Native IPv6 via xDSL - how to tweak your LNS

DENOG #2: November 4, 2010

SwiNOG #21: November 11, 2010

RIPE #61: November 16, 2010

Fredy Künzler
kuenzler at init7 dot net
AS13030

Init Seven AG
Elias-Canetti-Strasse 7
CH-8050 Zürich, Switzerland

www.init7.net
www.blogg.ch
www.bgp-and-beyond.com

Init Seven AG / Init7

- Internet Service Provider, based in Zurich / Switzerland

- privately owned company

- in Switzerland: mainly Business- and Wholesale Customers

- in Europe: MPLS and IP connectivity (full and partial table)

- own international dual-stacked v4 and v6 backbone (AS13030), mostly 10gig enabled (two rings: Zurich-Geneva-London-Amsterdam-Frankfurt-Zurich and Zurich-Basel-Frankfurt(2)-Zurich)

- connected to 20+ internet exchanges, close to 1000 BGP peers / customers

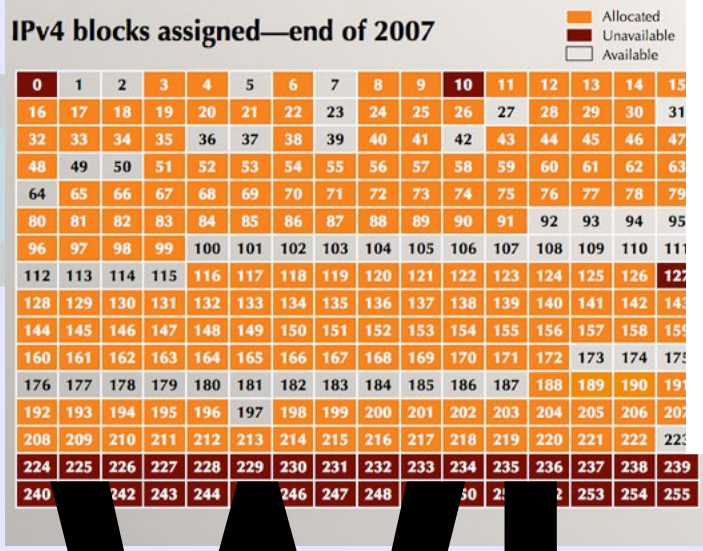
IPv4 blocks assigned—end of 2007
Allocated
Unavailable
Available

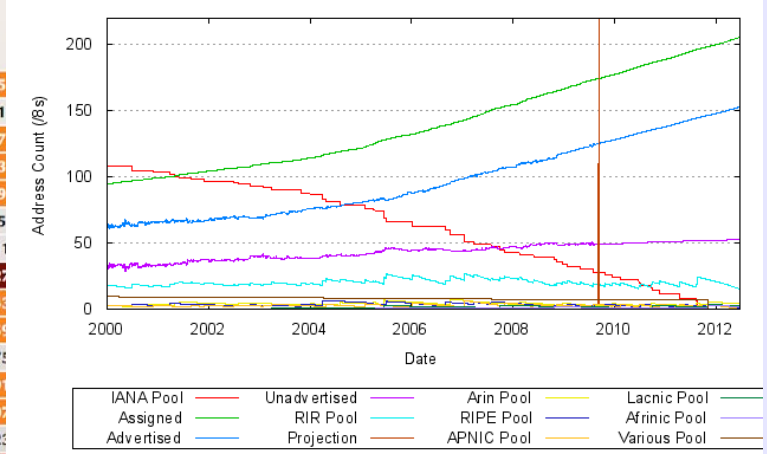
Address Count (/8s)
Date
IANA Pool
Assigned
Advertised
Unadvertised
RIR Pool
Projection
Arin Pool
RIPE Pool
APNIC Pool
Lacnic Pool
Afrinic Pool
Various Pool

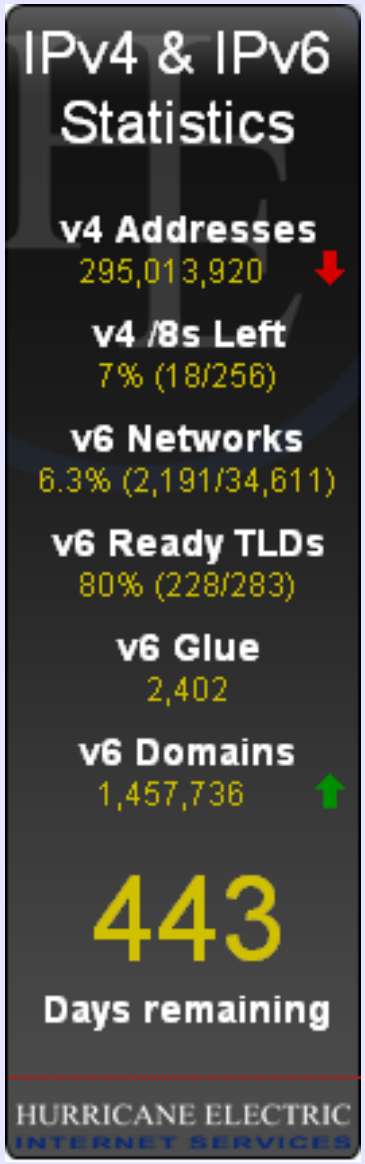
IPv4 & IPv6 Statistics
v4 Addresses
295,013,920
v4 /8s Left
7% (18/256)
v6 Networks
6.3% (2,191/34,611)
v6 Ready TLDs
80% (228/283)
v6 Glue
2,402
v6 Domains
1,457,736
443
Days remaining
HURRICANE ELECTRIC
INTERNET SERVICES

Why?

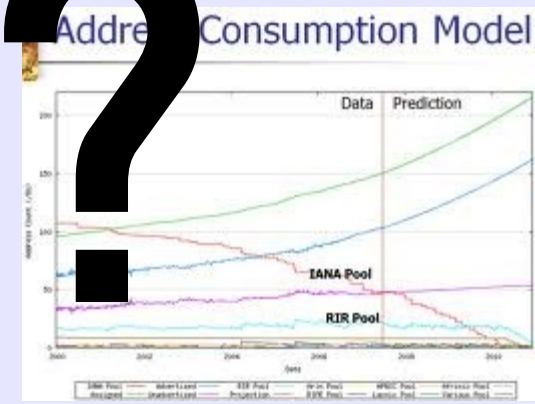
Addre Consumption Model
Data
Prediction
IANA Pool
RIR Pool

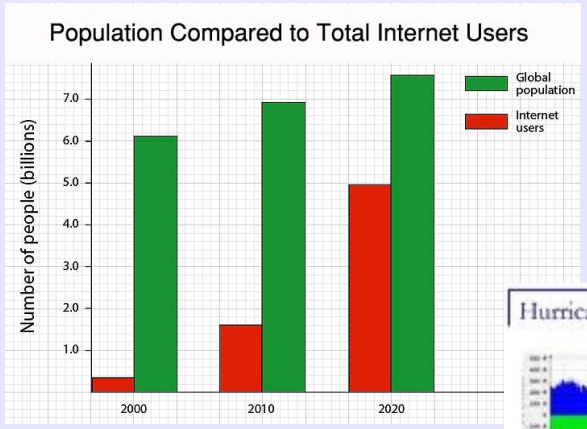
Population Compared to Total Internet Users
Number of people (billions)
Global population
Internet users
2000
2010
2020


Cogent (As 174)
Pleas IPv6 Peer With Us
xoxox
HURRICANE ELECTRIC
(AS6939)

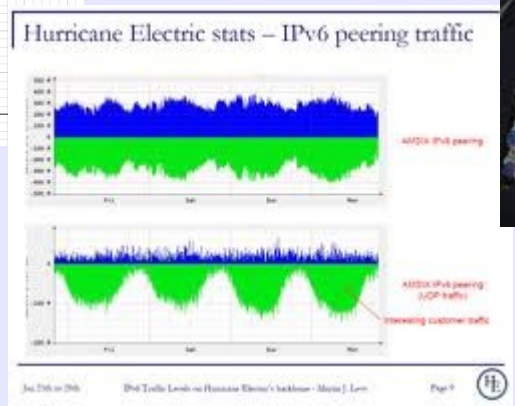
Hurricane Electric stats – IPv6 peering traffic

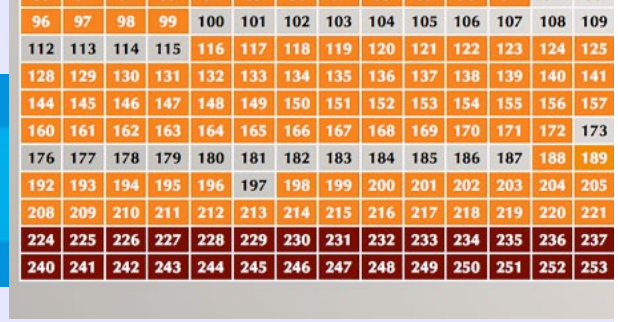

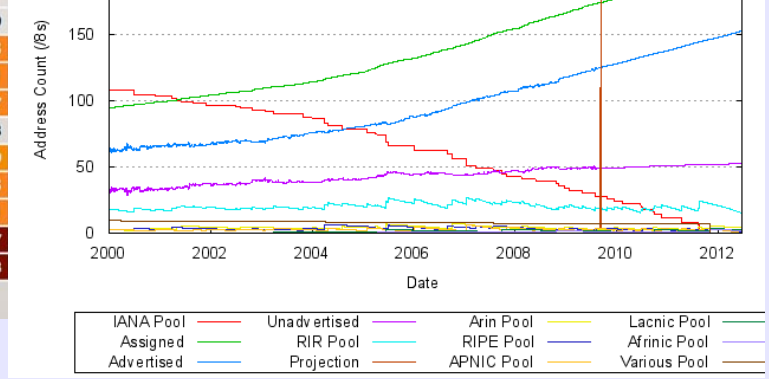

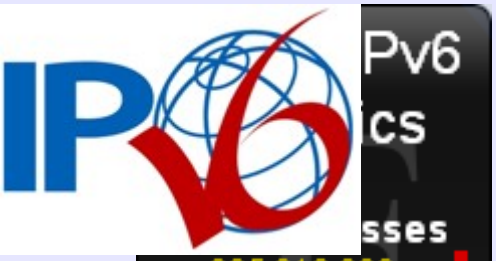

Because these three elderly men keep saying it for years and years and years...

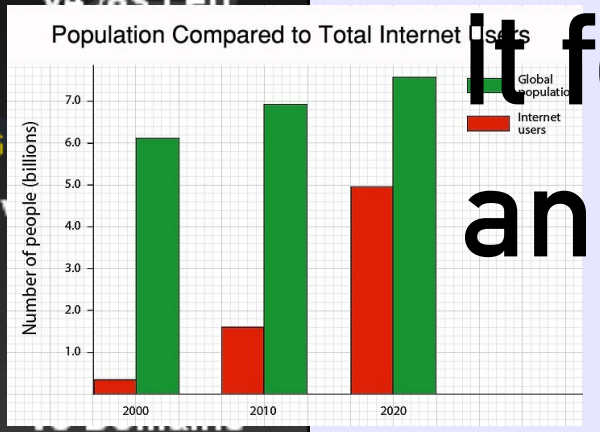

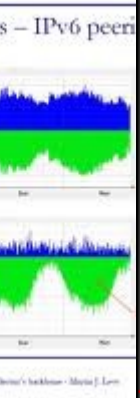

Disclaimer:

These slides show how Init7 implemented native IPv6 for it's xDSL customers. The examples show how it worked for Init7, they are not considered as “correct” or “right”.

IPv6 addressing and deployment can be argued. The examples may give an indication but there is no guarantee it will work in your environment, and there are plenty of different ways to implement it .

Init7 is also not able to support your infrastructure, but is interested in corrections and comments nevertheless.

The 'chicken or egg' discussion is over, whether access or content provider need to implement IPv6 first … because the chicken has laid an egg:

```
#ping6 -n www.heise.de
PING www.heise.de(2a02:2e0:3fe:100::7) 56 data bytes
64 bytes from 2a02:2e0:3fe:100::7: icmp_seq=1 ttl=57 time=8.18 ms
64 bytes from 2a02:2e0:3fe:100::7: icmp_seq=2 ttl=57 time=7.92 ms
64 bytes from 2a02:2e0:3fe:100::7: icmp_seq=3 ttl=57 time=7.85 ms
```

It's now time for the access provider to make a move and dual-stack their customers connections …

Unfortunately not every customer has a simple ethernet plug like those in the datacenter… there are nasty things such as L2TP and Radius authentication…

Prerequisites – assuming that

… you are a LIR with a /32 or bigger IPv6 address block assigned from RIPE (ARIN, APNIC …)

… your backbone is IPv6 ready (dual stack) – if not, please do your homework first. See the slides „IPv6 deployment for the IPv4 clueful“:
http://www.blogg.ch/uploads/IPv6-deployment-for-the-IPv4-clueful-heise-ipv6-conference.pdf

… you are a wholesale partner of the incumbent with a working xDSL aggregation equipment (L2TP) and the LNS is Cisco hardware.

then these slides are for you.

The slides may also be helpful if your ISP provides native IPv6 via xDSL – it gives some hints how to configure the CPE device.

IPv6 Address Plan for xDSL customers #1:

Big discussion in the community whether an end user should get a /48 or a /56. Latest proposal, kind of a compromise is a /52... these slides are not going to argue the pros and cons for /48 or /56.

Init7 decided for a /48 per xDSL customer, but it works with a /56 too … make your choice. Either way the customer has plenty of space which lasts forever.

Init7's example how to assign /48 address blocks to xDSL customers is just one working example – there are of course many other possibilities, which work as good as this example.

IPv6 Address Plan for xDSL customers #2:

Init7 has several IPv4 C-classes reserved to xDSL customers, mostly assigned statically by the Radius login credentials. For simplicity the IPv6 address scheme follows the IPv4 address scheme.

For each C-class one /40 IPv6 block is reserved …

(1 C-class => 256 IPv4 addresses => 256 customers => 256* /48 => requires one /40 address block)

IPv6 Address Plan for xDSL customers #3:

Example:
192.168.201.0/24 - xDSL range #1
...
192.168.209.0/24 - xDSL range #9
192.168.210.0/24 - xDSL range #10
…

and the IPv6 ranges for these customers would be:
2001:db8:100::/40 – xDSL range #1
...
2001:db8:900::/40 – xDSL range #9
2001:db8:a00::/40 – xDSL range #10
…

etc.

IPv6 Address Plan for xDSL customers #4:

The last octet of the IPv4 address will be converted into HEX and used for the IPv6 address assignment.

Example #1:

IPv4 client with 192.168.201.10 gets 2001:db8:10a::/48

Explanation #1:

'10' is the last octet of the IPv4 address. In HEX it's 'a' (equals '0a', and as we previously gave the 2001:db8:100::/40 range to the customers within 192.168.201.0/24, this customer will get 2001:db8:10a::/48.

IPv6 Address Plan for xDSL customers #5:

Example #2:

IPv4 client with 192.168.210.227 gets 2001:db8:ae3::/48

Explanation #2:

'227' is the last octet of the IPv4 address. In HEX it's 'e3', and as we previously gave the 2001:db8:a00::/40 range to the customers within 192.168.210.0/24, this customer will get 2001:db8:ae3::/48.

IPv6 Address Plan for xDSL customers #6:

Point-to-Point IPv6 Link – an IPv6 Address for the „WAN"-Port of the CPE device is required.

Again, the address scheme is adapted from the existing IPv4 address.

For the PtP IPv6 links of all customers we reserve another /48 range, in this example:

2001:db8:4::/48

IPv6 Address Plan for xDSL customers #7:

The PtP link addressing is calculated as follows:

2001:db8:4:xxyy::/64

xx: 3rd octet of the IPv4 WAN address in hex
yy: 4th octet of the IPv4 WAN address in hex

Example:

192.168.201.10/32 → 2001:db8:4:c90a::/64

(201d → 0xc9, 10d → 0x0a)

xDSL aggregation network scheme:

LNS (Layer 2 Network Server), Radius Server: operated by the ISP

LAC (Layer 2 Access Concentrator): operated by the incumbent

CPE (Customer Premises Equipment): Router of the Enduser, can be managed by the ISP

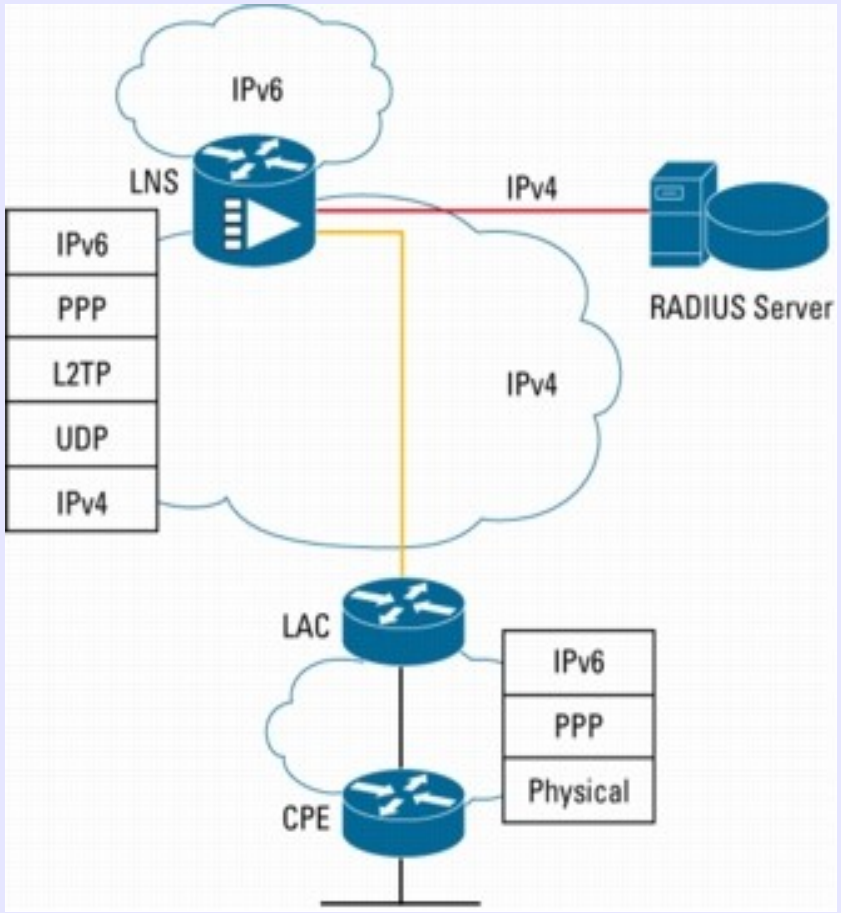

The good news!

IPv6 in such an environment can be implemented without interaction or support from the incumbents helpline (“IPv6? You mean IPTV? Of course, would you like the premium soccer package as well?”) … this fact will speed up the deployment of native IPv6 on xDSL for decades …

The L2TP is completely transparent – just leave the interconnection parameters of the LNS untouched.

IOS Version known to work

… schedule a maintenance window if required.

```
#sh version
Cisco IOS Software, 2800 Software (C2800NM-ADVIPSERVICESK9-
M), Version 12.4(13b), RELEASE SOFTWARE (fc3)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Tue 24-Apr-07 16:18 by prod_rel_team

ROM: System Bootstrap, Version 12.4(13r)T, RELEASE SOFTWARE
(fc1)

lns uptime is 1 year, 20 weeks, 2 days, 19 minutes
System returned to ROM by power-on
System restarted at 05:54:21 CET Mon Jun 22 2009
System image file is "flash:c2800nm-advipservicesk9-mz.124-
13b.bin"
```

IPv6 on LNS Deployment #1:

Enable IPv6 on the LNS:

Cisco:

```
!
ip cef    !required by ipv6 cef
!
ipv6 unicast-routing
ipv6 cef
!
```

IPv6 on LNS Deployment #2:

Configure Interfaces on LNS:

Cisco:

```
!
interface lo0
 ipv6 enable
 ipv6 address 2001:db8:1::101/128
!
interface gi0/0
 ipv6 enable
 ipv6 address 2001:db8:2::52/126
!
```

IPv6 on LNS Deployment #3:

Configure IPv6 OSPF on LNS:

Cisco:

```
!
ipv6 router ospf 1
 passive-interface default
 no passive-interface x/y
!
interface x/y
 ipv6 ospf 1 area 0
 ipv6 ospf cost 1
!
```

Make sure that you are not talking OSPF to external interfaces (passive-interface default).

IPv6 on LNS Deployment #4:

Check OSPF adjacencies and the IPv6 routing table:

Cisco:

```
# sh ipv6 ospf neighbor

# sh ipv6 route [ospf|connected|static]
```

Ensure that IPv6 connectivity is established.

IPv6 on LNS Deployment #5:

Enable IPv6 on DHCP / AAA:

Cisco:

```
!
aaa authorization configuration v6PREFIXDELEGATION group radius
!
ipv6 dhcp pool v6DHCPPOOL
 prefix-delegation aaa method-list v6PREFIXDELEGATION
!
```

IPv6 on LNS Deployment #6:

Enable IPv6 on Virtual Interface:

Cisco:

```
!
interface Virtual-Template1
 ipv6 enable
 no ipv6 nd suppress-ra
 ipv6 dhcp server v6DHCPPOOL
!
```

IPv6 on LNS Deployment #7:

Redistribute IPv6 address ranges from Virtual Access Interfaces:

Cisco:

```
!
route-map Virtual-Access-only deny 10
 match interface [all except Virtual-Access]
!
route-map Virtual-Access-only permit 20
!
ipv6 router ospf 1
 redistribute connected metric-type 1 route-map Virtual-Access-only
!
```

IPv6 on LNS Deployment #8:

[Redistribute IPv6 address ranges from Virtual Access Interfaces]

DISCLAIMER:

This solution doesn't scale. Redistributing every customers /48 into OSPF is not what we want if the IPv6 customer base is above a few dozen – aggregation and distribution of the corresponding /40 will be required.

We keep this item on the 'to do list' (at this time Init7 has some 10+ active IPv6 clients on the xDSL infrastructure).

IPv6 on LNS Deployment #9:

Radius Database – Example #1:

(IPv4 parameters are for reference only)

```
UserName            Attribute             op    Value
(Static IPv4 Address)
'username1@domain', 'Framed-IP-Address', ':=', '192.168.201.10'

(IPv6 Netblock)
'username1@domain', 'cisco-avpair',      '+=', 'ipv6:route#1=2001:db8:10a::/48'

(P2P IPv6 Link)
'username1@domain', 'cisco-avpair',      '+=', 'ipv6:prefix#1=2001:db8:4:c90a::/64'
```

Note: Init7 uses FreeRadius.

IPv6 on LNS Deployment #10:

Radius Database – Example #2:

(IPv4 parameters are for reference only)

```
UserName              Attribute              op    Value
(Static IPv4 Address)
'username2@domain', 'Framed-IP-Address', ':=', '192.168.210.227'

(IPv6 Netblock)
'username2@domain', 'cisco-avpair',      '+=', 'ipv6:route#1=2001:db8:ae3::/48'

(PtP IPv6 Link)
'username2@domain', 'cisco-avpair',      '+=', 'ipv6:prefix#1=2001:db8:4:d2e3::/64'


(optional, if the client also gets a /29 Ipv4 routed)
'username@domain', 'Framed-Route',      '+=', '192.168.17.232/29 0.0.0.0 1'
```

CPE sample IPv6 config (Example #1)

… known to work on Cisco 887V

```
!
ipv6 unicast-routing
ipv6 cef
!
! ADVANCED IP Services NEEDED on 880 Plattform
!
license udi pid CISCO887V-K9 sn xxxxxxxxxx
license boot module c880-data level advipservices
!
controller VDSL 0
!
interface Ethernet0
 no ip address
 pppoe-client dial-pool-number 1
!
interface FastEthernet0
 switchport access vlan XX
!
interface VlanXX
 ip address 192.168.21.1 255.255.255.0
 ip nat inside
 ip virtual-reassembly
 ipv6 address 2001:db8:10a:1:1/64
 ipv6 enable
!
interface Dialer1
 ip address negotiated
 ip nat outside
 ip virtual-reassembly
 encapsulation ppp
 dialer pool 1
 ipv6 address 2001:db8:4:c90a::1/64
 ipv6 address autoconfig default
 ipv6 virtual-reassembly
 ppp authentication chap callin
 ppp chap hostname username1@domaim
 ppp chap password 0 password
 ppp ipcp dns request accept
 ppp ipcp route default
 no cdp enable
!
ip nat inside source list ACL4-NATSRC interface Dialer1 overload
!
ip access-list extended ACL4-NATSRC
 permit ip 192.168.1.0 0.0.0.255 any
!
```

Native IPv6 on xDSL. Now!

Questions?

Fredy Künzler, AS13030
kuenzler at init7 dot net

www.init7.net
www.blogg.ch
www.bgp-and-beyond.com