Supporting IPv6

Communication Systems group 40th years anniversary
1st November 2023
Edoardo Martelli

IPv6 and CERN: a long story started before IPv6 was born

CERN datacentre, 1972

1990s: IPv4 starts showing its limits

In the late 80s, Internet started expanding faster and faster

IPv4, although it counts 2^{32} (~4 billions) addresses, had a fixed subnetting scheme (the A, B, C classes) which led to wasteful allocations and the fear of a premature exhaustion of the available address space

1993: RFC 1338 acks IPv4 exhaustion

Network Working Group — V. Fuller (BARRNet), T. Li (Cisco)
Request for Comments: 1338 — J. Yu (MERIT), K. Varadhan (OARnet)
June 1992

Supernetting: an Address Assignment and Aggregation Strategy

Abstract

This memo discusses strategies for address assignment of the existing IP address space with a view to conserve the address space and stem the explosive growth of routing tables in default-route-free routers run by transit routing domain providers

1. Problem, Goal, and Motivation

As the Internet has evolved and grown over in recent years, it has become painfully evident that it is soon to face several serious scaling problems. These include:

1. Exhaustion of the class-B network address space.
2. Growth of routing tables in Internet routers beyond the ability of current software (and people) to effectively manage.
3. **Eventual exhaustion of the 32-bit IP address space.**

It has become clear that the first two of these problems are likely to become critical within the next one to three years. [...] **It does not attempt to solve the third problem, which is of a more long-term nature, but instead endeavors to ease enough of the short to mid-term difficulties to allow the Internet to continue to function efficiently while progress is made on a longer term solution.**

1993: Call for a new IP protocol

In 1993, although some workaround to extended the life of IPv4 were being developed (namely variable prefix length (VLSM, CIDR) and Network address translation (NAT)) the IETF (Internet Engineering Task Force) calls for a new IP protocol to replace IPv4

1993: RFC 1550 call for IPng proposals

```
Network Working Group                                        S. Bradner
Request for Comments: 1550                           Harvard University
Category: Informational                                       A. Mankin
                                                                    NRL
                                                          December 1993

          IP: Next Generation (IPng) White Paper Solicitation

1. Introduction
```

The IP: next generation (IPng) area in the IETF is soliciting white papers on topics related to the IPng requirements and selection criteria.

```
   All interested parties are invited to submit white papers detailing
   any specific requirements that they feel an IPng must fulfill or any
   factors that they feel might sway the IPng selection.  An example of
   the former might be a submission by a representative of a utility
   company detailing the scaling and addressing features which would be
   required to service future inclusion of utility meters on the
   network.  An example of the other case might be a paper outlining the
   potential effect on IPng of some sections of the future network
   connectivity being provided via wireless networks.
```

1993-95: Several proposal come forward

Several protocols were proposed for IPng, the next generation IP.

In 1995 they converged in three main proposals:
- **CATNIP**, which tried to include several of the existing protocols (IPv4, IPX, CLNP)
- **SIPP**, a complete new protocol with 64 bits addresses
- **TUBA**, a reworked ISO/OSI CNLP

1995: RFC 1752 converge on three proposals

```
Network Working Group                                    S. Bradner
Request for Comments: 1752                       Harvard University
Category: Standards Track                                 A. Mankin
                                                                ISI
                                                       January 1995


        The Recommendation for the IP Next Generation Protocol

Abstract
   This document presents the recommendation of the IPng Area Directors
   on what should be used to replace the current version of the Internet
   Protocol.  This recommendation was accepted by the Internet
   Engineering Steering Group (IESG).

Table of Contents
[...]
   7.       IPng Proposals . . . . . . . . . . . . . . . . . . . . . . 11
     7.1    CATNIP. . .  . . . . . . . . . . . . . . . . . . . . . . . 11
     7.2    SIPP. . . .  . . . . . . . . . . . . . . . . . . . . . . . 12
     7.3    TUBA. . . .  . . . . . . . . . . . . . . . . . . . . . . . 13
```

CN-CS position on IPng

Brian Carpenter:
“I knew of OSI since 1982. During networking studies for LEP. C. Piney, a member of the CERNET software team, and I wrote a paper, called 'Site-Wide Datagrams at CERN', advocating that we should adopt the ISO/OSI standard for 'connectionless datagrams' as the lowest common denominator for network traffic
across the CERN site (https://cds.cern.ch/record/155952/files/198504345.pdf)

I also did a rough implementation of CLNP over the controls network at the PS, where I worked at that time, and I think at least one application using it went into production.

So personally I was an advocate for CLNS, even before DECnet Phase V became a factor.”

1994: RFC 1670 CERN prefers CLNS

```
Network Working Group                                    D. Heagerty
Request for Comments: 1670                                      CERN
Category: Informational                                  August 1994

                Input to IPng Engineering Considerations
Abstract
   This document was submitted to the IETF IPng area in response to RFC
   1550.  Publication of this document does not imply acceptance by the
   IPng area of any ideas expressed within.

Summary
   This white paper expresses some personal opinions on IPng engineering
   considerations, based on experience with DECnet Phase V transition.
   It suggests breaking down the IPng decisions and transition tasks
   into smaller parts so they can be tackled early by the relevant
   experts.

Author's Address
   Denise Heagerty
   Communications Systems Group
   Computing and Networks Division
   EMail: denise@dxcoms.cern.ch
```

1994: RFC 1671 CERN recommendations

```
Network Working Group                                  B. Carpenter
Request for Comments: 1671                                     CERN
Category: Informational                                 August 1994

       IPng White Paper on Transition and Other Considerations

Summary
   This white paper outlines some general requirements for IPng in selected areas.
   It identifies the following requirements for stepwise transition:
   A) Interworking at every stage and every layer.
   B) Header translation considered harmful
   C) Coexistence.
   D) IPv4 to IPng address mapping.
   E) Dual stack hosts.
   F) DNS.
   G) Smart dual-stack code.
   H) Smart management tools.

 Transition and deployment
   It is clear that the transition will take years and that every site
   will have to decide its own staged transition plan.
```

1994: RFC 1671

Disclaimer and Acknowledgements

This is a personal view and does not necessarily represent that of my employer.

CERN has been through three network transitions in recent years (IPv4 renumbering managed by John Gamble, AppleTalk Phase I to Phase II transition managed by Mike Gerard, and DECnet Phase IV to DECnet/OSI routing transition managed by Denise Heagerty). I could not have written this document without having learnt from Them. [...]

Author's Address

Brian E. Carpenter
Group Leader, Communications Systems
Computing and Networks Division
CERN
European Laboratory for Particle Physics
1211 Geneva 23, Switzerland

Phone: +41 22 767-4967
Fax: +41 22 767-7155
Telex: 419000 cer ch
EMail: brian@dxcoms.cern.ch

CN-CS is involved in the IPng decision

Brian Carpenter:
“When the IETF was just starting on the IPng problem, the IAB unilaterally proposed CLNP as IPng, which led of course to TUBA.
But the IETF did not accept that 'diktat' at all, and **I was invited to join the IPng Directorate which led to the IPv6 decision in 1994**.
I was certainly viewed as a TUBA supporter.”

But TUBA didn’t make it

1995: RFC 1883 - A modified SIPP is chosen

```
Network Working Group                                  S. Deering, Xerox PARC
Request for Comments: 1883                         R.  Hinden, Ipsilon Networks
Category: Standards Track                                         December 1995
```

Internet Protocol, Version 6 (IPv6)
Specification

Abstract

This document specifies version 6 of the Internet Protocol (IPv6), also sometimes referred to as IP Next Generation or IPng.

1. Introduction

IP version 6 (IPv6) is a new version of the Internet Protocol, designed as a successor to IP version 4 (IPv4) [RFC-791]. The changes from IPv4 to IPv6 fall primarily into the following categories:

o Expanded Addressing Capabilities

CN-CS reaction to IPv6

Brian Carpenter:
“The one thing I did immediately after the decision in 1994 (I mean literally, within one day of the announcement) was to start work on RFC1888, because we might need a mapping between DECnet V and IPv6. In fact that was a waste of time, as it turned out. So the actual work on IPv6 at CERN started many years later.”

1996: RFC 1888 – CERN prepares for IPv6

```
Network Working Group                                          J. Bound
Request for Comments: 1888             Digital Equipment Corporation
Category: Experimental                                     B. Carpenter
                                                                   CERN
                                                          D. Harrington
                                        Digital Equipment Corporation
                                                        J. Houldsworth
                                                   ICL Network Systems
                                                              A. Lloyd
                                                Datacraft Technologies
                                                           August 1996

                          OSI NSAPs and IPv6

Abstract
   This document recommends that network implementors who have planned
   or deployed an OSI NSAP addressing plan, and who wish to deploy or
   transition to IPv6, should redesign a native IPv6 addressing plan to
   meet their needs.  However, it also defines a set of mechanisms for
   the support of OSI NSAP addressing in an IPv6 network.
```

2001: IT-CS starts testing IPv6

IT-CS-EN started testing IPv6 around the year 2002. Daniel Davids (IT-CS-EN) was in charge.

A testbed was built in the computer centre in B513, but with no connection to the IPv6 Internet

CERN Computer Centre around 2001

2003: public IPv6 prefix assigned to CERN

To get an official IPv6 block, it was necessary to become a member of RIPE (Réseaux IP Européens), the entity that manages IP addresses for the European region.

CERN was a founder of RARE (Réseaux Associés pour la Recherche Européenne, which later became TERENA (Trans-European Research and Education Networking Association)).
RARE provided the initial legal framework for the RIPE NCC.
But CERN was not member of RIPE because the IPv4 addresses it used were allocated before RIPE came to existence

The IT-CS-EN section took care of the full process

	Contributor	RIPE NCC
Place	Genève	Amsterdam
Date	10/12/2004	23 August 2004
Name of authorised person	T. Lagrange	Axel Pawlik
Function	Head of Purchasing Service	Managing Director
Signature		

Please return a signed copy of the RIPE NCC Standard Service Agreement and a copy of an extract from the Commercial Trade Register or an equivalent document proving the registration of Contributor's business with the national authorities.

2003: IPv6 allocation request to RIPE

#[IPV6 ALLOCATION REQUEST]#

#[REGISTRY ID]#

X-NCC-RegID: ch.cern

#[OVERVIEW OF ORGANISATION]#

Q1: Please provide a short description of the organisation requesting the IPv6 address space.

CERN's primary mission is to provide facilities for high energy particle physics experiments. CERN is open to scientists from its 20 member states and from all other countries of the world. This makes CERN one of the largest sources of numerical scientific data in the world. Computer networking, and in particular Internet connectivity, is therefore a mission-critical requirement. CERN operates an IXP in order to facilitate the exchange of Internet traffic in the region and to maximise its own Internet connectivity.

Q2: Is this IPv6 request for the entire organisation or will other parts of the organisation also be requesting an allocation?

This request is for CERN and its istitutional activities. CIXP, the Internet Exchange Point managed by CERN, is going to send a request for a /48 allocation.

Q3: Will your organisation also be requesting (or do you already have) an IPv6 allocation from another Regional Internet Registry (RIR)?

CERN will only request an IPv6 allocation from the RIPE NCC.

2003: IPv6 allocation request to RIPE

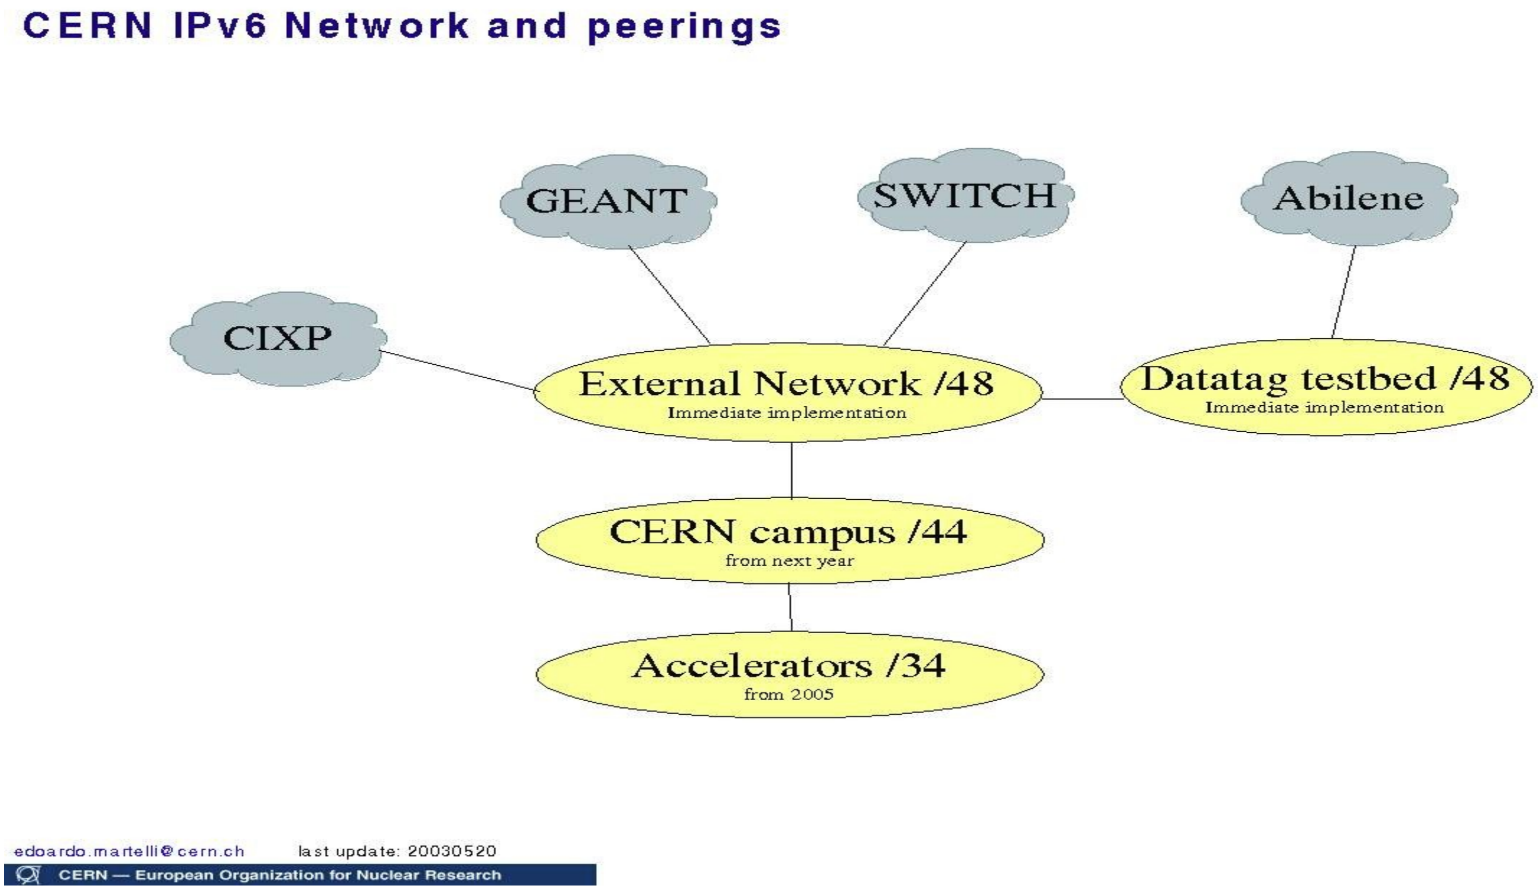

2003, June: IPv6 allocation from RIPE

From: RIPE NCC Hostmaster <hostmaster@ripe.net>
Sent: Wednesday, 11 June 2003 16:28
To: Edoardo Martelli <Edoardo.Martelli@cern.ch>
Cc: Paolo Moroni; Edoardo Martelli <Edoardo.Martelli@cern.ch>; Cern Extip; Paolo Moroni; Edoardo Martelli <Edoardo.Martelli@cern.ch>
Subject: NCC#2003054349 Internet Network Numbers 2001:1458::/32 Allocated

Dear Edoardo Martelli,
This is to inform you that we have allocated the following range of IPv6 address space to your registry ch.cern:

2001:1458::/32

The database shows the following information:

```
inet6num:     2001:1458::/32
netname:      CH-CERN-20030611
descr:        PROVIDER LOCAL REGISTRY
descr:        CERN
descr:        European Organization for Nuclear Research
country:      CH
admin-c:      OM10
tech-c:       PM1829-RIPE
tech-c:       EM9228
notify:       extip@cern.ch
status:       ALLOCATED-BY-RIR
mnt-by:       RIPE-NCC-HM-MNT
mnt-lower:    CERN-MNT
changed:      hostmaster@ripe.net 20030611
source:       RIPE
```

2003: Prefix announced to R&E networks

In September 2003, IPv6 is deployed in the CERN External Network.

The CERN prefix 2001:1458::/32 could then be announced to the Research and Education networks.

At this point there was only a partial visibility of the CERN prefix on the global IPv6 Internet

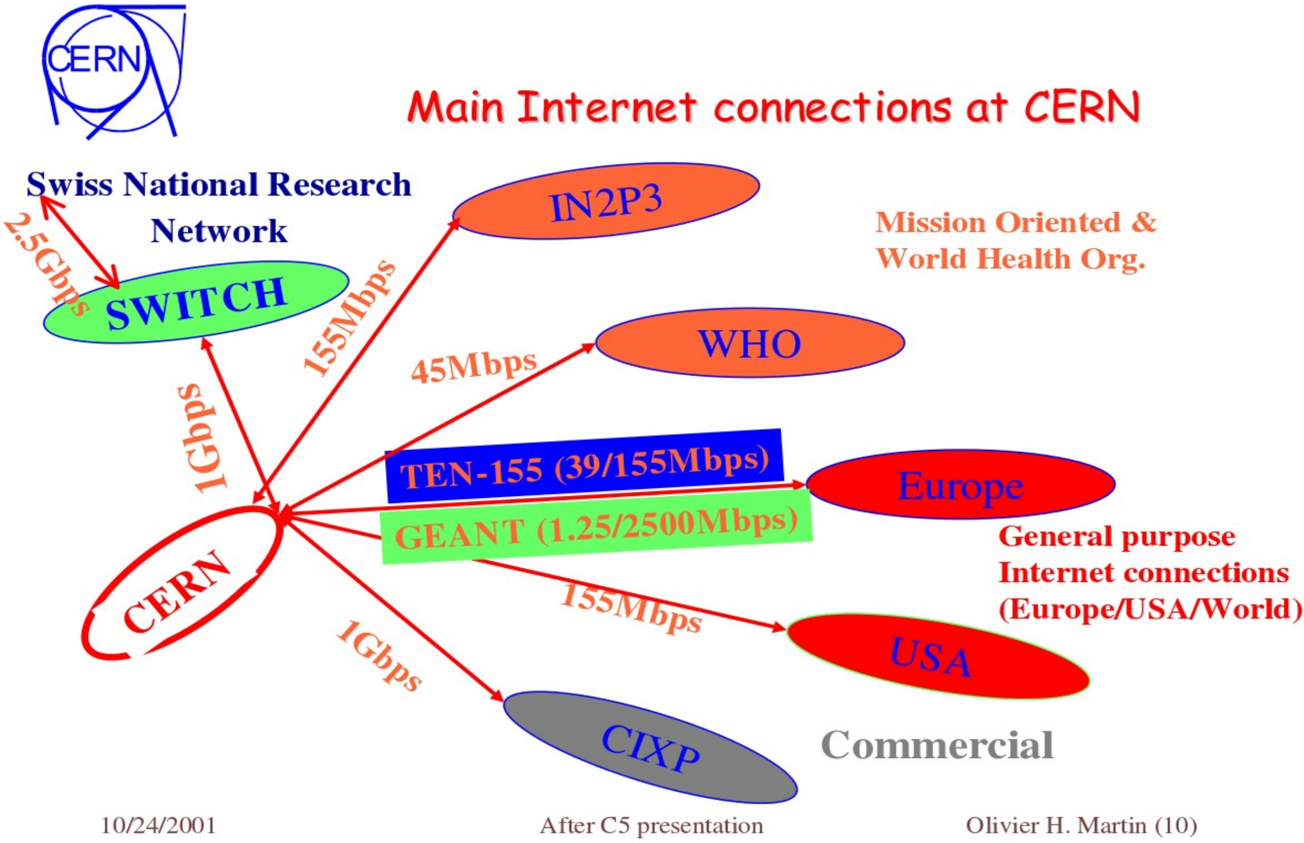

2003: Land Speed Record

In November 2003, the members of the DataTAG project in IT-CS-EN and the Caltech network team at CERN achieved one of the IPv6 Land Speed records

From the adjudication: “On November 18, 2003, Caltech and CERN transferred 560 Gigabytes of data in 20 minutes between Geneva and the Caltech stand at SuperComputing 2003 in Phoenix, through the LHCnet/DataTag, Abilene and SciNet backbones, using a single TCP/IPv6 stream”

Distance: 11539 km
Data transferred: 560 Gbytes
Average speed (over 20 minutes):
    4.00 Gbps
Record submitted:
    46,156,000,000,000,000 meters-bits/sec

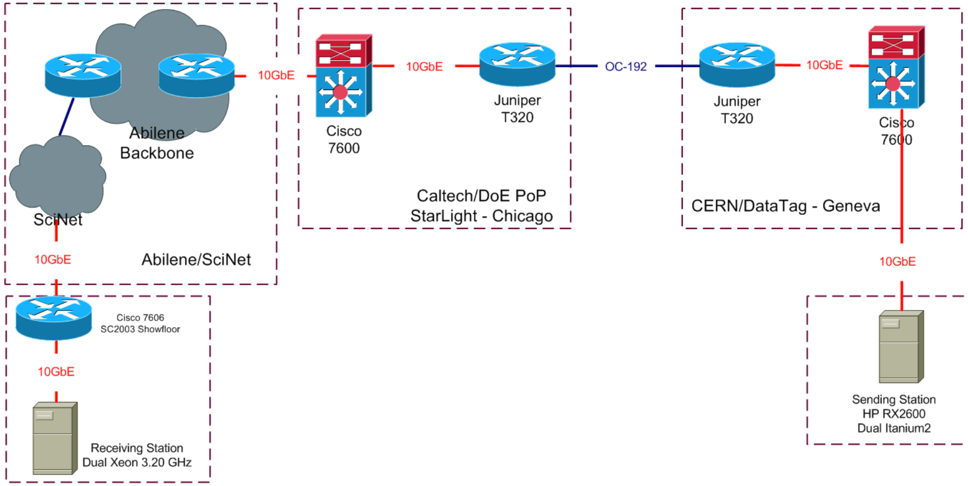

2009: Full Internet reachability

In 2009, Sunrise was selected as one of the new Commercial IPv4 Internet upstreams for CERN

It happened that Sunrise was part of the global IPv6 Internet

Following our informal request, they accepted to peer with CERN also in IPv6

In November 2009, the CERN prefix 2001:1458::/32 became reachable to the global IPv6 Internet

2011: IPv6 deployment project

At CERN, the use of Virtual Machines started around 2010. Because of their success, It was soon **planned to have 130,000 VMs with public IP addresses** ready for LHC Run2
The CERN IPv4 address pools (5x /16) were already mostly fully used at that time, so in order to not be stopped by a possible IPv4 shortage, it was decided to start a full deployment of IPv6

The IPv6 deployment project was approved in Q1 2011
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

Objective: to be ready for production by 2013

Work plan

Software tools:

- New IPv6-compatible Network Database schema
- IPv6 Address assignments in Network Database
- Development of NMS tools (Network device configurations, DHCPv6, DNS…)
- Network-Database interface for end-users (webreq)

Network:

- Testing on installed base
- Configuration of Network devices
- Deployment of Network services (DNS, DHCPv6, Radius, NTP)
- Training of Support-Lines and Advanced Users

IPv6 service definition

- Dual Stack
- At least one IPv6 address for every IPv4 address in use
- Identical performance as IPv4, no degradation
- Common provisioning tools for IPv4 and IPv6
- Same network service portfolio as IPv4
- Common security policies for IPv4 and IPv6

LANDB

Main tasks:
- Addition of IPv6 network tables and records
- Address plan
- Assignment of IPv6 subnet/addresses to all IPv4 entries
- Web interfaces for engineers and users

Challenges:
- New schema compatible with all legacy queries
- Consistently populate IPv6 records

Limitations:
- IPv4 still needed

Network devices configuration

Main tasks:

- Test IPv6 functionalities and performance on all running devices
- Development of NMS tools for automatically generated configurations and ACLs
- Finally add commands for IPv6 addressing and routing on all the devices

Challenges:

- Translate firewall ACLs
- DHCPv6 configuration and functionalities

Limitations:

- Management and Monitoring still over IPv4
- IPv6 Policy Base Routing (used in HTAR) not line rate

IPv6 project 2011-12

2011, January: IPv6 deployment project approved

2011, February: IPv6 address plan issued

2011, March: Development of LANDB schema including IPv6 information started

2011, July: IPv6 connectivity in part of LCG, CORE and GPN backbones

2011, July: Prototype of DNS servers

2011, August: IPv6 Internet connectivity via dedicated firewall

2011, August: Pilot IPv6 services for LCG and GPN users

2012, March: LANDB (Network database) with IPv6 tables in production

2012, March: CSDWEB support of IPv6 information

2012, June: cfmgr router compiler can generate IPv6 configuration

IPv6 project 2012-14

2012, June: cfmgr router compiler can generate IPv6 configuration
2013, March: all routers in the Computer Centre 513 support IPv6 for end-users
2013, April: DHCPv6 for static devices (IP services)
2013, April: All LCG routers have dual-stack services
2013, June: GPN NTP servers are dual-stack
2013, September: DHCPv6 for portable devices
2013, September: DNS replies over IPv6 from DNS servers
2013, October: Gates software ready
2013, October: firewall rules translated to IPv6
2013, October: DNS automaticaly configured from LANDB information
2013, November: All GPN routers have dual-stack services
2013, November: WEBREQ shows IPv6 information to any user
2014, January: Automatic IPv6 configuration in the central firewall

IPv6 project 2014-18

2014, April 1st: DHCPv6 leases to any device in the IT datacentre in building 513

2014, May 6th: DHCPv6 leases to any registered device connected to a portable socket or WIFI

2014, May 8th: dual-stack lxplus instance available at lxplus-ipv6.cern.ch

2014, May 12th: imap, pop, smtp, ldap services dual stack

2014, June 3rd: DHCPv6 leases to any static device in GPN. DHCPv6 deployment completed.

2018, August: IPv6 HTAR

Questions?

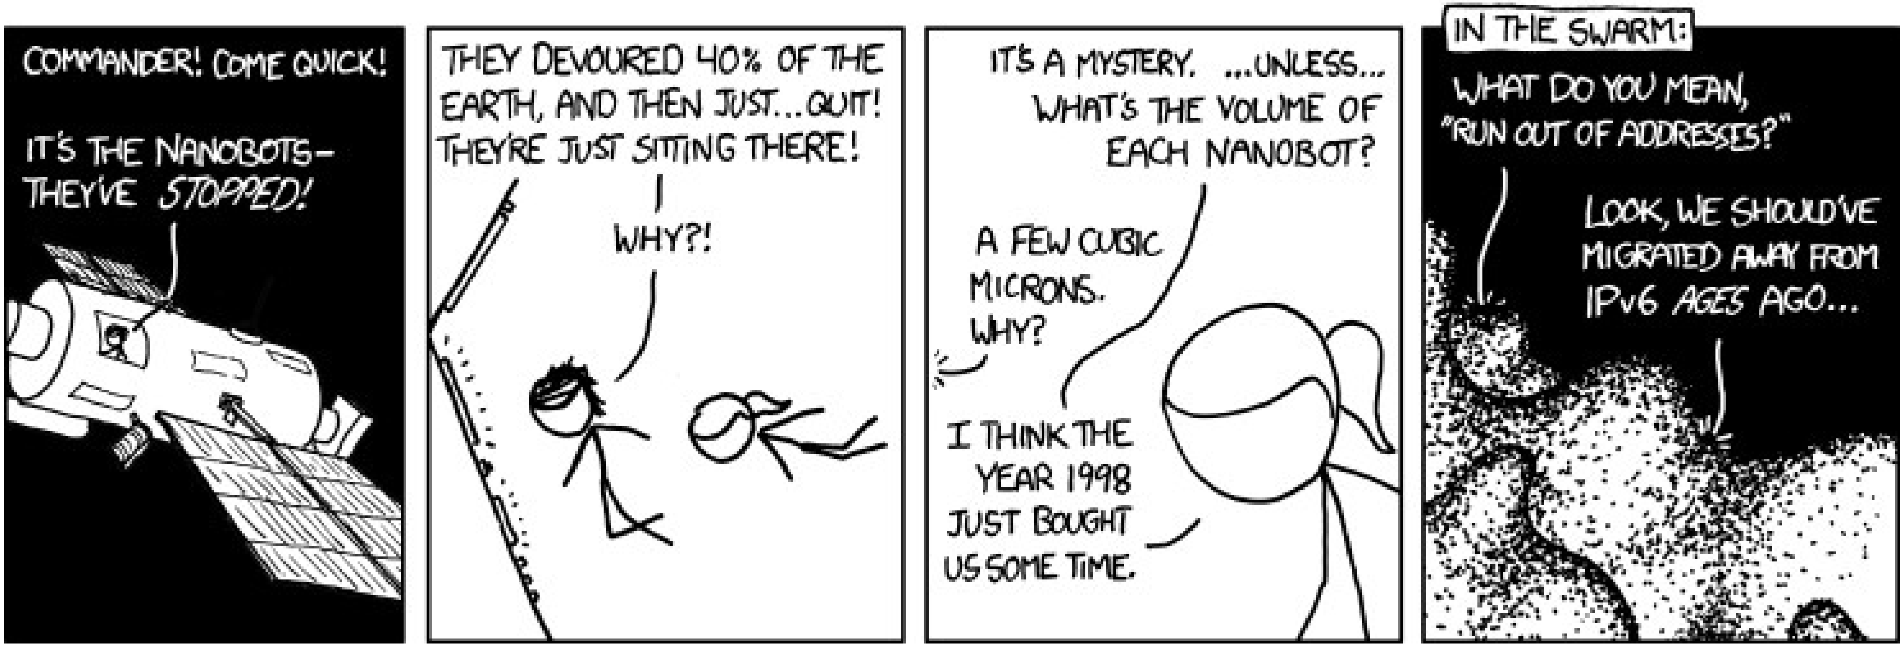
COMMANDER! COME QUICK!
IT'S THE NANOBOTS– THEY'VE STOPPED!
THEY DEVOURED 40% OF THE EARTH, AND THEN JUST...QUIT! THEY'RE JUST SITTING THERE!
WHY?!
IT'S A MYSTERY. ...UNLESS... WHAT'S THE VOLUME OF EACH NANOBOT?
A FEW CUBIC MICRONS. WHY?
I THINK THE YEAR 1998 JUST BOUGHT US SOME TIME.
IN THE SWARM:
WHAT DO YOU MEAN, "RUN OUT OF ADDRESSES?"
LOOK, WE SHOULD'VE MIGRATED AWAY FROM IPv6 AGES AGO...