

# LHCOPN infrastructure status

**Roma - 4<sup>th</sup> April 2006**

[edoardo.martelli@cern.ch](mailto:edoardo.martelli@cern.ch)

## **LCG backbone**

The LCG backbone is now fully operational with six Force10 E1200 in production.

Two of those E1200 are dedicated to the connections to the Tier1s. Primary and Backup links are distributed between the two.

Two 10Gbps links connect the LCG backbone directly to the the CERN's external network for the L3 connections via ESnet and Geant2-IP.

**ASGC:** Still connected with 2x1Gbps. L3 backup via CERN-ASnet peering at Starlight.

**BNL:** Still connected at L3 via Esnet, but the deployment of the 10G connections has started: ESnet and USLHCnet will soon provide two links. There is also a plan for a third one via Chicago using ESnet's MPLS network.

**CNAF:** Connected with one dedicated 10G lightpath provided by Geant2-E2E. Since March they have been using their full LHCOPN prefix (131.154.128.0/17). L3 backup via Geant2-IP.

**FNAL:** Connected with primary and backup 10Gbps connections provided by USLHCnet. L3 backup via CERN-ESnet peering at Starlight and CERN-USLHCnet-ESnet. Plan for a third one via NY using ESnet's MPLS network.

**FZK:** GN2 has provided the lightpath from Geneva to Frankfurt. DFN is going to provide the remaining part up to Karlsruhe. At the moment they are still connected at L3 via DFN - Geant2-IP.

**IN2P3:** Connected with 10Gbps lightpath provided by Renater. L3 backup via Renater - Geant2-IP

**NDGF:** still connected at L3 via Sunet-Nordunet-Geant2-IP.

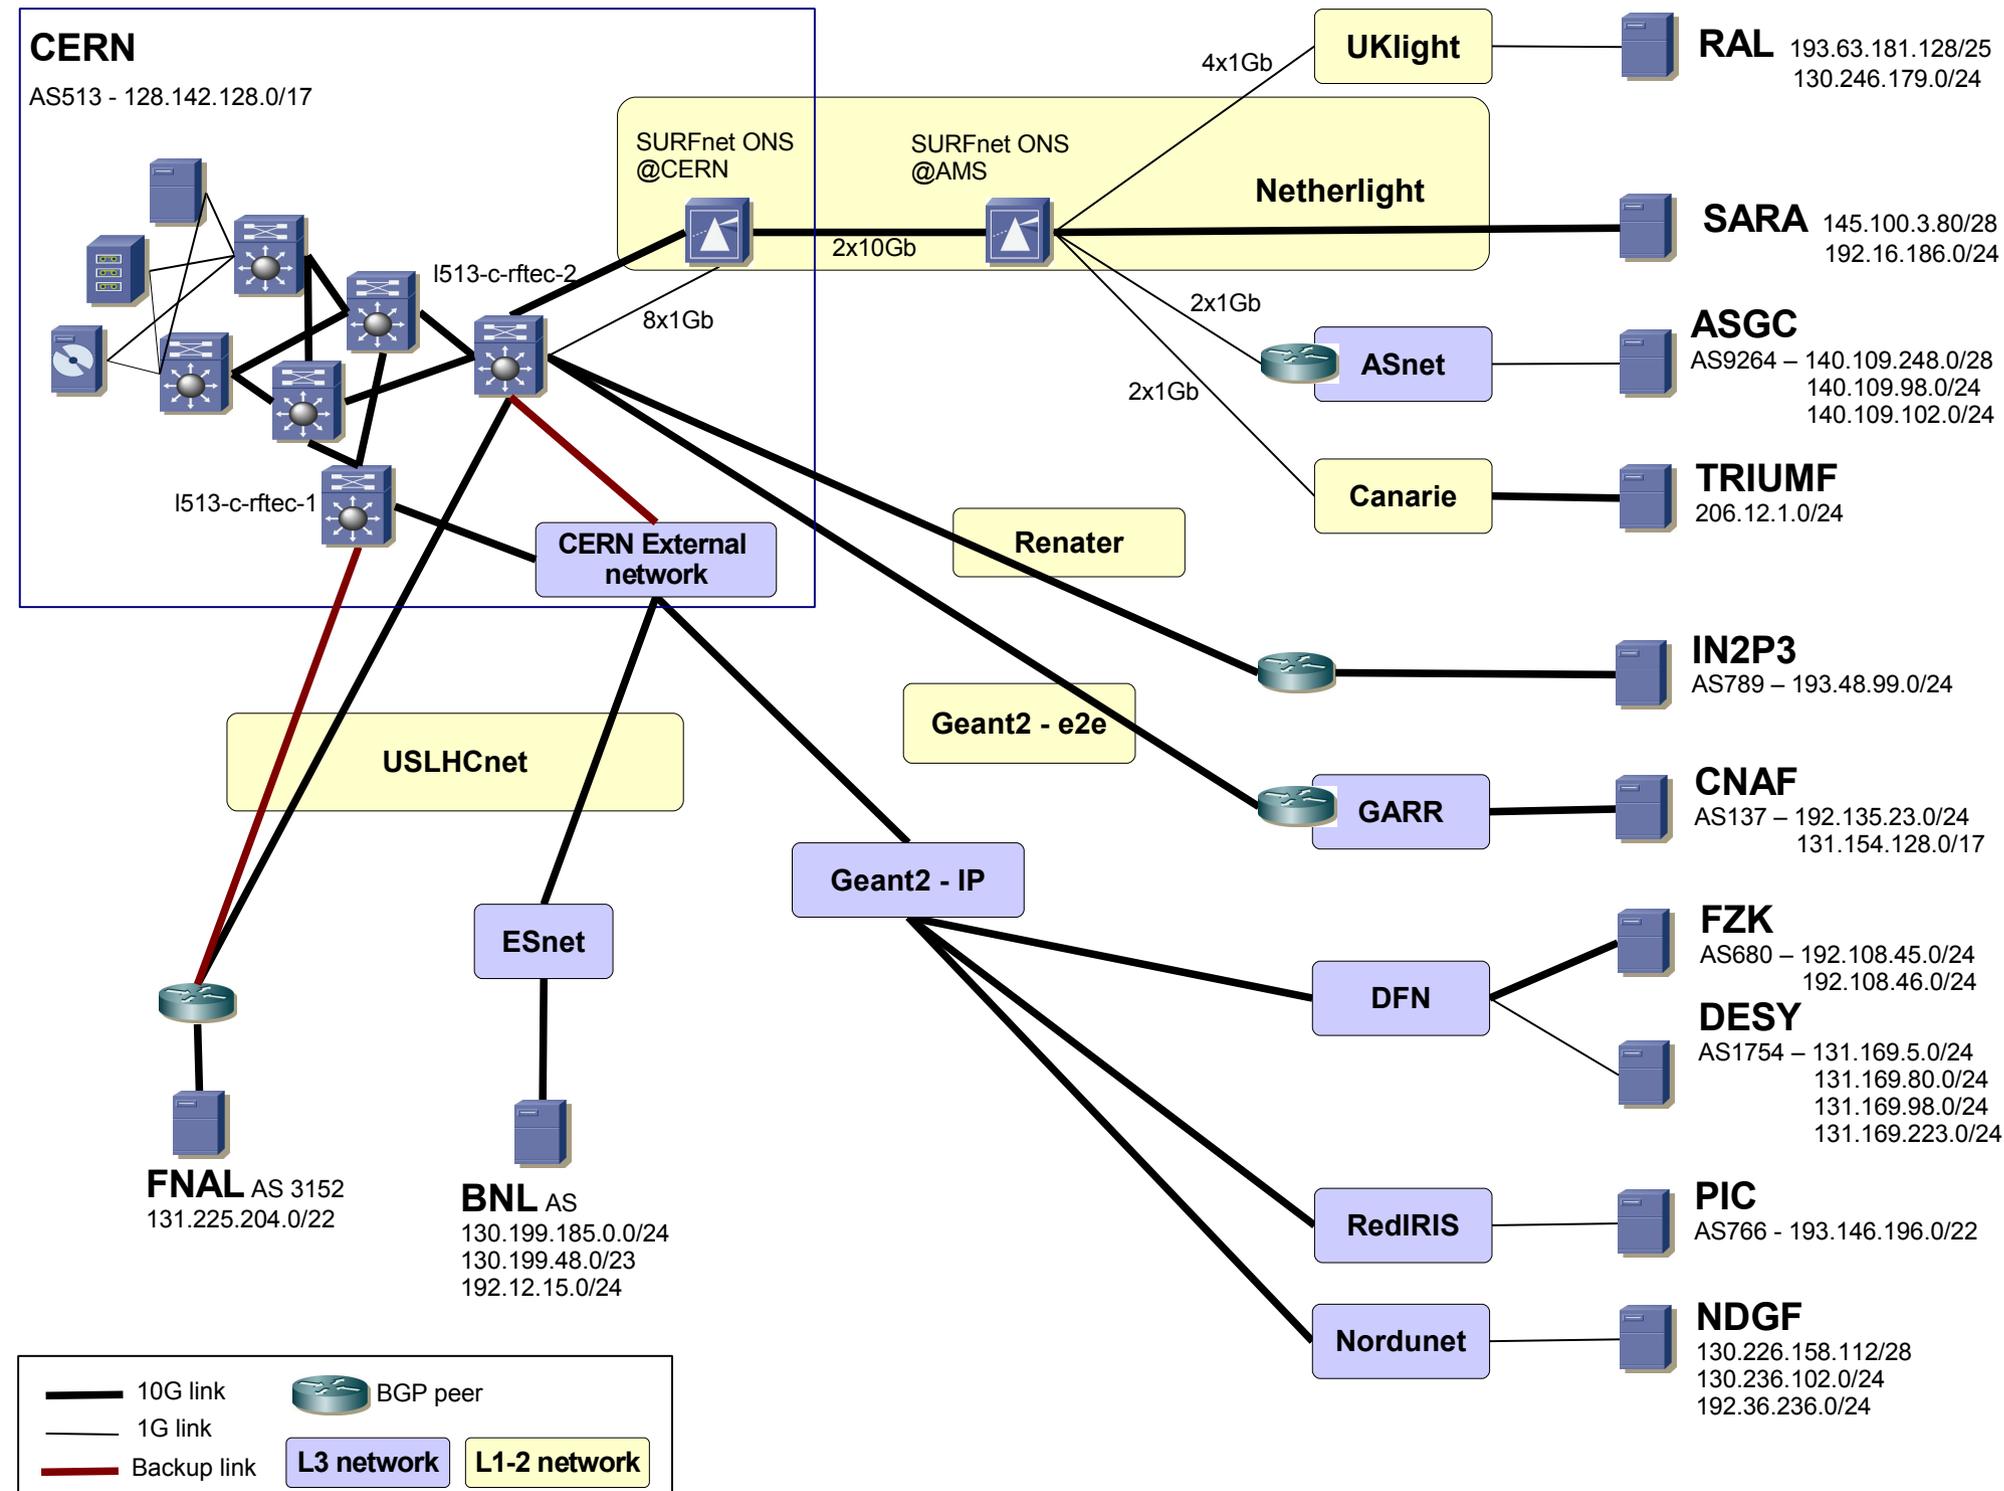
**PIC:** still connected at L3 via Rediris-Geant2-IP.

**RAL:** Connection upgraded to four 1Gbps lightpaths provided by Janet and Surfnet, but still no BGP configured. No automatic backup.

**SARA:** 10Gbps lightpath provided by Surfnet, but still no BGP configured. No automatic backup.

**TRIUMF:** still connected with two 1Gbps connections provided by Canarie and Surfnet, still no BGP configured. No automatic backup.

# LHCOPN – current status





January 2006



March 2006

All the new machines deployed are directly connected to the LCG backbone

## LHCOPN Twiki:

[Edit](#) [Attach](#) [Printable](#) LHCOPN.WebHome r1.18 - 07 Mar 2006 - 13:07 - [EdoardoMARTELLI](#) [topic end](#)

### LHC Optical Private Network

- LHCOPN architecture document: [LHCopnArchitecture](#)

#### Routing

- Routing working group document: [LHCopnRoutingDoc](#)
- LHCOPN at a glance: [LHCopnTables](#)
- [ImplementationDetails](#)
- [Traffic Statistics](#)
- [LHCOPN network map](#)
- LHCOPN prefixes: [RS-LHCOPN](#)
- [Routing WG Mailing List archive](#)

#### Operations

- NOCs information: [LHCopnOperations](#)
- [LHCOPN traffic statistics](#)
- [Operations WG Mailing List archive](#)

#### Security

- [LHCOPN security policy document](#)
- [Security WG Mailing List archive](#)

#### Monitoring

- [Monitoring WG Mailing List archive](#)

#### Other links:

- [LCG network activity page](#)
- [LGC Service Challenge TWiki](#)

<https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome>  
or jump from <http://lhcopn.cern.ch/>

The Routing Working Group document with the recommendations that complement the original LHCOPN Architecture document is on the LHCOPN twiki:  
<https://twiki.cern.ch/twiki/bin/view/LHCOPN/LHCopnRoutingDoc>

Edit Attach Printable

LHCOPN.LHCopnRoutingDoc r1.2 - 21 Mar 2006 - 14:19 - [EdoardoMARTELLI](#) [topic end](#)

## LHCOPN - IP addressing and routing

### Scope

This document presents some guidelines for the implementation of the LHCOPN network. It focuses on IP addressing and routing aspects; it uses and completes the instructions already provided by the [LCG Network Architecture document](#); it collects all the decisions taken by the LHCOPN routing working group.

### IP addressing

- **Use of public IP addresses:** Every Tier must allocate publicly routable IP address space to the machines that need to be reached over the T0-T1 links (the "LHC prefixes"). LHC prefixes cannot be from RFC1918.
- **Aggregate in few prefixes:** every Tier must aggregate the address space dedicated to the LHCOPN traffic into one or few CIDR blocks.
- **Addresses for the T0-T1 links:** the prefix 192.16.166.0/24 has been allocated for the addressing of the point-to-point links between T0 and the Tier1s. Refer to [LHCopnTables#AnchorASnumbersLHCprefixes](#) for the allocation of the /30.
- **Security:** for security reason, only packets with source and destination IP address that belong to one of the LHCOPN prefixes can transit on the LHCOPN. Refer to LINK-TO-BE-PROVIDED for more information about security in the LHCOPN.
- **LHCOPN prefixes repository:** the list of the LHCOPN prefixes is saved in the RIPE route-set object [RS-LHCOPN](#). The Object is maintained by CERN, and all the request for modification must be sent to [extip@cernNOSPAM.ch](mailto:extip@cernNOSPAM.ch).

### Routing

- **Traffic flows:** the LHCOPN is essentially a star with the T0 at the centre. Traffic will normally flow from the T0 to every T1 and back.

The list of all the NOCs involved is in:

<https://twiki.cern.ch/twiki/bin/view/LHCOPN/LHCopnOperations>

*The mailing list that collect all of them has not been created: is it necessary?*

- [SARA NOC](#)

### Time zones

- [Time zones](#)
- [Current time at Tier sites](#)

### ASGC

Provider	Section	NOC email	NOC telephone	Availability
ASGC	Taipei	<a href="mailto:noc@twgridNOSPAM.org">noc@twgridNOSPAM.org</a>	+886.9.18215162	Phone: 24x7 for Emergency Only. Email: 8x5 UTC+8
ASnet	Taipei-Amsterdam	<a href="mailto:noc@twgridNOSPAM.org">noc@twgridNOSPAM.org</a>	+886.9.18215162	Phone: 24x7 for Emergency Only. Email: 8x5 UTC+8
<a href="#">SARA</a>	Amsterdam-Geneva	Details <a href="#">here</a>		
<a href="#">CERN</a>	Geneva	Details <a href="#">here</a>		

### BNL

Provider	Section	NOC email	NOC telephone	Availability
BNL	Brookhaven	<a href="mailto:itdhelp@bnlNOSPAM.gov">itdhelp@bnlNOSPAM.gov</a>	+1-631.344.5522	8x5 EST (UTC-5)
ESnet	Brookhaven-MANLAN	<a href="mailto:trouble@esNOSPAM.net">trouble@esNOSPAM.net</a>	+1 510-486-7607	24x7 PST (UTC-8)
USLHCnet	MANLAN-Geneva	<a href="mailto:noc@uslhcnetsNOSPAM.org">noc@uslhcnetsNOSPAM.org</a>		8x5 CET (UTC+1)
<a href="#">CERN</a>	Geneva	Details <a href="#">here</a>		

### CERN

Provider	Section	NOC email	NOC telephone	Availability
<a href="#">CERN</a>	Geneva	Details <a href="#">here</a>		

### CNAF

Some sites have already implemented the suggested architecture and use BGP to route the traffic.

/30 from the 192.16.166.0/24 are used for the addressing of the direct lightpaths. Removal of the legacy /30 is still on going.

Only T0-T1 traffic is allowed at the moment; no T1-T1 transit.

Pairs of T1s can use their direct link to provide mutual backup.

Security ACLs in place on the T0's interfaces facing the Tier1s.

ACLs are at IP level for the time being (src-dst ip addresses are LHCOPN prefixes).

The list of IP prefixes allowed in the LHCOPN is stored in the RIPE database in the route-set object **RS-LHCOPN**.

To retrieve it:

<http://www.ripe.net/perl/whois?&searchtext=RS-LHCOPN>

CERN maintains the object. Please contact [extip@cern.ch](mailto:extip@cern.ch) for any request.

On the twiki there is an example script that build ACLs and filters using this object.

Implement BGP routing for all the sites: static routing and extended LAN don't allow automatic backups, are difficult to debug, require lot of reconfiguration effort in case of prefix change.

Keep deploying the direct lightpaths as soon as they are available. BNL and GRIDKA should be soon ready.