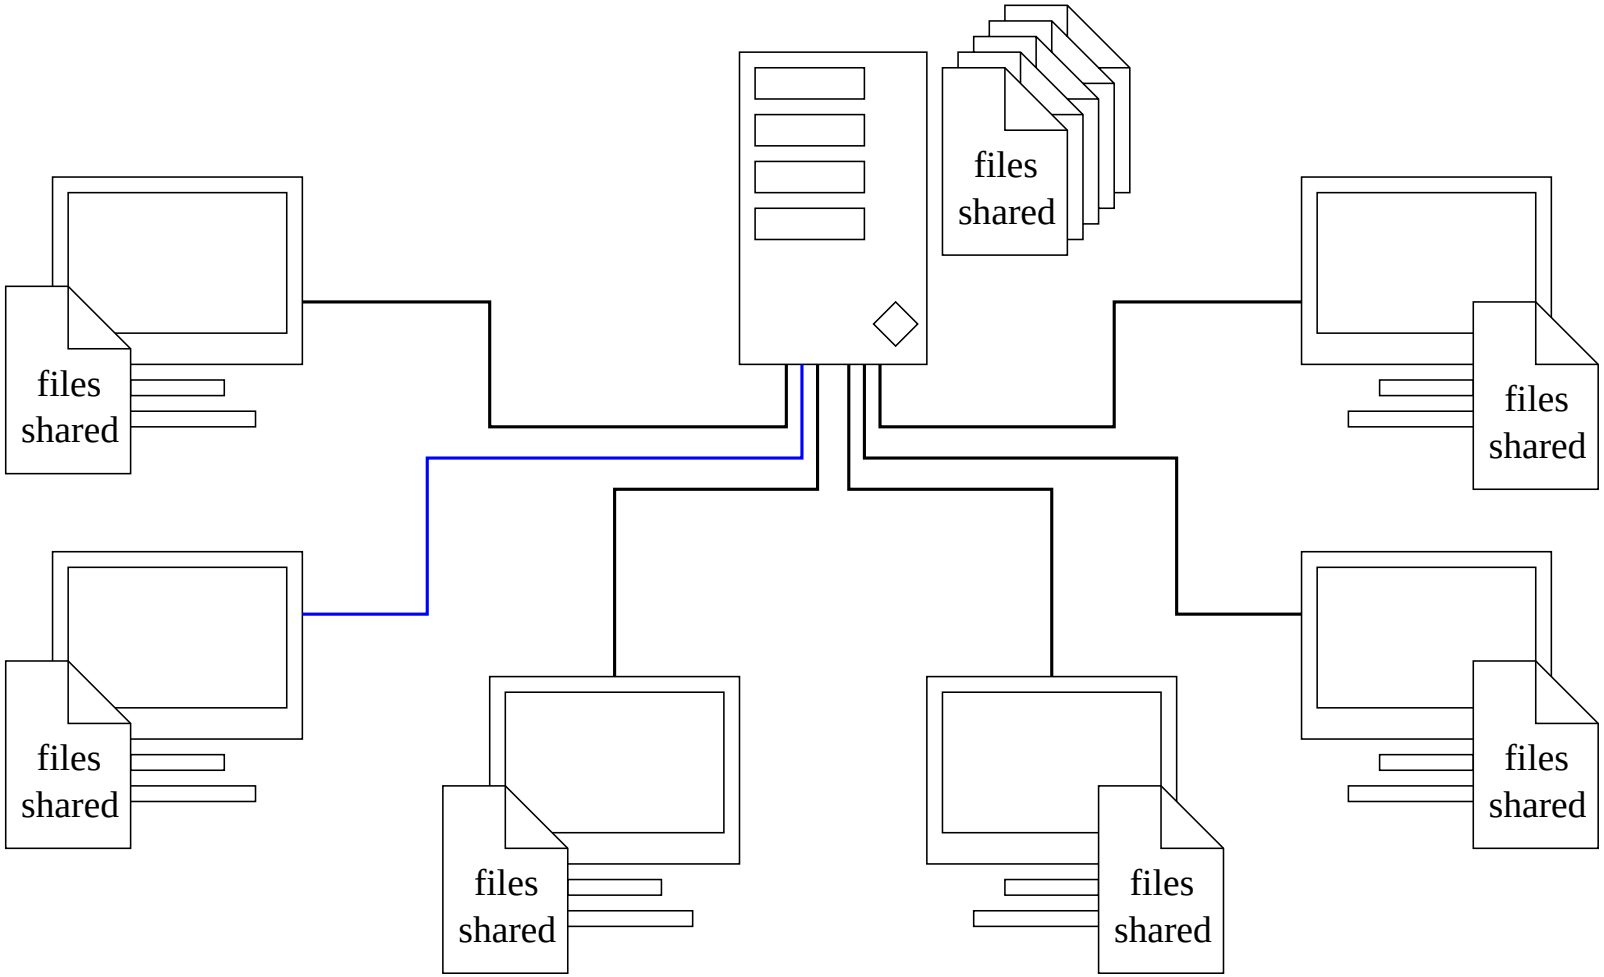

Suchanfrage

Servergestütztes Peer2Peer

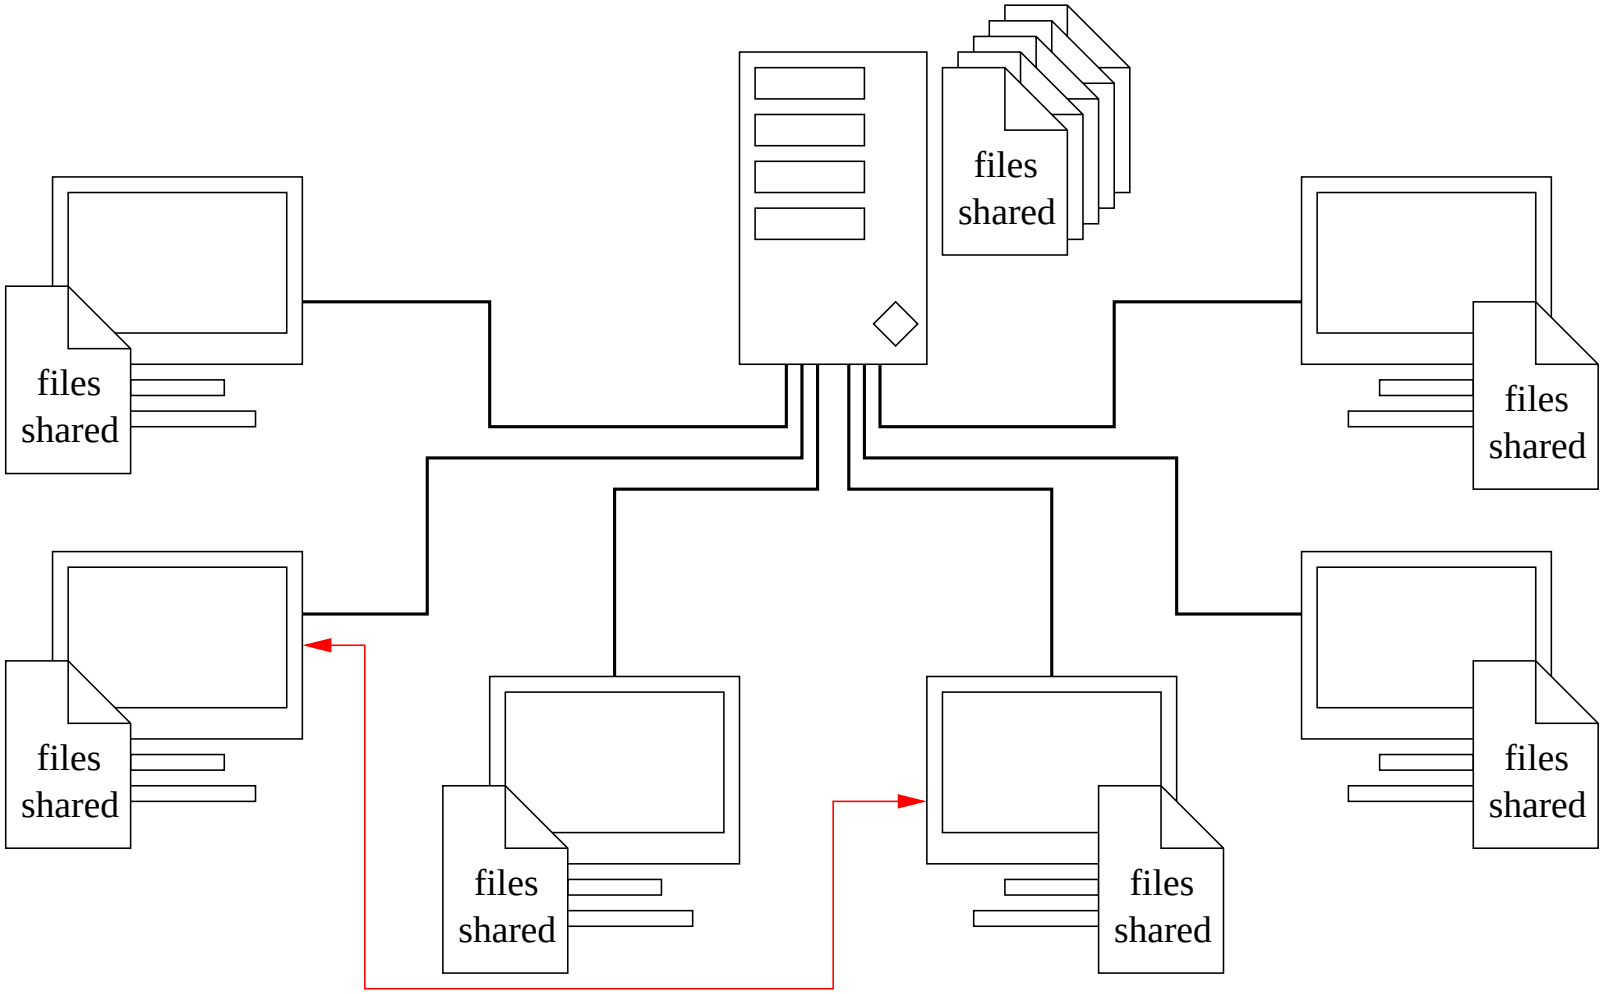

Filetransfer

Reines Peer2Peer

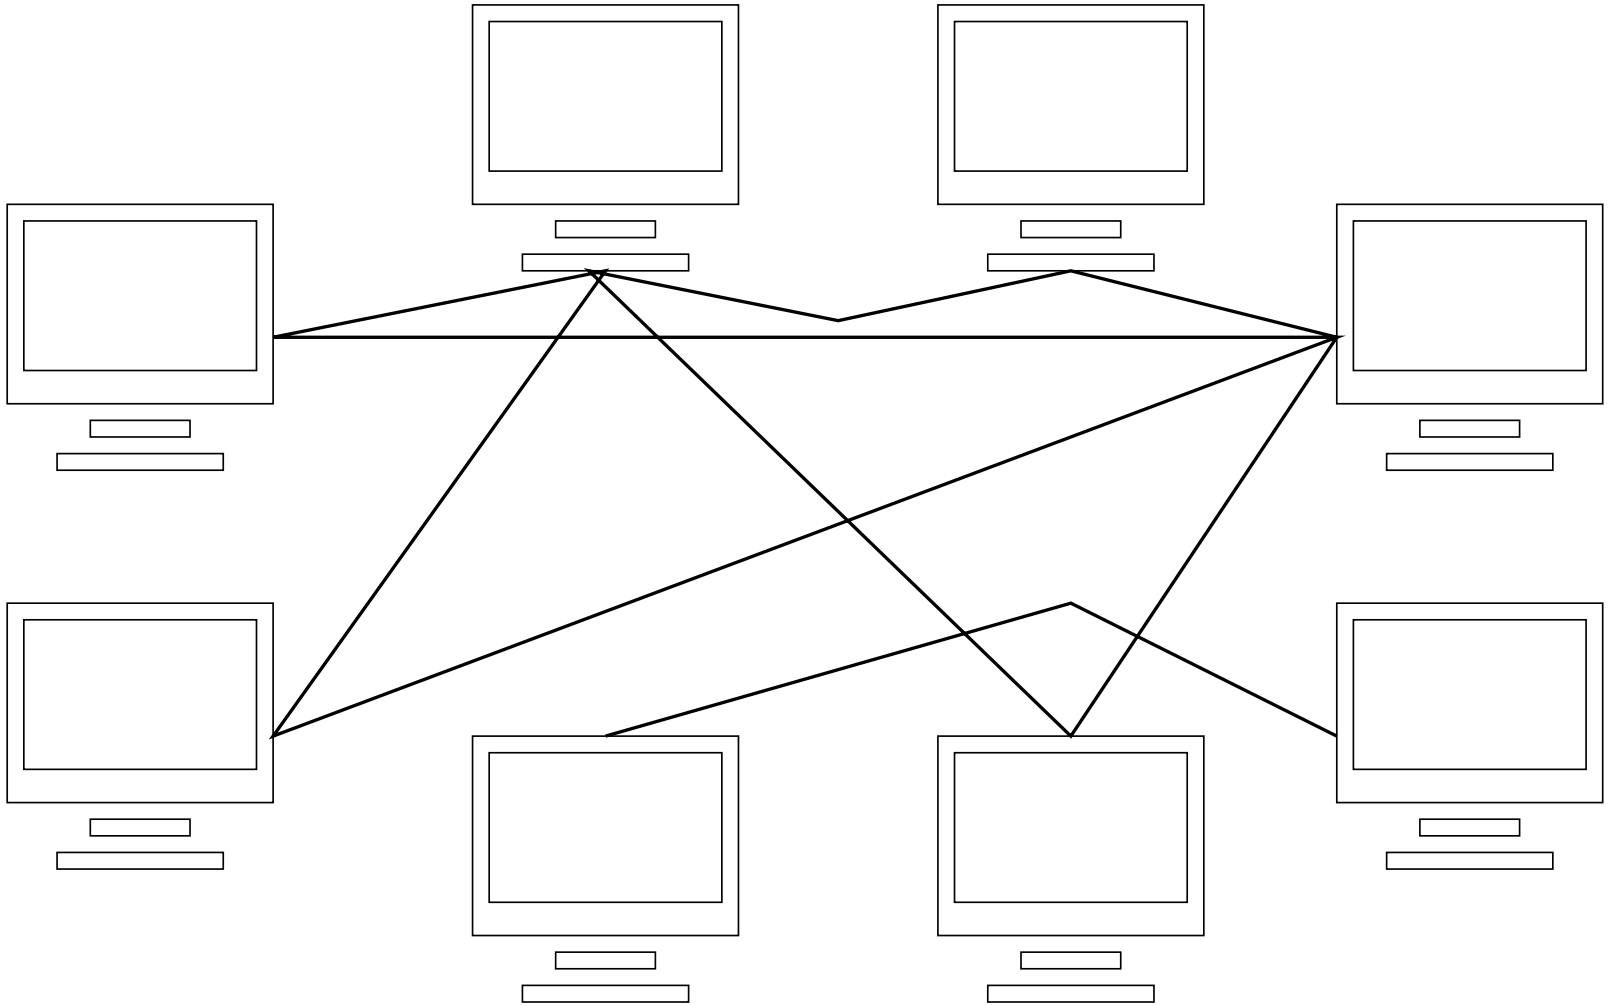

Reines Peer2Peer

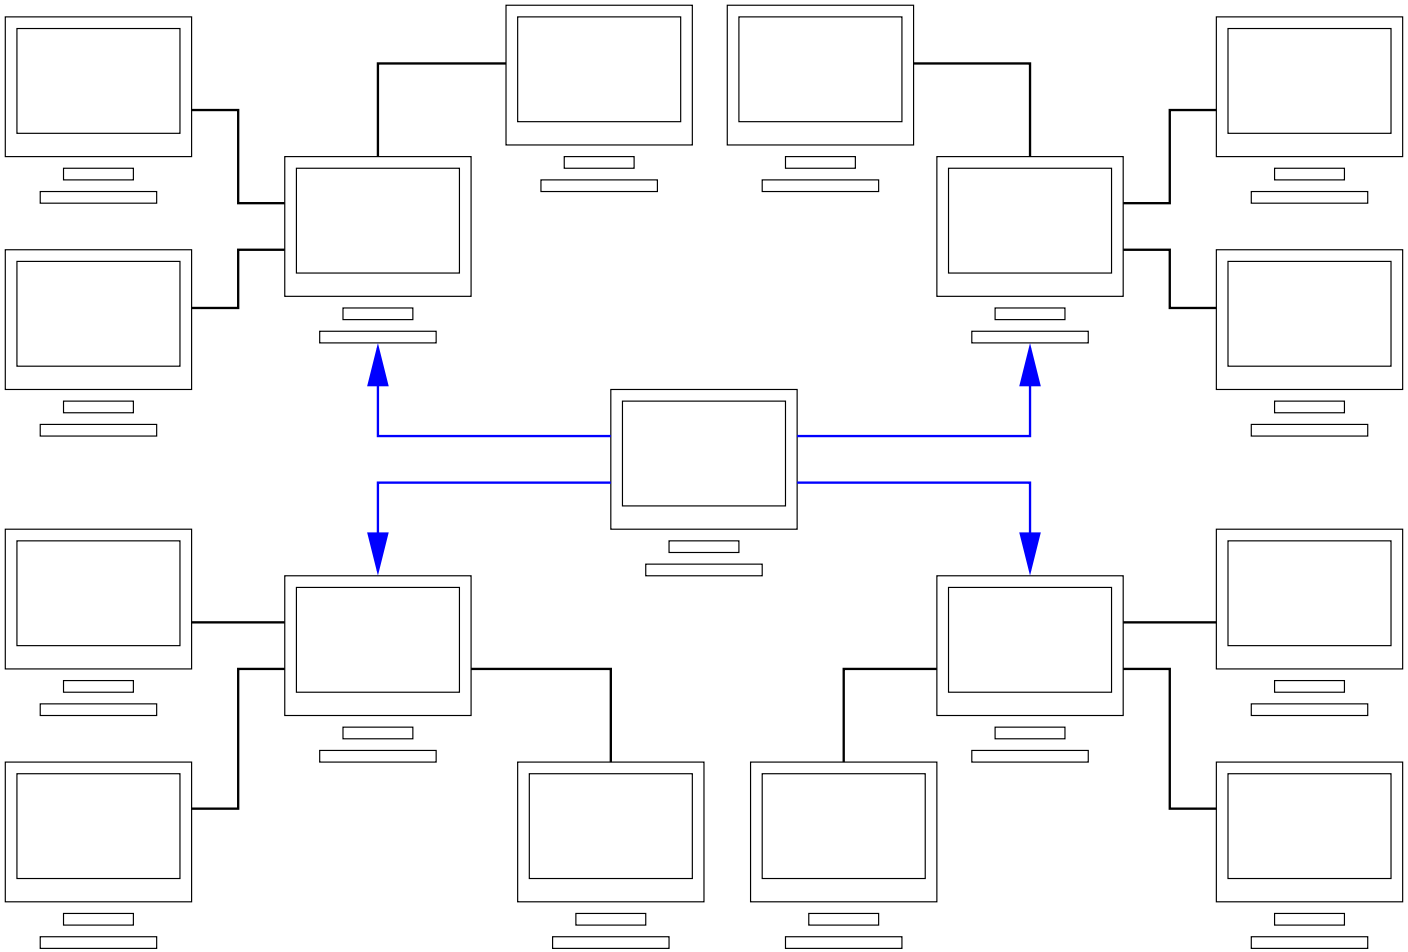

Suchanfrage

Reines Peer2Peer

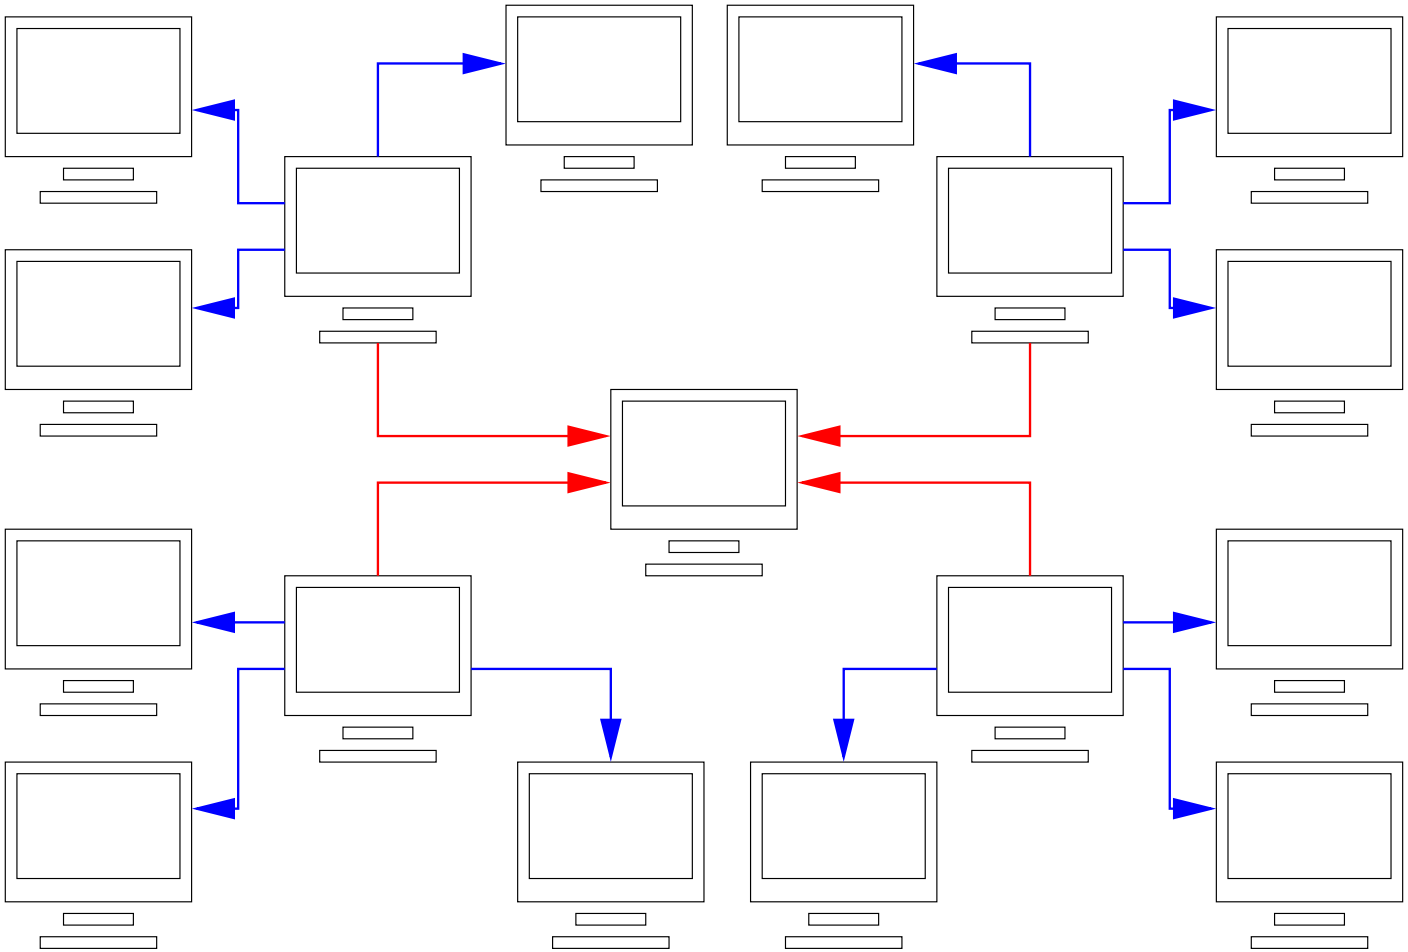

Suchanfrage / Antwort

Reines Peer2Peer

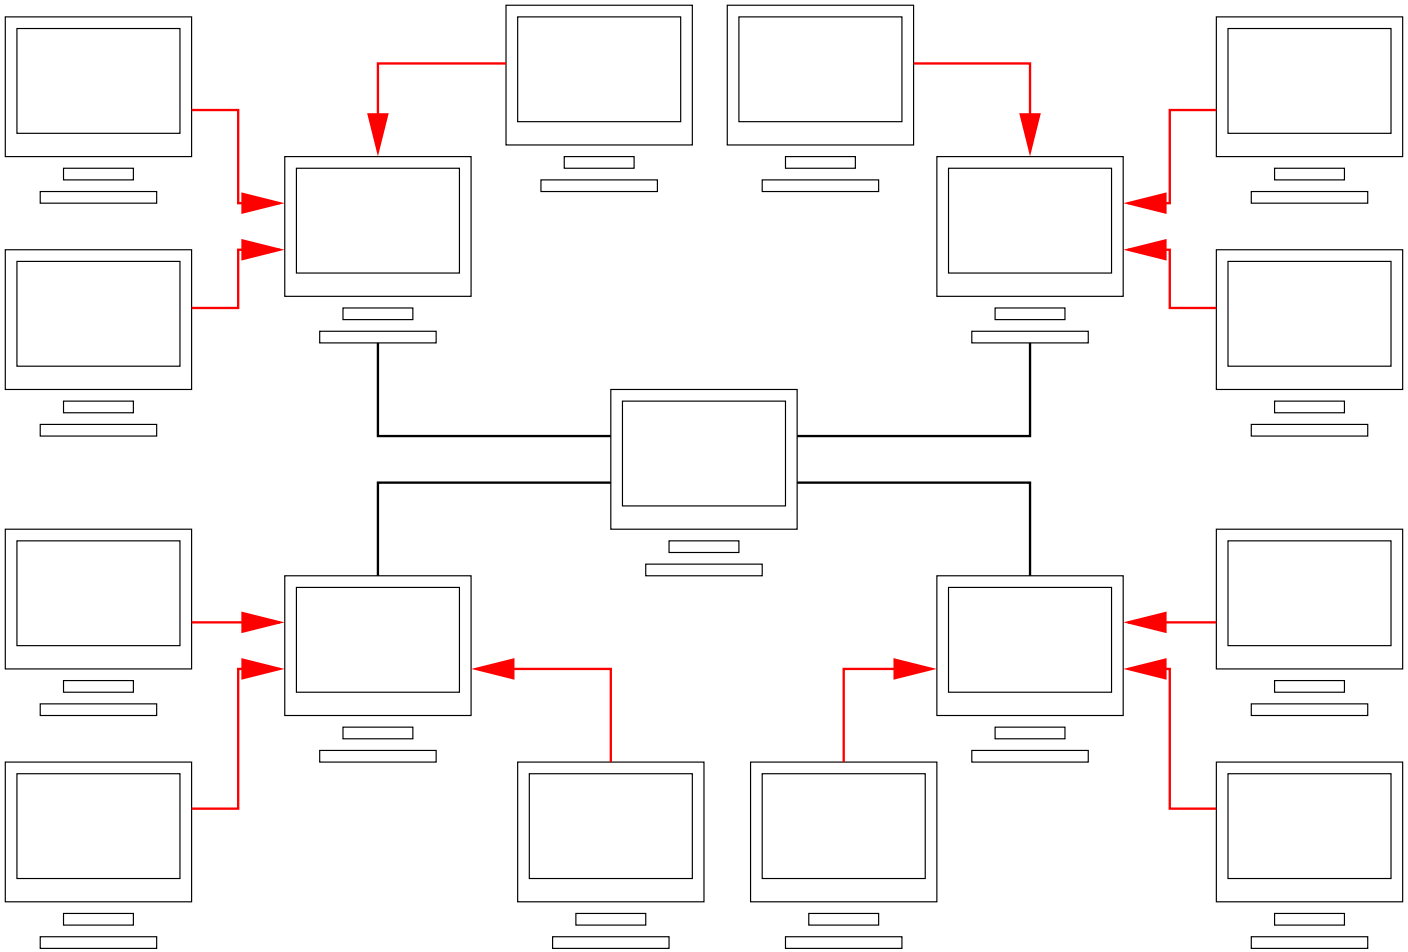

Suchanfrage / Antwort

Reines Peer2Peer

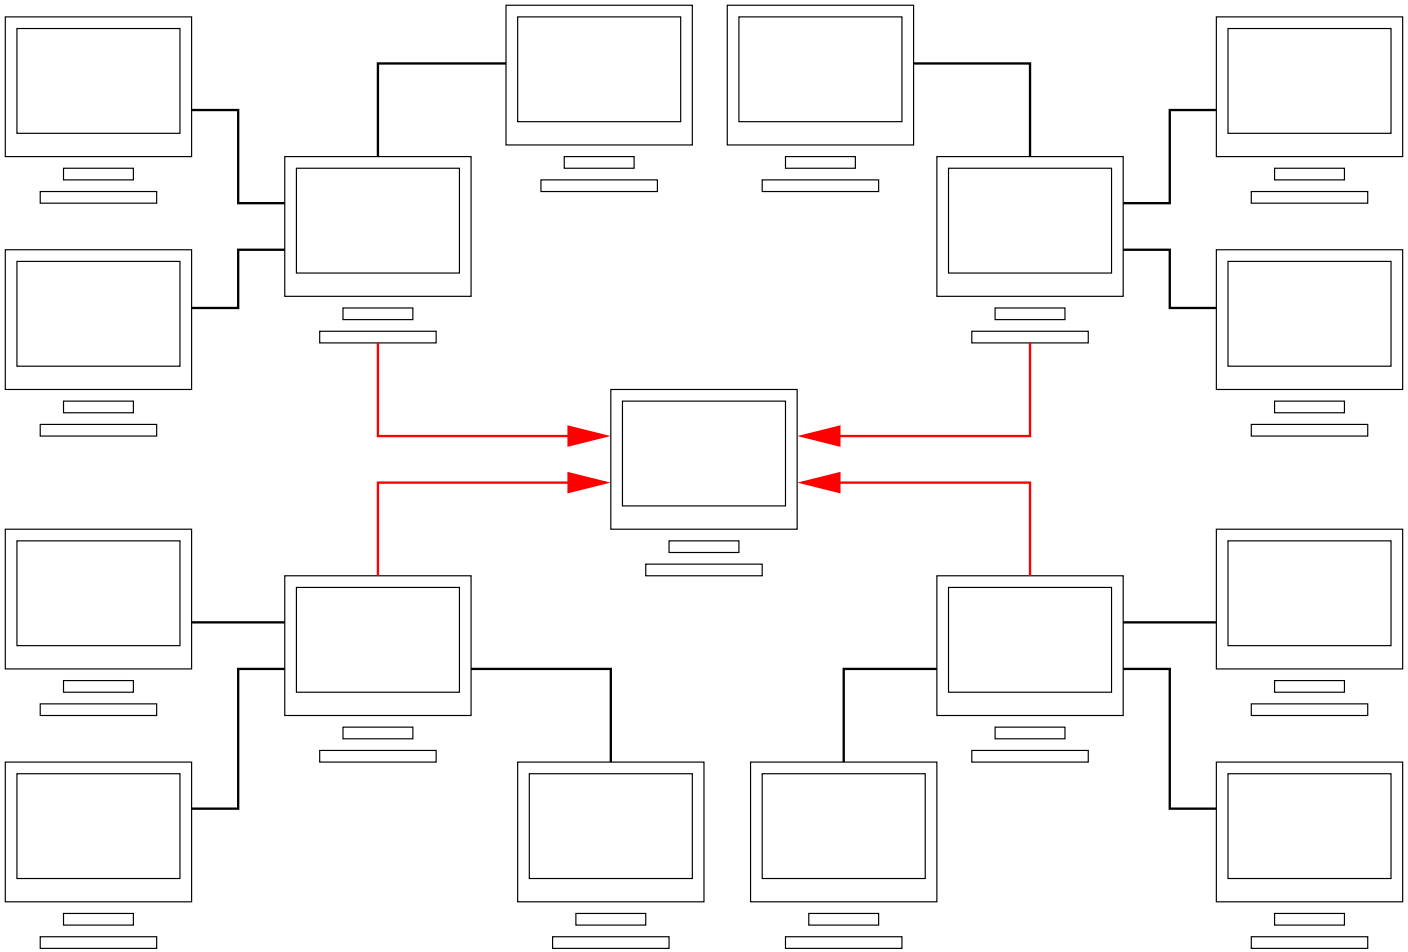

Suchanfrage / Antwort

Firewallproblem

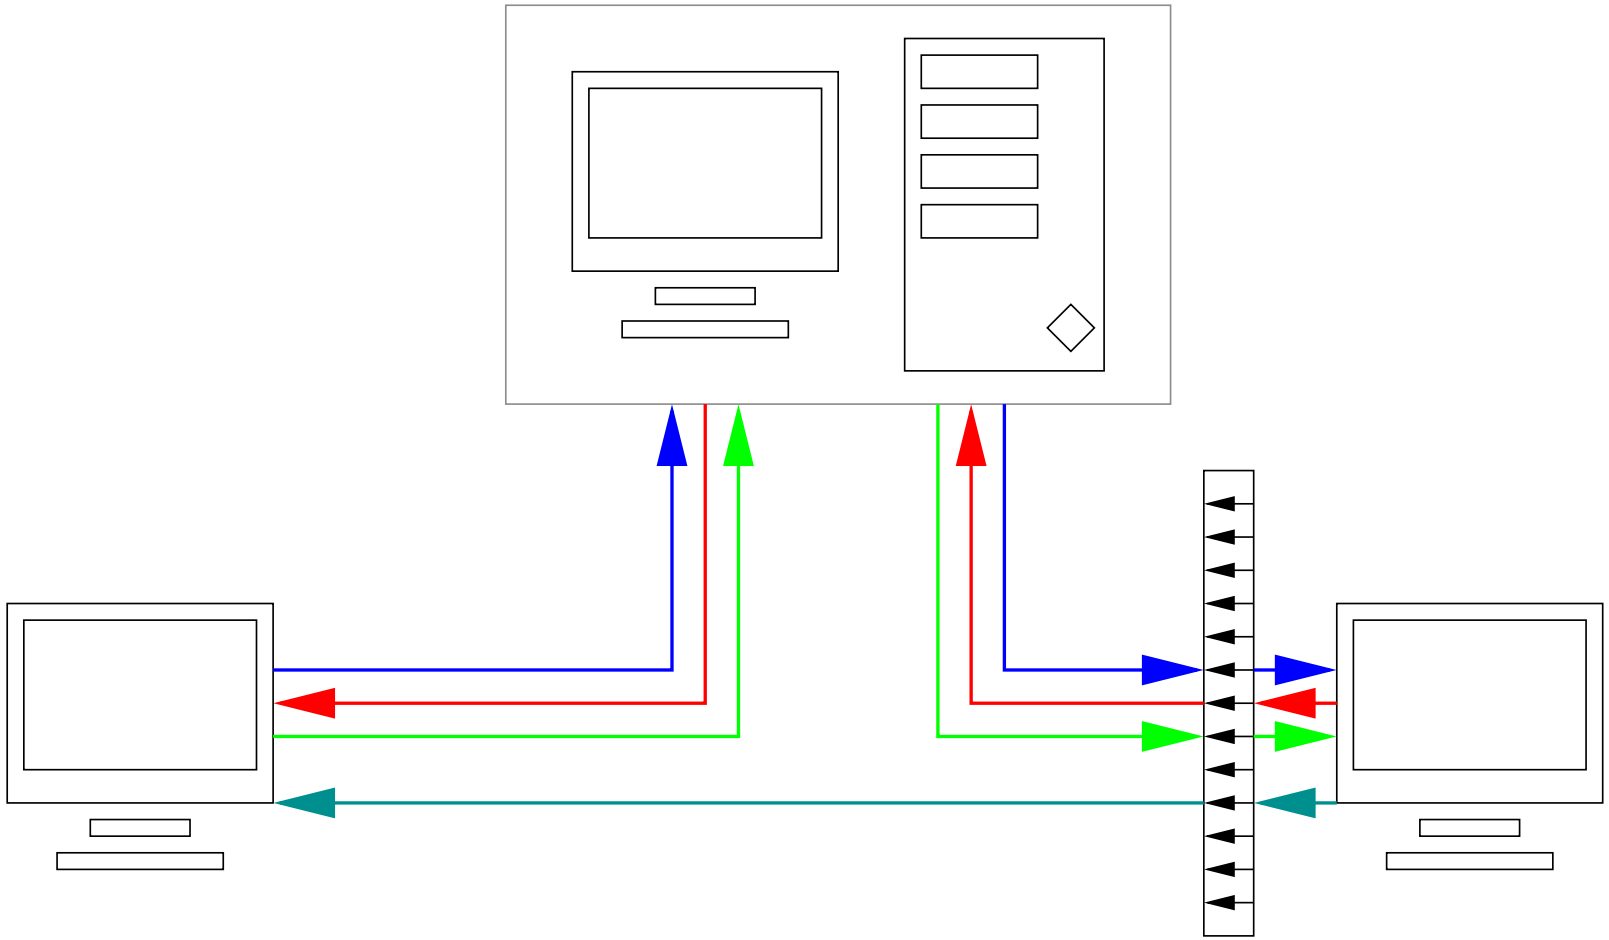

Suchanfrage / Antwort / Push-Request

Content

Napster:

- nur mp3-Dateien
- Suchbegriffe aus ID3-Tags

Gnutella

- alle Dateienformate
- Suchbegriffe aus Dateinamen

Gnutella-Protokoll

Das Gnutella-Protokoll verfügt über folgende Nachrichtentypen:

- PING
- PONG (PING-Response)
- Query
- Push-Request
- Hits(Query-Response)
- Request (via HTTP)

Gnutella Protokoll

Header:

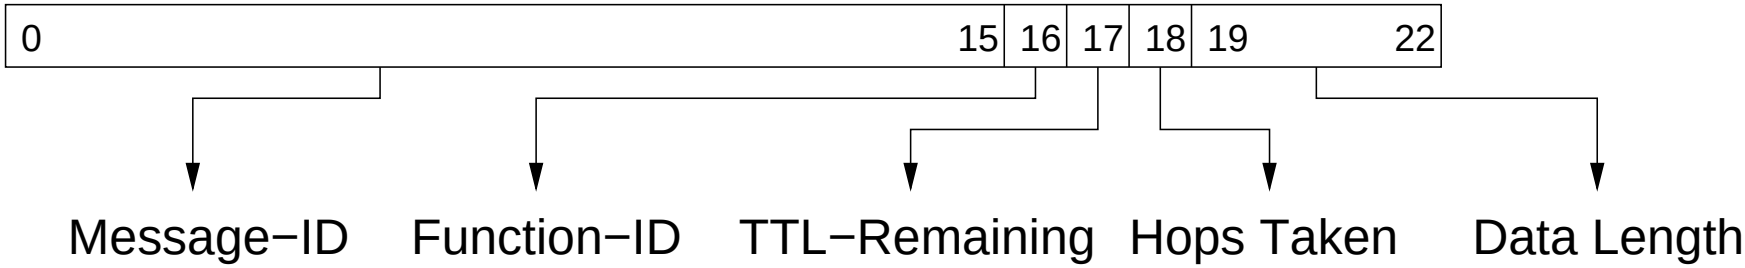

Ping:

kein Nachrichten-Body

Pong:

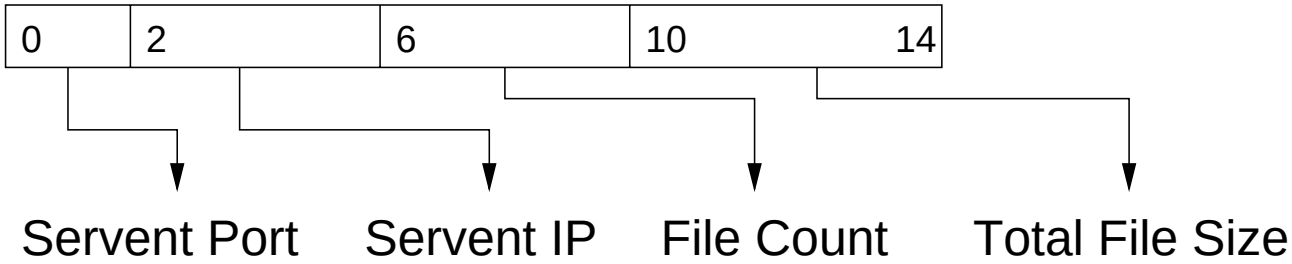

Gnutella Protokoll

Query:

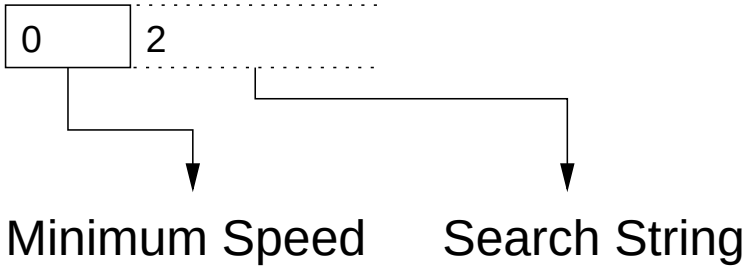

Pushrequest:

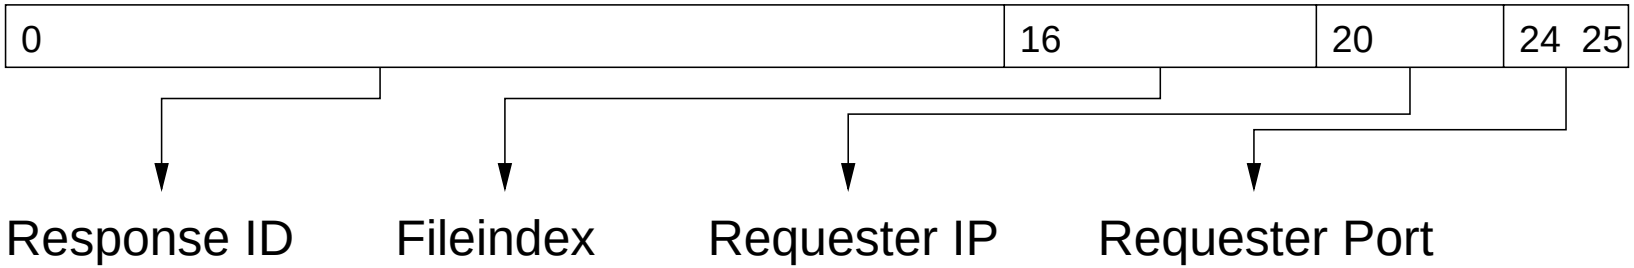

Gnutella Protokoll

Hits:

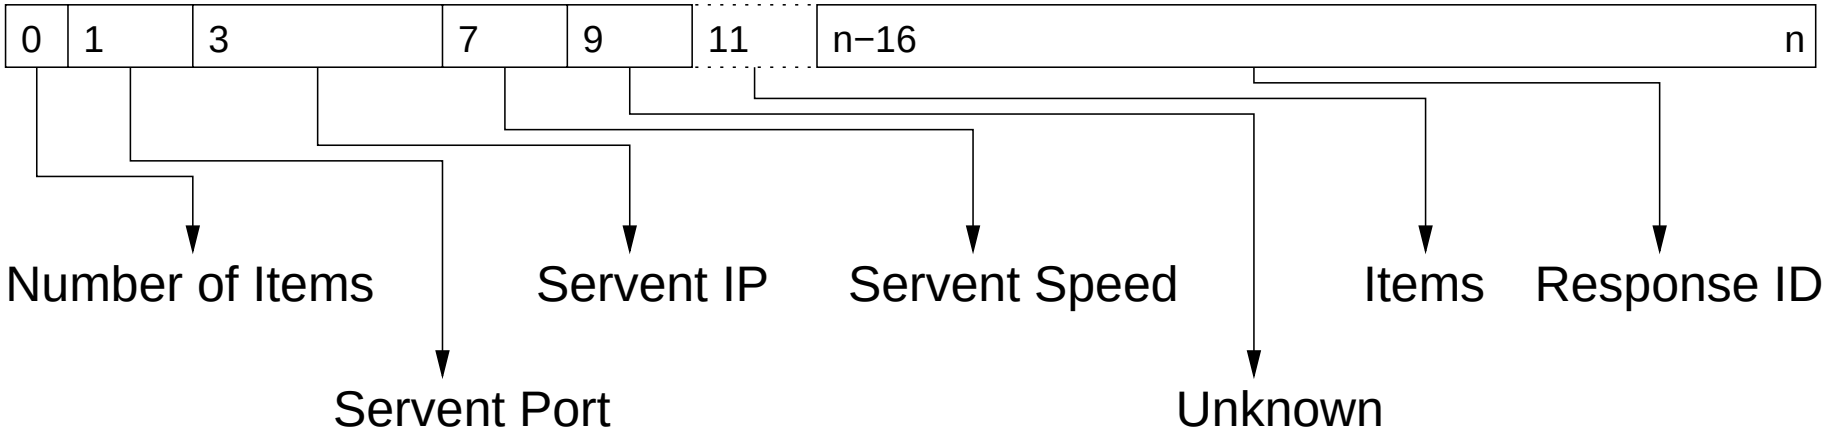

Items:

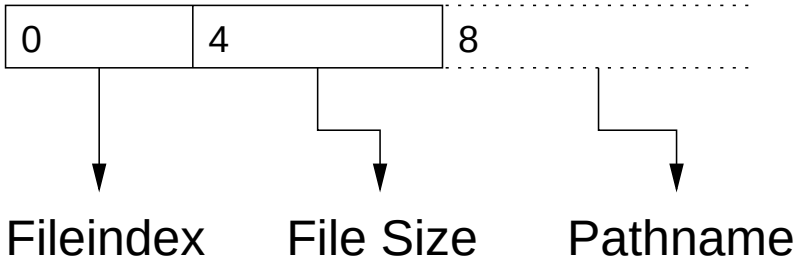

Erreichbare Benutzer

Annahme: Jeder Nutzer habe N offene Verbindungen, $T =$ Time to Live der Query-Pakete. Die erreichbaren Benutzer P errechnen sich folgendermaßen:

$$P(N,T) = N * f(N-1, T-1)$$

Mit:

$$f(n,t) = n^t + f(n-1, t-1)$$

Erreichbare Benutzer

N\T	1	2	3	4	5	6	7	8
2	2	4	6	8	10	12	14	16
3	3	9	21	45	93	189	381	765
4	4	16	52	160	484	1.456	4.372	13.120
5	5	25	105	425	1.705	6.825	27.305	109.225
6	6	36	186	936	4.686	23.436	117.186	585.936
7	7	49	301	1.813	10.885	65.317	391.909	2.351.461
8	8	64	456	3.200	22.408	156.864	1.098.056	7.686.400

N = Anzahl ausgehender Verbindungen je Knoten

T = TTL (Time To Live für gesendete Pakete)

Paketgrösse: Suchanfrage

IP-Header	20 bytes
TCP-Header	20 bytes
Gnutella-Header	23 bytes
Minimum Speed	1 byte
Search String	α bytes
Gesamtgrösse	$64 + \alpha$ bytes

Sei die Suchanfrage "*The Offspring .mp3*"

$\Rightarrow \alpha = 18$ (17 Zeichen und $\backslash 0$)

$\Rightarrow$ 83 bytes Gesamtgrösse.

Datentransfer der Suchanfrage

N \ T	1	2	3	4	5	6	7	8
2	166	332	498	664	830	996	1.13kB	1.30kB
3	249	747	1.7kB	3.65kB	7.54kB	15.32kB	30.88kB	62.01kB
4	332	1.3kB	4.21kB	12.97kB	39.23kB	118.02kB	354.37kB	1.04MB
5	415	2.03kB	8.51kB	34.45kB	138.20kB	553.20kB	2.16MB	8.65MB
6	498	2.92kB	15.08kB	75.87kB	379.82kB	1.86MB	9.28MB	46.38MB
7	581	3.97kB	24.40kB	146.95kB	882.28kB	5.17MB	31.02MB	186.13MB
8	664	5.19kB	36.96kB	259.38kB	1.77MB	12.42MB	86.92MB	608.42MB

Ausgehender Datentransfer in Byte (Suchanfrage mit $17 + 1$ Bytes)

$\Rightarrow$ max 637 971 200 Byte $=\sim$ 608.42MB für eine Suchanfrage

Paketgrösse: Antwort

IP-Header	20 bytes
TCP-Header	20 bytes
Gnutella-Header	23 bytes
Number of Hits	1 byte
Port	1 byte
IP-Address	4 bytes
Speed	3 bytes
Result Set	$\beta * (8 + \gamma + 2)$ bytes
Servent Identifier	16 bytes
Gesamtgrösse	$88 + (\beta * (10 + \gamma))$ bytes

Mit $\beta =$ Anzahl der Dateien, $\gamma =$ Länge des Dateinamen

Datentransfer der Antwort

Der Datentransfer der Anwort lässt sich berechnen wie folgt:

$$D_{Antwort} = P * (\delta * \varepsilon) * (88 + (\beta * (10 + \gamma)))$$

Mit:

$\delta =$ Prozentualer Anteil der “Sharing Peers“

$\varepsilon =$ Prozentualer Anteil der antwortenden Peers.

Datentransfer der Antwort

Mit $\beta = 5, \gamma = 10, \delta = .3$ und $\varepsilon = .5$:

N \ T	1	2	3	4	5	6	7	8
2	101.4	405.6	912.6	1.6kB	2.5kB	3.6kB	4.9kB	6.3kB
3	152.1	912.6	3.1kB	8.9kB	23.0kB	56.1kB	132.0kB	303.0kB
4	202.8	1.6kB	7.7kB	31.7kB	119.8kB	432.5kB	1.48MB	5.07MB
5	253.5	2.5kB	15.6kB	84.2kB	422.1kB	1.98MB	9.24MB	42.25MB
6	304.2	3.6kB	27.6kB	185.4kB	1.13MB	6.80MB	39.66MB	226.65MB
7	354.9	4.9kB	44.7kB	359.1kB	2.63MB	18.95MB	132.65MB	909.57MB
8	405.6	6.3kB	67.7kB	633.8kB	5.42MB	45.51MB	371.65MB	2.90GB

Datentransfer der Antwort

Mit $\beta = 10, \gamma = 60, \delta = .3$ und $\varepsilon = .4$:

N \ T	1	2	3	4	5	6	7	8
2	189.1	756.5	1.7k	3.0k	4.6k	6.6k	9.0k	11.8k
3	283.7	1.7k	5.8k	16.6k	42.9k	104.7k	246.3k	565.1k
4	378.2	3.0k	14.4k	59.1k	223.5k	806.7k	2.76M	9.47M
5	472.8	4.6k	29.1k	157.0k	787.2k	3.69M	17.24M	78.80M
6	567.4	6.6k	51.5k	345.7k	2.11M	12.68M	73.97M	422.72M
7	661.9	9.0k	83.4k	669.7k	4.91M	35.34M	247.39M	1.66G
8	756.5	11.8k	126.3k	1.15M	10.10M	84.88M	693.15M	5.42G

Angaben in xB.

Datentransfer komplett

Der Datentransfer der Anwort lässt sich berechnen wie folgt:

$$D_{komplett} = (P * (64 + \alpha)) + (P * (\delta * \varepsilon) * (88 + (\beta * (10 + \gamma))))$$

Datentransfer komplett

Mit $\alpha = 17, \beta = 10, \gamma = 60, \delta = .3$ und $\varepsilon = .4$:

N \ T	1	2	3	4	5	6	7	8
2	355.1	1.1k	2.1k	3.6k	5.4k	7.6k	10.2k	13.1k
3	532.7	2.4k	7.5k	20.3k	50.5k	120.0k	277.2k	627.2k
4	710.2	4.3k	18.6k	72.1k	262.7k	924.7k	3.11M	10.50M
5	887.8	6.6k	37.6k	191.4k	925.4k	4.23M	19.40M	87.44M
6	1065.4	9.6k	66.6k	421.6k	2.48M	14.54M	83.25M	469.09M
7	1242.9	13.0k	107.8k	816.6k	5.77M	40.51M	278.42M	1.84G
8	1420.5	17.0k	163.3k	1.41M	11.88M	97.29M	780.07M	6.01G

Angaben in xB.

Bandbreite

Im Normalbetrieb, so DSS2 Clip, werden pro Sekunde 3-5 Suchanfragen gestartet, zur Zeit der *Napster-Flut* 10 Anfragen pro Sekunde

$$D_{bandbreite} = D_{komplett} * \zeta$$

Mit:

$\zeta =$ qps (Queries per Second).

Datentransfer komplett

Mit $\alpha = 17, \beta = 10, \gamma = 60, \delta = .3$, $\varepsilon = .4$ und $\zeta = 3$:

N \ T	1	2	3	4	5	6	7	8
2	1.0k	3.2k	6.4k	10.8k	16.3k	22.9k	30.6k	39.4k
3	1.6k	7.2k	22.6k	60.8k	151.4k	360.1k	831.5k	1.84M
4	2.1k	12.8k	55.9k	216.2k	788.1k	2.71M	9.32M	31.51M
5	2.6k	19.9k	112.8k	574.3k	2.71M	12.70M	58.19M	262.33M
6	3.1k	28.7k	199.8k	1.24M	7.45M	43.61M	249.75M	1.37G
7	3.6k	39.1k	323.4k	2.39M	17.31M	121.54M	835.25M	5.52G
8	4.2k	51.0k	489.9k	4.22M	35.63M	291.88M	2.29G	18.03G

Angaben in xB/s.

Datentransfer komplett

Mit $\alpha = 17, \beta = 10, \gamma = 60, \delta = .3$, $\varepsilon = .4$ und $\zeta = 5$:

N \ T	1	2	3	4	5	6	7	8
2	1.7k	5.3k	10.7k	18.0k	27.1k	38.1k	50.9k	65.6k
3	2.6k	12.0k	37.6k	101.3k	252.4k	600.2k	1.35M	3.06M
4	3.5k	21.3k	93.1k	360.3k	1.28M	4.52M	15.53M	52.52M
5	4.3k	33.2k	188.0k	957.2k	4.52M	21.17M	96.99M	437.22M
6	5.2k	47.8k	333.0k	2.06M	12.42M	72.68M	416.25M	2.29G
7	6.1k	65.1k	538.9k	3.99M	28.85M	202.56M	1.36G	9.19G
8	6.9k	85.0k	816.4k	7.04M	59.39M	486.46M	3.81G	30.05G

Angaben in xB/s.

Datentransfer komplett

Mit $\alpha = 17, \beta = 10, \gamma = 60, \delta = .3$, $\varepsilon = .4$ und $\zeta = 10$:

N \ T	1	2	3	4	5	6	7	8
2	3.5k	10.6k	21.5k	36.0k	54.3k	76.2k	101.8k	131.2k
3	5.2k	23.9k	75.2k	202.7k	504.8k	1.17M	2.71M	6.12M
4	6.9k	42.5k	186.2k	720.7k	2.57M	9.03M	31.06M	105.04M
5	8.7k	66.4k	376.0k	1.87M	9.04M	42.33M	193.98M	874.45M
6	10.4k	95.7k	666.0k	4.12M	24.84M	145.36M	832.50M	4.58G
7	12.1k	130.2k	1.05M	7.97M	57.70M	405.12M	2.72G	18.38G
8	13.9k	170.1k	1.59M	14.08M	118.77M	972.92M	7.62G	60.09G

Angaben in xB/s.

Verbesserungen

- Suchanfragen und Antworten zwischenspeichern
- Geschwindigkeitsgesteuerte Subnetze
- Subnetze je nach Interessengebiet (Audio, Video, etc...)
- Reflektoren und Supernodes einbinden

Quellen

- `www.gnutelladev.com`

  Spezifikation von Gnutella

- “Why Gnutella can’t scale. No really!“ [2000 Ritter, Jordan]

- `www.heise.de/tp`

  Telepolis: Diverse Artikel und Berichte zu Gnutella