

LHC[OPN/ONE] → IPv6 → history + status

Bruno Hoefft / DE-KIT

STEINBUCH CENTRE FOR COMPUTING - SCC



- First meeting 28 April 2011
- Motivation → IPv4 address space limits (RIPE/CERN)
- Starting with testbeds
(FNAL/DESY/KIT/FZU/CERN/NDGF/PIC/...)
- A small group of volunteers investigating and analysing currently deployed HEP applications
- Several issues preventing enabling of IPv6 at production env.
 - E.g.: Router, Firewall (IPv6 in software only), ...
 - HEP Applications

Grid-[Middle-]Ware


- Collection of Grid-applications (52 different apps)
- analyse of IPv6 readiness

url: <http://hepix-ipv6.web.cern.ch/wlwg-applications>
 IPv6 compliance of WLCG applications

Software Component	Type	Used by Experiment	Version	IPv6 Compliance
AliEN	LHC Experiment Application	ALICE		
ARC CE	Middleware	ATLAS, CMS		YES
ARGUS	Middleware	ALICE, ATLAS, CMS, LHCb		YES
BDII	Middleware	ATLAS, CMS, LHCb	EMI 2	YES
BestMAN	Middleware	ATLAS, CMS		YES
CASTOR	Middleware	ALICE, ATLAS, CMS, LHCb		NO
...				
glideinWMS	Middleware	CMS		Claimed
...				
WMAgent	LHC Experiment Application	CMS		Unknown
xrootd	Middleware	ALICE, ATLAS, CMS, LHCb	4.0.0	YES with caveats
xrootd	Middleware	ALICE, ATLAS, CMS	3.1.0	NO

Xrootd → caveats

url: <http://hepix-ipv6.web.cern.ch/content/xrootd>



Home

xrootd

Submitted by [Andrea Sciaba](#) on Mon, 11/25/2013 - 11:24

Contact details for Supplier and/or Developer:

Andy Hanushevsky

Version:

4.0.0

Does developer claim IPv6 compliance?:

Yes

Did you try running on dual stack system?:

Yes

If so, did it break the application?:

No

Does it bind to both stacks?:

Yes

Is IPv6 preferred by default?:

Yes

Can it be configured to prefer IPv6 or IPv4?:

Don't know/not tested

Type:

Middleware

Used by Experiment:

ALICE

ATLAS

CMS

LHCb

How does it behave with a private IPv4 and a public IPv6?:

It will always advertise any IPv4 in preference of any IPv6 address, even when initially connected to with IPv6.

IPv6

- Request
 - all Tier-1s offer IPv6 peering to LHCOPN
 - provide a dual stack PerfSONAR machine
 - by **April 2015.**
- Request
 - Tier-2s offer IPv6 peering to LHCONE
 - provide a dual stack PerfSONAR machine
 - by **August 2015.**

IPv6 Status 2015

Tier-1	LHCOPN IPv6 Peering	LHCONE IPv6 Peering	Perfsonar via IPv6
ASGC	--	--	
BNL	not on their priority list		
CH-CERN	✓	✓	LHC[OPN/ONE]
DE-KIT	✓	✓	LHC[OPN/ONE]
FNAL	✓	✓	LHC[OPN/ONE] <small>not yet in dashboard</small>
FR-CCIN2P3	✓	✓	LHCONE
IT-INFN-CNAF	--	✓	LHCONE
NGDF	✓	✓	LHC[OPN/ONE]
ES-PIC	✓	✓	LHCOPN
KISTI	startet but no IPv6 Peering implemented		
NL-T1	no peering implemented		
Triumf	IPv6 peering planned at end of 2015		
RRC-KI-T1	--	--	

proposal (steered by Alastair Dewhurst) agreed by MB in 2016 activities/requests → require IPv6 readiness

- By April 1st 2017
 - Sites can provide IPv6-only CPUs if necessary
 - Tier-1's must provide dual-stack storage access with sufficient performance and reliability
 - At least in a testbed setup
 - Stratum-1 service at CERN must be dual-stack
 - A dedicated ETF infrastructure to test IPv6 services must be available
 - ATLAS and CMS must deploy all services interacting with WNs in dual-stack
 - All the above, without disrupting normal WLCG operations
- By April 1st 2018
 - Tier-1's must provide dual-stack storage access in production with increased performance and reliability
 - Tier-1's must upgrade their Stratum-1 and FTS to dual-stack
 - The official ETF infrastructure must be migrated to dual-stack
 - GOCDB, OIM, GGUS, BDII should be dual-stack

■ Request

- all Tier-1s offer some IPv6 services
- by **April 2017.**

■ Request

- Tier-1s offer IPv6 services (~95%)
- by **April 2018.**

IPv6 tier-1 site readiness ticket

status at Tuesday 02. Feb 2017

GGuS Tickets created,
 Responce:

IPv6 peering + dual stack perfonar

work in progress

no response



IPv6 Status 2017

Tier-1				
	LHCOPN IPv6 Peering	LHCONE IPv6 Peering		dual stack Perfsonar
ASGC	--	--		
BNL	not on their priority list			
CH-CERN	✓	✓		LHC[OPN/ONE]
DE-KIT	✓	✓		LHC[OPN/ONE]
FNAL	✓	✓		LHC[OPN/ONE]
FR-CCIN2P3	✓	✓		LHCONE
IT-INFN-CNAF	✓	✓		LHC[OPN/ONE]
NGDF	✓	✓		LHC[OPN/ONE]
ES-PIC	✓	✓		LHC[OPN/ONE]
KISTI	startet but no IPv6 Peering implemented			
NL-T1	✓	✓		LHC[OPN/ONE]
RAL	--	--		IPv6 only - GPI
RRC-KI-T1 + JINR	--	--		
Triumf	IPv6 peering planned and started at end of 2015/2016 ?			

Name	Type	LHCOPN IPv6 peering	LHCONE IPv6 peering	LHCONE IPv6 peers	NREN IPv6 peers	IPv6 LAN	dualstack perfSONAR	dualstack storage percentage by 1st April 2017	dualstack storage percentage by 31st May 2017	dualstack storage percentage by 31st July 2017	Network Statistics	Notes	date last update
UKI-LT2-IC-HEP	Tier2		Yes	Janet	Janet	Yes		100					02/17/2017
TRIUMF	Tier1	Yes	Yes	Canarie	BCNET	Yes	Yes	0				TRIUMF MW readiness storage is 100% running on dual stack since April 26. And data transfer is primarily using ipv6 now.	05/05/2017
RAL	Tier1	Yes			JANET	Yes	Yes	0	20	100		IPv6 storage connectivity refers to ECHO/Ceph, Castor will not (ever) be dual stack.	05/23/2017
pragueicg2 (FZU)	Tier2	No	Yes	CESNET	CESNET	Yes	Yes	100			http://netreport.cesnet.cz/netreport/hep-cesnet-experimental-facility2/		02/17/2017
NL-T1-Nikhef	Tier1	Yes	No	GEANT	SURFnet	Yes	No	0	0	100	http://beer.nikhef.nl/cgi-bin/grapher.cgi?target=%2Fparkwachter.ipmi.nik...	Full dual stack coming soon; direct peering with CERN is up.	04/07/2017
NL-T1 SARA-MATRIX	Tier1	Yes	Yes	GEANT	SURFnet	Yes	Yes	0	100	100		Production dual stack enabled May 1 2017.	05/04/2017
FR-CCIN2P3	Tier1	Yes	Yes	RENATER	RENATER	Yes	Yes	50	100	100	N/A		03/23/2017
FNAL	Tier1	Yes	Yes	ESnet	ESnet	Yes	Yes	0	0			Dual-stack storage configuration in progress...	02/17/2017
ES-PIC	Tier1	Yes	Yes	RedIRIS	RedIRIS	Yes	Yes	100	100	100	N/A		05/08/2017
DESY	Tier2		Yes	DFN	DFN	Yes	Yes	0					02/16/2017
DE-KIT/GridKa	Tier1	Yes	Yes	DFN	DFN	Yes	Yes	2	5	60	N/A	new dCache storage is currently deployed in dual-stack	02/16/2017

IPv6 Status 2018

Tier-1	LHCOPN IPv6 Peering	LHCONE IPv6 Peering	dual stack Perfsonar
ASGC	✓	✓	LHC[OPN/ONE]
BNL	✓	✓	LHC[OPN/ONE]
CH-CERN	✓	✓	LHC[OPN/ONE]
DE-KIT	✓	✓	LHC[OPN/ONE]
FNAL	✓	✓	test LHC[OPN/ONE]
FR-CCIN2P3	✓	✓	LHC[OPN/ONE]
IT-INFN-CNAF	✓	✓	LHC[OPN/ONE]
NGDF	✓	✓	LHC[OPN/ONE]
ES-PIC	✓	✓	LHC[OPN/ONE]
KISTI	✓	✓	LHC[OPN/ONE]
NL-T1	✓	✓	LHC[OPN/ONE]
RAL	✓		LHC[OPN]
RRC-KI-T1	✓	✓	??
JINR	✓	✓	??
Triumpf	✓	✓	LHC[OPN/ONE]

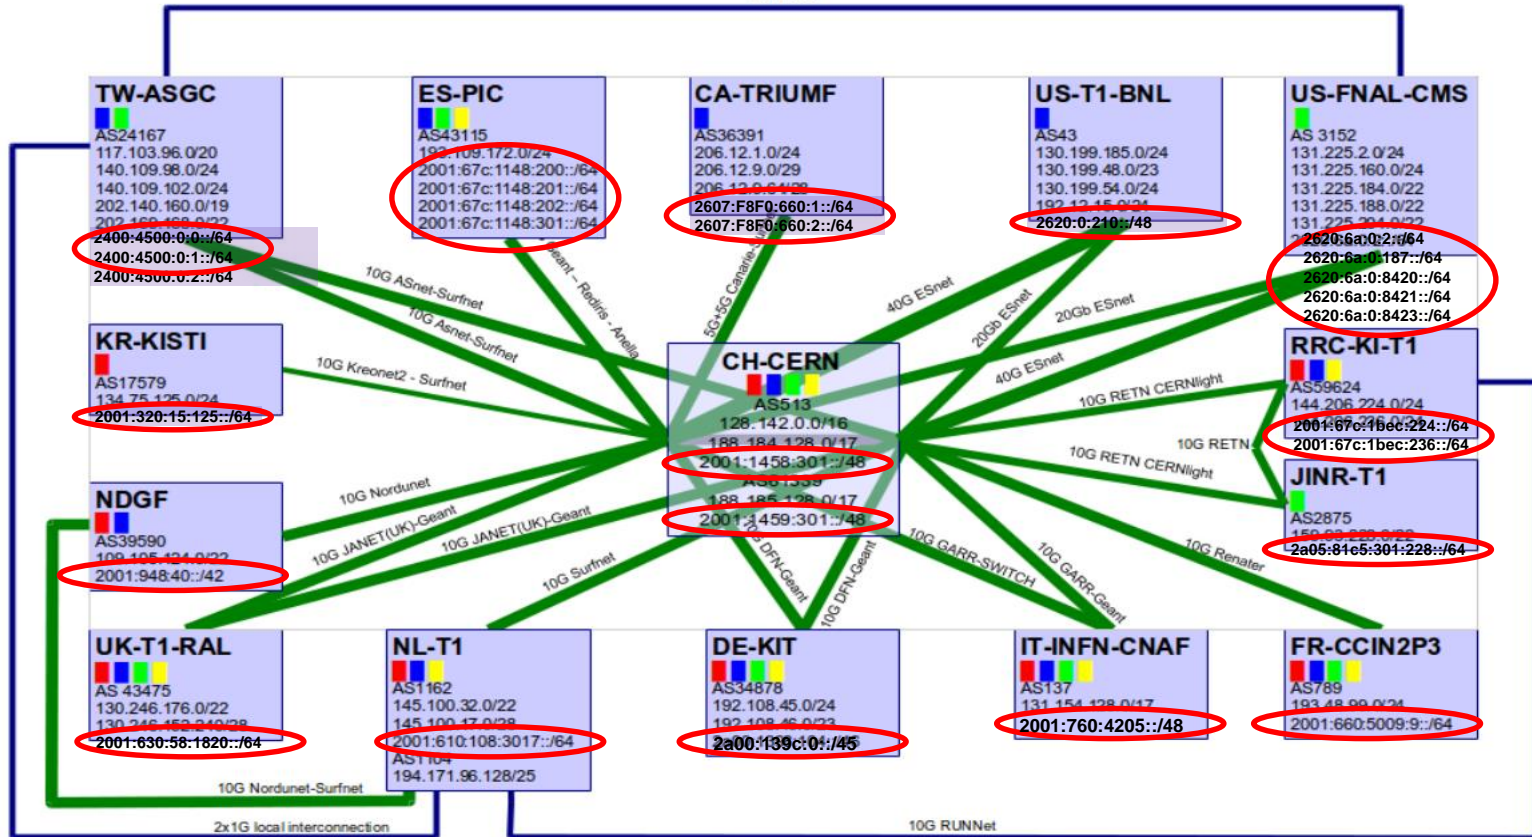
IPv6 Status 2019

Tier-1	LHCOPN	LHCONE	IPv6 Perfsonar
CA-TRIUMF	OK	OK	LHC[OPN/ONE]
CH-CERN (Tier-0)	OK	OK	LHC[OPN/ONE]
DE-KIT	OK	OK	LHC[OPN/ONE]
ES-PIC	OK	OK	LHC[OPN/ONE]
FR-CCIN2P3	OK	OK	LHC[OPN/ONE]
IT-INFN-CNAF	OK	OK	LHC[OPN/ONE]
KR-KISTI	OK	OK	LHC[OPN/ONE]
NGDF	OK	OK	LHC[OPN/ONE]
NL-T1 - NIKHEF	OK	OK	--
NL-T1 - Sara-Matrix	OK	OK	LHC[OPN/ONE]
RRC-KI-T1	OK	OK	--
RRCC-JINR-T1	OK	OK	LHC[OPN/ONE]
TW-ASGC	OK	OK	LHC[OPN/ONE]
UK-T1-RAL	OK	OK	LHC[OPN]
US-T1-BNL	OK	OK	LHC[OPN/ONE]
US-T1-FNAL	OK	OK	LHC[OPN/ONE]

Tier	AS Number	v4 p2p link address with CH-CERN	LHC v4 prefixes	LHC v6 prefixes
CA-TRIUMF	36391 (TRIUMF)	192.16.166.80/30	206.12.1.0/24	2607:F8F0:660:1::/64
		192.16.166.84/30	206.12.9.0/29 206.12.9.128/25	2607:F8F0:660:2::/64 2607:F8F0:660:3::/64
CH-CERN	513 (CERN)	-	128.142.0.0/16 188.184.128.0/17 188.185.48.0/20 188.185.128.0/17	2001:1458:301::/48 2001:1458:302::/48 2001:1458:303::/48 2001:1459:301::/48 2001:1459:303::/48
DE-KIT	58069 (FZK)	192.16.166.32/30	192.108.45.0/24	2a00:139c::/45
		192.16.166.36/30	192.108.46.0/23 192.108.68.0/24	
ES-PIC	43115 (PIC)	192.16.166.56/30 192.16.166.60/30	193.109.172.0/24	2001:67c:1148:200::/64 2001:67c:1148:201::/64 2001:67c:1148:202::/64 2001:67c:1148:301::/64
FR-CCIN2P3	789 (IN2P3)	192.16.166.40/30 192.16.166.44/30	193.48.99.0/24	2001:660:5009:9::/64
IT-INFN-CNAF	137 (GARR)	192.16.166.16/30 192.16.166.192/29	131.154.128.0/17	2001:760:4205::/48
KR-KISTI	17579 (Kreonet)	192.16.166.164/30	134.75.125.0/24	2001:320:15:125::/64
NDGF	39590 (NDGF)	192.16.166.48/30 192.16.166.52/30	109.105.124.0/22	2001:948:40::/42
NL-T1	1162 (SARA)	192.16.166.80/30	145.100.17.0/28 (perfSONAR MP) 145.100.32.0/22 (storage cluster)	2001:610:108:203a::/64 2001:610:108:3017::/64
		1104 (NIKHEF)	194.171.96.128/25 (storage cluster)	2a07:8504:120:e060::/64
RRC-KI-T1 JINR	2875	192.16.166.156/30	159.93.228.0/22	2a05:81c5:301:228::/64
RRC-KI-T1 KIAE	59624	192.16.166.156/30	144.206.224.0/24 144.206.236.0/24	2001:67c:1bec:224::/64 2001:67c:1bec:236::/64
		24167 (ASGCnet)	192.16.166.0/30 192.16.166.4/30 192.16.166.96/30	117.103.96.0/20 140.109.98.0/24 140.109.102.0/24 202.169.168.0/22 202.140.160.0/19
UK-T1-RAL	43475 (RAL)	192.16.166.64/30 192.16.166.68/30	130.246.176.0/22 130.246.152.240/28	2001:630:58:1820::/64
US-FNAL-CMS	3152 (FNAL)	192.16.166.24/30	131.225.2.0/24	2620:6a:0:2::/64
		192.16.166.28/30	131.225.160.0/24	2620:6a:0:187::/64
		192.16.166.6/30	131.225.184.0/22	2620:6a:0:8420::/64
			131.225.188.0/22 131.225.204.0/22	2620:6a:0:8421::/64 2620:6a:0:8423::/64
US-T1-BNL	43 (BNL)	192.16.166.8/30 192.16.166.12/30 192.16.166.88/30	130.199.185.0/24 130.199.48.0/23 192.12.15.0/24	2620:0:210::/48

LHCOPN

2G ASnet



— T0-T1 and T1-T1 traffic
— T1-T1 traffic only
--- Not deployed yet
— (thick) >= 10Gbps
--- (thin) < 10Gbps
■ = Alice ■ = Atlas
■ = CMS ■ = LHCb
 p2p prefix: 192.16.166.0/24 - 2001:1458:302::/48
 edoardo.martelli@cern.ch 20 160322

Still issues

- Some solved:
 - DNS + DHCP + site dependend IPv6 enabling
 - Early dCache versions needed special configuration
 - EOS had caveats
 - dCache does not support IPv6 dual-home
redesign of network deployment
 - Francesco Prelz enabled dCache code to support dual-home
code not jet committed by dCache developers
 - Network equipment does not support ...

- Some still pending:
 - FNAL → FTS server is still not dual-stack enabled
 - RRC-KI-T1 → dual-stack is not enabled

Questions ?



LHC IPv6 Working Group

