IPv6!

Utrecht, vrijdag 12 september 2008

Iljitsch van Beijnum

1969: ARPANET

- I.o.v. US Defense Advanced Research Projects Agency
- Om computers research-instellingen te verbinden
- Eerste pakketgeschakelde netwerk
- Voorheen alleen circuitgeschakeld
  - vgl. telefoonnetwerk

1981: TCP/IP

- Na 10 jaar ervaring nieuwe protocollen. Werking:
  - TCP hakt communicatie in pakketten
  - IP verstuurt pakketten adhv IP adres
    - ieder pakket staat op zich
  - ontvangende TCP zet pakketten weer achter elkaar

1986: IETF

- Internet Engineering Task Force
- Mensen die:
  - het netwerk runnen
  - de apparatuur en software bouwen
  - communicatieprotocollen ontwerpen
  - RFCs schrijven

Ca. 1992: Adressen op?

- IP adres is 32 bits lang:
  - 4294,97 miljoen mogelijke adressen
    - 3706,65 miljoen bruikbare adressen
- Gebruik verdubbelde iedere 5 jaar
- Prognose: op in 2005

IETF-prognose (± 1993)

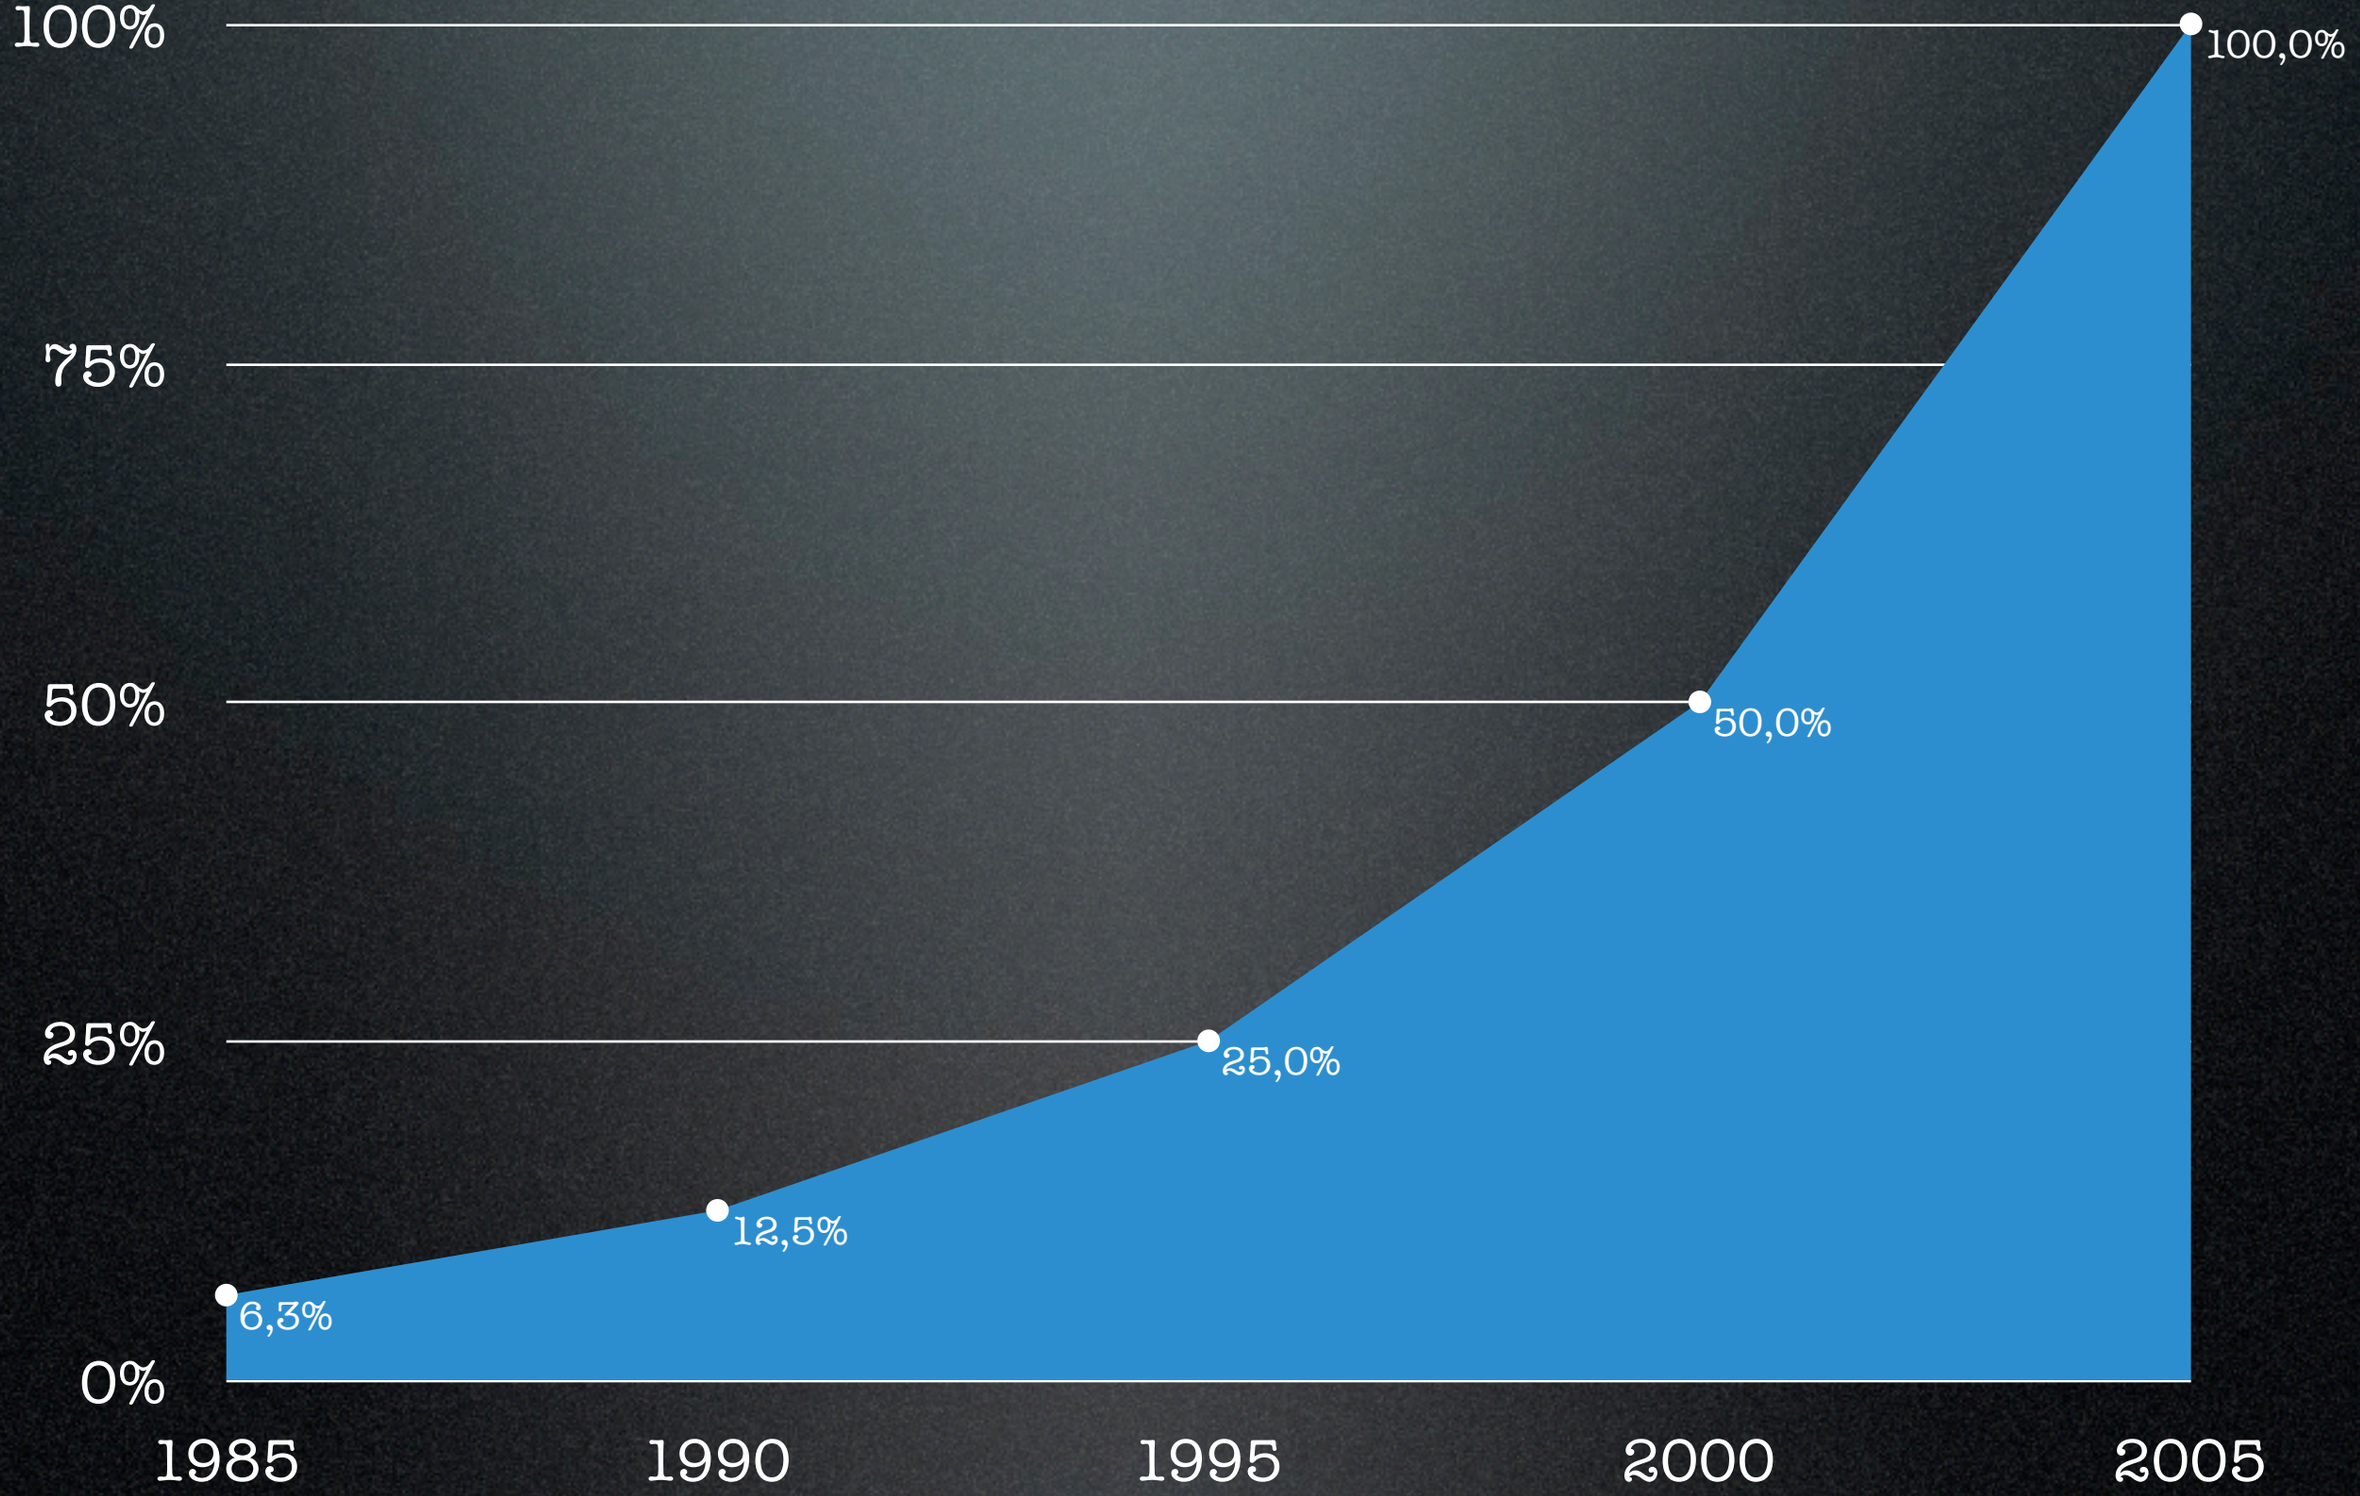

Legenda

- ★ Niet bruikbaar
- ★ Uitgegeven aan eindgebruiker
- ★ "Various registries"
- ★ RIPE NCC (Europa + aanhang)
- ★ ARIN (Noord-Amerika)
- ★ APNIC (Azië, Australië en Pacific)
- ★ LACNIC (Latijns Amerika + Caribbean)
- ★ AfriNIC (Afrika)

1993: 97 /8s (43.9%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

1997: 109 /8s (49.3%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	523	254	255

2000: 118 /8s (53.4%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2001: 125 /8s (56.6%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2002: 129 /8s (58.4%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2003: 134 /8s (60.6%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2004: 143 /8s (64.7%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2005: 156 /8s (70.6%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2006: 166 /8s (75.1%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2007: 178 /8s (80.5%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

2008: 179 /8s (80.1%)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Sep 2008: 182 /8s, 82.4%

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

"IANA global pool"

(Adresruimte die nog niet aan Regional Internet Registry's uitgegeven is)

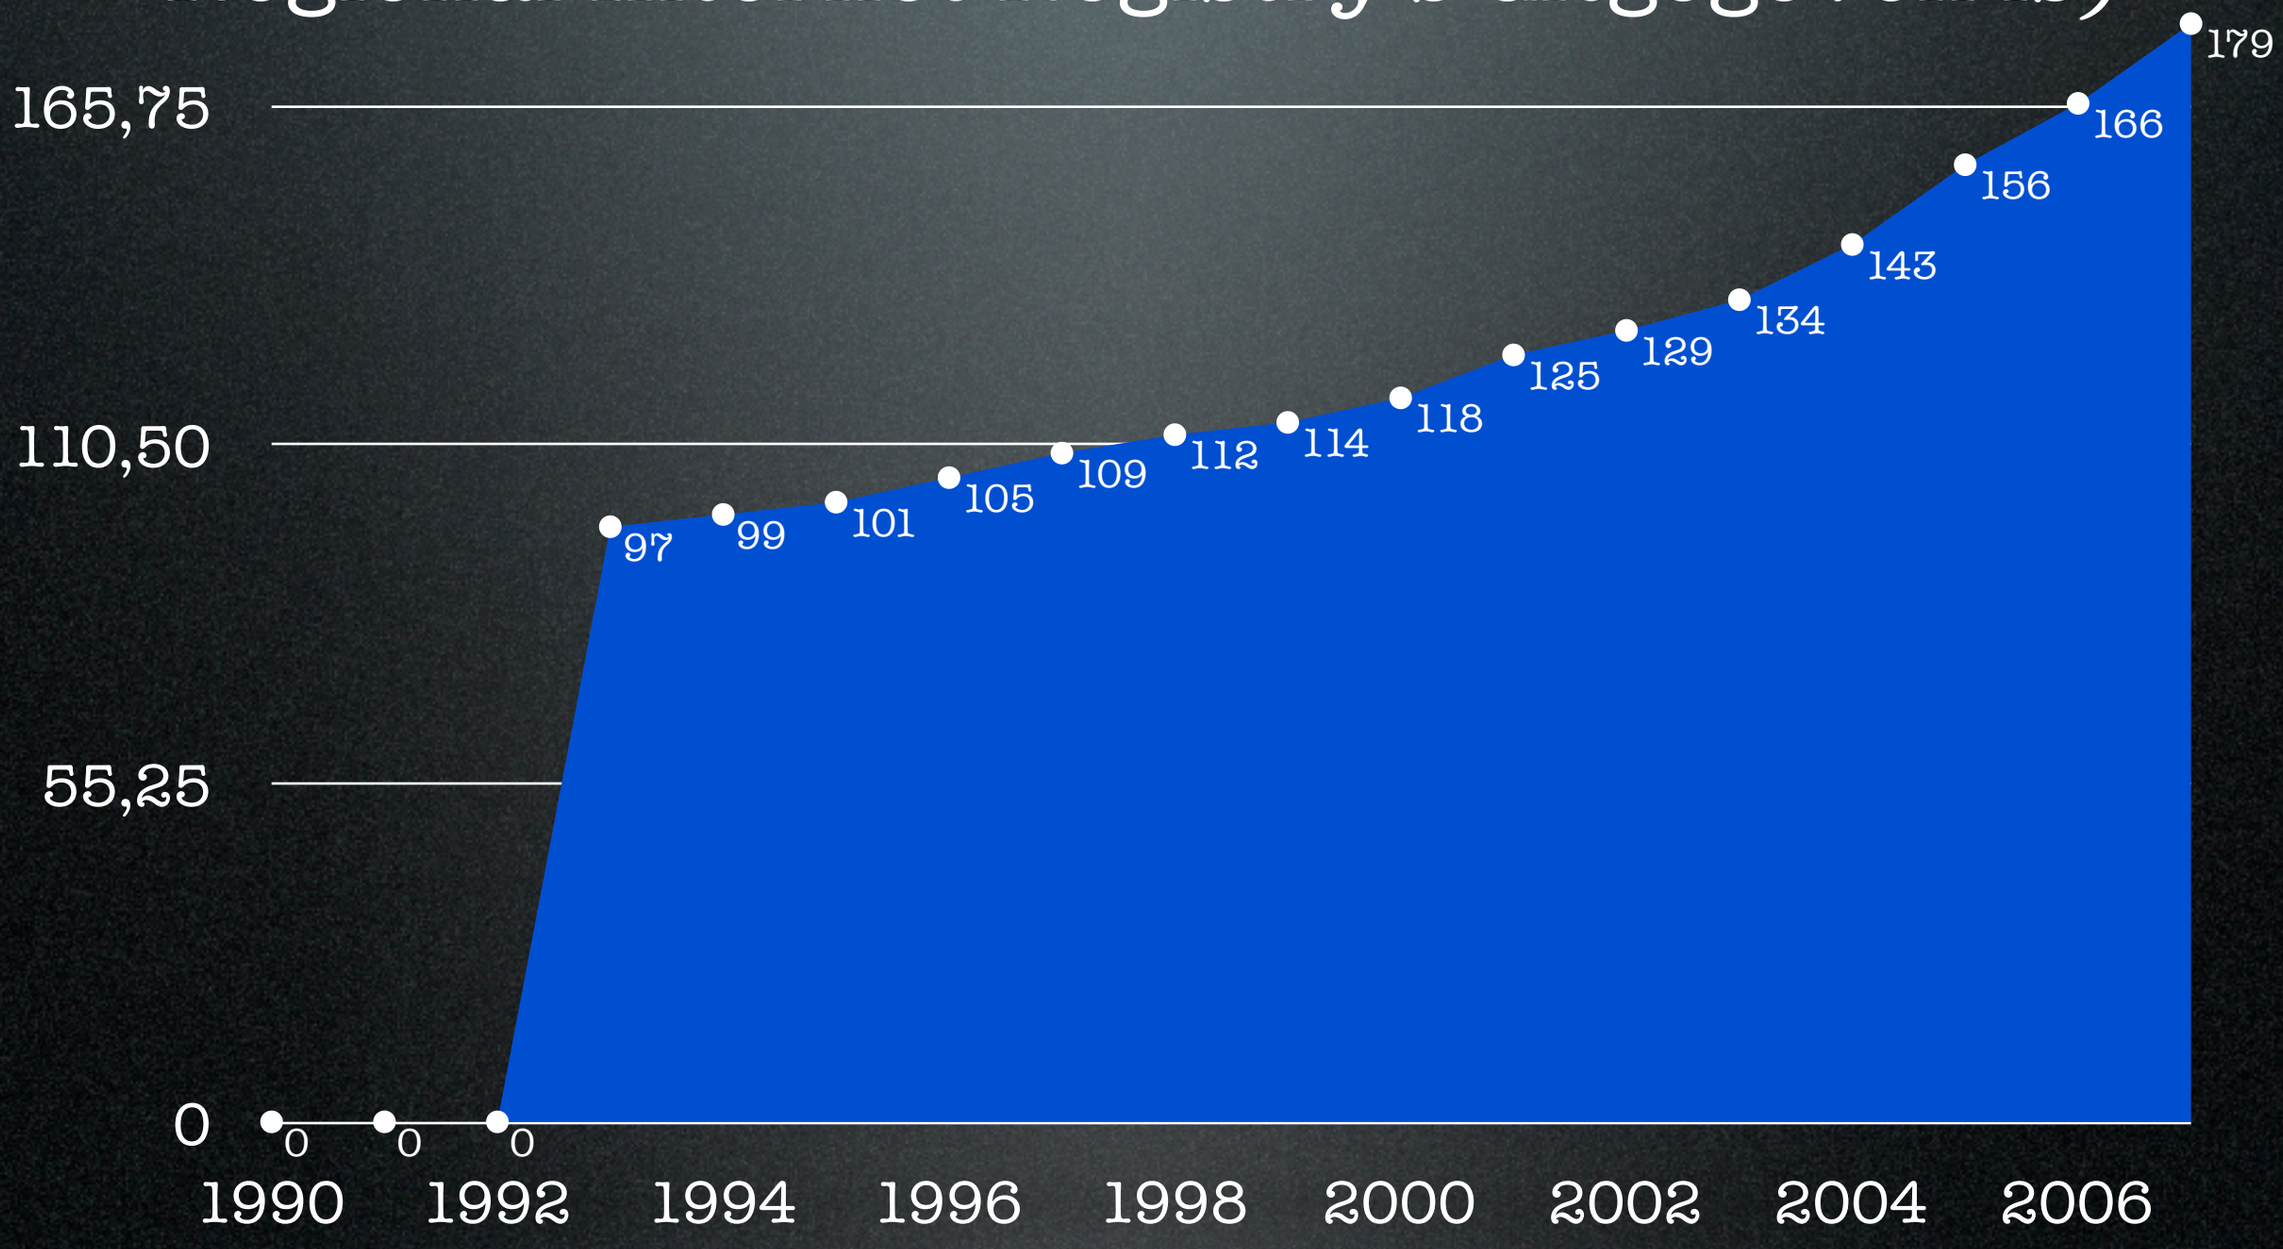

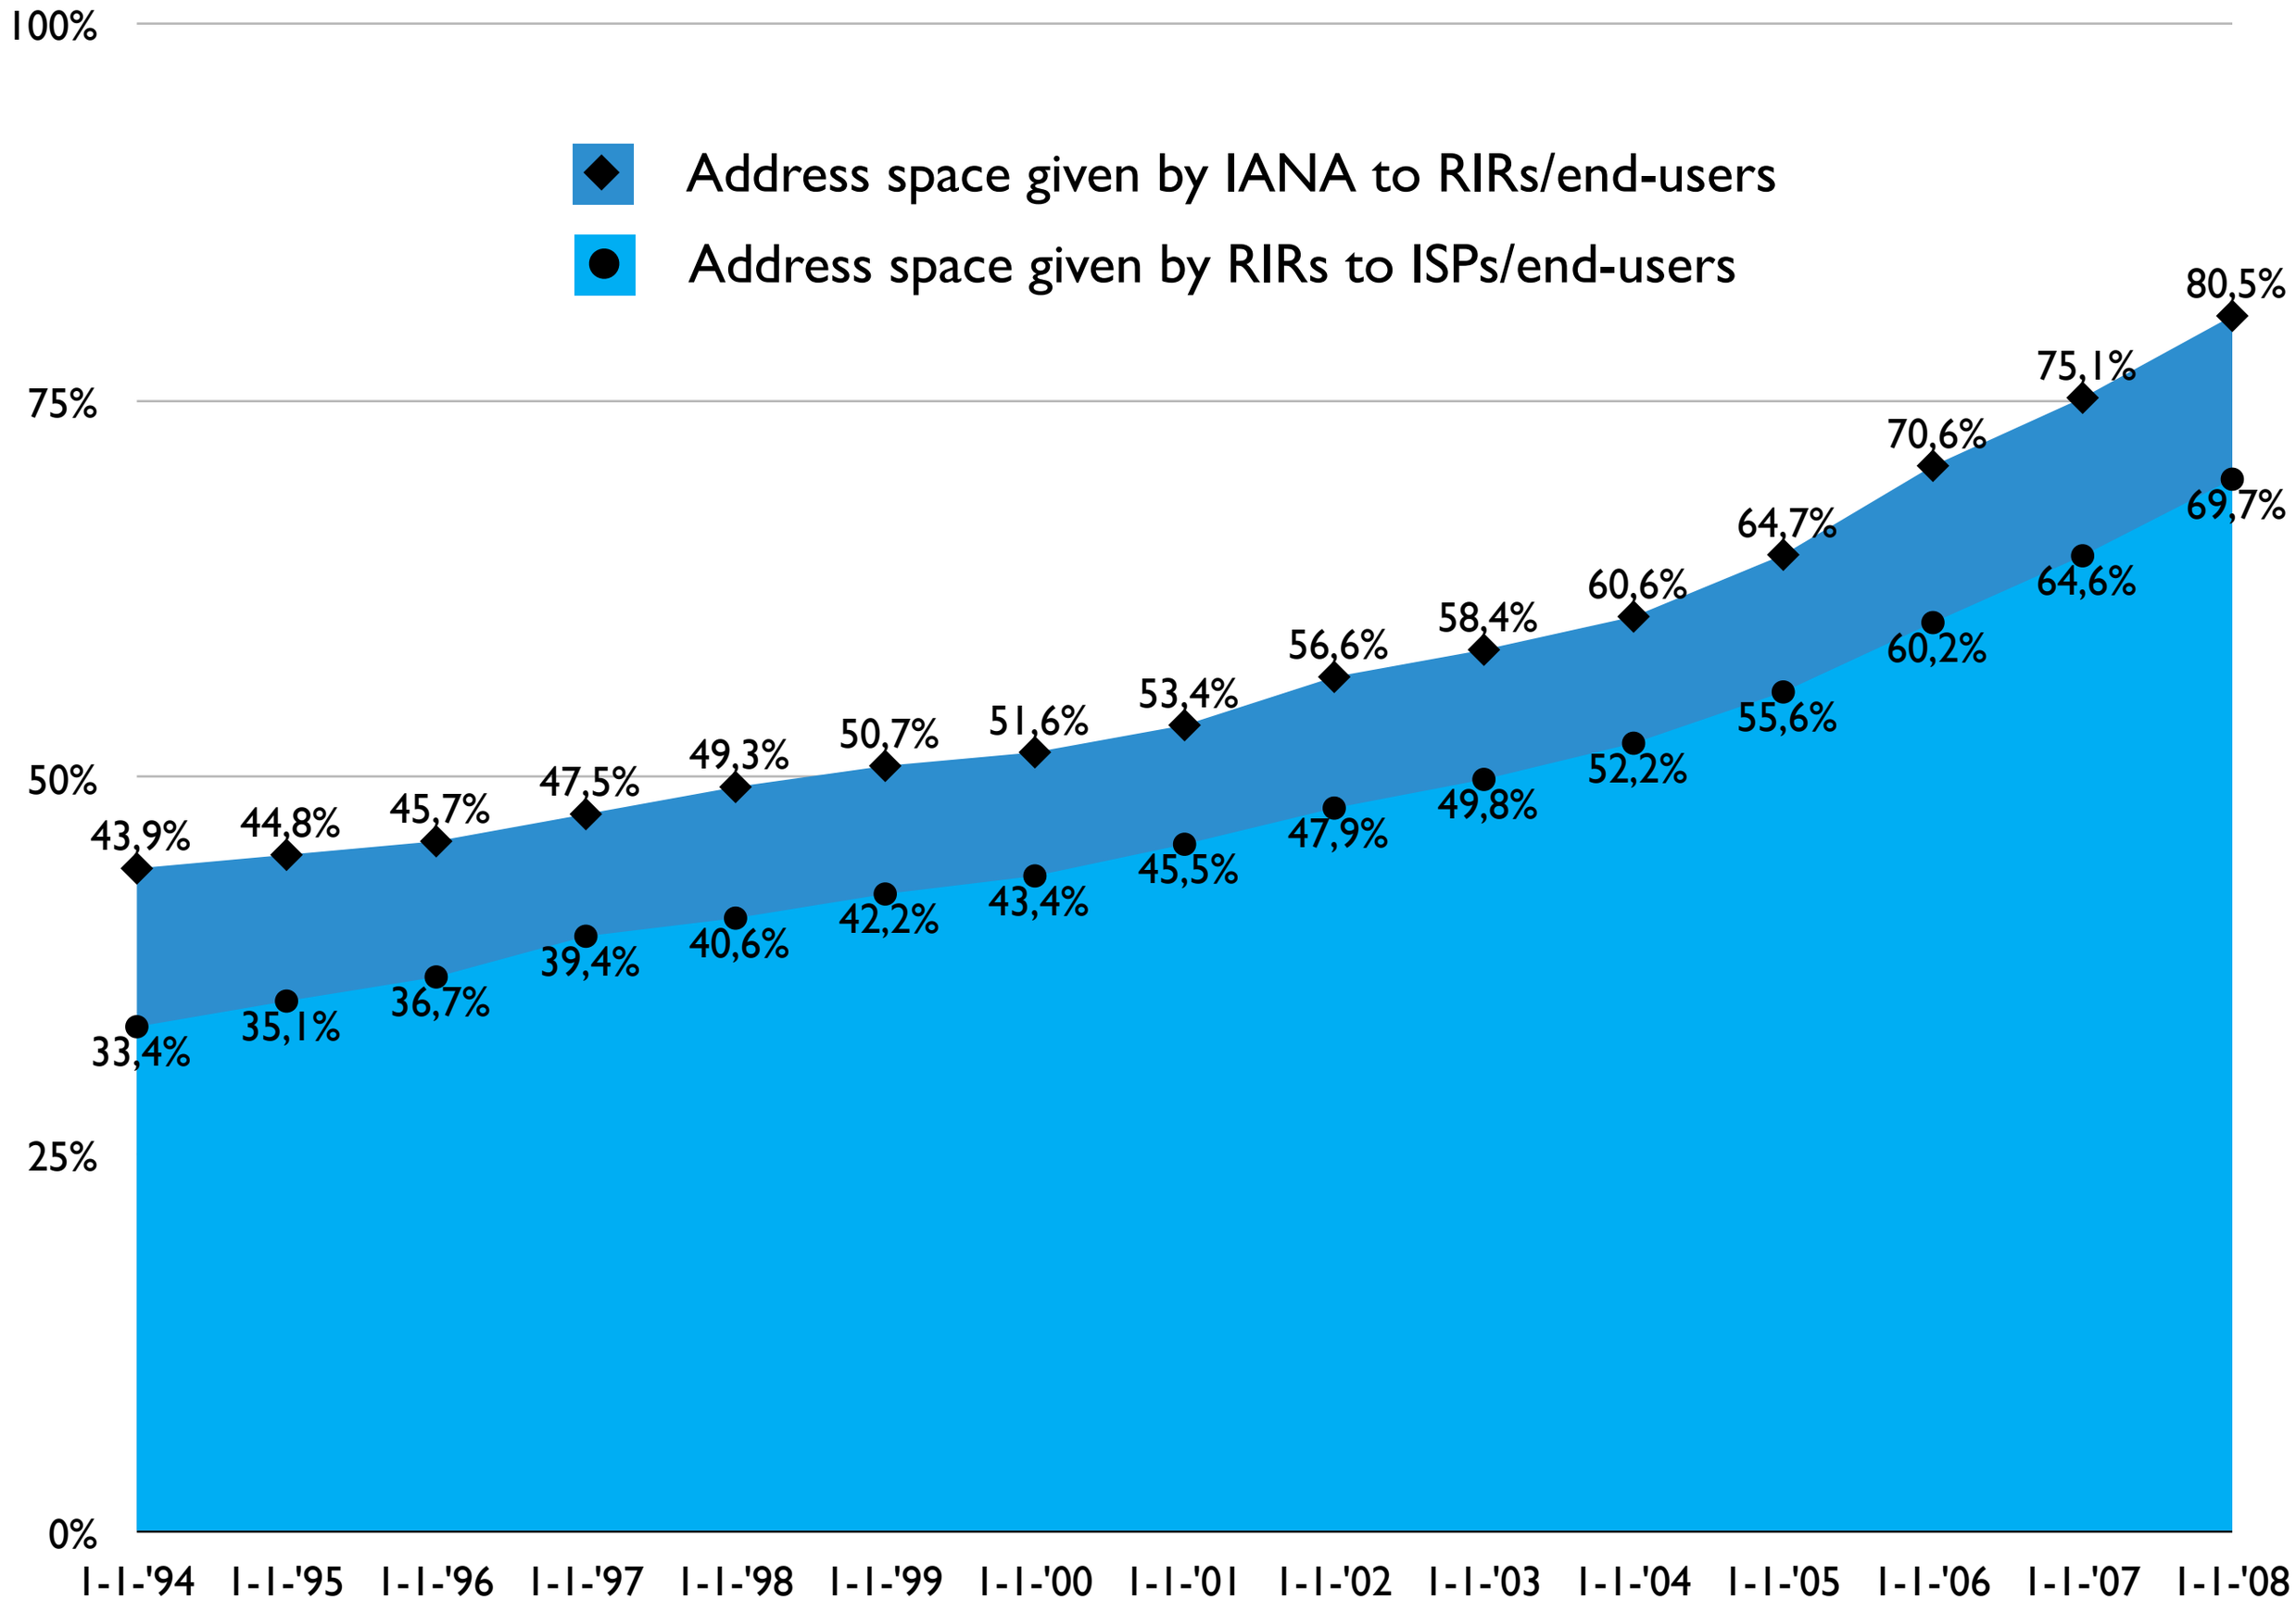
100%
75%
50%
25%
0%
Address space given by IANA to RIRs/end-users
Address space given by RIRs to ISPs/end-users
43,9%
44,8%
45,7%
47,5%
49,3%
50,7%
51,6%
53,4%
56,6%
58,4%
60,6%
64,7%
70,6%
75,1%
80,5%
33,4%
35,1%
36,7%
39,4%
40,6%
42,2%
43,4%
45,5%
47,9%
49,8%
52,2%
55,6%
60,2%
64,6%
69,7%
1-1-'94
1-1-'95
1-1-'96
1-1-'97
1-1-'98
1-1-'99
1-1-'00
1-1-'01
1-1-'02
1-1-'03
1-1-'04
1-1-'05
1-1-'06
1-1-'07
1-1-'08

Grootste adresgebruikers

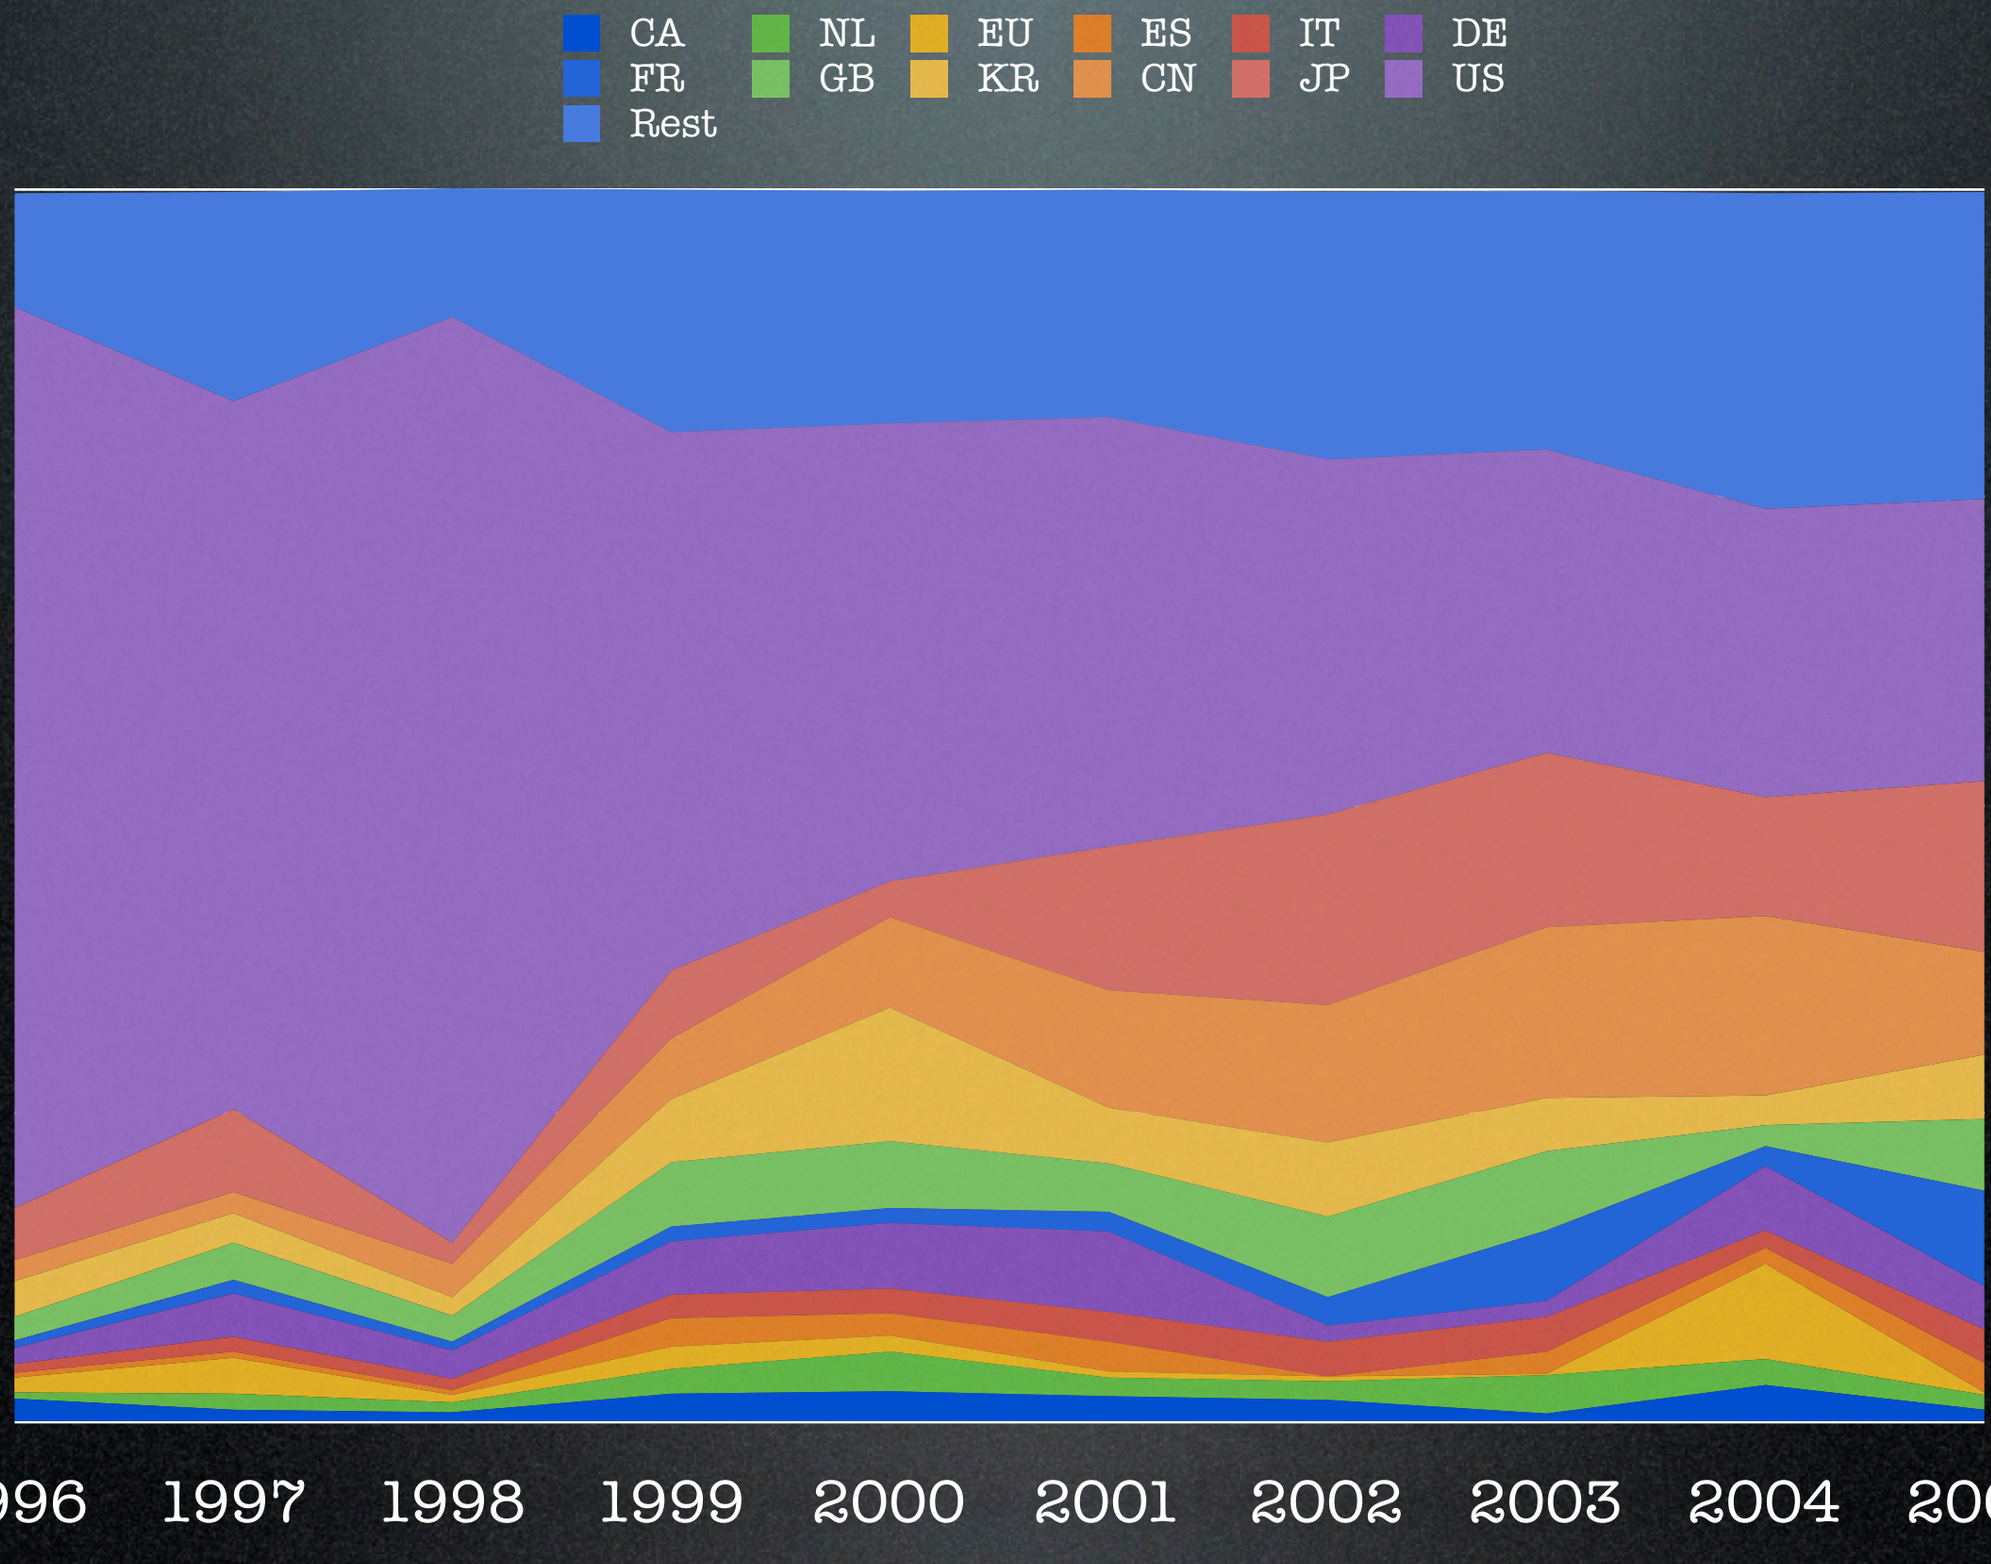

IPv4 Address Lifetime Expectancy

- This was an IETF activity starting as part of the Routing and Addressing (ROAD) activity in the early 1990's
- The objective was to understand the rate of allocation of IPv4 addresses and make s predictions as to the date of eventual exhaustion of the allocated address po
- This is a re-visiting of this activity w consideration of additional dat the characteristics of the B

Geoff Huston (2003)

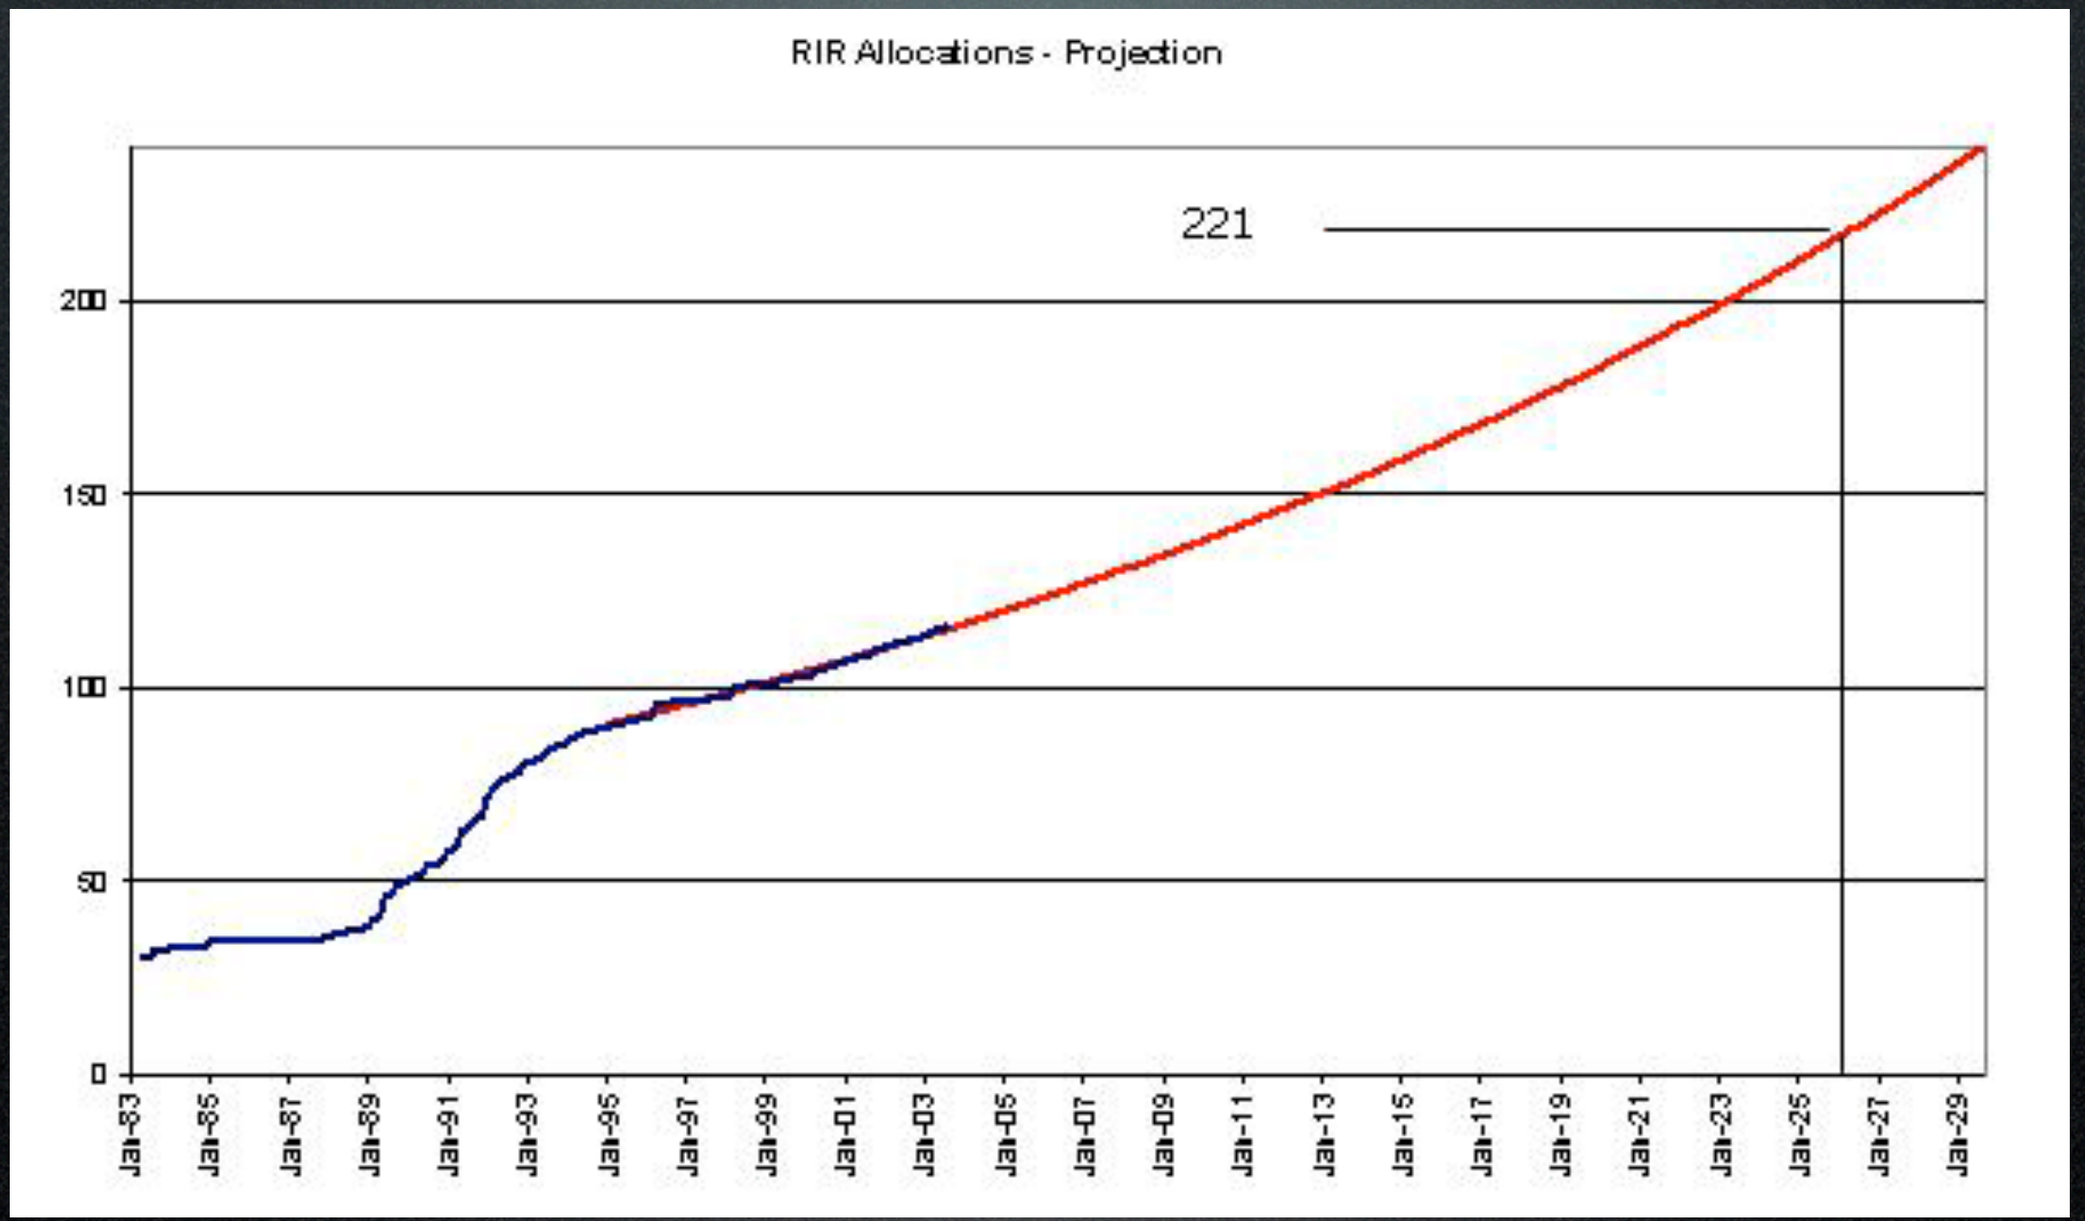

Geoff Huston (2005)

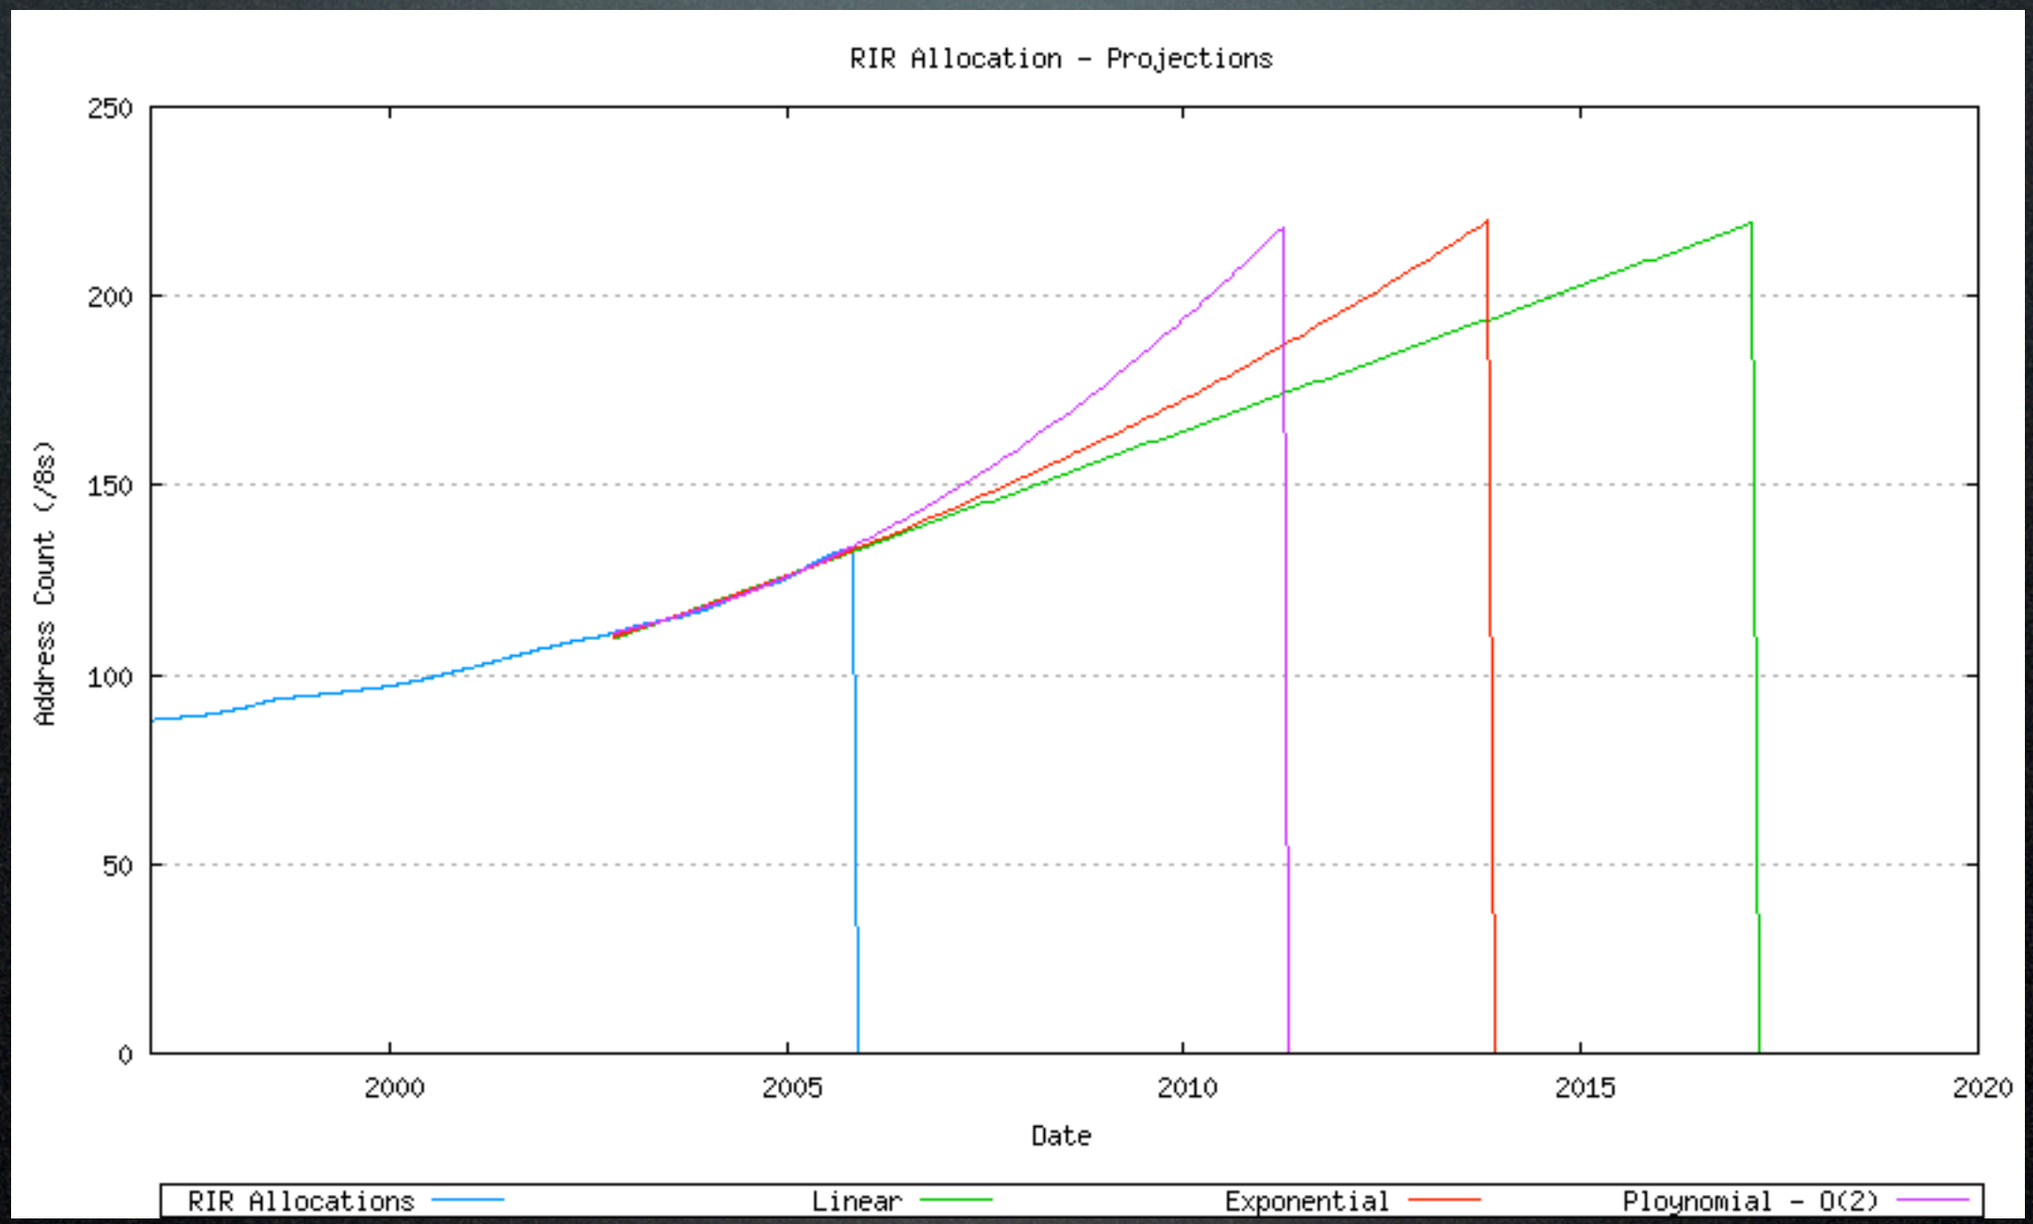

Tony Hain (2005)

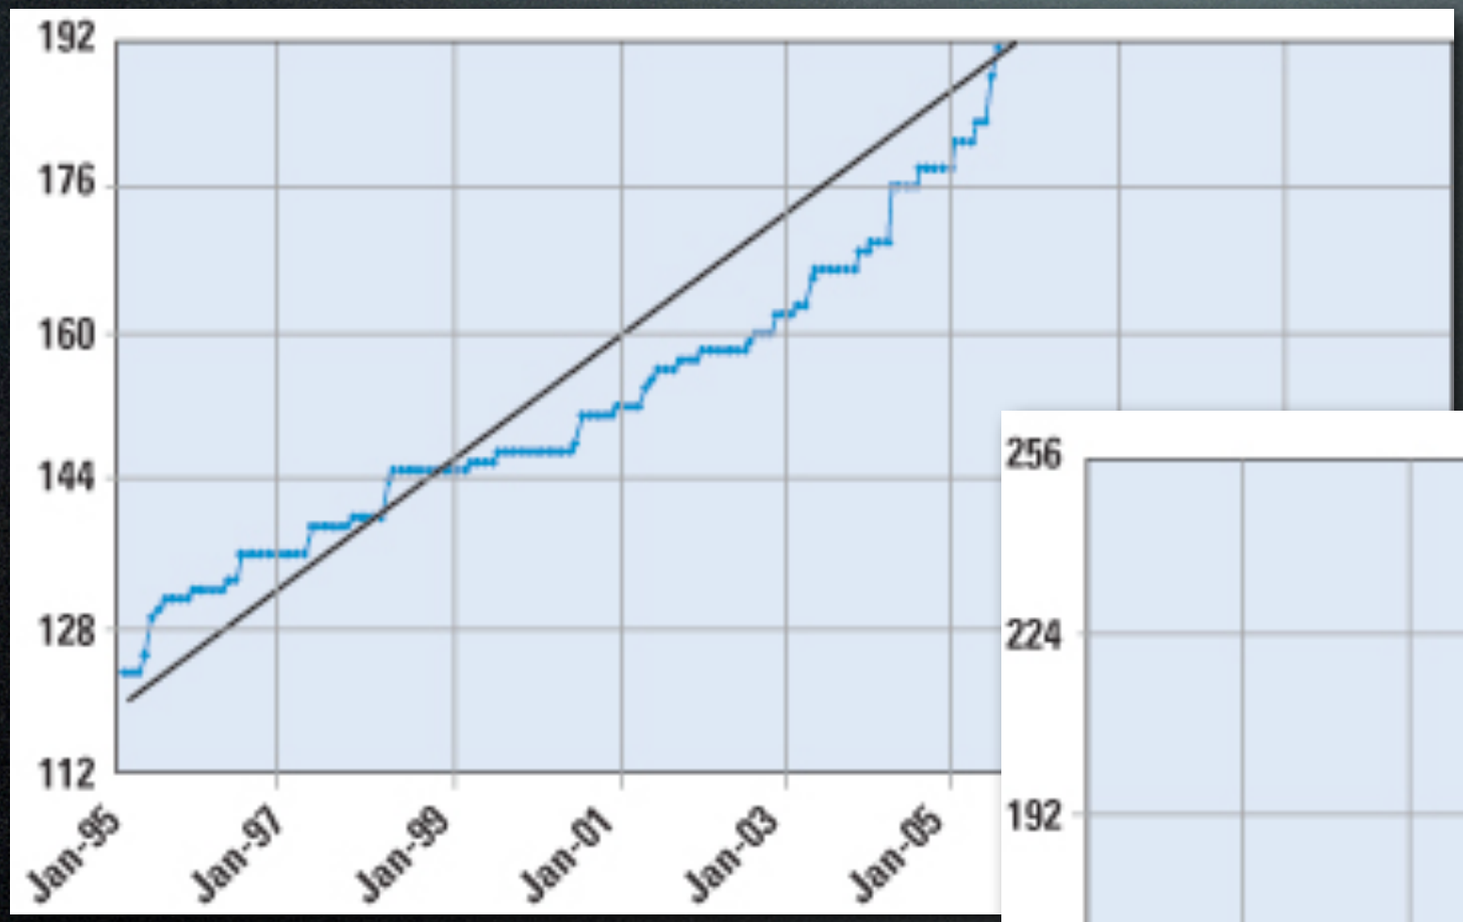

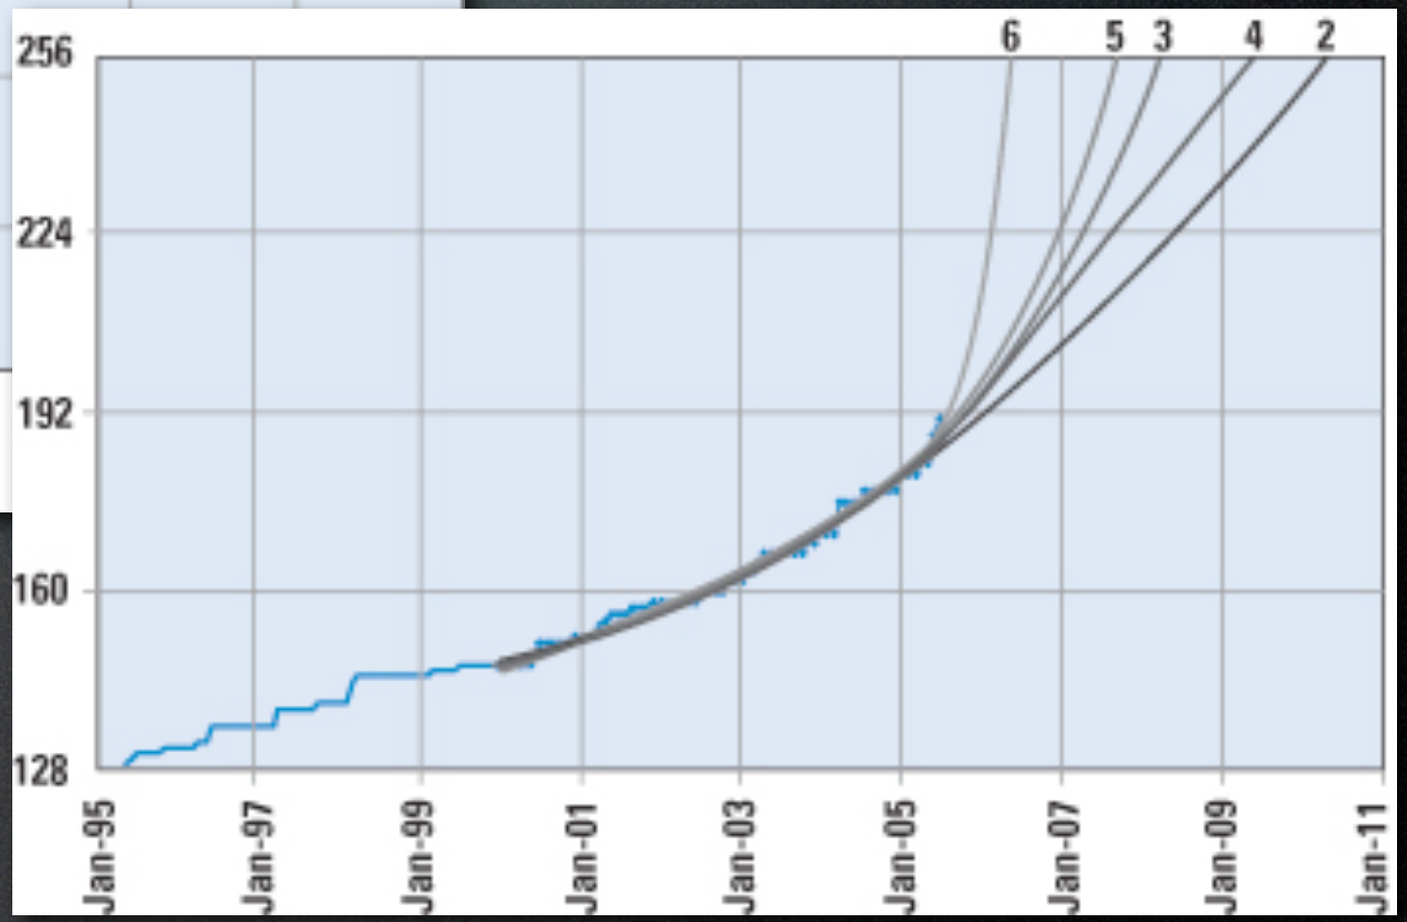

Dus?

- 3.7 miljard bruikbare adressen, 982 miljoen vrij, vorig jaar 196 miljoen uitgegeven
  - tot nu toe in 2008: 141.93 miljoen
- Geoff Huston, 2003: laatste ± 2028
- Geoff Huston, 2005: ± 2012 (update)
- Tony Hain, 2005: ± 2009
- Iljitsch van Beijnum, 2006: 2010-2015

IPng → IPv6

- IP next generation project binnen IETF:
  - los (vooral) het adresprobleem op
- Ca. 1995 keuze voor IPv6:
  - zo min mogelijk radicale wijzigingen
  - + voor de hand liggende verbeteringen
- Huidige IP heeft versienummer 4 (IPv4), het nieuwe is versie 6 (IPv6)

Adreslengte

Aantal adressen in IPv4:

4.294.967.296

32 bits (10 cijfers)

IPv6:

340.282.366.920.938.463.4
63.374.607.431.768.211.456

128 bits (39 cijfers)

Voordelen IPv6

- Beveiliging:
  - verplichte IPsec, nuttig in praktijk?
  - willekeurig scannen niet zinnig meer
- Quality of service: niet echt anders
- Geen ingewikkelde adresplannen meer
- Stateless autoconfig: veel simpeler
- Meer opties toekomstige verbeteringen

Voordelen IPv6 (2)

- Geen Network Address Translation (NAT) meer
  - maakt alles simpeler
    - wel stateful firewall nodig
- Unique Site Local als privé-adressen
  - in plaats van iedereen 10.x.x.x
  - scheelt hoofdpijn bij fusies

Het verschil

IPv4

https://noc.sixxs.net/main/

SixXS

Main | About | Contact | News | Home | PoPs | Presentations | FAQ | Forum | Misc | Links

SixXS - IPv6 Deployment & Tunnel Broker

SixXS (*Six Access*) is a free, non-profit, non-cost service for Local Internet Registries (LIR's). The main target is to create a common portal to help company engineers find their way with IPv6 networks deploying IPv6 to their customers in a rapid and controllable fashion. To reach these targets we are providing a whitelabel IPv6 Tunnel Broker and Ghost Route Hunter, an IPv6 route monitoring tool and various other services to help out where needed.

This service accommodates an enhanced version of the IPng IPv6 tunnel broker to assist LIR's in having a rapid IPv6 deployment in their organisation by providing a whitelabeled IPv6 Tunnel Broker, giving IPv6 access to their clients without the burden of developing their own software allowing full and easy integration into their existing systems.

SixXS offers the RIPE, ARIN, APNIC, LacNIC and AfriNIC communities pre-production deployment expertise based on the experience gathered while running the IPng IPv6 tunnel broker since 1999 combined with it's successor, SixXS, totaling in more than 5 years of massive experience.

Most of the PoP's are publicly available and provide free IPv6 connectivity to endusers allowing them to enjoy quality IPv6 connectivity. SixXS also aims to keep the IPv6 routing system stable and clean making sure that users experience a quality IPv6 internet. To assure this we provide Latency monitoring and the so called Ghost Route Hunter (GRH) service which determines the anomalies in the IPv6 routing tables. Full public statistics for this project are available from the Misc page.

Contact information
Points of Presence
AICCU
10 Easy St...
Signup for new users
Login for existing users
Ghost Route Hunter
Frequently Asked Questions
SixXS for ISP's
SixXS for endusers
SixXS Archive

Not logged in.
SSL IPv4 connection from 82.192.90.27

©2001-2006 SixXS - IPv6 Deployment & Tunnel Broker
Last revision: Fri 24 Mar 2006 16:58:42 CET

IPv6

https://noc.sixxs.net/main/

SixXS

Main | About | Contact | News | Home | PoPs | Presentations | FAQ | Forum | Misc | Links

SixXS - IPv6 Deployment & Tunnel Broker

SixXS (*Six Access*) is a free, non-profit, non-cost service for Local Internet Registries (LIR's). The main target is to create a common portal to help company engineers find their way with IPv6 networks deploying IPv6 to their customers in a rapid and controllable fashion. To reach these targets we are providing a whitelabel IPv6 Tunnel Broker and Ghost Route Hunter, an IPv6 route monitoring tool and various other services to help out where needed.

This service accommodates an enhanced version of the IPng IPv6 tunnel broker to assist LIR's in having a rapid IPv6 deployment in their organisation by providing a whitelabeled IPv6 Tunnel Broker, giving IPv6 access to their clients without the burden of developing their own software allowing full and easy integration into their existing systems.

SixXS offers the RIPE, ARIN, APNIC, LacNIC and AfriNIC communities pre-production deployment expertise based on the experience gathered while running the IPng IPv6 tunnel broker since 1999 combined with it's successor, SixXS, totaling in more than 5 years of massive experience.

Most of the PoP's are publicly available and provide free IPv6 connectivity to endusers allowing them to enjoy quality IPv6 connectivity. SixXS also aims to keep the IPv6 routing system stable and clean making sure that users experience a quality IPv6 internet. To assure this we provide Latency monitoring and the so called Ghost Route Hunter (GRH) service which determines the anomalies in the IPv6 routing tables. Full public statistics for this project are available from the Misc page.

Contact information
Points of Presence
Signup for new users
Frequently Asked Questions

Not logged in.
SSL IPv6 connection from 2001:1af8:6:0:20a:95ff:fef5:246e

...yment & Tunnel Broker
...Mar 2006 16:58:42 CET

Wijzigingen voor IPv6

- ISP-routers ✓
- Operating systems ✓
- Domain Name System (DNS) ✓
- Managementsystemen
- Groot deel applicaties
- Kabel/ADSL-modems/routers ?
- Firewalls/load balancers ?

Transition

- Een IPv4-systeem kan niet met een IPv6-systeem praten!
- Oplossingen:
  - in eerste instantie "dual stack": gebruik zowel IPv4 als IPv6
  - tunnels verbinden "IPv6 eilanden" (= stop IPv6-pakket in IPv4-pakket)
  - NAT-PT (PT = protocol translation)

"Flag day"

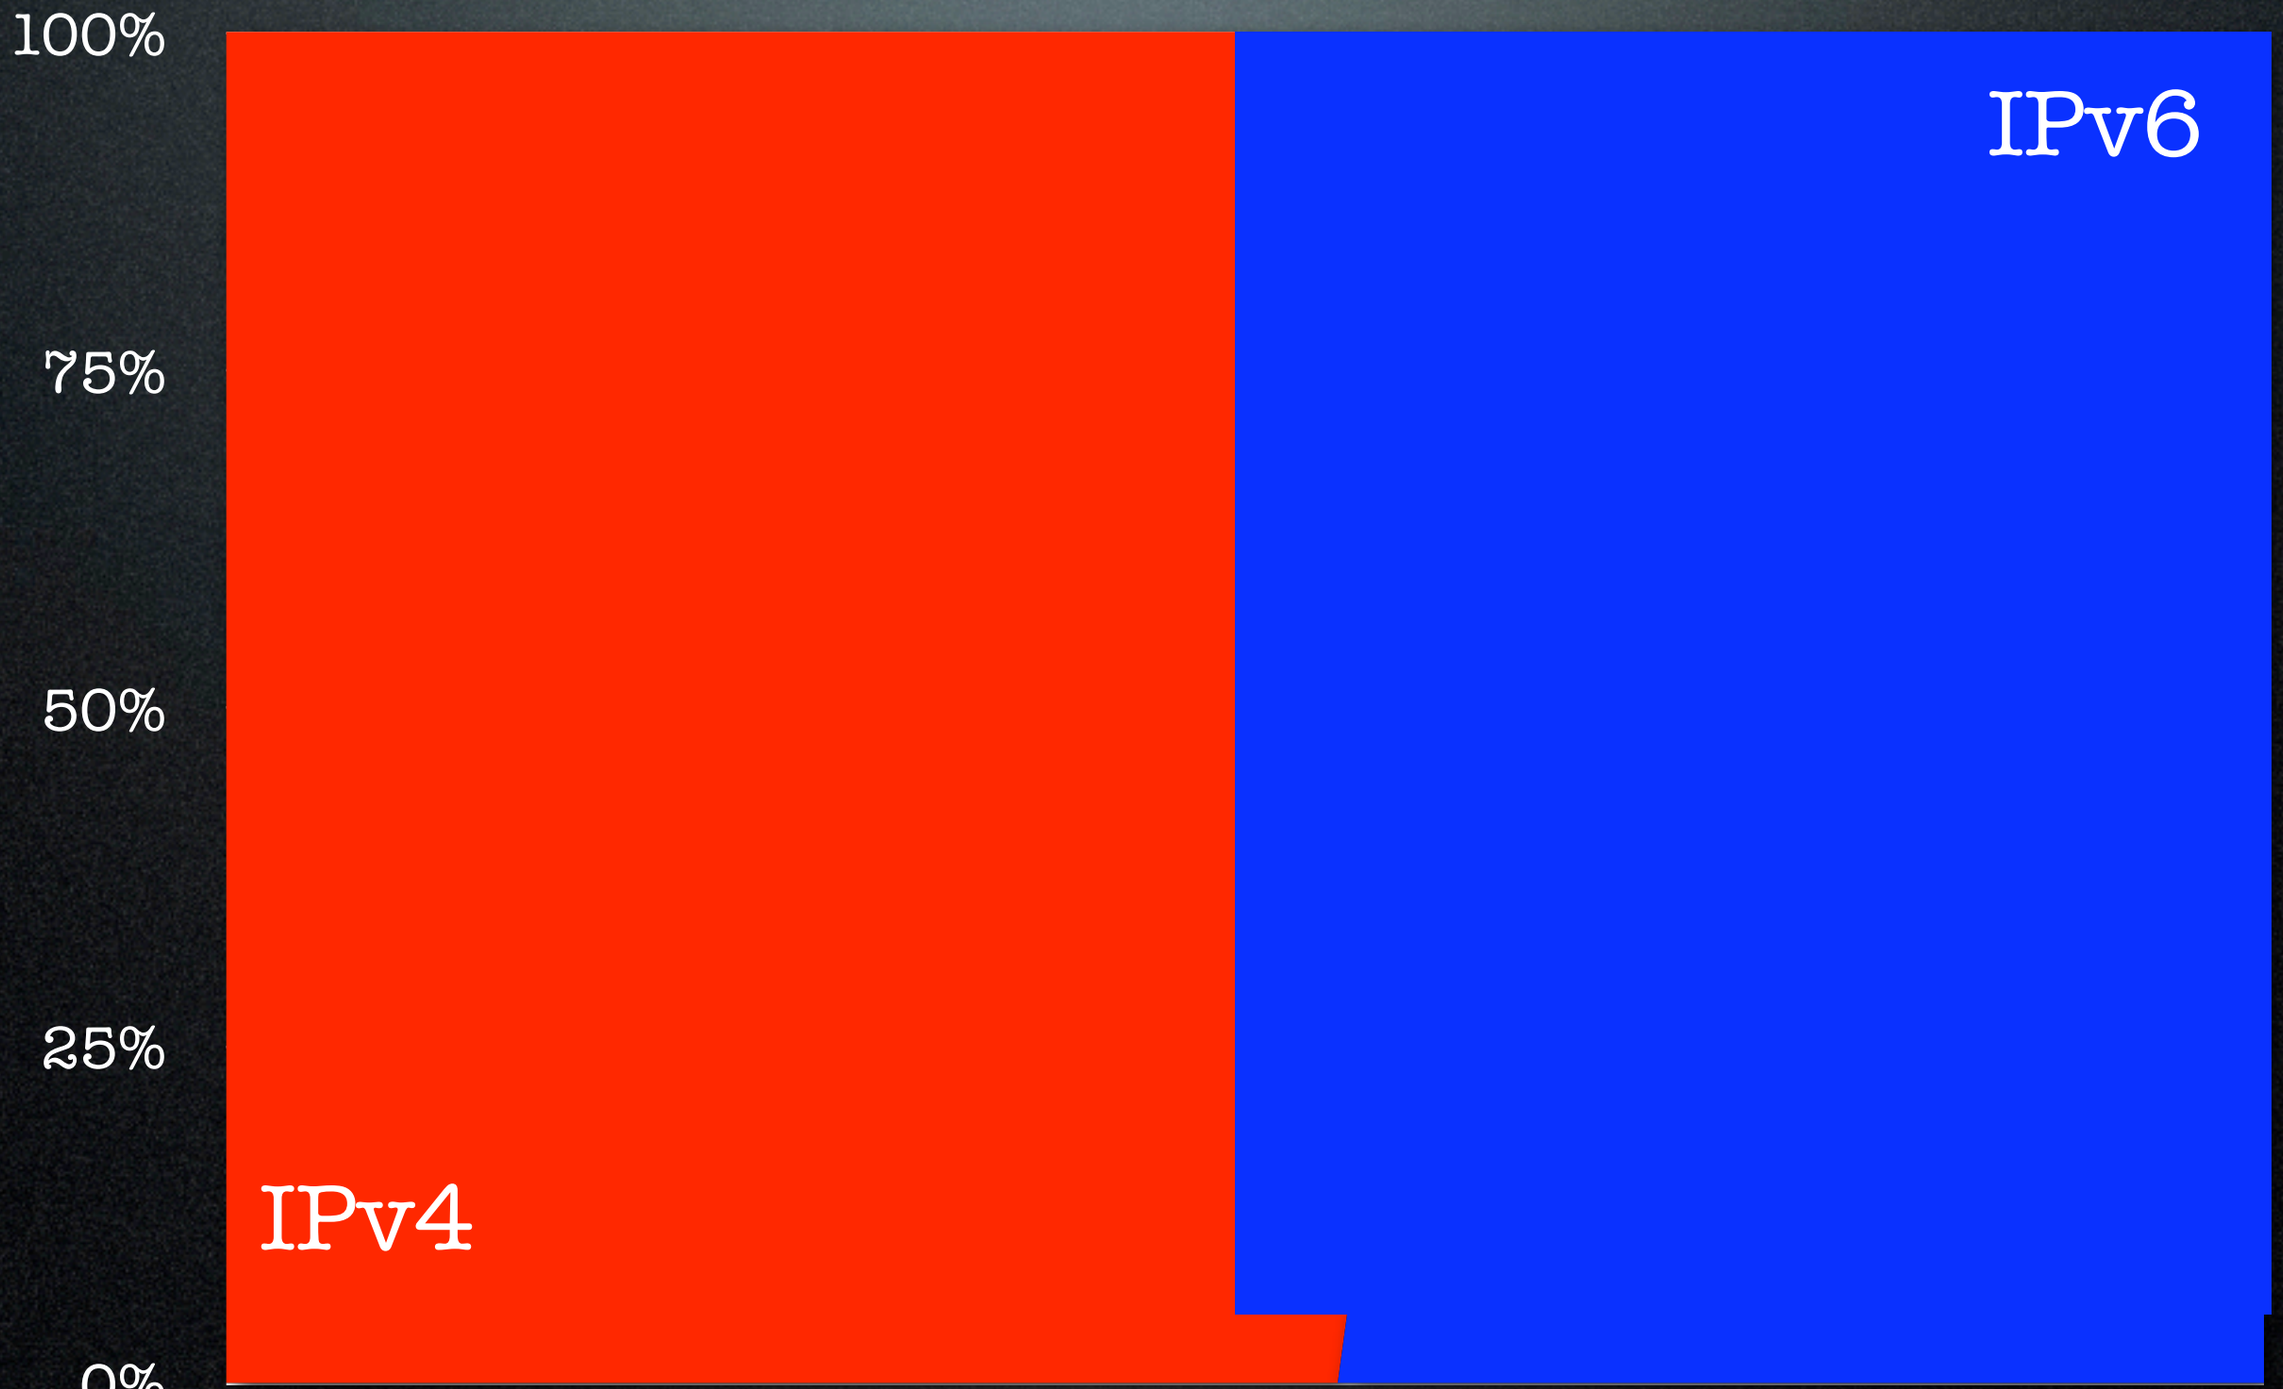

Dual stack

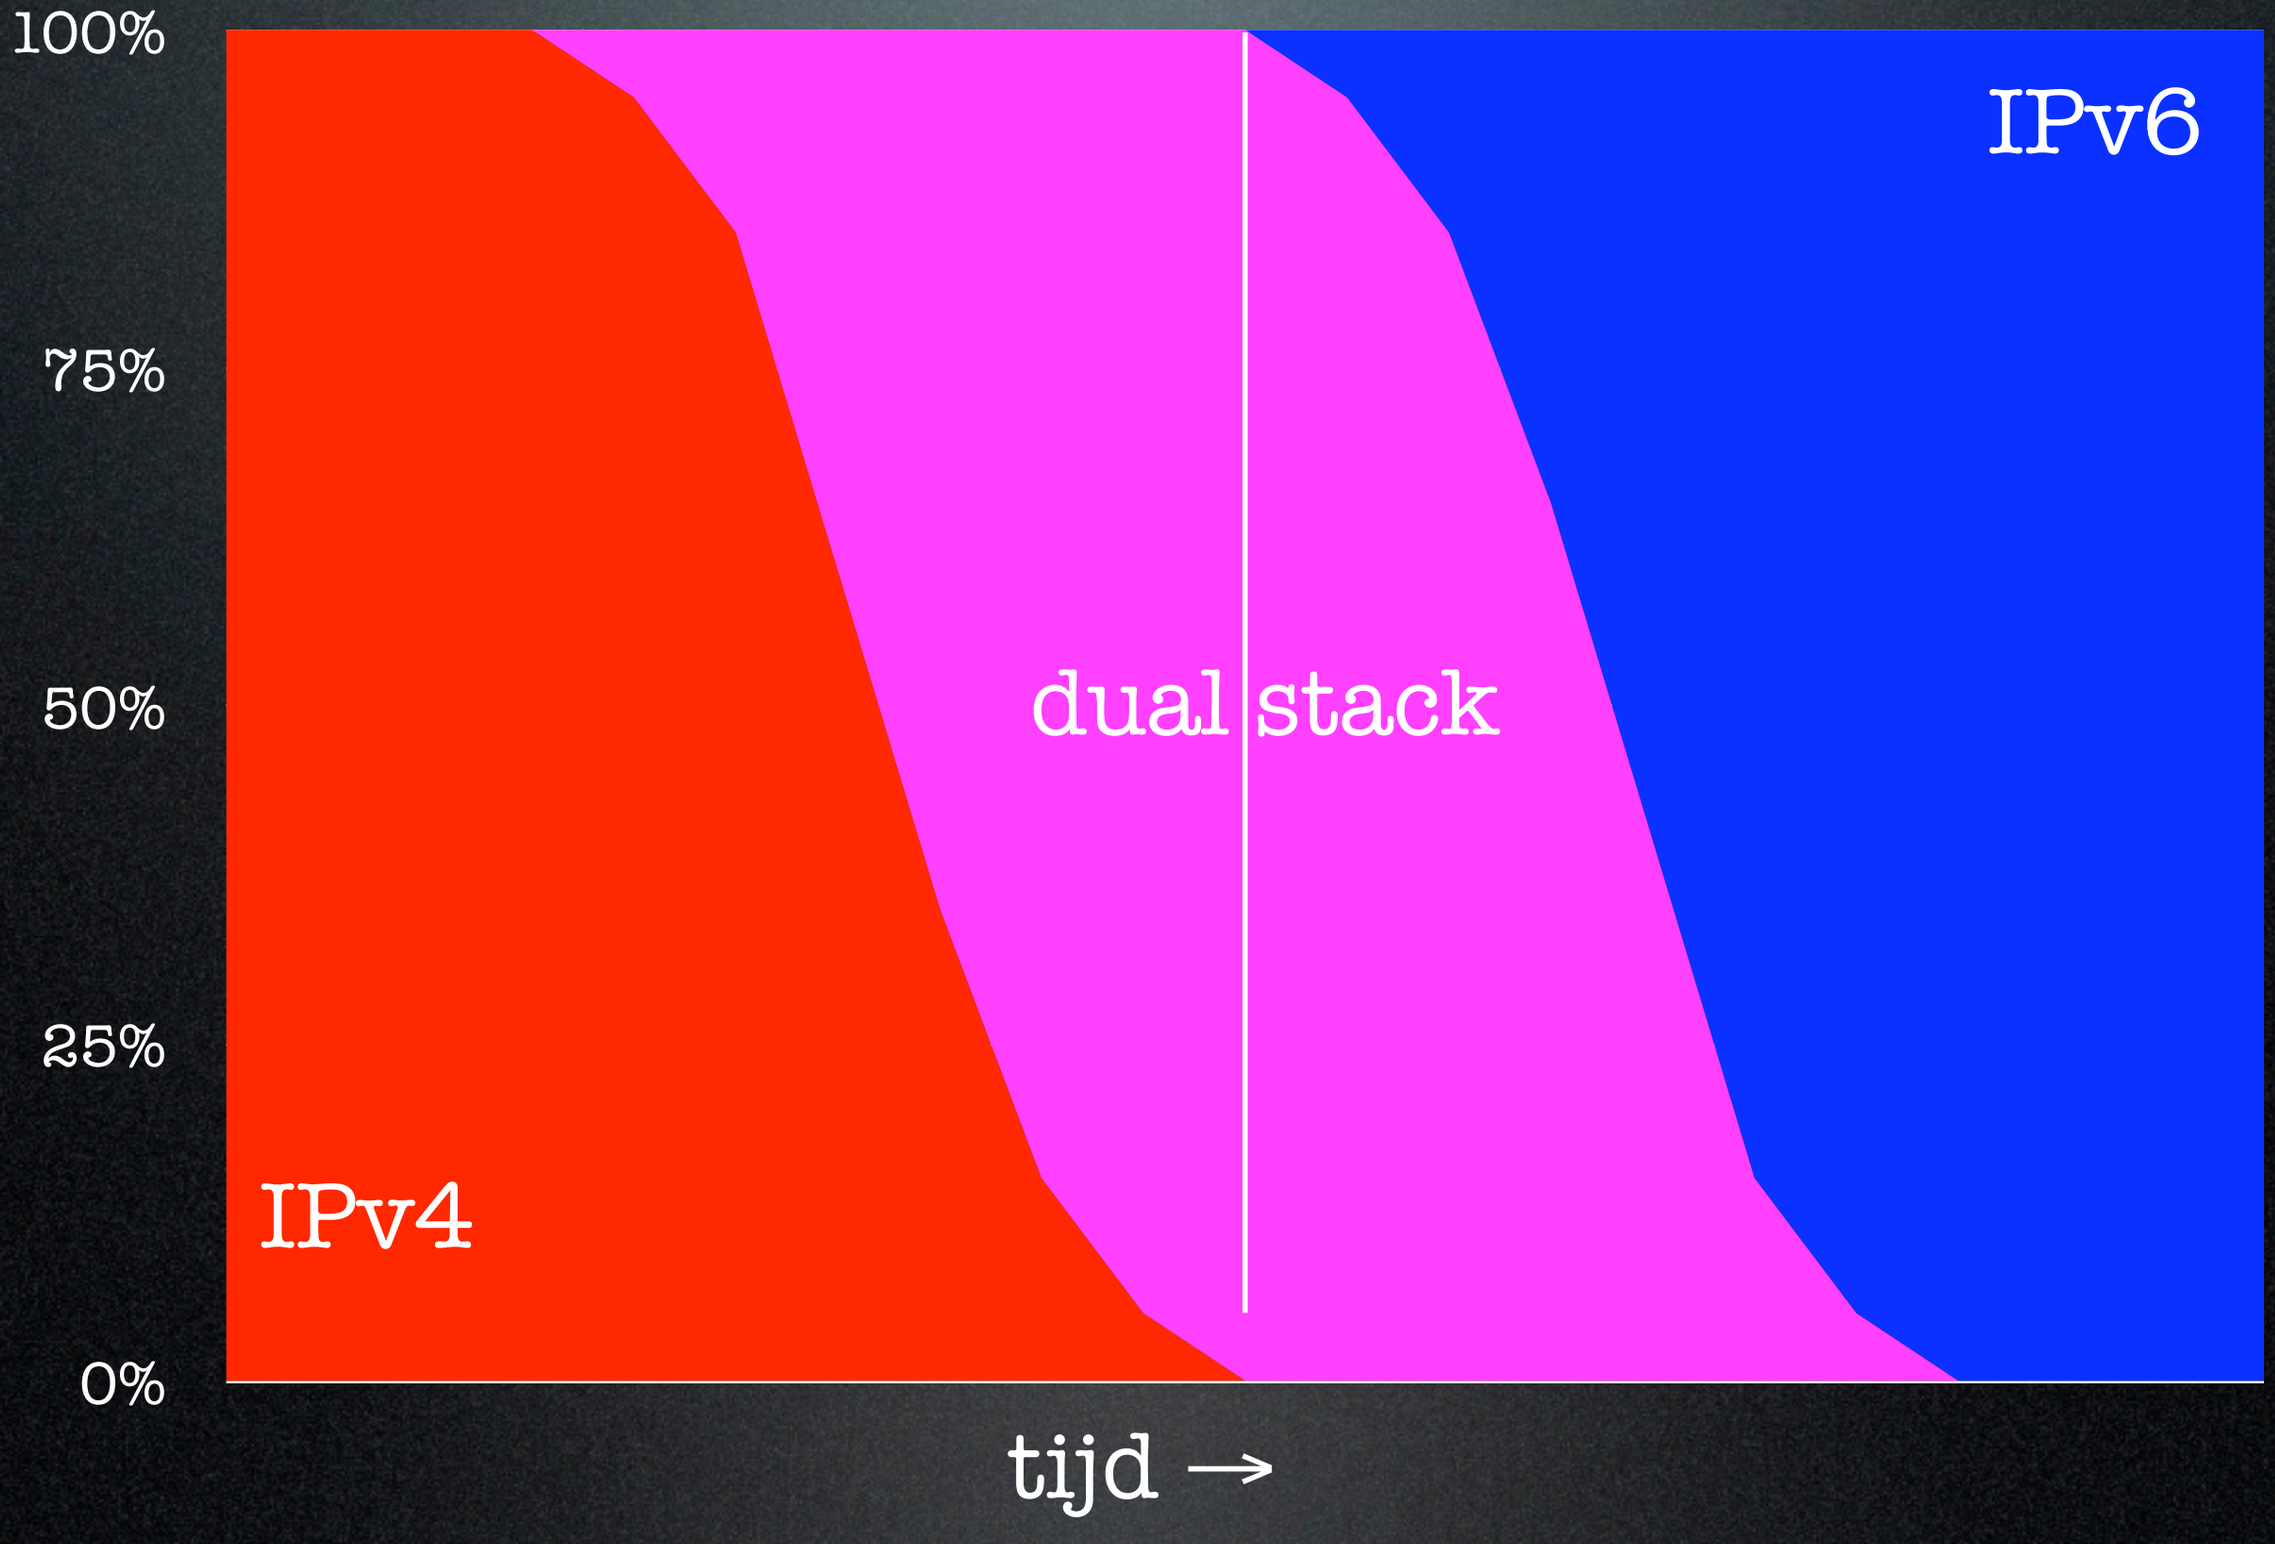

De werkelijkheid?

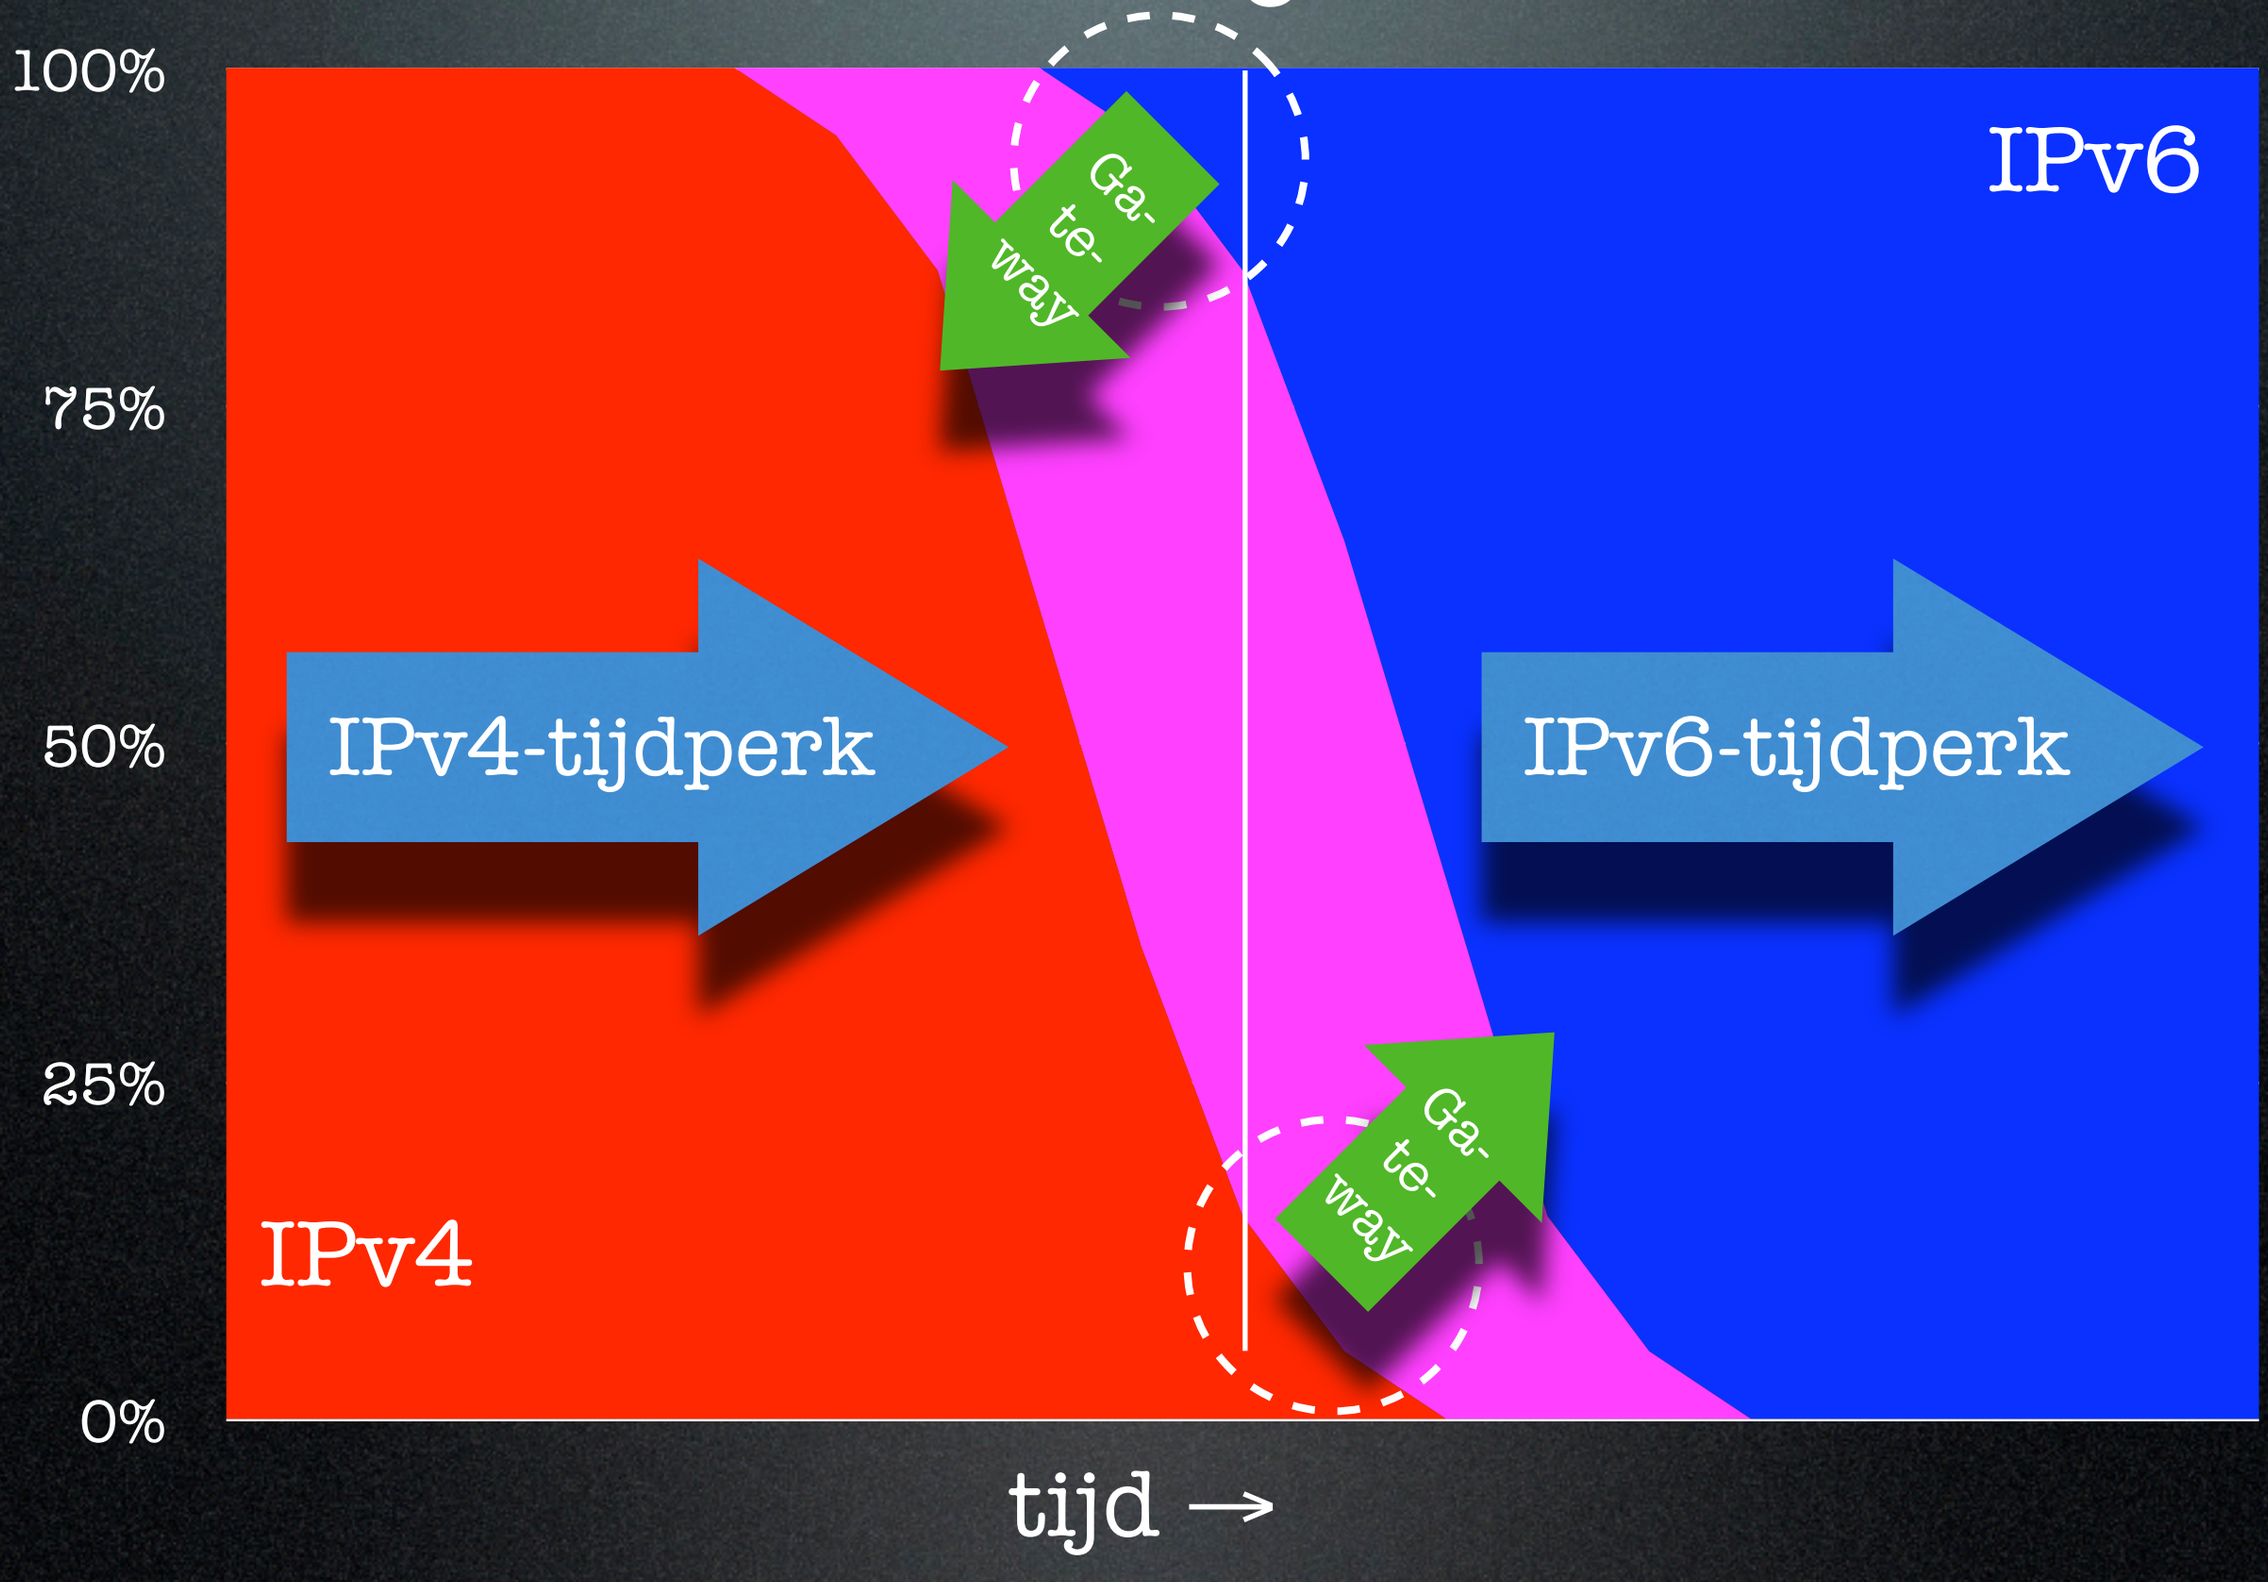

IPv4 Header

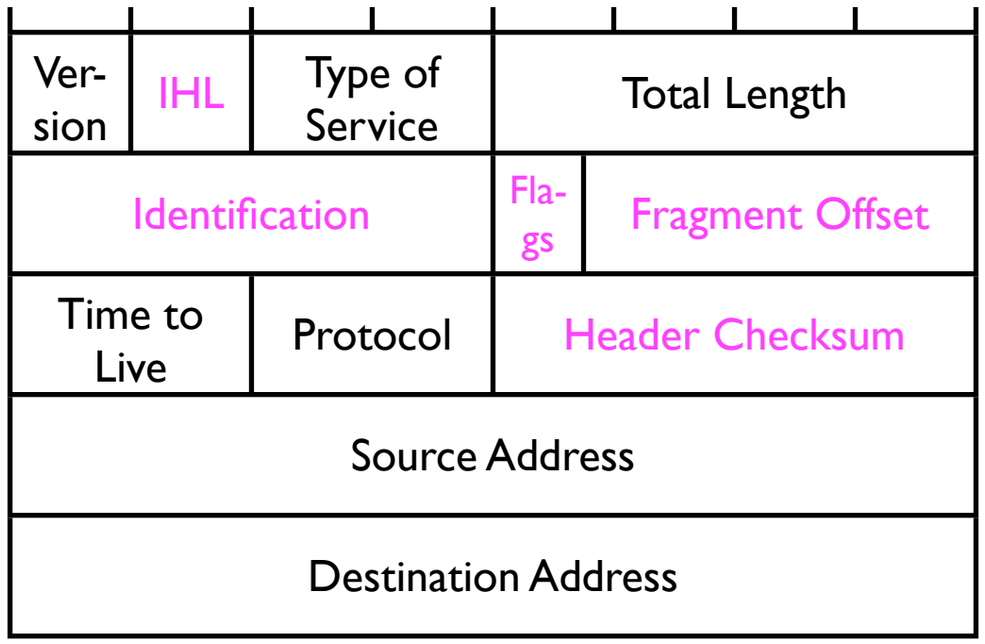

IPv6 Header

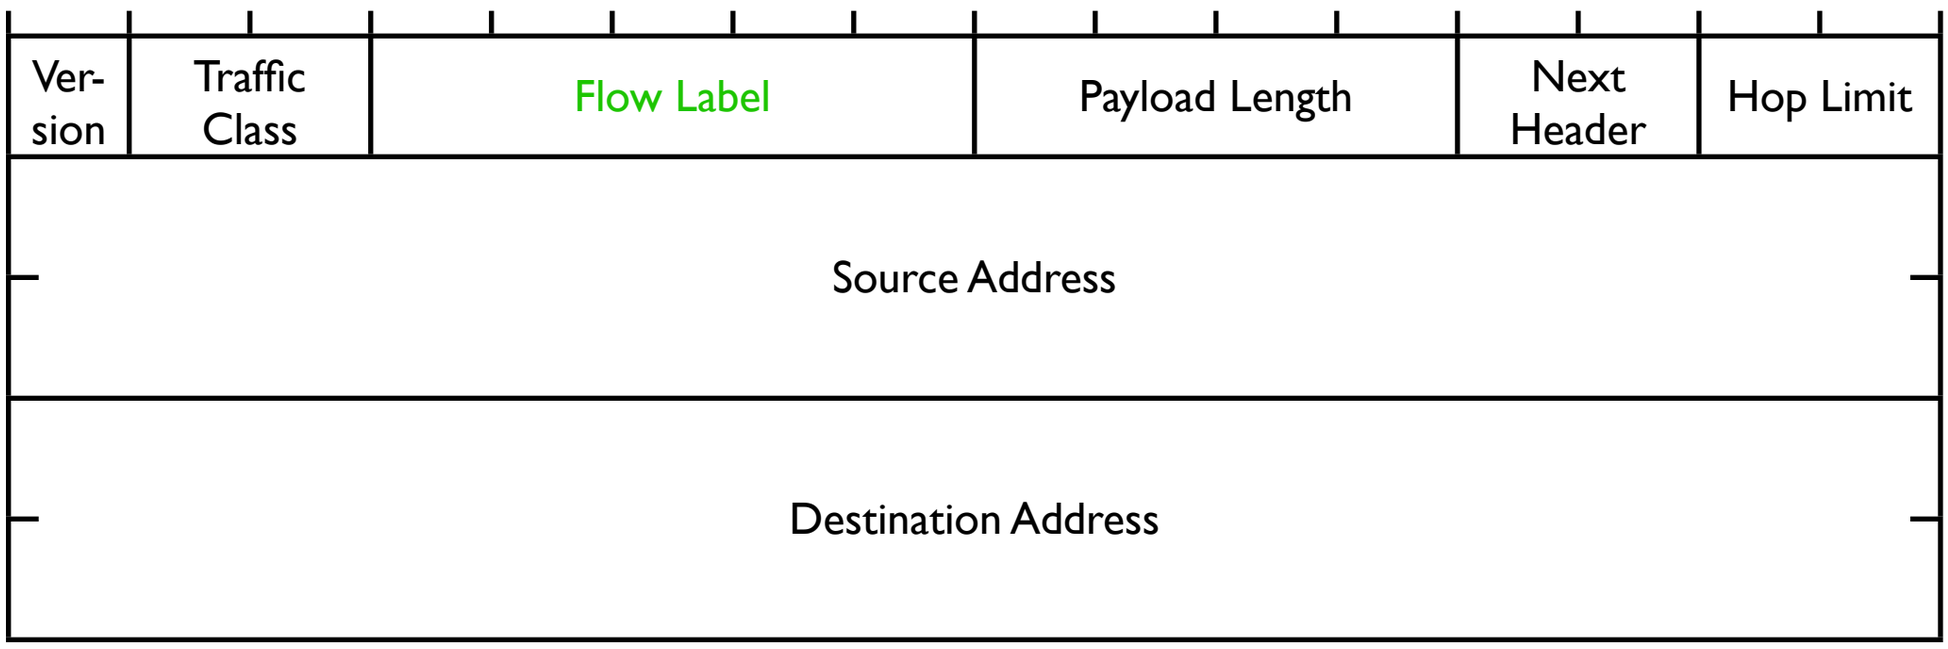

Protocol Chain

- IPv6 header is fixed length
- So options must be implemented in a header of their own!
- For instance, fragmentation header sits between IPv6 header and UDP header
  - (so "next header" rather than "protocol")
- Extra effort to find ultimate payload type...

Address Notation

IPv4:

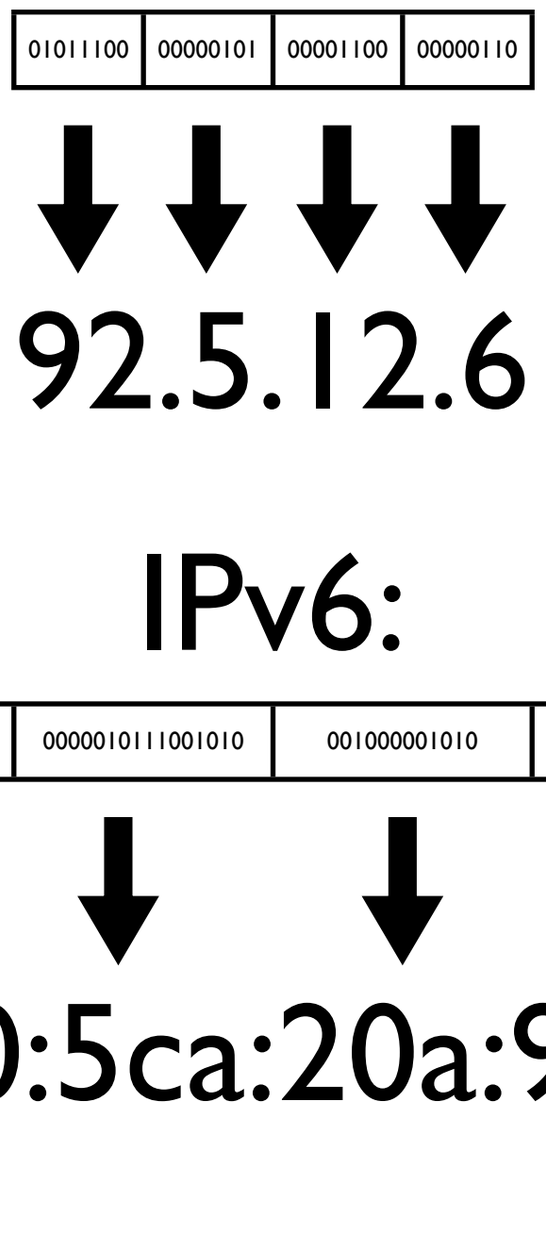

92.5.12.6

IPv6:

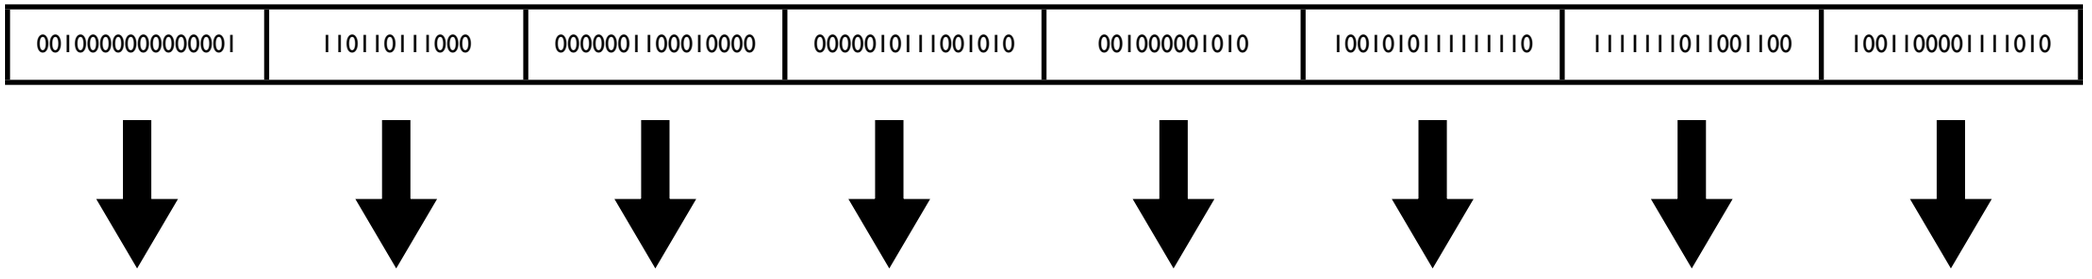

2001:db8:310:5ca:20a:95ff:fecd:987a

Address Notation

- Zero compression for one contiguous range of zeros:
  - 2001:db8:31:0:0:0:0:1 → 2001:db8:31::1
  - 2001:db8:0:0:ff:0:0:1 → 2001:db8::ff:0:0:1
    or 2001:db8:0:0:ff::1, *not* 2001:db8::ff::1
- Localhost: 0:0:0:0:0:0:0:1 → ::1
- Unspecified address: 0:0:0:0:0:0:0:0 → ::

Address Types

- Global unicast: for regular world-wide use
- Site local unicast: within own "site"
- Link local unicast: on a "link" (subnet), won't pass a router
- Multicast: one-to-many
- Anycast: one-to-any (shares unicast space)

IPv6 Address Space

First Bits	Prefix	Purpose
000	::/3	Special uses
001	2000::/3	Global unicast
01 - 1111 1011	4000::/2 - fb00::/8	Reserved
1111 1100	fc00::/8	Unique site local
1111 1110 10	fe80::/10	Link local
1111 1110 11	fec0::/10	Old site local
1111 1111	ff00::/8	Multicast

Enabling IPv6

Windows XP

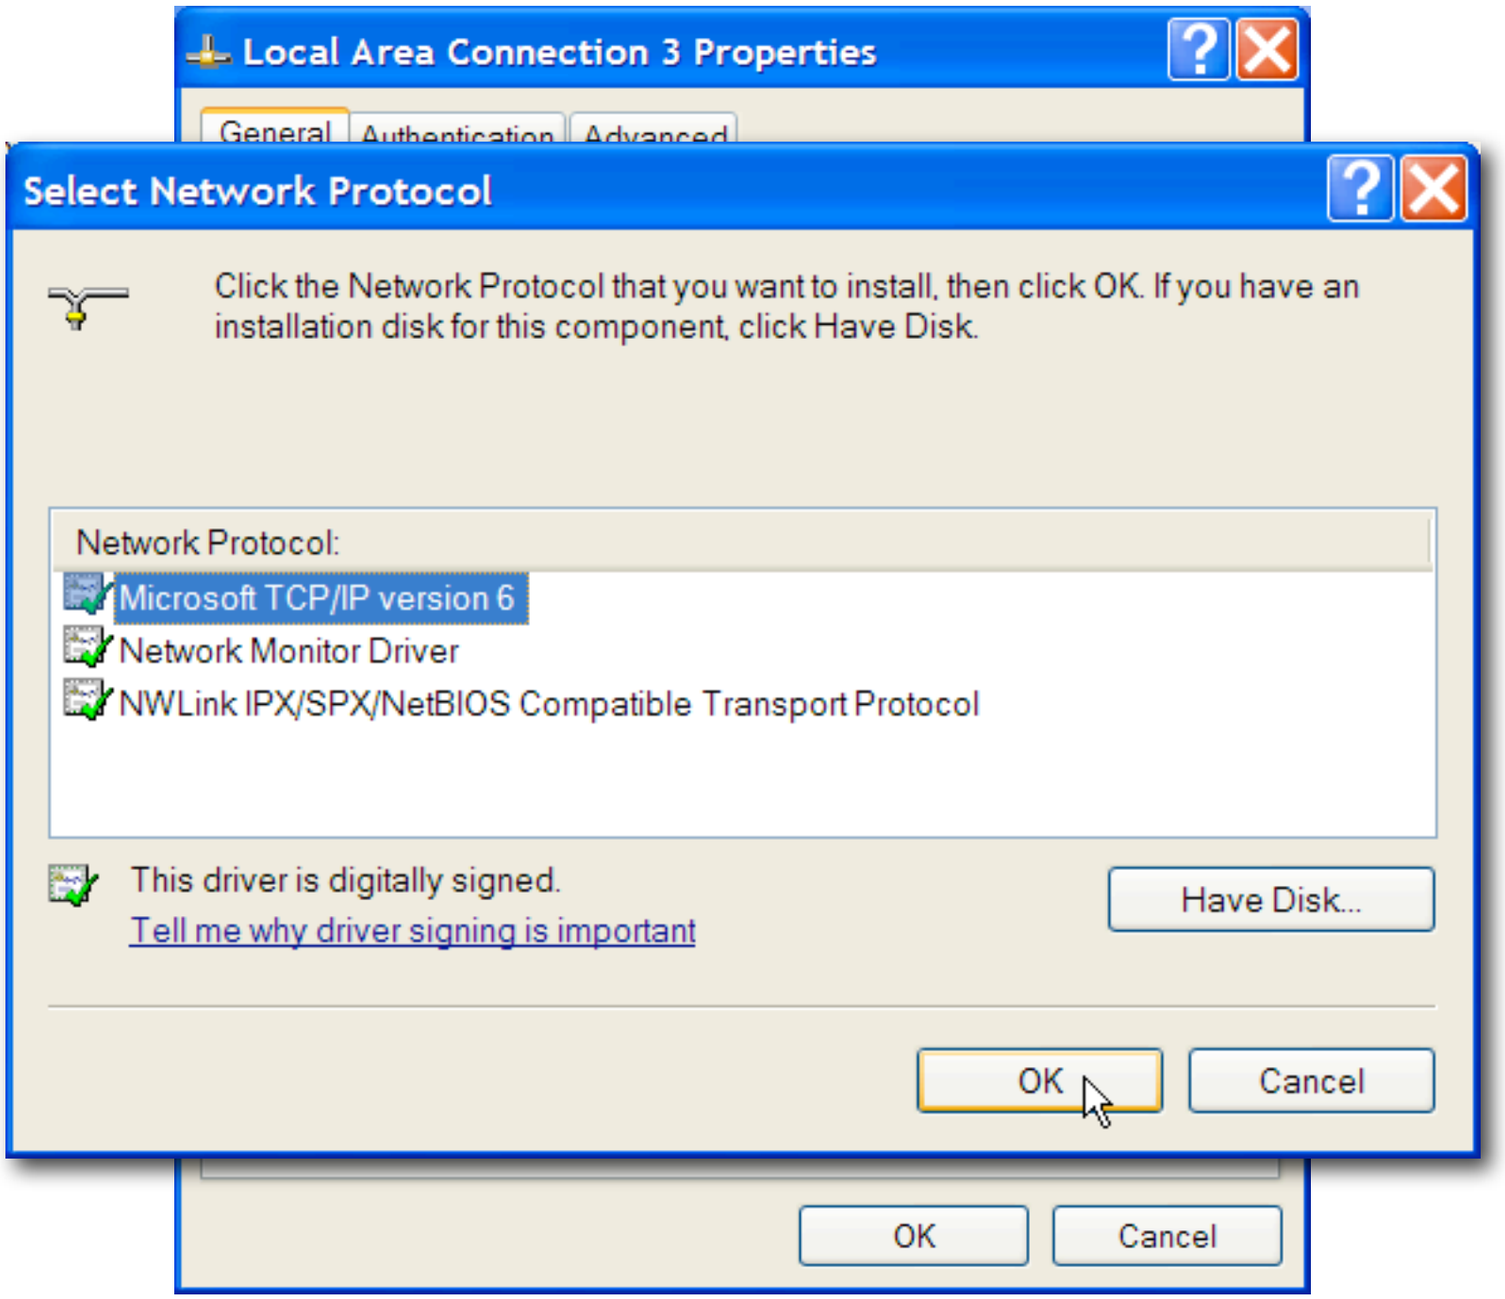

- "ipv6 install"
- Or as of SP1 in GUI
- will then do autoconf
- 6to4 and/or Teredo when no native IPv6 connectivity

MacOS X

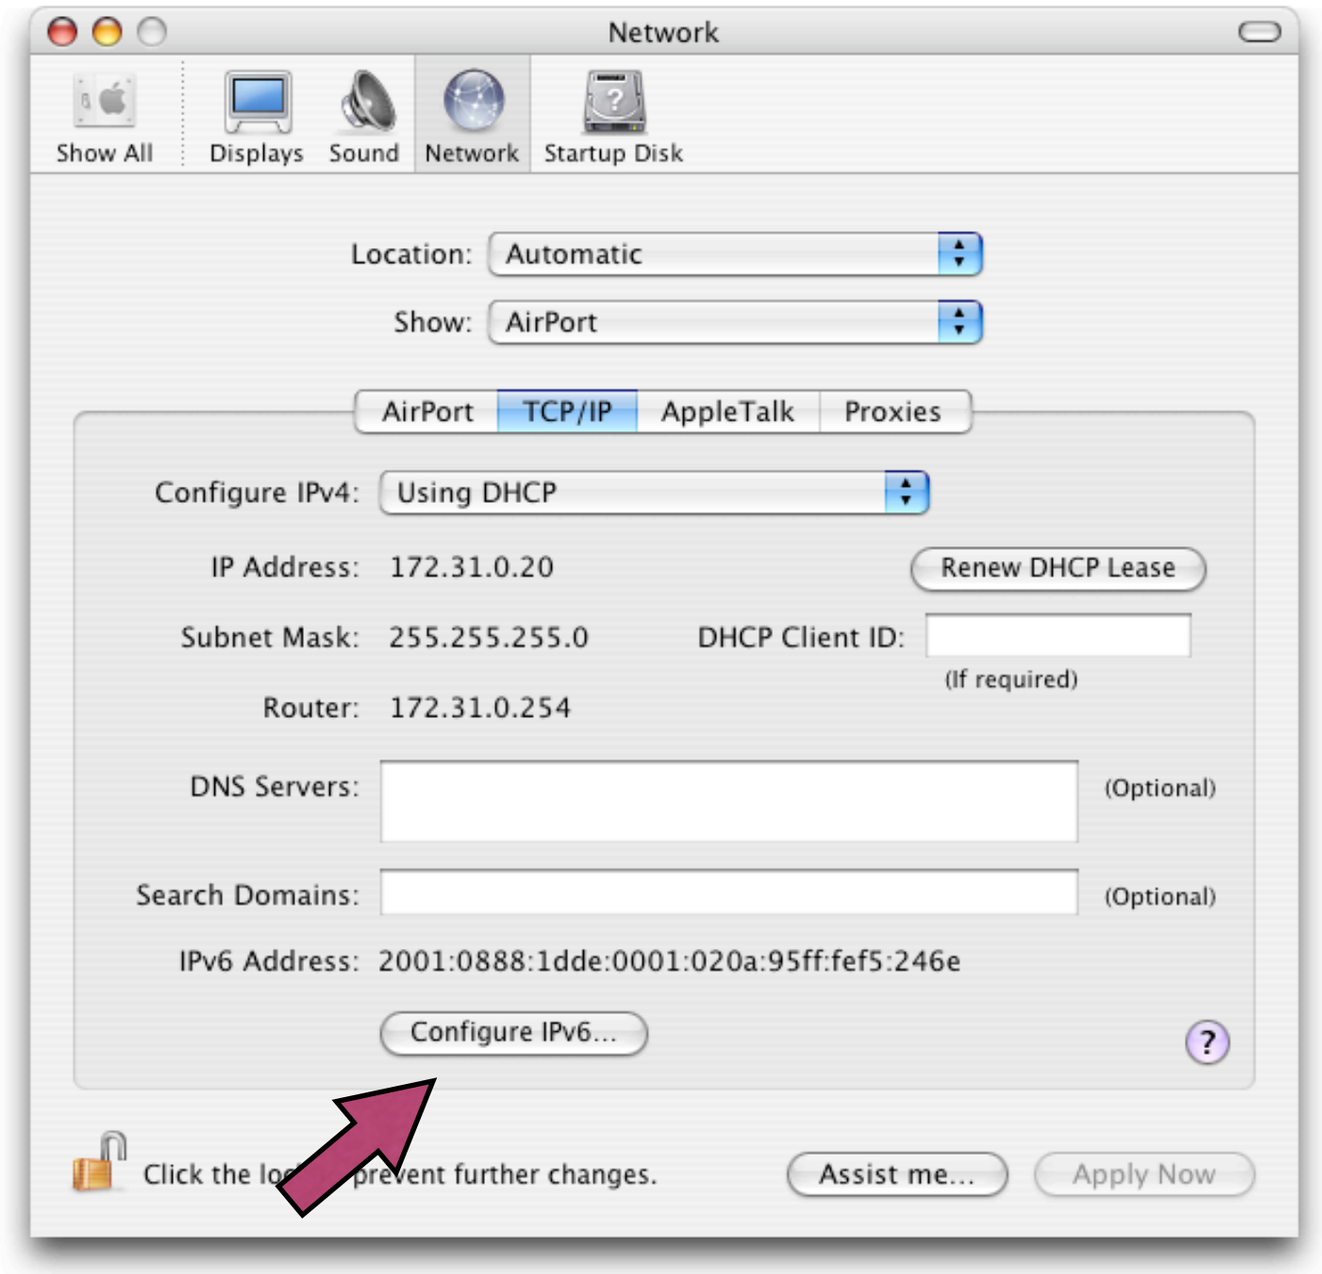

- On and uses auto-conf by default
- Disable/manual setup in Network Preferences
- 6to4 must be enabled manually
- Otherwise Mac networking very much like BSD

FreeBSD

- Usually present in kernel: link-local addresses, but no stateless autoconfig
  - configure in /etc/rc.conf:

```
ipv6_enable="YES"
```

  - (without further configuration) enables stateless autoconfiguration
  - `ipv6_enable="NO"` doesn't do anything

FreeBSD

- Manual configuration (no autoconf):
  - in /etc/rc.conf:

```
ipv6_enable="YES"
ipv6_network_interfaces="em0"
ipv6_ifconfig_em0="2001:1af8:2:5::2/64"
ipv6_defaultrouter="2001:1af8:2:5::1"
```

- Additional options in /etc/defaults/rc.conf

Linux

- As with most stuff: depends on distribution
- Generally IPv6 enabled, autoconf also enabled if no IPv6 forwarding
- In many distros, IPv6 startup scripts poorly documented and buggy
  - may want to working around them

Linux (2)

- Lots of information on:

  http://www.bieringer.de/linux/IPv6/

  - But not always helpful, often outdated
- Have a look at the "ip" package
  - "ip help"
- (BSD is so much easier... ☺)

Red Hat / Fedora

- Usually enabled by default, otherwise:
  - in /etc/sysconfig/network:

    `NETWORKING_IPV6="yes"`
    - will then do autoconfig
    - (again, `"no"` has no effect)
- Additional settings in `/etc/sysconfig/network-scripts/ifcfg-eth0` or similar

Stateless Autoconfiguration

IPv6 Address Creation

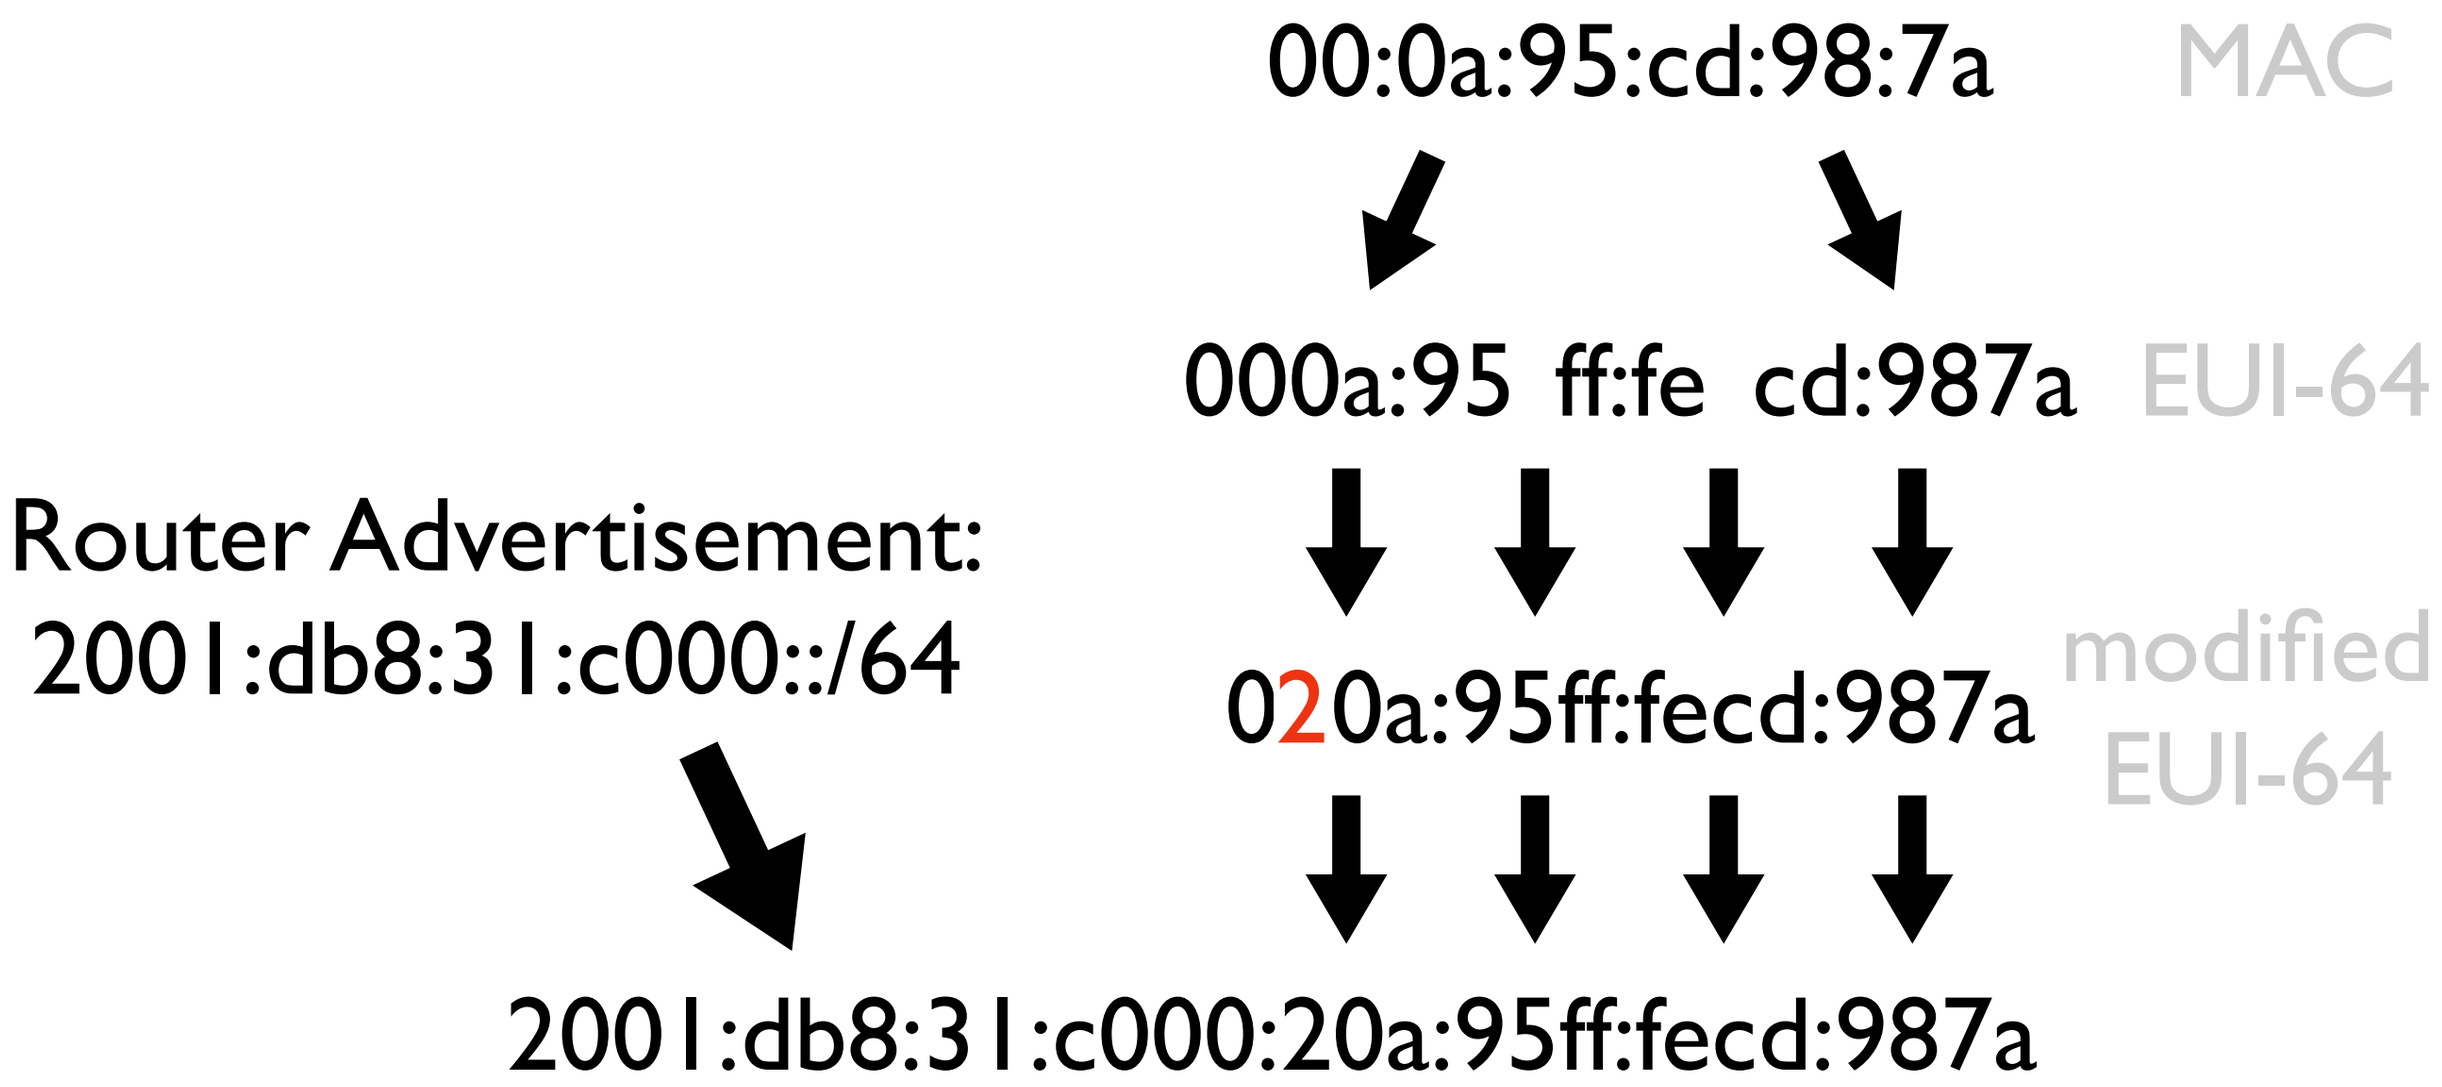

Address Privacy

- Ugh, when you move around people can recognize your MAC address!
- RFC 3041: temporary addresses
  - use random number to generate address
  - generate new one every 24 hours
  - Default for outgoing sessions in Win XP
  - Next level: random MAC address in Vista?

Address Privacy (2)

- Not enabled by default on other OSes
- BSD/MacOS:
  - net.inet6.ip6.use_tempaddr=1
- Linux:
  - net.ipv6.conf.default.use_tempaddr=2
  - (1 = outgoing sessions use MAC-derived)

IPv6 and the Domain Name System

AAAA and ip6.arpa

- A record holds IPv4 address
- AAAA (quad A) record holds IPv6 address
- Reverse like d.e.a.d.b.e.e.f.ip6.arpa
- Old stuff (forget this):
  - A6 records, bitlabels, ip6.int

Signal IPv6 Capability

- Adding AAAA records for a server opens it for IPv6 business
- AAAA and broken IPv6: timeouts or worse
- So may want to use alternative names:
  - www.example.com: A and AAAA
  - www.ipv4.example.com: A
  - www.ipv6.example.com: AAAA

Configuring Resolvers

- Usually no DHCP in IPv6, implementations not very mature
  - it is in Vista and some Linuxes
- Resolver addresses not in router advertisements (yet: RFC 5006)
- Can use IPv4 resolver if dual stack
- Or /etc/resolv.conf

Failure Modes

- "Server fail" for AAAA queries → timeout
  - used to be common, rare now
  - Dumb residential gateways (DLINK?) → timeout for AAAA query
  - (Many OSes/applications do AAAA lookup even when there is no IPv6 connectivity)

Lookup Tools

```
# host -t aaaa www.isc.org

# nslookup -query=aaaa www.isc.org

# dig www.isc.org aaaa
```

- Dig has -4 and -6 flags to force IPv4 or IPv6 transport
- Query types: A, AAAA and many others, ANY returns them all

Applications and API Issues

Socket API

- BSD-style socket API most widely used
- Predates DNS so works with addresses
- Not suitable for 128-bit IPv6 addresses
- So extensions for IPv6 (basic + advanced)
- *All software that uses the socket API must be updated to support IPv6*

Other APIs

- Higher-level APIs don't require changes:
  - Java as of J2SDK/JRE 1.4 `InetAddress` is superclass of `Inet4Address`/`Inet6Address`
  - Stuff like PHP `fsockopen()` works on domain names, so no changes necessary
- When protocols embed IP addresses changes required regardless of API...

Multiple Addresses

- It's normal for IPv6 hosts to have multiple addresses on one interface
- If only because most are dual stack: IPv4 + IPv6
- So one DNS name → multiple addresses
- Applications *should* try all addresses until one works

The Transition

When?

- Some people happy to go to IPv6 now/soon
- Some people very much against it
- Most users: huh?
  - depend on vendors / service providers
- Vendors in reasonable shape
- Service providers: stick with v4 to the end

The End of IPv4

- Small address users: pretty much never
- Large address users: end around 2012, then:
  - Existing large users: fairly light NAT
  - New large users: very heavy NAT
- Heavy Network Address Translation / multiple NATs bad for peer-to-peer

NAT Crunch

- VoIP, BitTorrent, personal servers etc. harder and harder
- IPv6 to bypass NAT
- IPv6 will be promoted by service providers with few IPv4 addresses to be competitive
- People with adequate IPv4 will add IPv6 to talk to others behind NAT

Perception is Reality

- When end of IPv4 address space is undeniable
  - some people will start adopting IPv6 without waiting until the end
  - (see US Department of Defense)
- Likely when available < 3 x yearly address use (now 982 vs 196 million = 5 x)

And Reality is...

- ... that it's easier to build an IPv6 network than an IPv4 network (see routing)
- Can proxy from IPv4 to IPv6 / IPv6 to IPv4
  - NAT-PT (NAT with Protocol Translation) so IPv6-only clients can access IPv4 services
    - deprecated by IETF in summer 2007 (but now trying to revive in some form)
    - Disadvantages of NAT, otherwise easy

Phases

1. Gaining IPv6 experience
   - not much risk, performance irrelevant
2. Adding some IPv6 support
   - do it where it's easy, possible to go back
3. Making IPv6 equal to IPv4
   - legacy and performance issues, no return
4. Turning off IPv4

Planning the Transition

- You need a plan!
- Step zero: don't buy new stuff that can't do IPv6 for sure (i.e., routers w/ IPv4 in ASICs)
- Goal: in 2050 IPv6-only, IPv4 is gone
- Apart from that, keep dates flexible:
  - IPv4 land rush: need to implement fast
  - IPv4 slump: lots of time between steps

The Plan

- Take the easy steps first, save the hard parts for later
- Don't set up IPv6 and then forget about it
- Obviously dependencies: need address space and connectivity
  - own address block easy if ISP
  - otherwise get addresses from ISP

Path MTU Discovery and Fragmentation

Fragmentation

- TCP adjusts to too bigs by reducing packet size
- UDP/ICMPv6 can't
- Solution: host fragments at the source
- Fragment header inserted between IPv6 header and UDP or other payload
  - (firewalls often don't understand this)

Fragmentation (2)

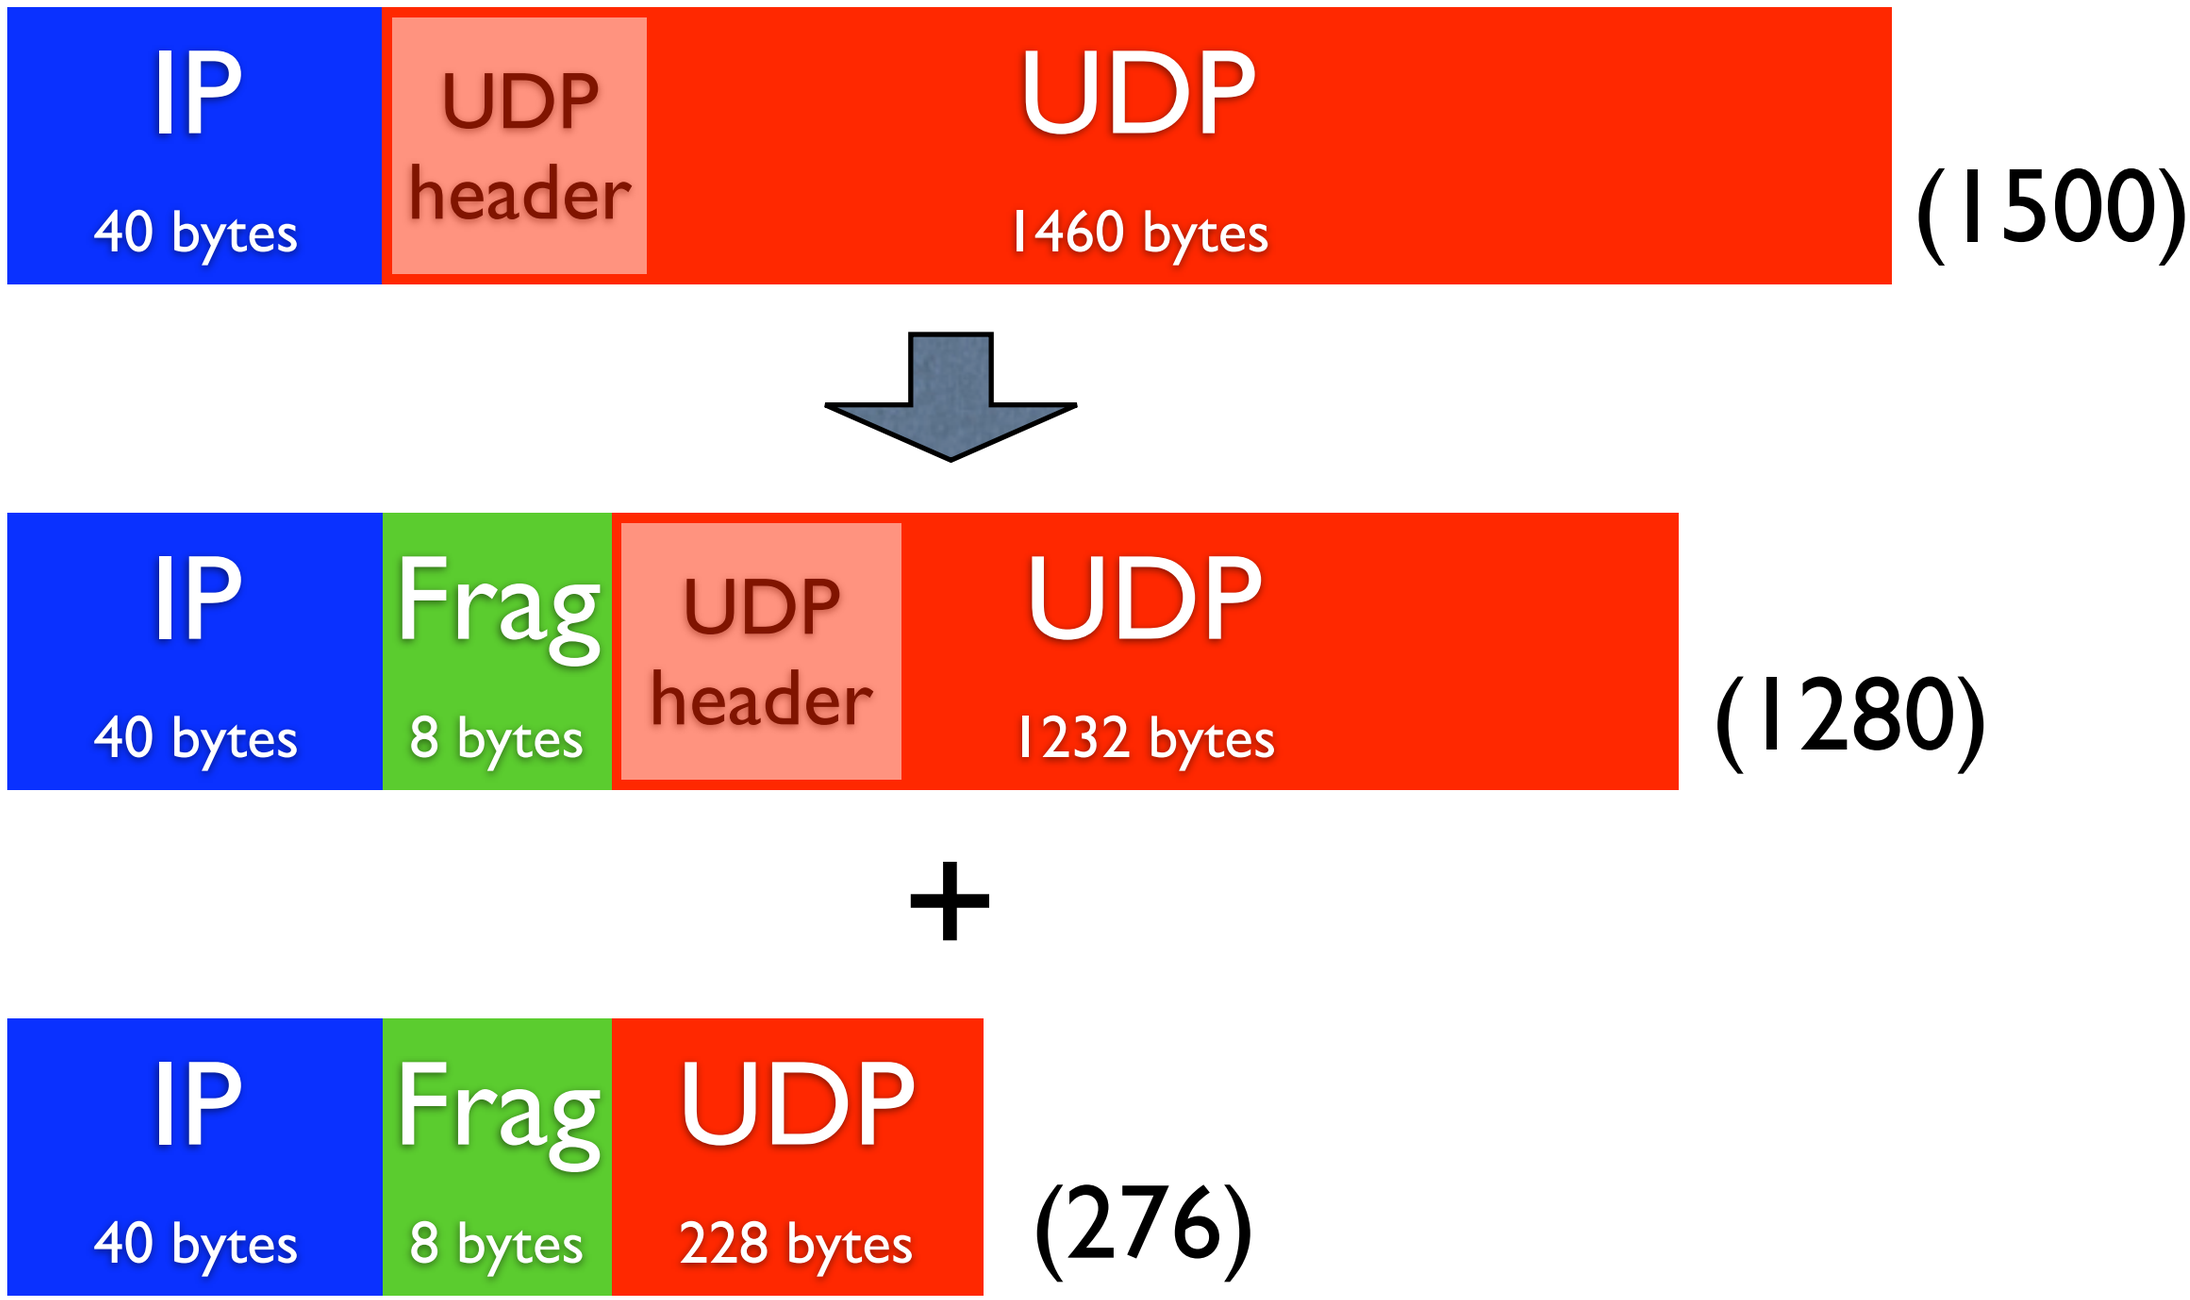

Multicast

Multicast

- No broadcasts in IPv6
- Multicast prefix: ff00::/8
- Ethernet group addresses are MAC addresses 33:33:x:x:x:x

Scope Identifier

- Multicast and link-local addresses not particular to a specific interface
- So need to identifier interface/link on which we send/receive them
- (Or use default interface for this)
- Often with %interface: ff15::1%x10
- Or with flag (especially on Linux)

Security

More Secure?

- IPv6 is supposed to support IPsec
- IPsec is very cool, but...
  - needs lots of config, not in general use
- Nifty hop limit = 255 trick to protect link-local protocols such as ICMPv6 redirect
- Large address space makes random scanning impossible (still targeted scanning)

Less Secure?

- Lots of automatic behavior to corrupt...
  - but generally need to be on local subnet
- Discovery with multicast ping
- Node information queries
- SEND: SEcure Neighbor Discovery, protects ND/RA with crypto (future...)

Less Secure? (2)

- When IPv6 is enabled, ALWAYS link local addresses!
    - possible attack vector on local subnet
- IPv4 filters generally ignore IPv6
    - make sure you aren't accidentally leaving services that are firewalled over IPv4 open over IPv6

Multicast Ping

- Find IPv6 systems on local subnet:
  - ping6 -c 2 -I en0 ff02::1
- All IPv6 systems except latest Windows respond to this

Node Information Query

- (KAME = BSD/MacOS only so far)
- Retrieve names:
  - ping6 -c 2 -I xl0 -w ff02::1
- Retrieve addresses:
  - ping6 -c 2 -I xl0 -a ag ff02::1
  - ping6 -c 1 -a ag www.kame.net

IPv6 Filtering

- Linux: ip6tables
- MacOS/FreeBSD: ip6fw
- FreeBSD: IPF
- OpenBSD/FreeBSD/more: PF
- Windows: netsh firewall
  - (but I don't get how it works)
- Cisco access lists

Careful!

- Don't forget to allow neighbor discovery!
  - (Cisco allows them before implicit deny)
- IPv6 "protocol train" concept where there can be a fragmentation header or other header between IP and TCP/UDP isn't supported by many filters

Tunnels

Tunnels

- Put IPv6 packet inside IPv4 packet to cross IPv4-only parts of the network

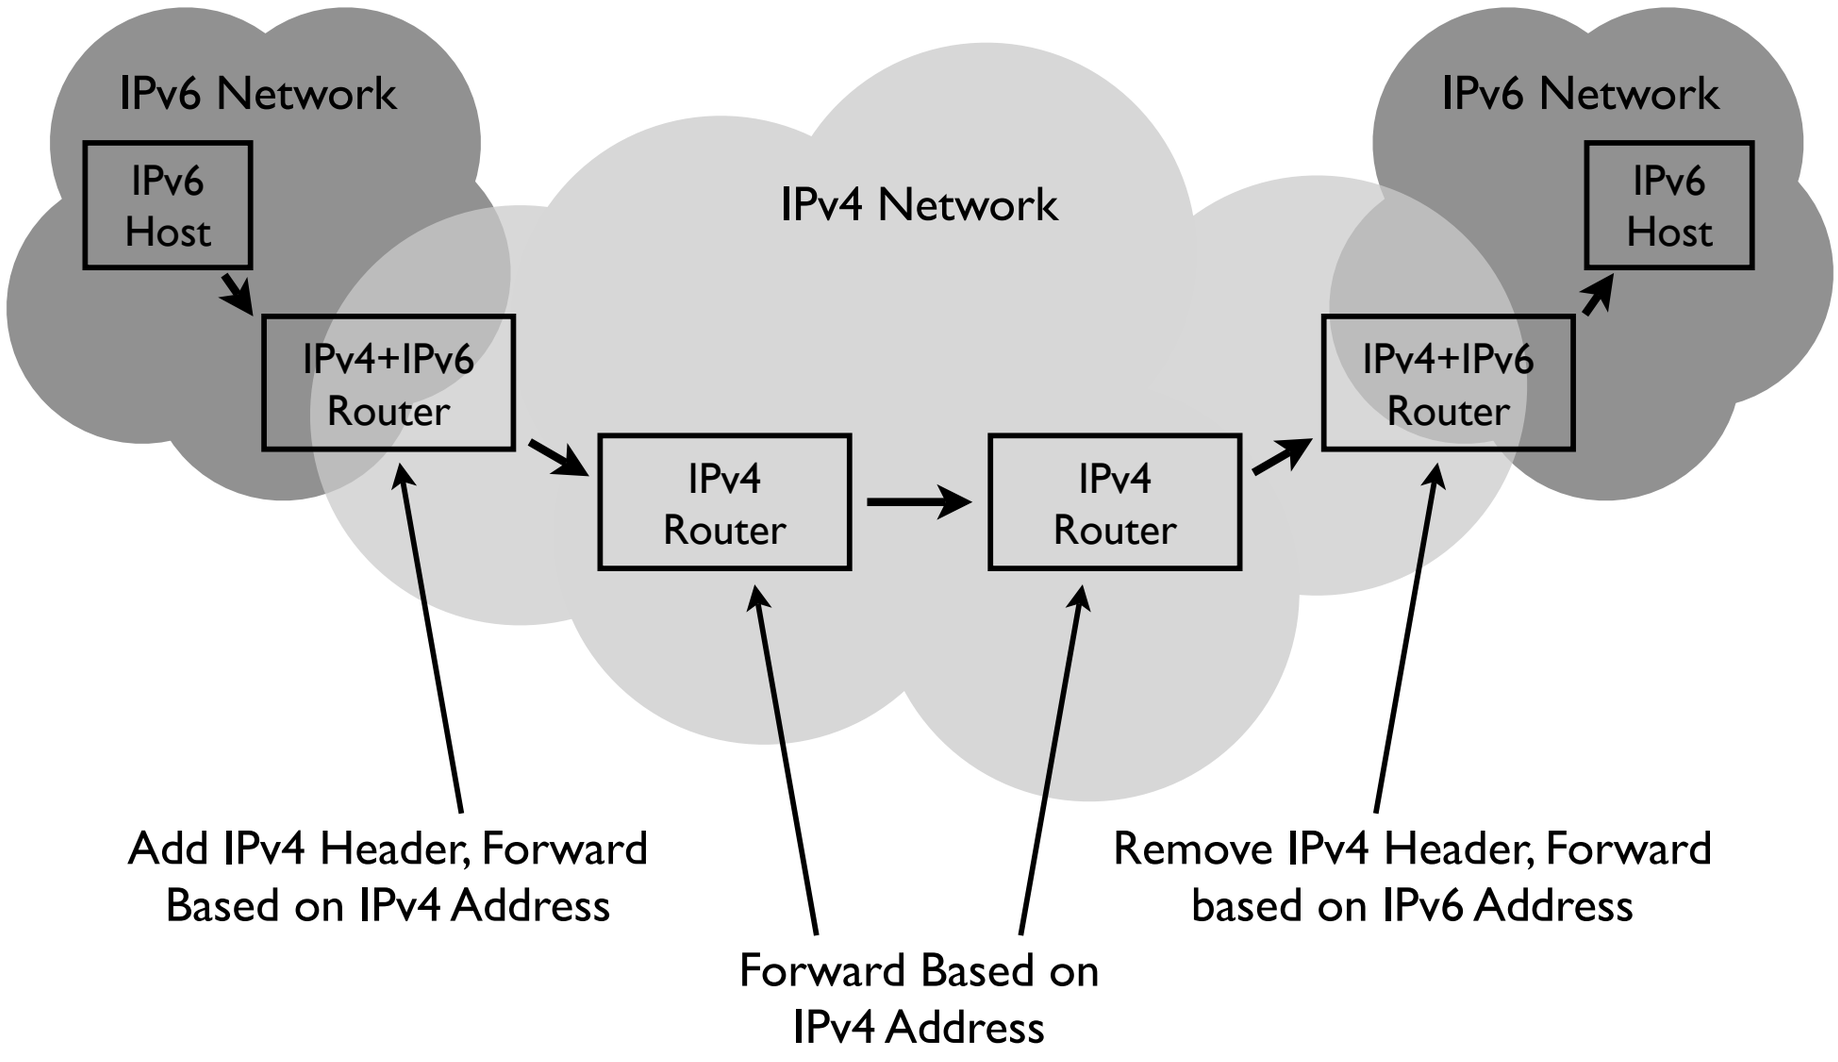

6to4

- (Ignore all automatic tunneling mechanisms everything except 6to4)
  - 6to4: create IPv6 prefix from IPv4 address
    - 6to4 addresses in 2002::/16

6to4 (2)

- Entire 2002::/16 prefix is "directly connected" to 6to4 tunnel interface
- Regular IPv6 addresses reachable through gateway
  - gateway anycast address: 192.88.99.1 (= 2002:c058:6301::)
  - gateways inject 2002::/16 into IPv6 BGP

Manual Tunnels

- Tunnel interface encapsulates IPv6 packet and transmits to remote tunnel endpoint
- Decapsulates incoming IPv6 packets
- Supports routing protocols and multicast
- Can be used in own network, to ISP, ...
- Or "tunnel broker" such as www.sixxs.net

Details

- 6to4 and regular IPv6-in-IPv4 tunnels both use protocol 41 (GRE protocol 47)
  - which may be firewalled / broken by NAT
- Tunnel MTU: often 1280 bytes (minimum maximum packet size in IPv6)
- Also often physical interface - 20 = 1480
- Mismatch can hinder Path MTU Discovery

Addressing

From Where?

- IANA: Addresses come from Internet Assigned Numbers Authority
- ICANN: IANA is now part of the Internet Corportation for Assigned Names and Numbers
- RIRs: IANA distributes to Regional Internet Registries

LIRs

- To keep IPv6 routing table small:
  - only ISPs ("Local Internet Registries") get an address block of their own
- But now a lot of push for provider independent addressing in some form
  - but even then few people qualify

End-Users

- Until now:
  - Need one address? /128
  - Need one subnet? /64
  - Need more? /48
- Talk of /56s (sigh...)
- But ISPs get much leeway

Site Local Addressing

- Like link locals, but for site
  - big mess: what is a site, ambiguity...
- Unique site locals
  - generate site local /48 prefix with random number
  - unique, but for local use
  - nice when merging networks!

Thanks for listening!

http://www.runningipv6.net/

iljitsch@bgpexpert.com