IPv6 Business Conference

Organized by SWISS IPv6 COUNCIL

2014 June, 17

Case Studies – an Overview

SWISS IPv6 COUNCIL

Taking Off

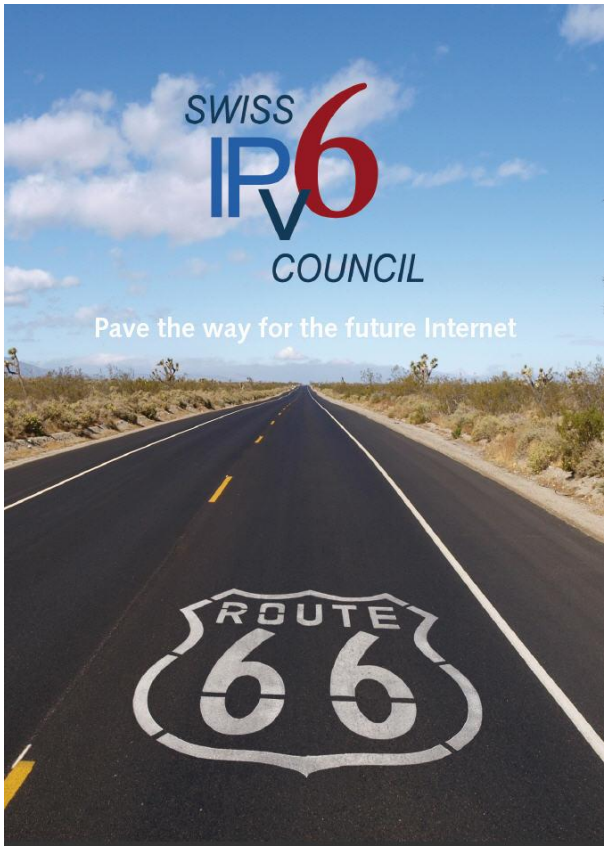

IPv6 World Day, June 8, 2011

Schweizerische Post

Wer sind wir
Projektteam, Projekt

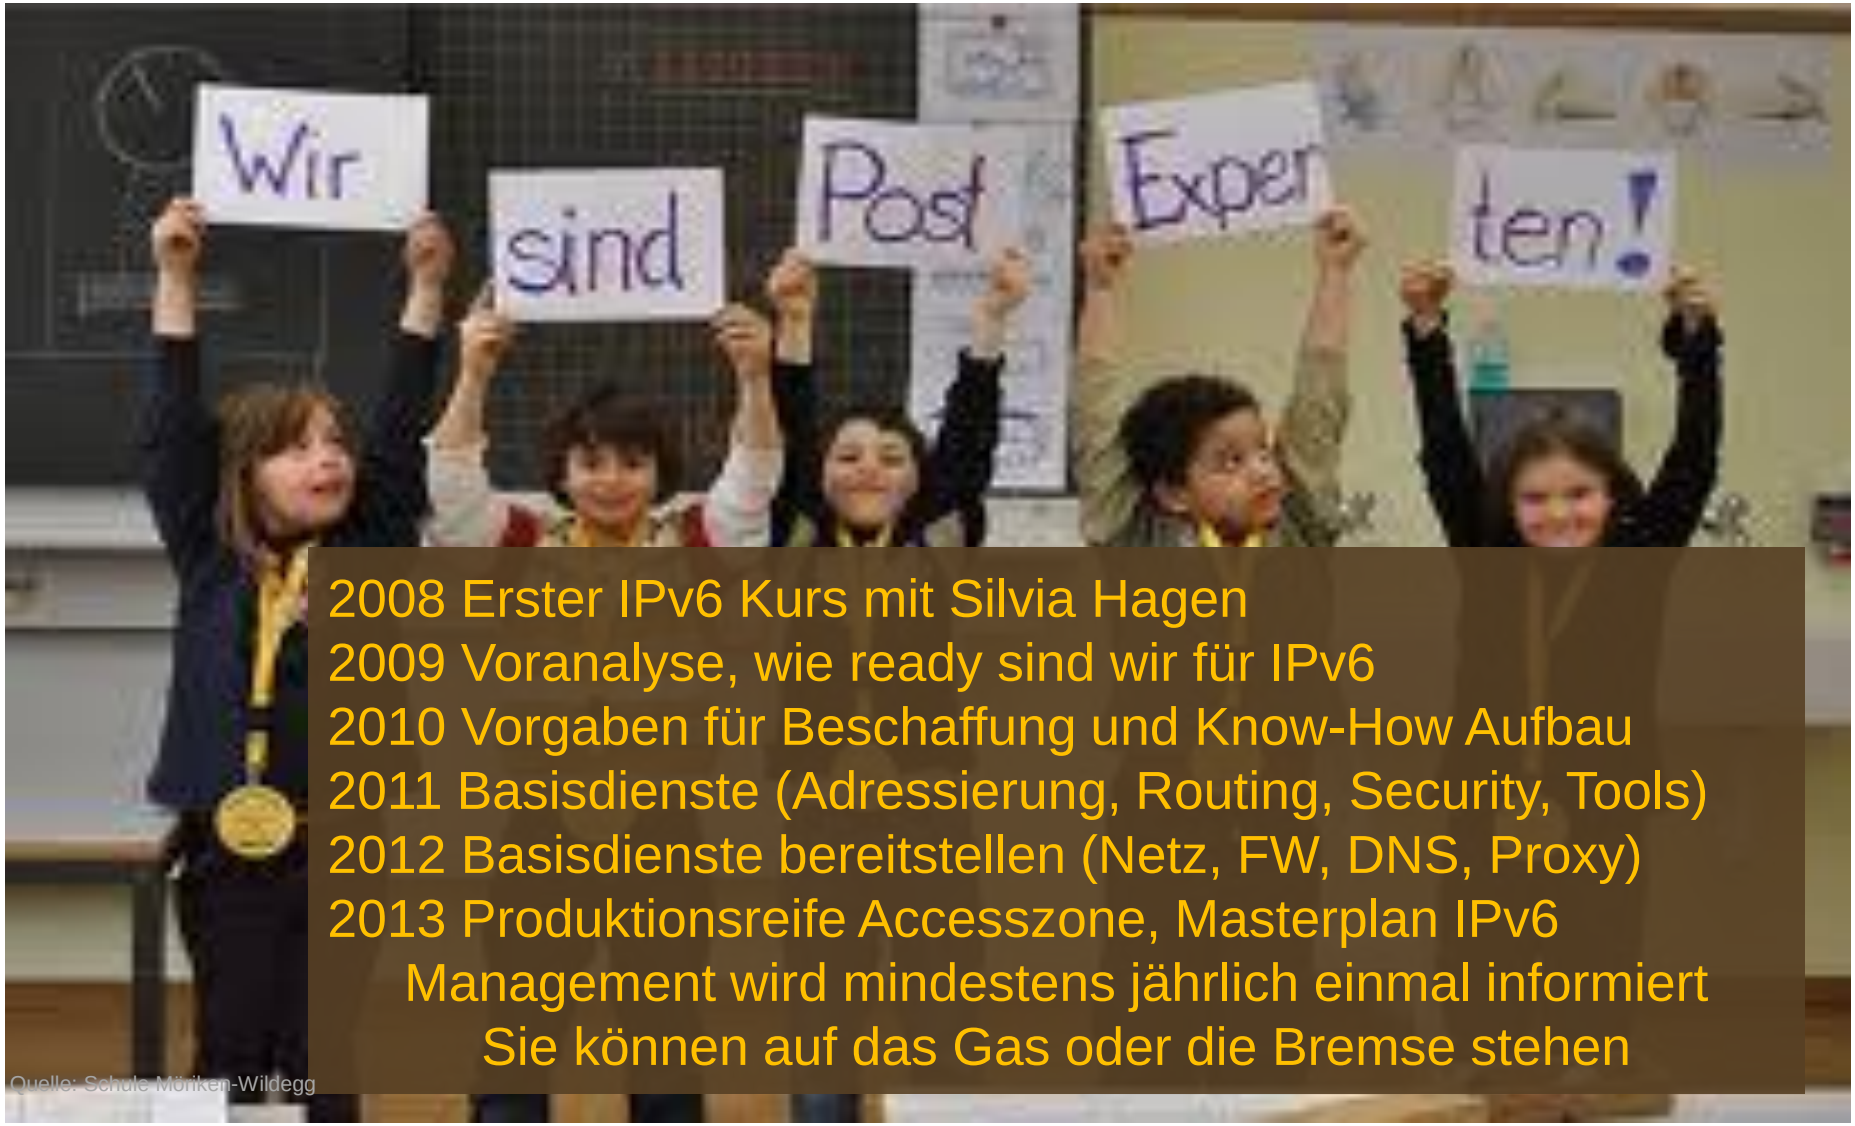

2008 Erster IPv6 Kurs mit Silvia Hagen
2009 Voranalyse, wie ready sind wir für IPv6
2010 Vorgaben für Beschaffung und Know-How Aufbau
2011 Basisdienste (Adressierung, Routing, Security, Tools)
2012 Basisdienste bereitstellen (Netz, FW, DNS, Proxy)
2013 Produktionsreife Accesszone, Masterplan IPv6
Management wird mindestens jährlich einmal informiert
Sie können auf das Gas oder die Bremse stehen

Quelle: Schule Möriken-Wildegg

Wo der Mischkonzern Post tätig ist
Vier Märkte

- **Logistikmarkt**

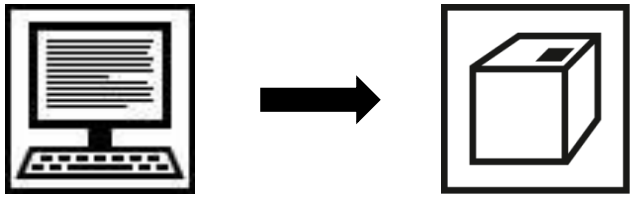

- **Retailfinanzmarkt**

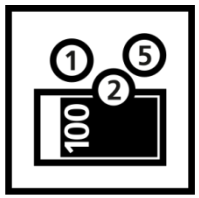

- **Markt für öffentlicher Personenverkehr**

- **Kommunikationsmarkt**

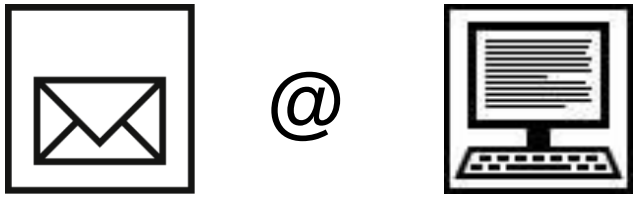

Milestones

- 2013 Post wird Aktiengesellschaft
- Most Domain entries had to be adjusted – this was used to add IPv6 records at the same time
- Saved a lot of money!
- Took more than a year.

Addressing

- Post Worldwide /32
  - Country Aggregation /40
- Post Switzerland /36
    - Region Aggregation /48
    - Local Aggregation /56
- Optimized for routing
- Based on geography. Each country has its own legal situation.
- Internet Services are offered over Internet only. Everything else is routed internally.
- Default Gateways (link-local) are standardized
- No identifiers in IPv6 address (e.g. phone, VLAN etc)
- No ULA's
- No decimal identifications in IPv6 address

Der Weg vom Home User zur Post
Wo stehen wir heute

User *(Home User)*
Ready OS, Browser

ISP Home User *(Internet Service Provider)*
Ready Schweiz ca. 10 %

ISP Post *(Internet Service Provider)*
Ready seit 2012

DDos Abwehr beim ISP
Ready Swisscom, Sunrise

Netz *(Access Zone / DMZ)*
Ready seit 2012

IPS / Firewall *(Intrusion Prevention System)*
Ready seit 2012

DNS *(Domain Name Server)*
Ready seit 2012

GSLB *(Global Server Load Balancer)*
Ready, Ersatz im Lifecycle

NTP (*Network Time Protocol)*
Ready seit September 2013

LSLB *(Local Server Loadbalancer)*
Ready seit 2012

Web Server *(Testserver)*
Ready seit 2012

XML GWY *(Extensible Markup Language)*
Ready seit 2013

Was es sonst noch braucht
Wo stehen wir heute

Adressverwaltung *(IPAM/Inventar)*
Ready seit 2013

Prozesse *(Workflow Tool)*
Ready seit September 2013

ADS *(Access Detection System)*
Ready seit 2012

NetViz *(Network Visualization)*
Ready seit 2013

MRTG/Cacti *(Messen, Statistik)*
Not Ready MIB's fehlen

Outgoing Proxies *(IPv6 Sites erreichen)*
Ready seit Herbst 2013

Mail (*Contentfilter)*
Not ready, geplant 2014

Homepage Post *(Internetauftritt Post)*
Not ready, Projekt für Ablösung

Management *(Server/Network)*
Not ready, zur Zeit auf IPv4

Schulung *(Know How)*
30 MA drei Tage Schulung

Issues

- Waiting for new NTP releae that can also work with off-link systems (worked only with link-locals, waiting 15 months for fix)
- Had to wait for update of loadbalancer services. Migration took quite some time.
- Addressmanagement: Want to work with standard products – not possible yet to fulfill all post-specific requirements, too many manual processes still needed, no full integration of inventory and address management
- A bug in Bluecoat prevented internal clients to reach dual-stack websites over IPv6, has now been solved (filtered on addressses instead of URLs)

Issues continued

- Mail/Contentscanner was announced for fall 2013. Tests are running. Now updating production systems, done by fall 2014.
- Design scenarios for downgrade in case of problems.
- Purchasing guidelines that have requirements at feature level are a MUST
- Threats:
  - Lack of vendor implementation especially in endpoint security
  - RA guard and Destination guard are important (Cisco announces support for both in Q2/2014

Status

- Backbone is IPv6 ready and can be activated
- Firewalls and network infrastructure are ready by 2015, maybe even by end of 2014
- IPv4 islands are identified (Paket- und Briefzentren, due to third party systems not IPv6-enabled, long lifecycles)
- Swiss Post will be reachable over IPv4 and IPv6 in the Internet (DNS, Mail and Website) by Quarter 3 2014
- Currently building testlab for applications
- We know what it takes to integrate IPv6 in the datacenter and in the intranet, Masterplan and roadmap defined

Fragen
bringen uns weiter

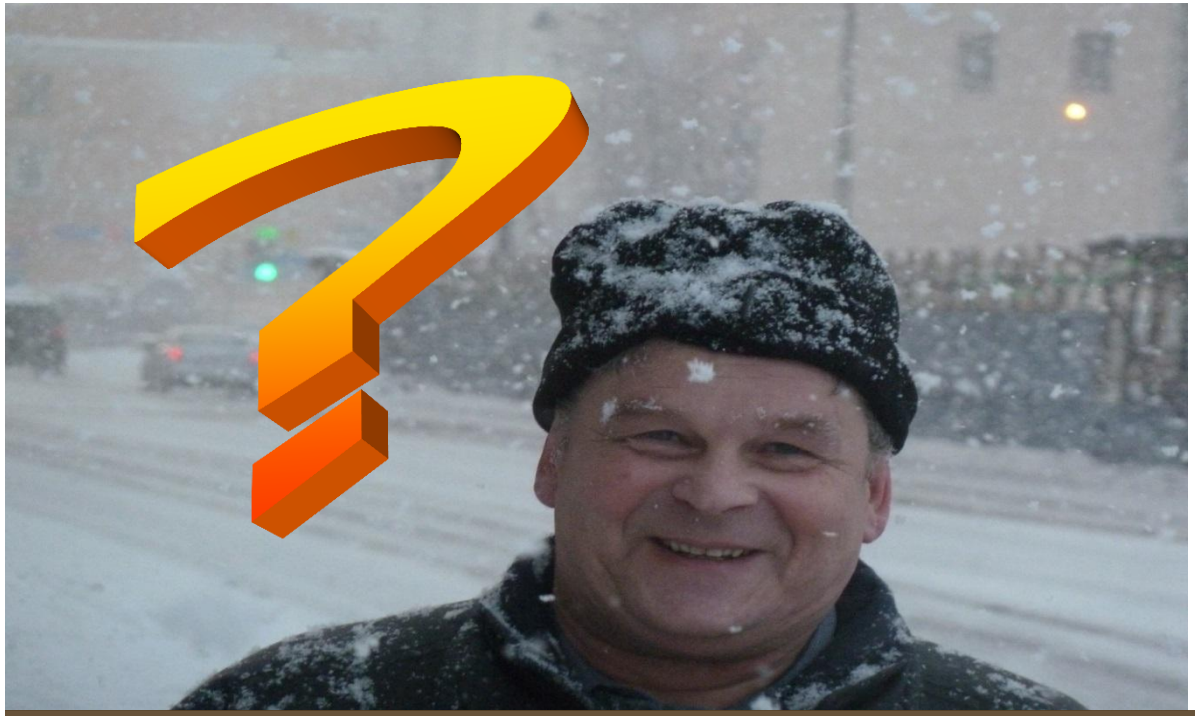

robert.buerk@post.ch

Full presentation can be downloaded at:
http://www.swissipv6council.ch/de/news/Download-Slides-IPv6-Einfuehrung-Post

Alle Photos: Robert Bürk

IPv6 @ SBB

SBB CFF FFS

- Founded in 1902

SWISS IPv6 COUNCIL

Reasons to plan for IPv6

SBB CFF FFS

- Lack of IPv4 addresses is real (Asia)
- IPv6 is the requirement for continued growth of the Internet
- RIPE only assigns small blocks of IPv4 addresses from the last /8
- Current operating systems enable IPv6 by default and prefer it over IPv4 communication
- IPv6 in an unmanaged environment can cause security issues
- IPv6 restores end-to-end connectivity, no more NATs needed.

SWISS IPv6 COUNCIL

Strategy

SBB CFF FFS

- General migration szenario is dual-stack. If needed IPv6 islands will be used.
- Dual-stack will be used for many years.
- First step is activation of website sbb.ch, planned for October 2014. First stage uses translation.
- Purchasing guidelines and application development guidelines have been updated to support IPv6
- RFP in 2014 for Datacom Next Generation, will be initialized in 2015. All detailed plans for migrating network will be done with new partner.
- Testlab has been implemented

Projekt

SBB CFF FFS

- Projektleiter:
- Marc Pauli, Plattform Manager Security

- Statement:
- The biggest challenge in our project is to get enough internal ressources. IPv6 still has not enough attention, and often gets forgotten behind daily business and short term thinking.

SWISS IPv6 COUNCIL

Aspectra

- Founded in 2000
- Hosting. Each customer has its dedicated zone and servers.
- Target customers: high availability, high security
- 2 Datacenters
- 23 employees

Finding a provider

- Asking for IPv6 we get the following answers:
  - The second largest referred us to sales (meaning: we don't do IPv6)
  - The largest said "we can, in 3 – 6 months, it'll cost xxxx
  - The smallest said "yes we can, no problem, here's the contract

Paradigm shift

- No hierarchy in address space
- Each project gets its own address space (project over customer, customer may have several projects – helps a lot in traceability – 12 bits)
- DC encoded in address, 4 bits
- Zone encoded, 12 bits

- Take your time to rethink current concepts

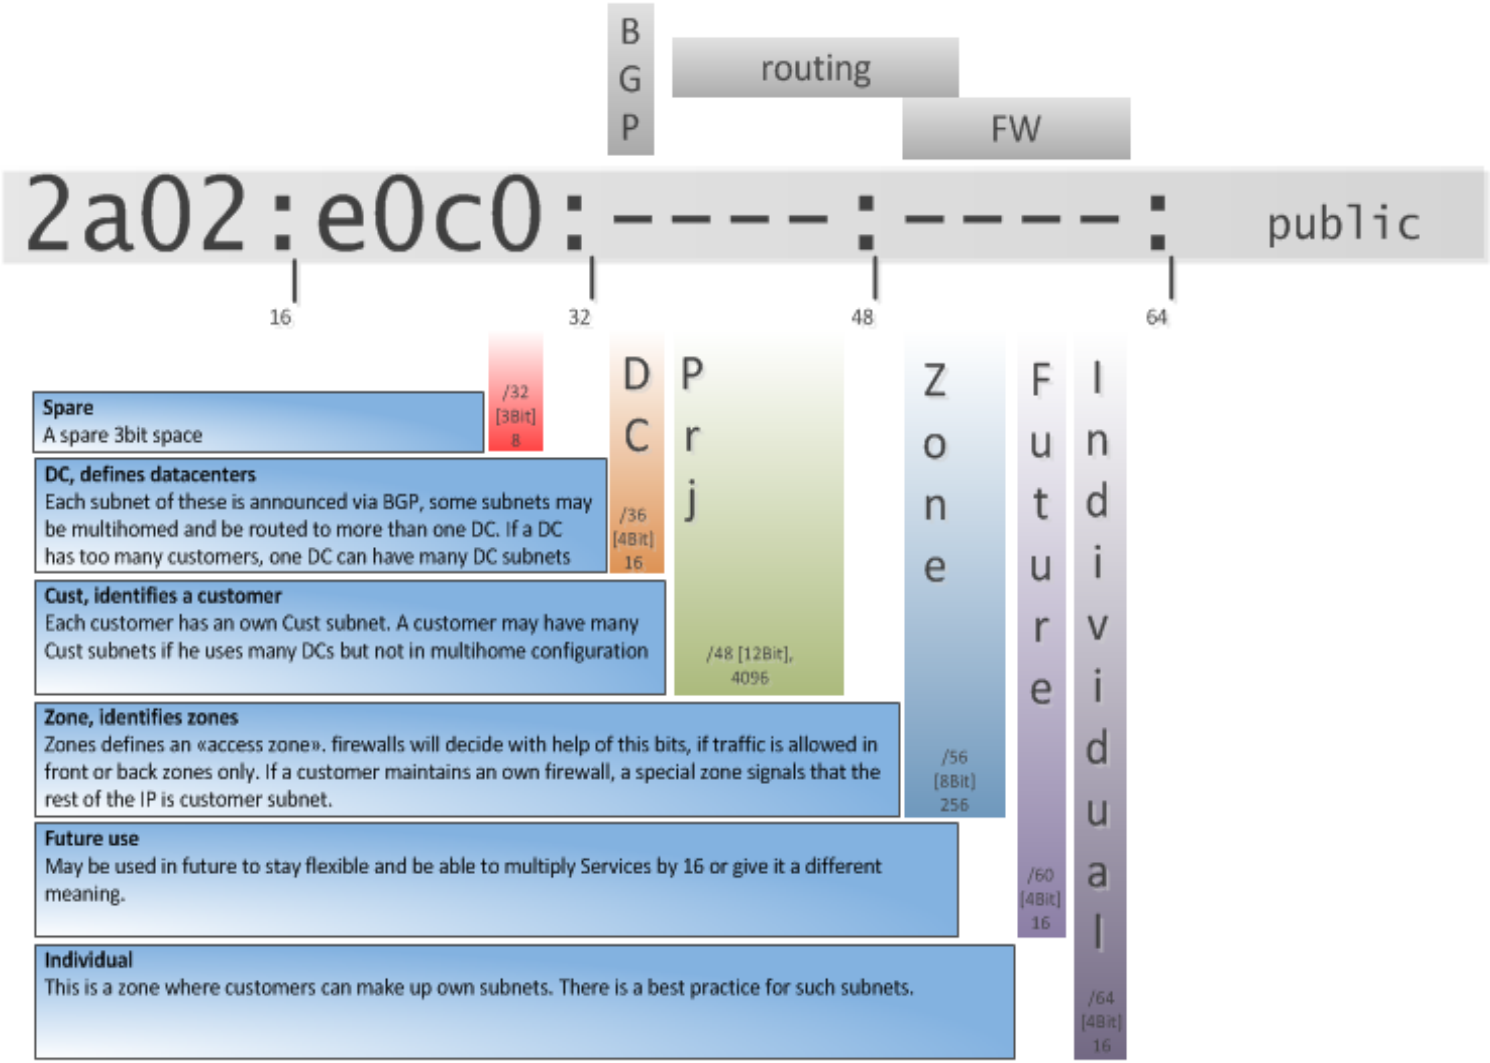

01.2013 – IP Konzept II

aspectra
hosting your future

Milestones and Timeline

- Started **August 2013** – 2nd site, BGP Announcement
- **September 2013** – Education all engineers, routing and firewalls 2nd site, first server implemented
- **October 2013** – DNS, Clients IPv6 range at office, tunnel to Datacenter
- **November 2013** – basic monitoring, DNS security and cleanup, start primary site
- **December 2013** – first vertical integration finished, christmas freeze, implement first services, webstats, homepage, reverse proxy, customer cockpit
- **January/February 2014** – DNS AAAA entry, activated reverse DNS at RIPE, glue record – **IPv6 goes production**!

Lessons learned – the Do's

- Use your own addresses, become LIR
- Find out if hardware can cope with extra load (dual-stack)
- Set a top level goal, define scopes, budget, roadmap
- Invest enough time for concept and deployment plan:
  - IPv6 addressing plan
  - Routing concept
  - DNS
- Test test test test test test test test test
- One thing at a time, isolated steps
- Expect the unexpected! (Monitoring, Microsoft AD, global catalogue)

Do not

- … think you can dual-stack in two weeks
- … overdefine your concepts, it will restrict you later
- … believe manufacturers are IPv6 ready. They will write anything on the box, just to sell it. Most products have limitations.
- … buy new hardware just because of IPv6. Use regular product life cycles
- … forget monitoring, logging, security: one for IPv4, one for IPv6
  - System/network performance
  - E2E scans
  - Firewall rules and logs

What it took

- 20% FTE
- Many maintenance windows on live systems (night work)
- External man power (for security and network)

- It's no rocket science
- Next step – activating customers

Who are they?

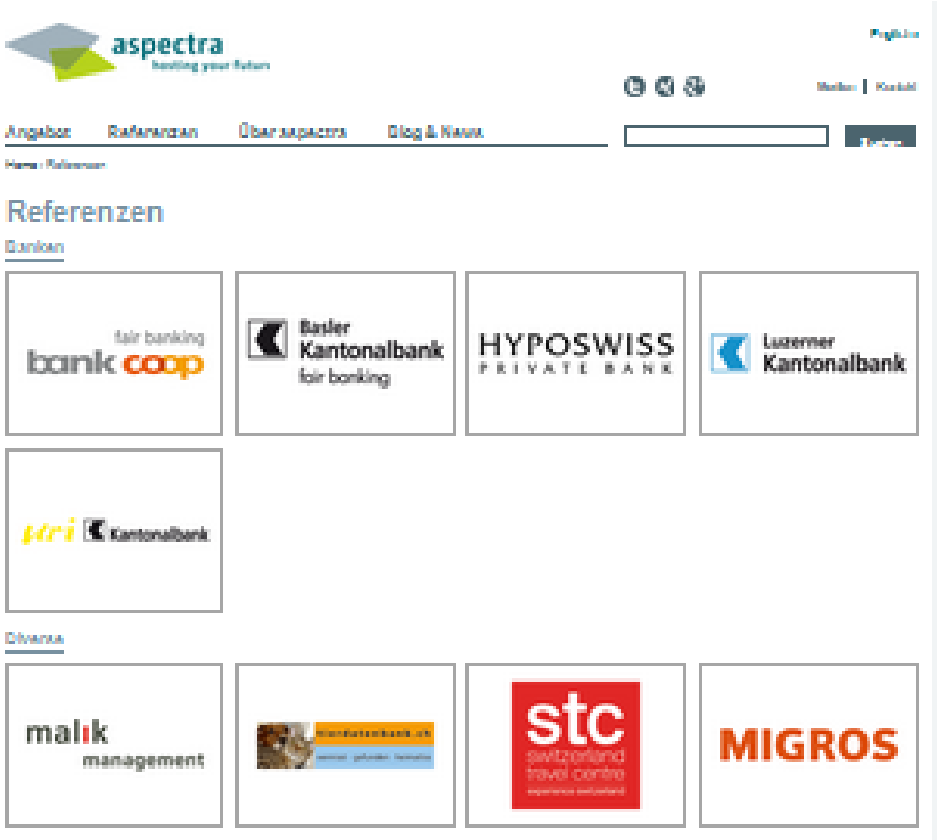

SWISS IPv6 COUNCIL

Holcim

- Founded in Switzerland in 1912
- Holcim is one of the world's leading suppliers of cement and aggregates (crushed stone, sand and gravel). We also supply ready-mix concrete and asphalt, and provide related services.

- 71'000 employees
- Production sites in 70 countries
- Market presence on every continent

IPv6 @ Holcim

- The main drivers to launch an IPv6 project at Holcim:
  - Be reachable from outside, especially for Asia, as Asia is an important region for Holcim
  - Create a foundation to be ready for new services in the area of the Internet of Things (sensor technologies).
  - Facilitate IP renumbering in the future (mergers and acquisitions)

What Holcim did

- Developed a high level concept and analyzed the address need
- Decided to become a LIR and acquired a /29
- Decided to go for a dual-stack strategy for the first phase
  - DMZ, public services first (plan for 2015)
  - Access networks and datacenters in a second stage
  - Communication and collaboration platform on Google, supports IPv6
- The concept has been approved on a global level
- Plan to invest in education and labs first
- HQ has to lead the way – other regions will follow

Recommendations

- Project initiator: Corporate IT
  - Massimo Fini

- Quote:
  There is no easy business case for IPv6,
  we have to promote it and fight for it ☺
- IPv6 address space management is a must

IPv6 Depoyment @ CERN

- RIPE Presentation by Edoardo Martelli

CERN Accelerators complex

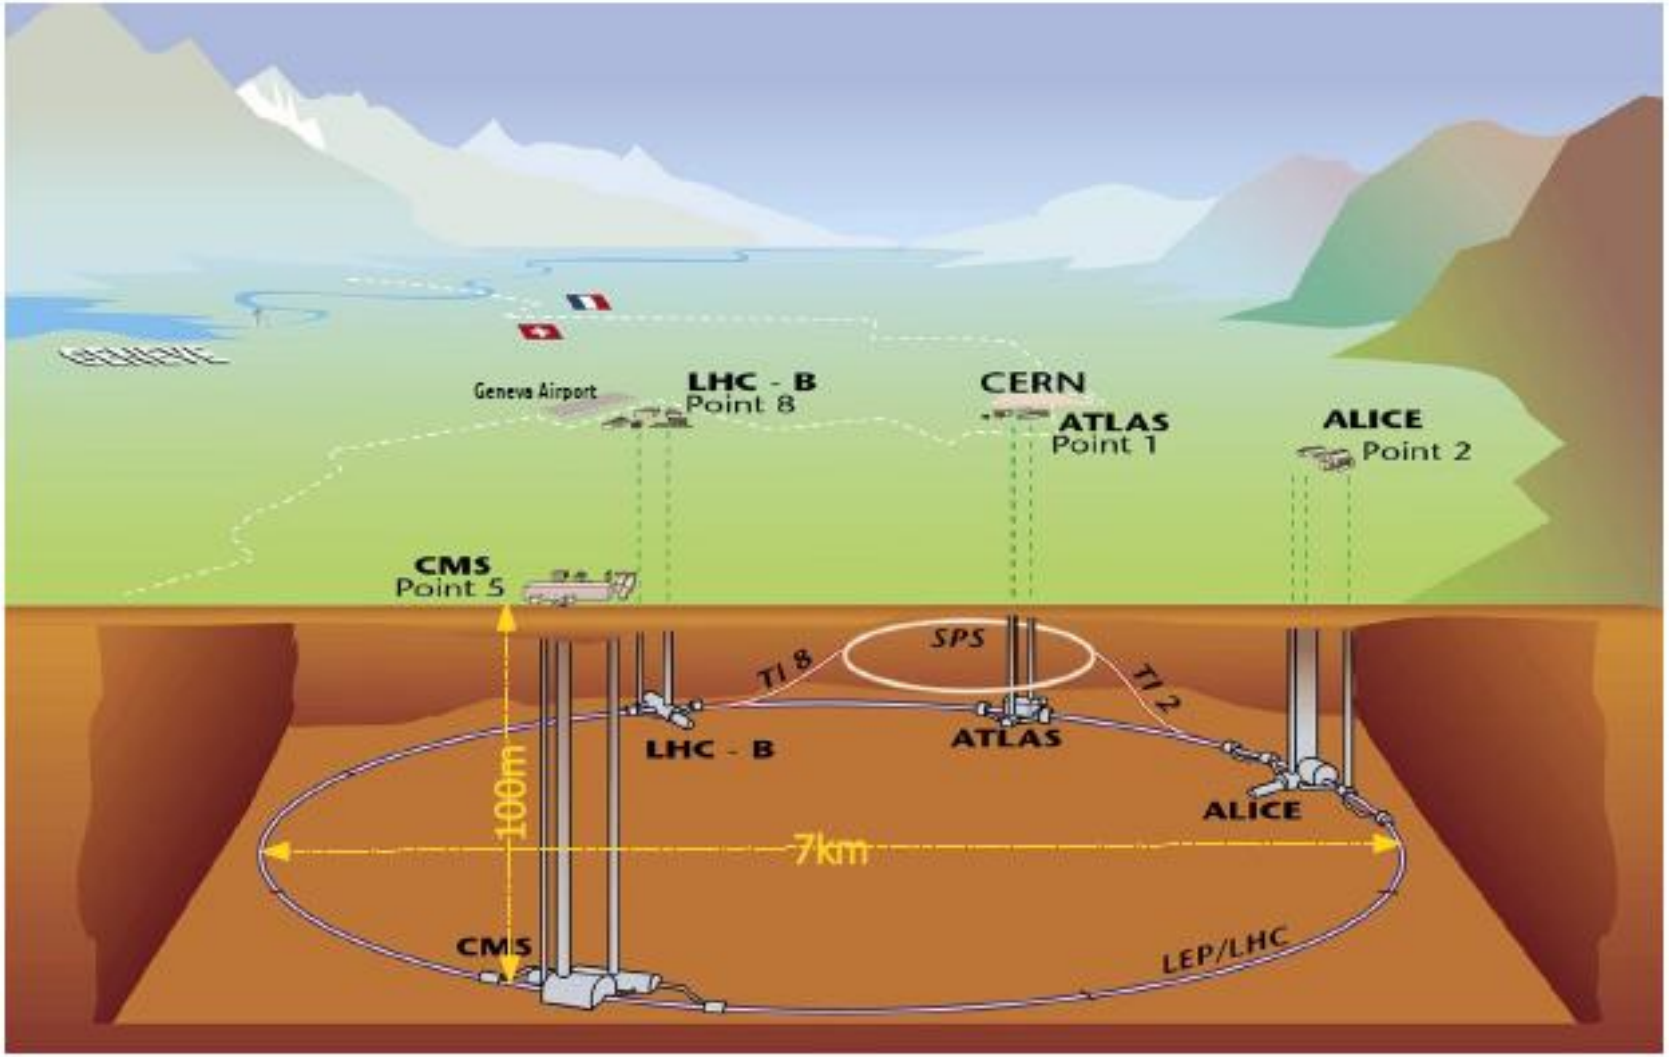

COUNCIL

CERN data network

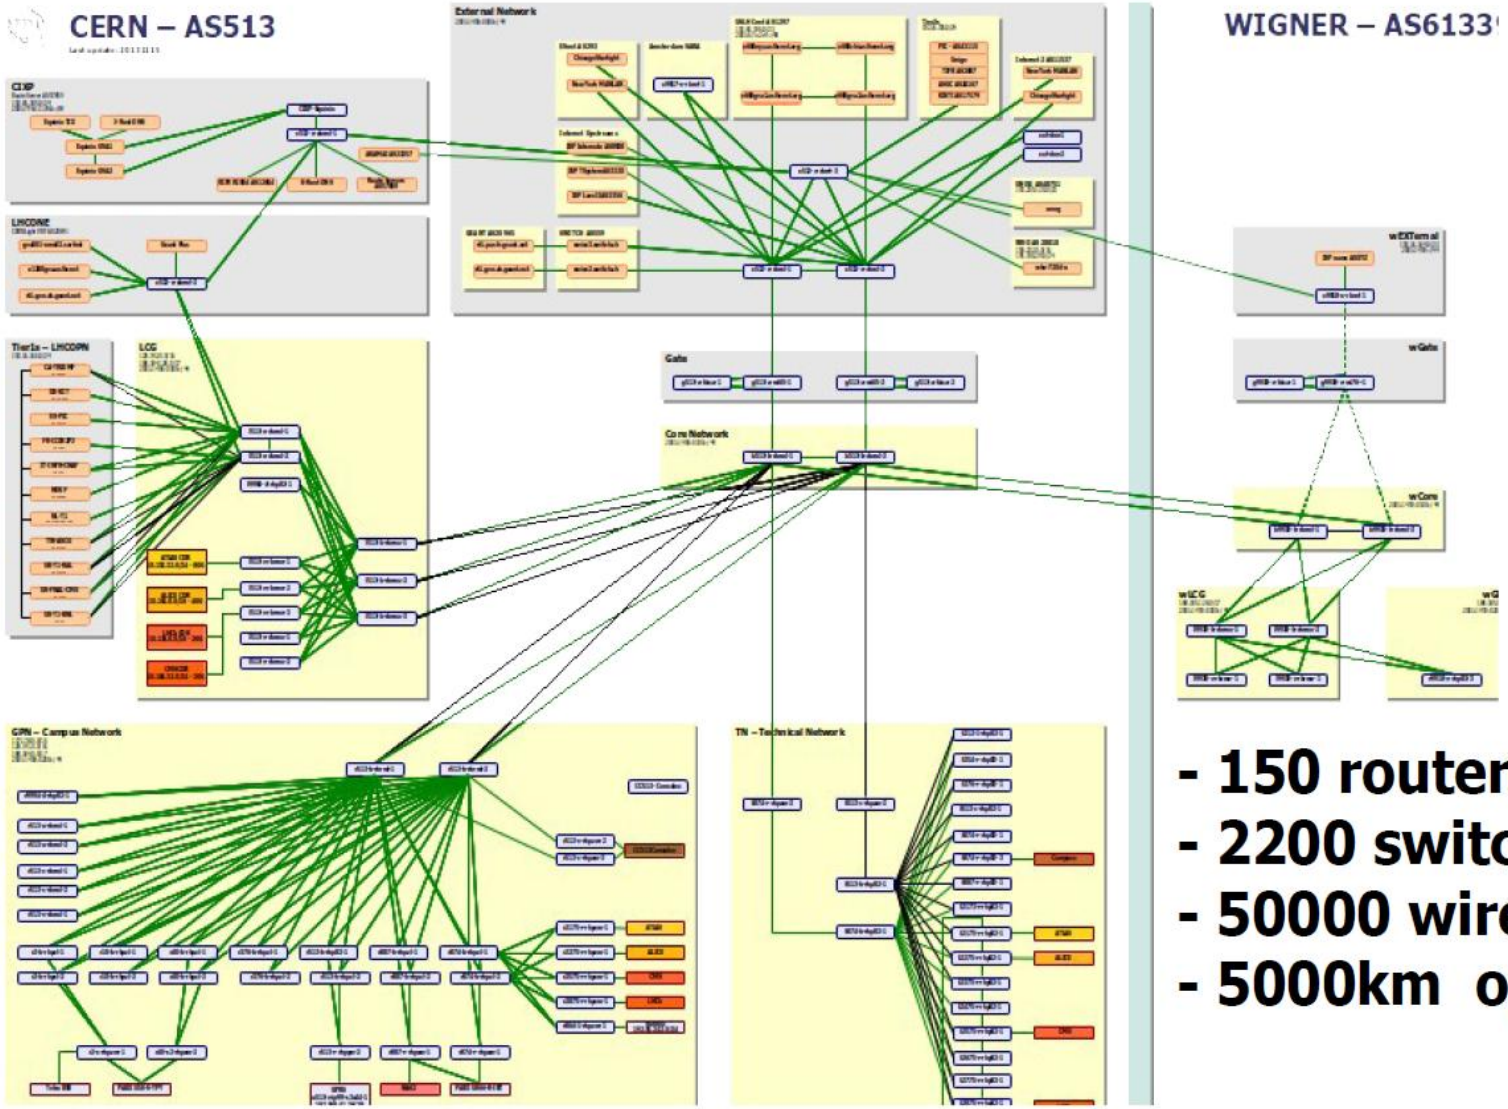

- 150 routers
- 2200 switches
- 50000 wired devices
- 5000km optical fibres

Network Provisioning and Management System

- 250 Database tables
- 100,000 Registered devices
- 50,000 hits/day on web user interface
- 1,000,000 lines of codes

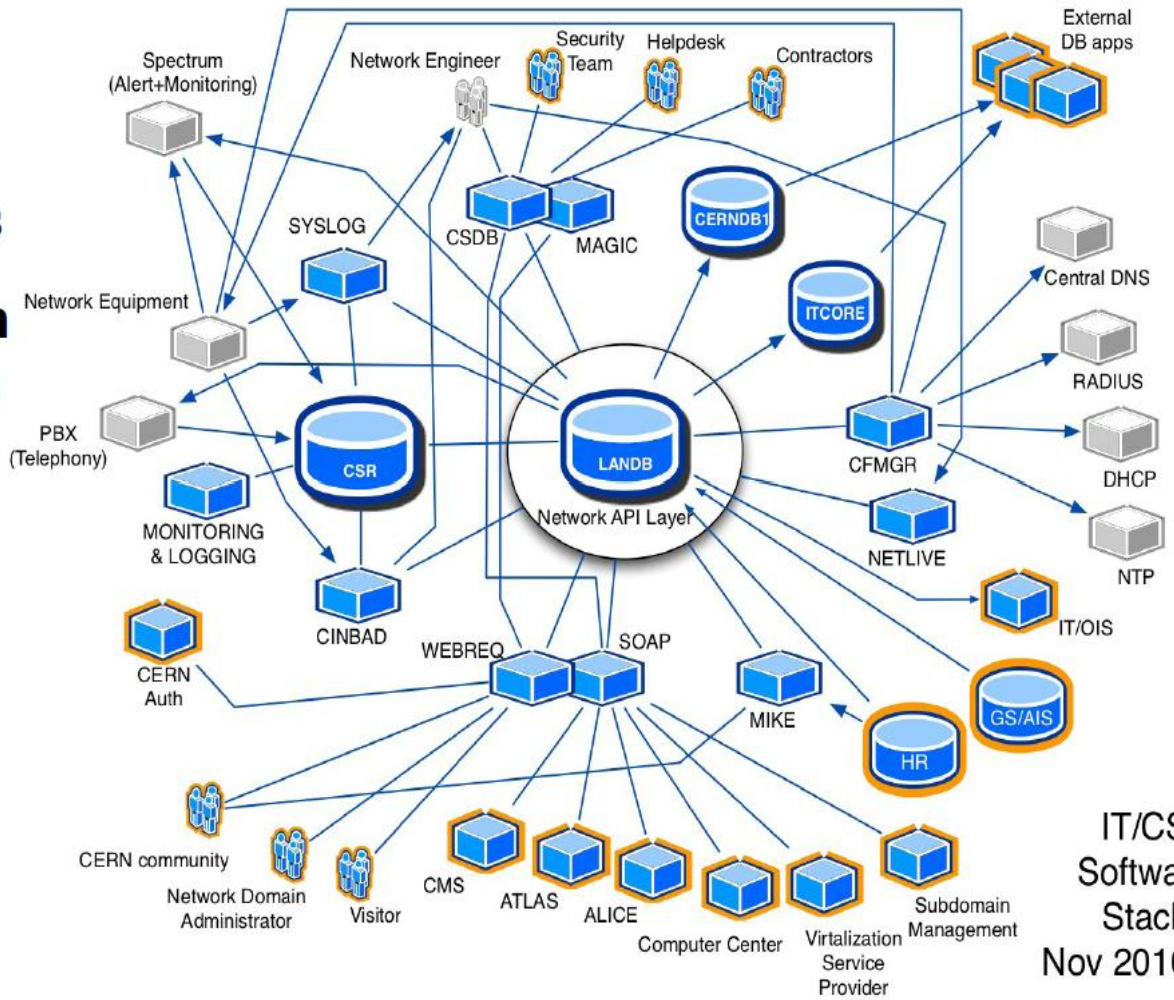

IT/CS Software Stack Nov 2010 V.3

IPv6 Service Definition

- Dual Stack
- At least one IPv6 address for every assigned IPv4 address
- Identical performance as IPv4, no degradation
- Common provisioning tools (NMS) for IPv4 and IPv6
- Same network services portfolio as IPv4
- Common security policies for IPv4 and IPv6

Driver

- CERN started playing with IPv6 in 2001, but for many years there was no reason for it.
- Large Virtual Machines deployment ramped up in 2010.
- Planned to have 130,000 VMs with public IP addresses to crunch the data from LHC after its upgrade in 2014. No NAT!

Addressing rules

- We split our address space to have reserve chunk. Reserved a lot of space for future use.
- Assigned /48 to every domain
- Subprefix for:
  - Routers, Network Devices, IP-phones, portable devices (for filtering in firewall)
- Using ULA for hosts that don't need to go out (e.g. isolated, printers), same addressing model like public prefix
- For DHCPv6 using RFC 6939, MAC-address registration

Approval and resources

- IPv6 deployment approved in Q1 2011 to be ready for production in 2013
- Allocated resources:
- Network design/testing/deployment:
  - 1x Network Engineer FTE for 2 years.
- Network database and NMS applications:
  - 2x Software Developers FTE for 2 years

How did it go?

Well ☺

Dual-Stack

"Dual stack" DONE!

Network database:
- IPv6 now main navigation data
- New schema compatible with all legacy queries
- IPv6 address tables fully populated

Network:
- All campus, data centres, firewalls, external interfaces are dual stack (except: LHC accelerator control network, LHC detectors data acquisition networks).

Same routing architecture (BGP and OSPF).

Addressing

"An IPv6 address to every IPv4 ones"

DONE!

- Every device with an IPv4 address has an IPv6 address assigned in the Network DB

- All assigned IPv6 addresses have a name in **ipv6.cern.ch**

```
# host ping.ipv6.cern.ch
ping.ipv6.cern.ch has IPv6 address 2001:1458:201:1c80::100:175

# host TELEPHONE-62470.ipv6.cern.ch
TELEPHONE-62470.ipv6.cern.ch has IPv6 address fd01:1458:204:27a::100:2e
```

- Dynamic (portable) devices get a name in **dyndns6.cern.ch**

```
# host myiphone.dyndns6.cern.ch
myiphone.dyndns6.cern.ch has IPv6 address 2001:1458:202:180::101:8a26
```

Performance

"Identical performance"

DONE!

Almost

All production network devices can forward IPv6 packets at wire speed

Only exception: policy base routing for statefull firewall bypass. Not a show stopper, because of low IPv6 traffic volume.

SWISS IPv6 COUNCIL

Provisioning

"Common provisioning tools" DONE!

NMS:
- routers configuration generator for all the vendors
- DHCPv6, DNS configurations from Network-DB
- ACL generator for firewalls from Network-DB

CSDBweb (Network-DB interface for engineers): IPv6 everywhere there is IPv4

WebReq (Network-DB interface for end-users): All IPv6 info visible together with IPv4, IPv6-ready flag settable

IPv6-ready Flag

IPv6-ready flag

Users can declare their own devices as "IPv6-ready"

IPv6-ready means:
- IPv6 connectivity is OK
- all running server applications are listening on both v4 and v6 sockets

Consequences:
- Firewall: IPv6 equivalent of IPv4 security openings applied to the central firewall
- DNS: DEVICENAME.cern.ch returns A and AAAA records, reverse returns DEVICENAME.cern.ch (and host certificates work)

Network Services

"Same network services as IPv4"

DONE!

DNS:
- queryable over IPv6
- announced in the DHCPv6 leases

NTP:
- reachable over IPv6 (ip-time-1.ipv6.cern.ch and ip-time-2.ipv6.cern.ch)

DHCPv6:
- Static devices: same servers and daemon of SHCP for IPv4
- Dynamic devices: different servers because running the very latest version (classes only works in 4.3.0)

Security Policies

"Common security policies" DONE!

Firewall rules database
- schema and management tools developed
- Most IPv4 rules automatically translated into IPv6
- IPv6 only rules manually created

ACL compiler generator
- Antispoofing ACLs applied to all router interfaces
- All firewalls managed by the NMS

What about the Accelerator?

Next maintenance window in 3 years

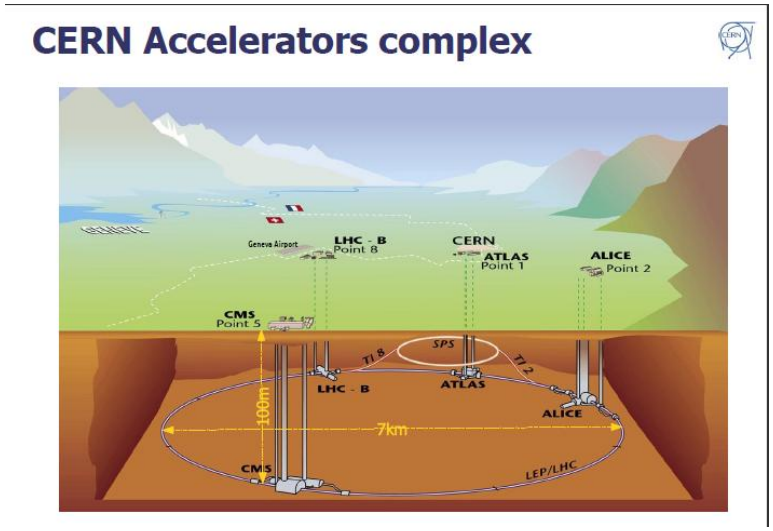

Recommendations and more info

- It takes time!
- Plan and think in advance
- Take your time for software planning and integration

- More Information:
- Find the full presentation at: https://ripe68.ripe.net/presentations/294-cern-ipv6-deployment.pdf
- Check the current status at http://cern.ch/ipv6/content/implementation-plan-and-status

Challenges

- Sizes of routing tables and ACLs have doubled
- New issues to be solved by Support lines
- DHCPv6 still in an early stage
- New security threats to take into account
- Legacy applications don't understand IPv6, and some will never do

Lessons learnt

- Catching up with 20 years of IPv4 experience and development takes a lot of time.
- The network is the easy part.
- DHCPv6 is definitely not DHCPv4.
- Don't rush. Have a staged deployment with a large variety of early adopters. And keep in touch with them: they may not report all the problems.
- Only the deployment on the live network will prove it can cope with the two protocols.
- Don't rely on support from applications developers: there are already enough bugs to fix without adding IPv6.

IPv6 @ face::b00c

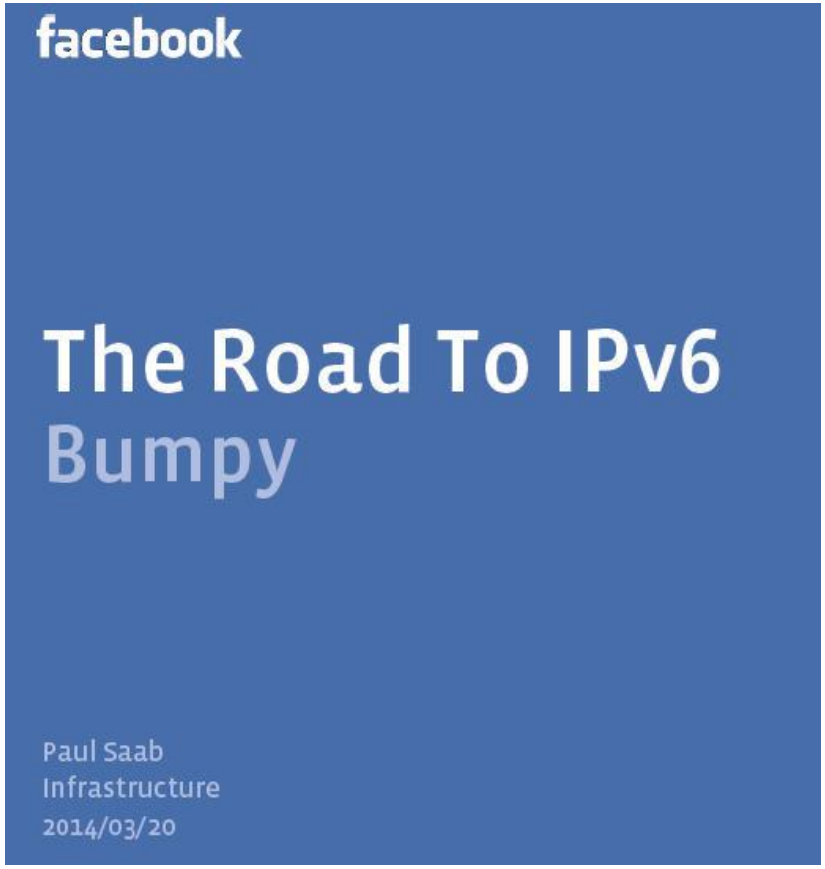

1. Motivation
2. Running out of RFC1918
3. IPv6 @ face:b00c
4. The good stuff
5. Where are we now?

SWISS IPv6 COUNCIL

Motivation

- We're out of RFC1918
- Why maintain two stacks in your entire network?
- Much easier to handle two stacks at the edge and one internally.
- It needs to be done someday, why not now while we're motivated?
- Stop engineers from continually writing IPv4 only code that will need to be fixed later.
- Push the industry to move faster and re-prioritize IPv6.

Running out of RFC 1918

- /24 for every rack – Genius!
- Math is easy
- Subnet mask is easy to remember
- Wastes a lot of space
  - 254 usable addresses
  - 80 addresses in a rack
- /25 is what we could re-number into
  - Not enough savings
- Do it if you want to get to IPv6 faster

IPv6 solution

- Each rack is a /64
- Each cluster is a /n (where n < 64)
- Layer 3
  - Core should not handle Neighbor Discovery
  - fping6 –g xxx::/64 confined to a rack
  - Just Do It

SLAAC vs. Static Assignment

- SLAAC
  - Great idea
  - Terrible for datacenter deployment
  - NIC changes, IP address changes

- Static Assignment
  - Avoid encoding IPv4 address in the IPv6 address
    - But it makes mapping back and forth easy!
    - What happens when you stop using IPv4?
    - Take the opportunity to have a clean slate with no dependencies

The Good

- We were able to get rid of a lot of technical debt
- IPAddress class
  - Death to strings and integers
- Rollout of traffic
  - Most services were able to slowly roll out IPv6 from 0-100%
  - Instantaneous rollback if needed
  - Problems may not show up at 1%, 5% or 10%, but they do at 100%
- Iterate Iterate Iterate
- Don't make IPv6 an all or nothing proposition. You will fail.

The Good

The Good – Neteng @ Facebook

- Backbone was upgraded a couple of years ago
- Clusters were converted to Layer 3
- IPv6 native to all cluster and rack switches after World IPv6 Day

SWISS IPv6 COUNCIL

Where are we now?

100%

of our hosts we care about respond on IPv6

- Hosts that are not IPv6 ready are going away

75%

of our internal traffic is now IPv6

- 100% Q3 2014 (or earlier)

98%

of traffic in & out of HHVM is IPv6

100%

of our memcache traffic is IPv6

100%

IPv6 only (no RFC1918) in 2-3 years

SWISS IPv6 COUNCIL

Where are we now (cont)

- New IPv6 traffic showing up daily
- Engineers asking if they can start writing IPv6-only code today
- Latency and other metrics show IPv6 to be the same as IPv4
- Plans for first IPv6 only cluster (no RFC1918) by end of 2014
- We will not remove RFC1918 from existing clusters

- Find full presentation at: http://www.swissipv6council.ch/sites/default/files/images/paul_saab_facebook_-_worldipv6congress-ipv6_lh_v2.pdf

Thank You For Your Attention!

IPv6 Grundlagen, Funktionalität, Integration

von Silvia Hagen, Deutsch
2. Auflage, Sunny Edition, 2009
ISBN 978-3-9522942-2-2

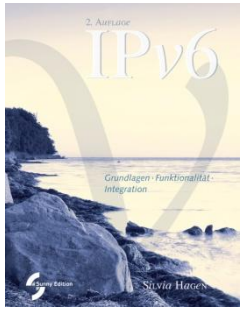

IPv6 Essentials

by Silvia Hagen, English
3rd Edition, O'Reilly, June 2014
ISBN 978-1-4493-1921-2

NEW

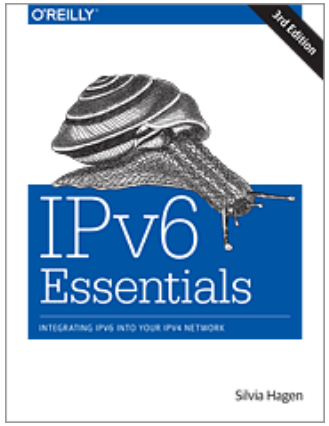

Planning for IPv6

by Silvia Hagen, English
O'Reilly, July 2011
ISBN 978-1-4493-0539-0
eBook 978-1-4493-0538-3

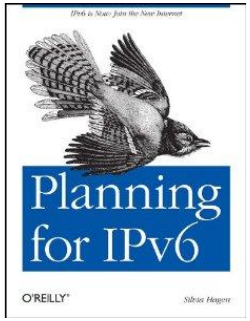

SWISS IPv6 COUNCIL