

Routing issues and backup effectiveness

Cambridge - 12 January 2007

edoardo.martelli@cern.ch

LHCOPN update

Gridka: 10Gbps link to CERN (GN2-DFN) fully operational.

Gridka-CNAF: 10Gbps link (DFN-SWITCH-GARR) deployed but not used yet.

Gridka-SARA: 10Gbps link (DFN-SURFNET) deployed but used yet.

NDGF: 10Gbps link to CERN (GN2) deployed but not used (routing issue).

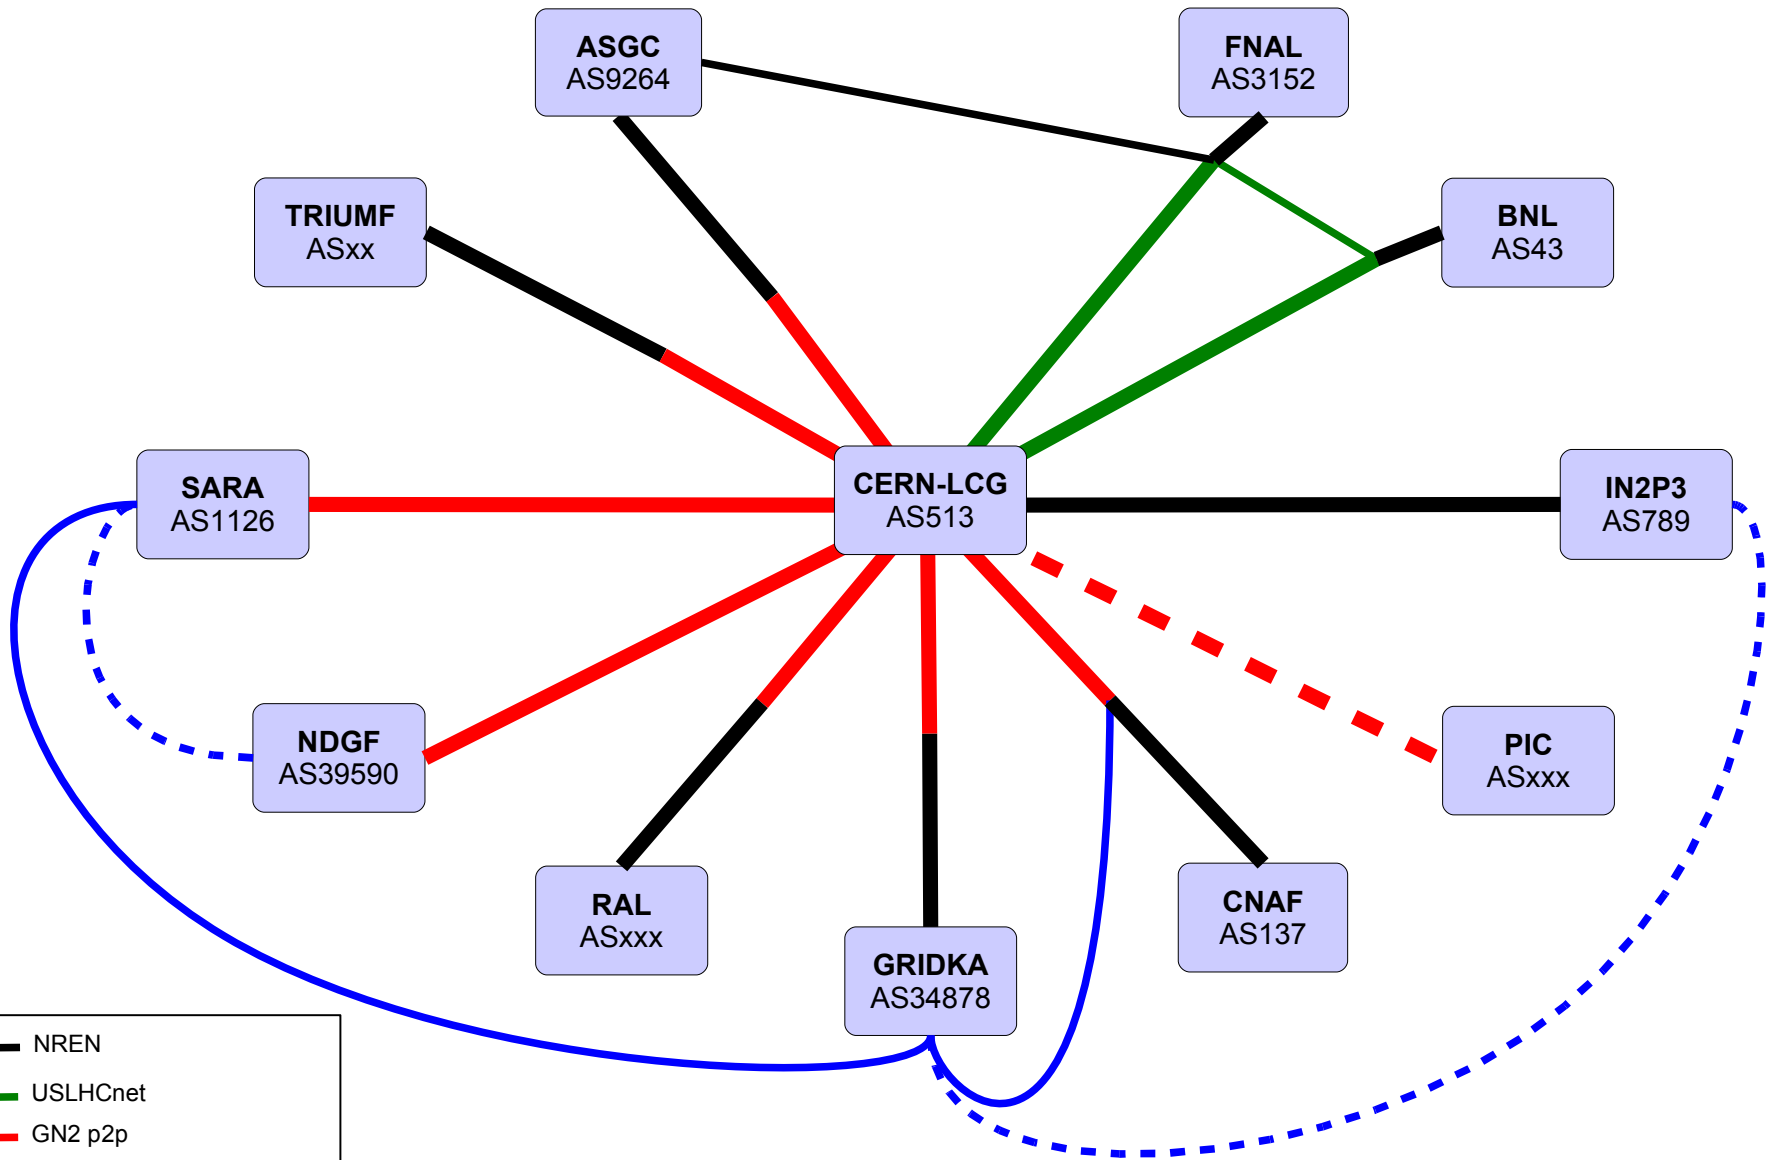
RAL: 10Gbps link (GN2-JANET) operational, but with static routing (public AS number still missing).

SARA: three GN2 lambdas Amsterdam-Geneva operational; at CERN they are connected to the recently installed SURFnet's Nortel OME6500.

SARA: 10Gbps link (GN2) deployed but not used (routing issue).

TRIUMF: 5Gbps link (Canarie-Surfnet-GN2) almost ready.

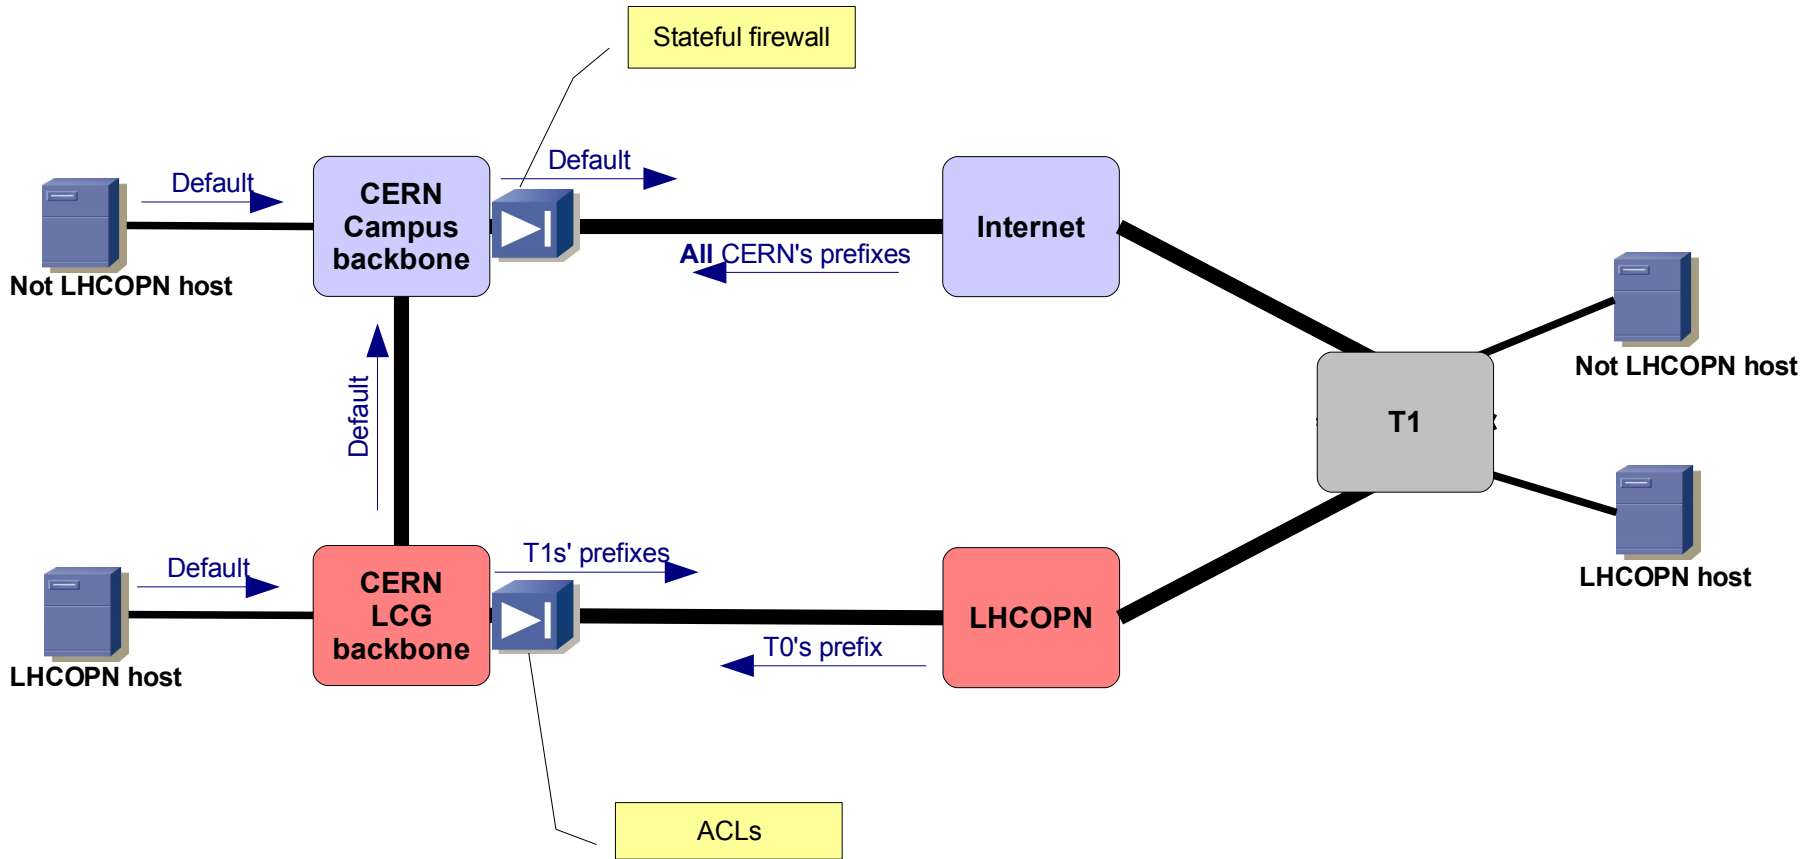
LHCOPN links (active and foreseen)



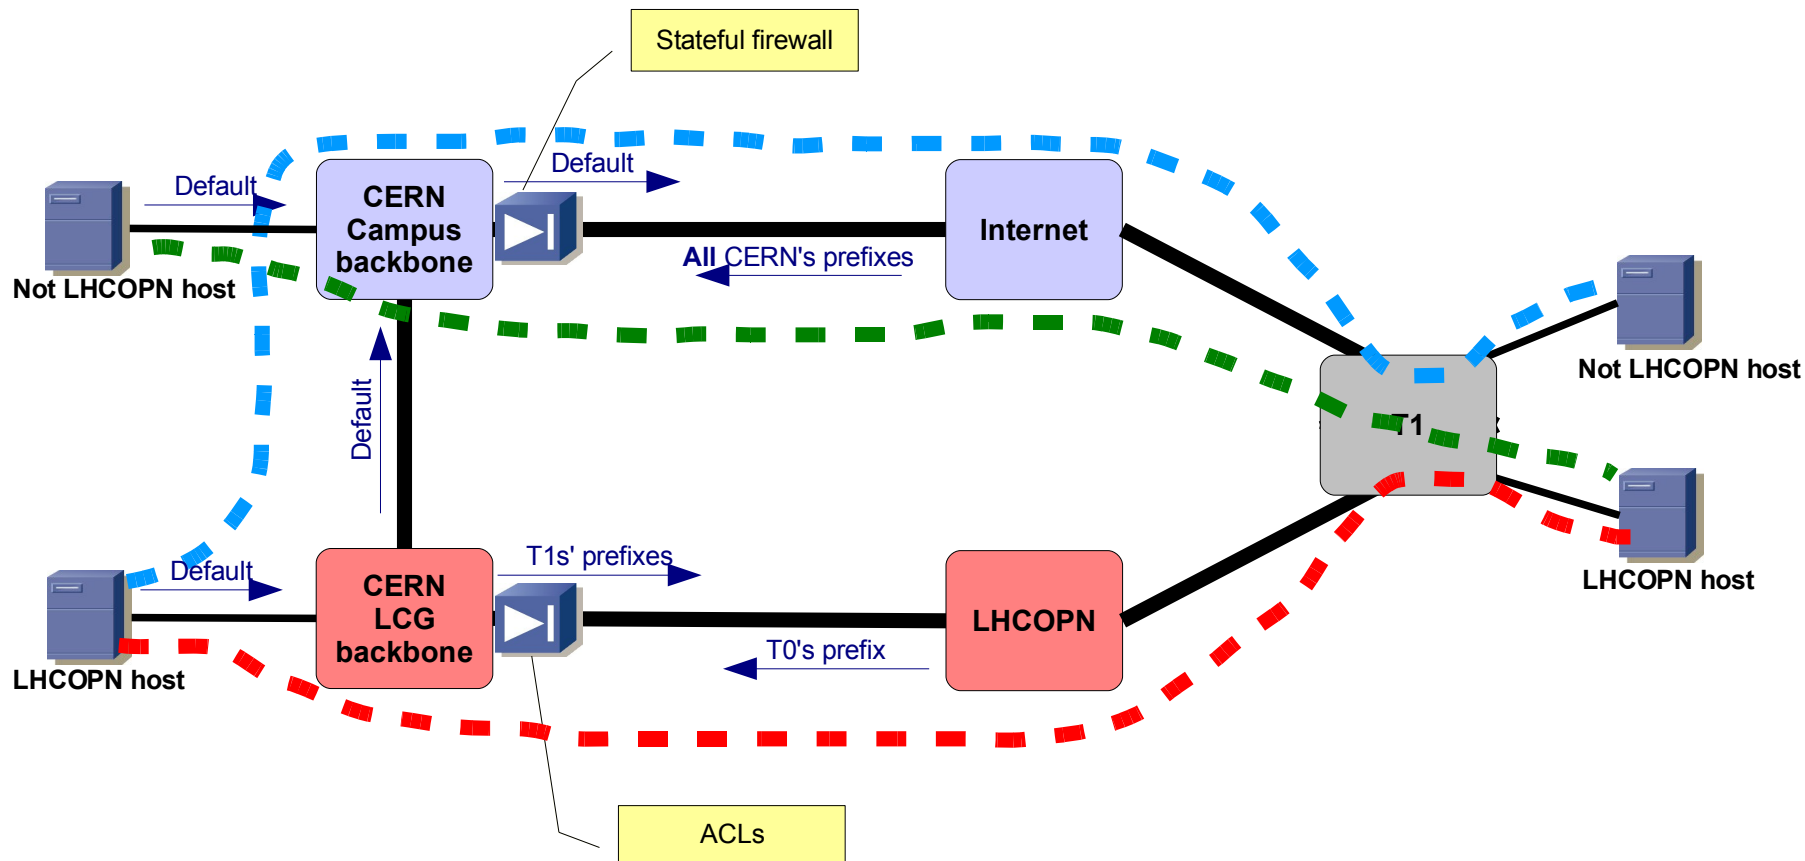
- NREN
- USLHCnet
- GN2 p2p
- Cross Border fibre
- - - Not deployed yet
- (thick) Primary path
- (thin) Backup path

Routing

CERN routing



Routing must be symmetric



Red dashed line: LHCOPN host to LHCOPN host

Blue dashed line: T0's LHCOPN host to T1's not LHCOPN host

Green dashed line: T0's not LHCOPN host to T1's LHCOPN host

Dedicated backbone/routers

(e.g. CERN, GRIDKA (?))

Policy Based routing (PBR)

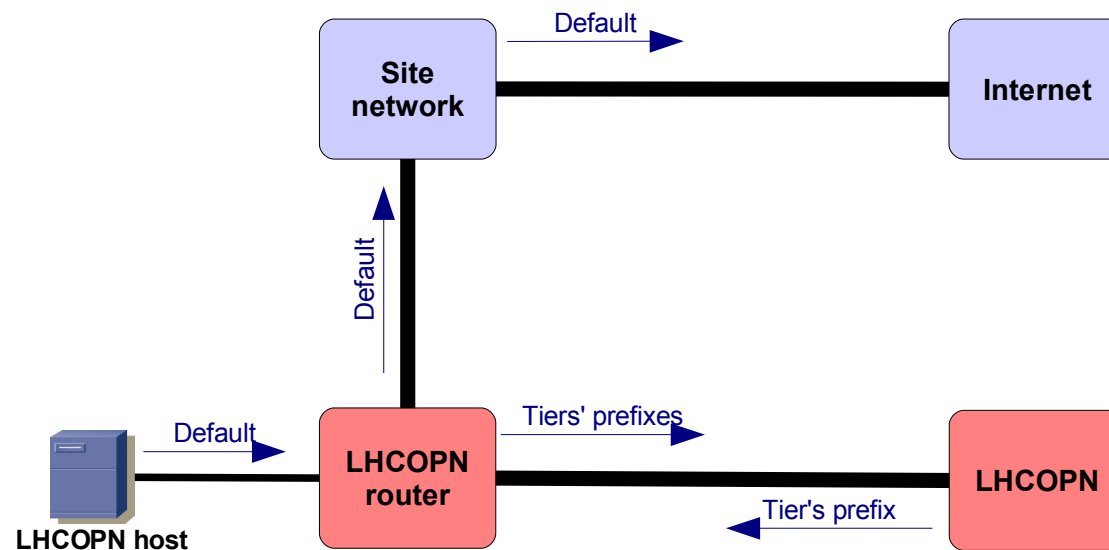
(e.g. BNL, FNAL, IN2P3, ASGC (?))

Separated routing instance

(e.g. CNAF-GARR)


The LHCOPN hosts are connected to routers dedicated to them and directly connected to the LHCOPN.

In the routing table of these routers there are only the LHCOPN prefixes, plus a default route towards a generic upstream.



LHCOPN hosts and generic hosts are connected to a shared backbone.

The router connected to the LHCOPN must policy route the LHCOPN traffic into the LHCOPN, basing its decisions on the source and destination addresses.

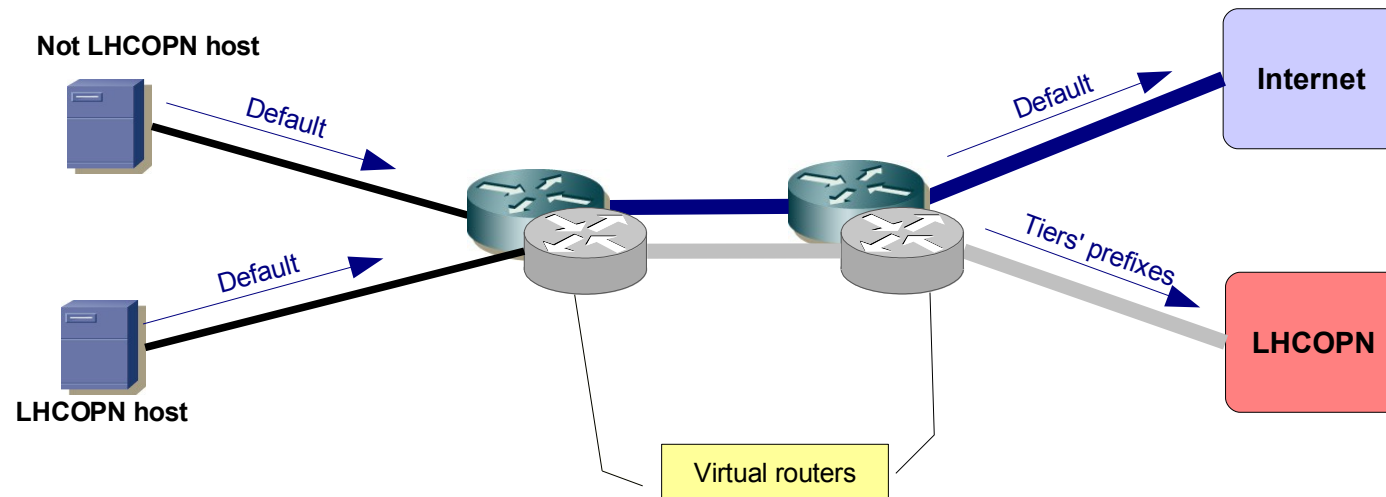


Separated Routing Instance

LHCOPN hosts and generic hosts are connected to a shared backbone, but to a different virtual routing instances (or VPN).

The virtual routing instance behaves like the dedicated backbone.

Virtual routers functionality not available on all the routers (Juniper and Cisco have it).



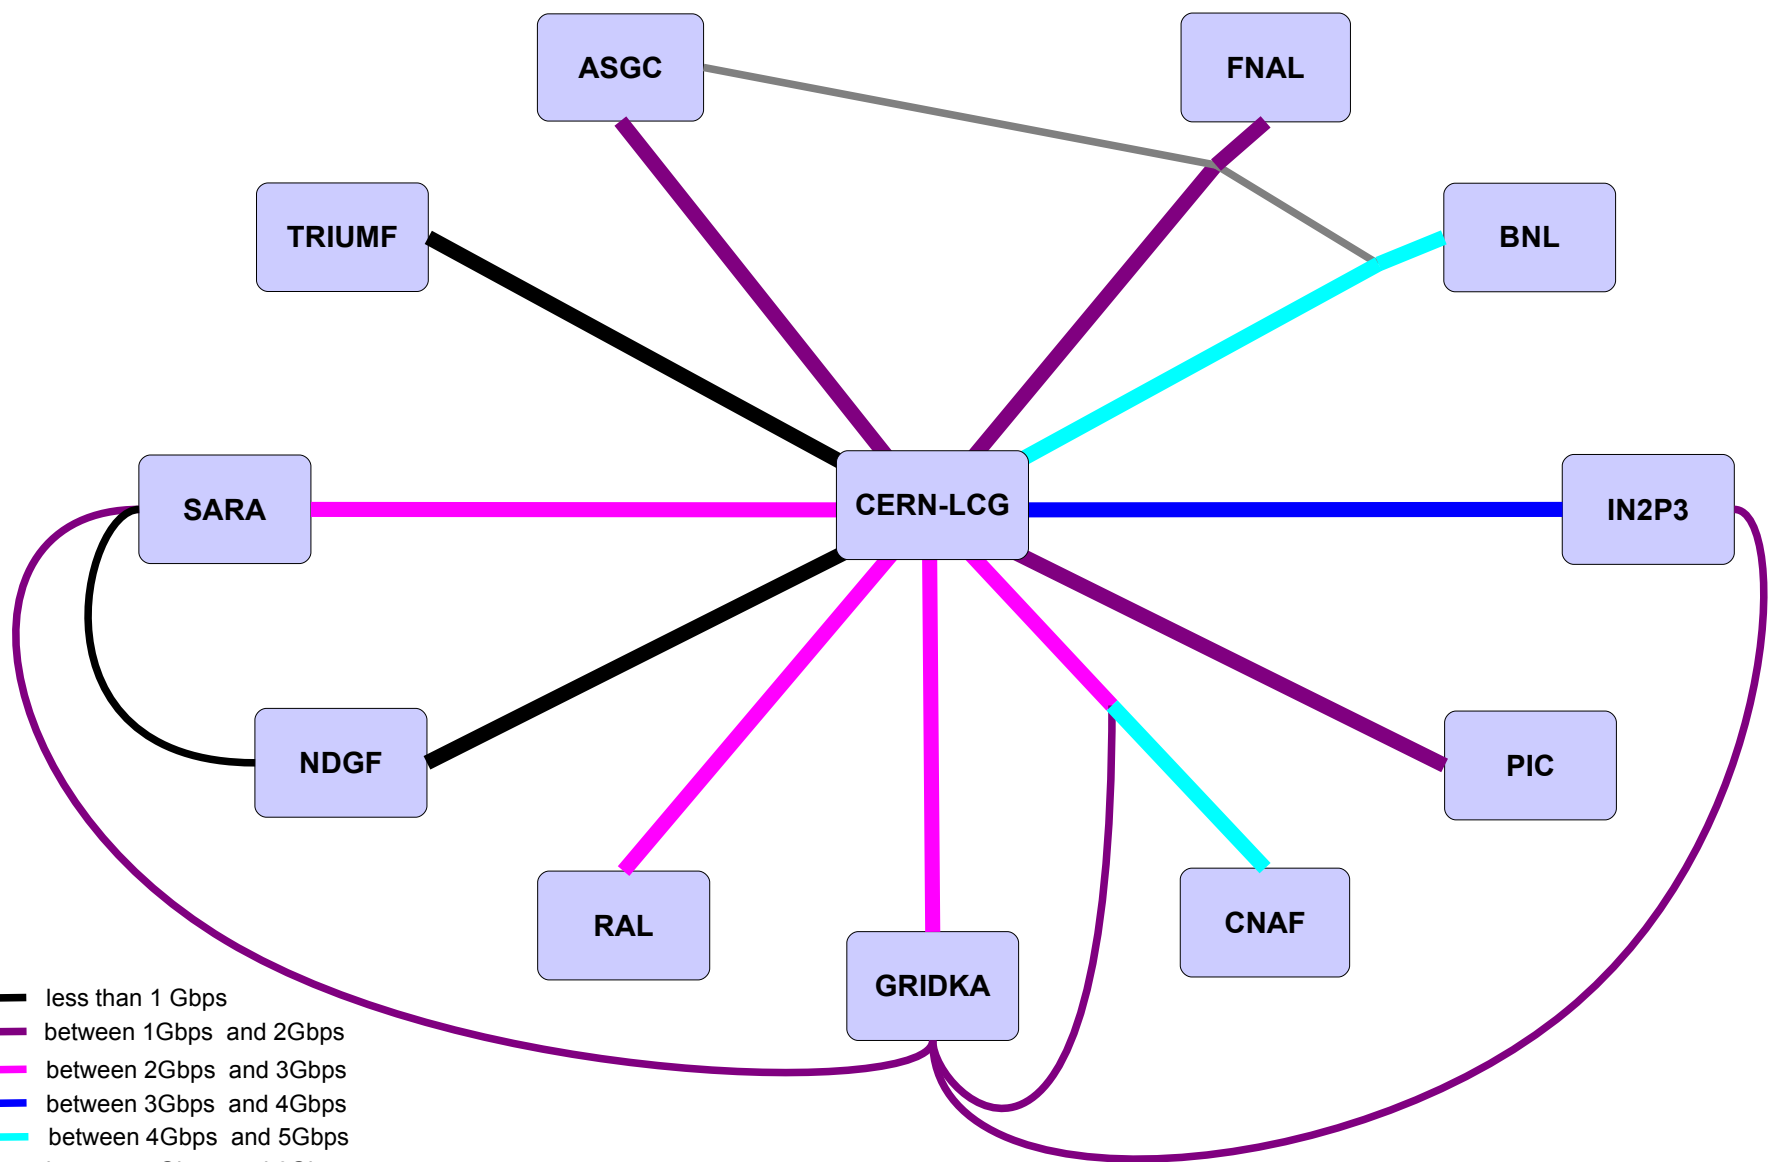
Backup effectiveness

Backup effectiveness

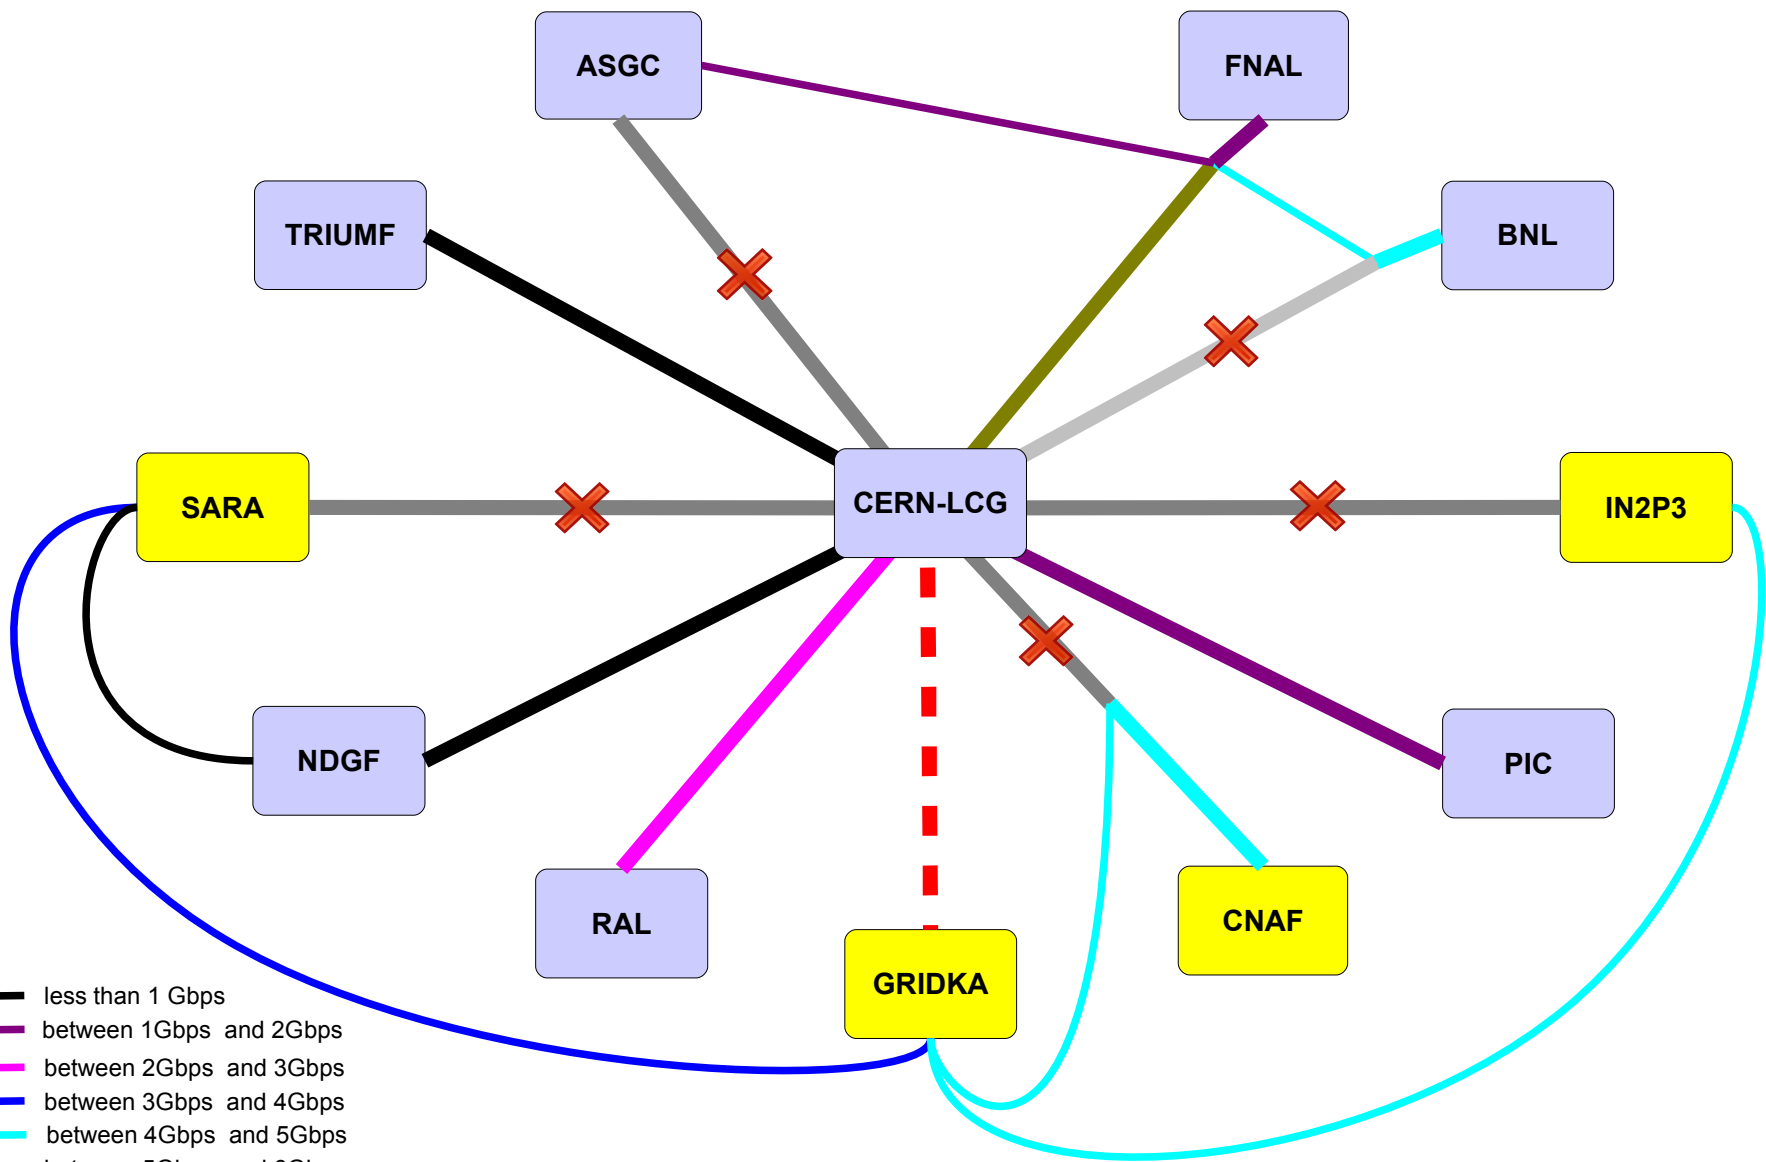


PART \ LINK	ASGC	ASGC - BU via Starlight	BNL	BNL - BU via Starlight	CNAF	CNAF - BU via Gridka	FNAL	FNAL - BU via MANLAN	GRIDKA	Gridka - BU via GARR	Gridka - BU via SARA	Gridka - BU via IN2P3	IN2P3	IN2P3 - BU via Gridka	NDGF	NDGF - BU via SARA	PIC	PIC - BU	RAL	RAL - BU	SARA	SARA - BU via GRIDKA	SARA - BU via NDGF	TRIUMF	TRIUMF - BU
CERN LCG router b513-crtec-1	x		x		x		x		x					x	x				x			x	x	x	
CERN LCG router b513-crtec-2		x		x	x		x			x	x	x	x			x						x			
Surfnets Nortel OME - Geneva	x																					x			x
Surfnets Nortel OME - Amsterdam	x										x											x			x
GN2 Alcatel switch - Geneva	x				x	x			x	x	x		x	x	x		x		x		x	x	x	x	x
GN2 Alcatel switch - Frankfurt	x					x			x		x		x	x	x						x	x	x	x	x
GN2 fibre Geneva-Frankfurt	x					x			x		x		x	x	x						x	x	x	x	x
GN2 fibre Frankfurt-Amsterdam	x										x										x				x
GN2 fibre Frankfurt-Copenhagen															x										
GN2 fibre Geneva-Milano					x					x															
GN2 fibre Geneva-Paris																				x					
GN2 fibre Paris-London																				x					
GN2 fibre Geneva-Barcelona-Madrid																									
Trench Geneva-Basel (Switch and GN2 fibres)	x				x	x			x	x	x		x	x	x						x	x	x	x	x
DFN fibre Karlsruhe-Frankfurt									x		x											x			
DFN fibre Karlsruhe-Renater						x				x	x		x												
Gridka router - Karlsruhe						x			x	x	x		x									x			
GARR router - Milano					x	x				x															
IN2P3 router - Lyon												x	x	x											
Nordunet router - Copenhagen															x	x									x
SARA router - Amsterdam											x					x						x	x	x	
USLHCnet router e600gva1 - Geneva			x					x																	
USLHCnet router e600gva2 - Geneva	x		x				x																		
USLHCnet router e600nyc - New York			x	x				x																	
USLHCnet router e600chi - Chicago	x		x					x	x																
USLHCnet link Geneva - New York			x						x																
USLHCnet link Chicago - New York				x					x																
USLHCnet link Geneva - Chicago	x		x					x																	

→ = Single point of failure for 2 or more T1s



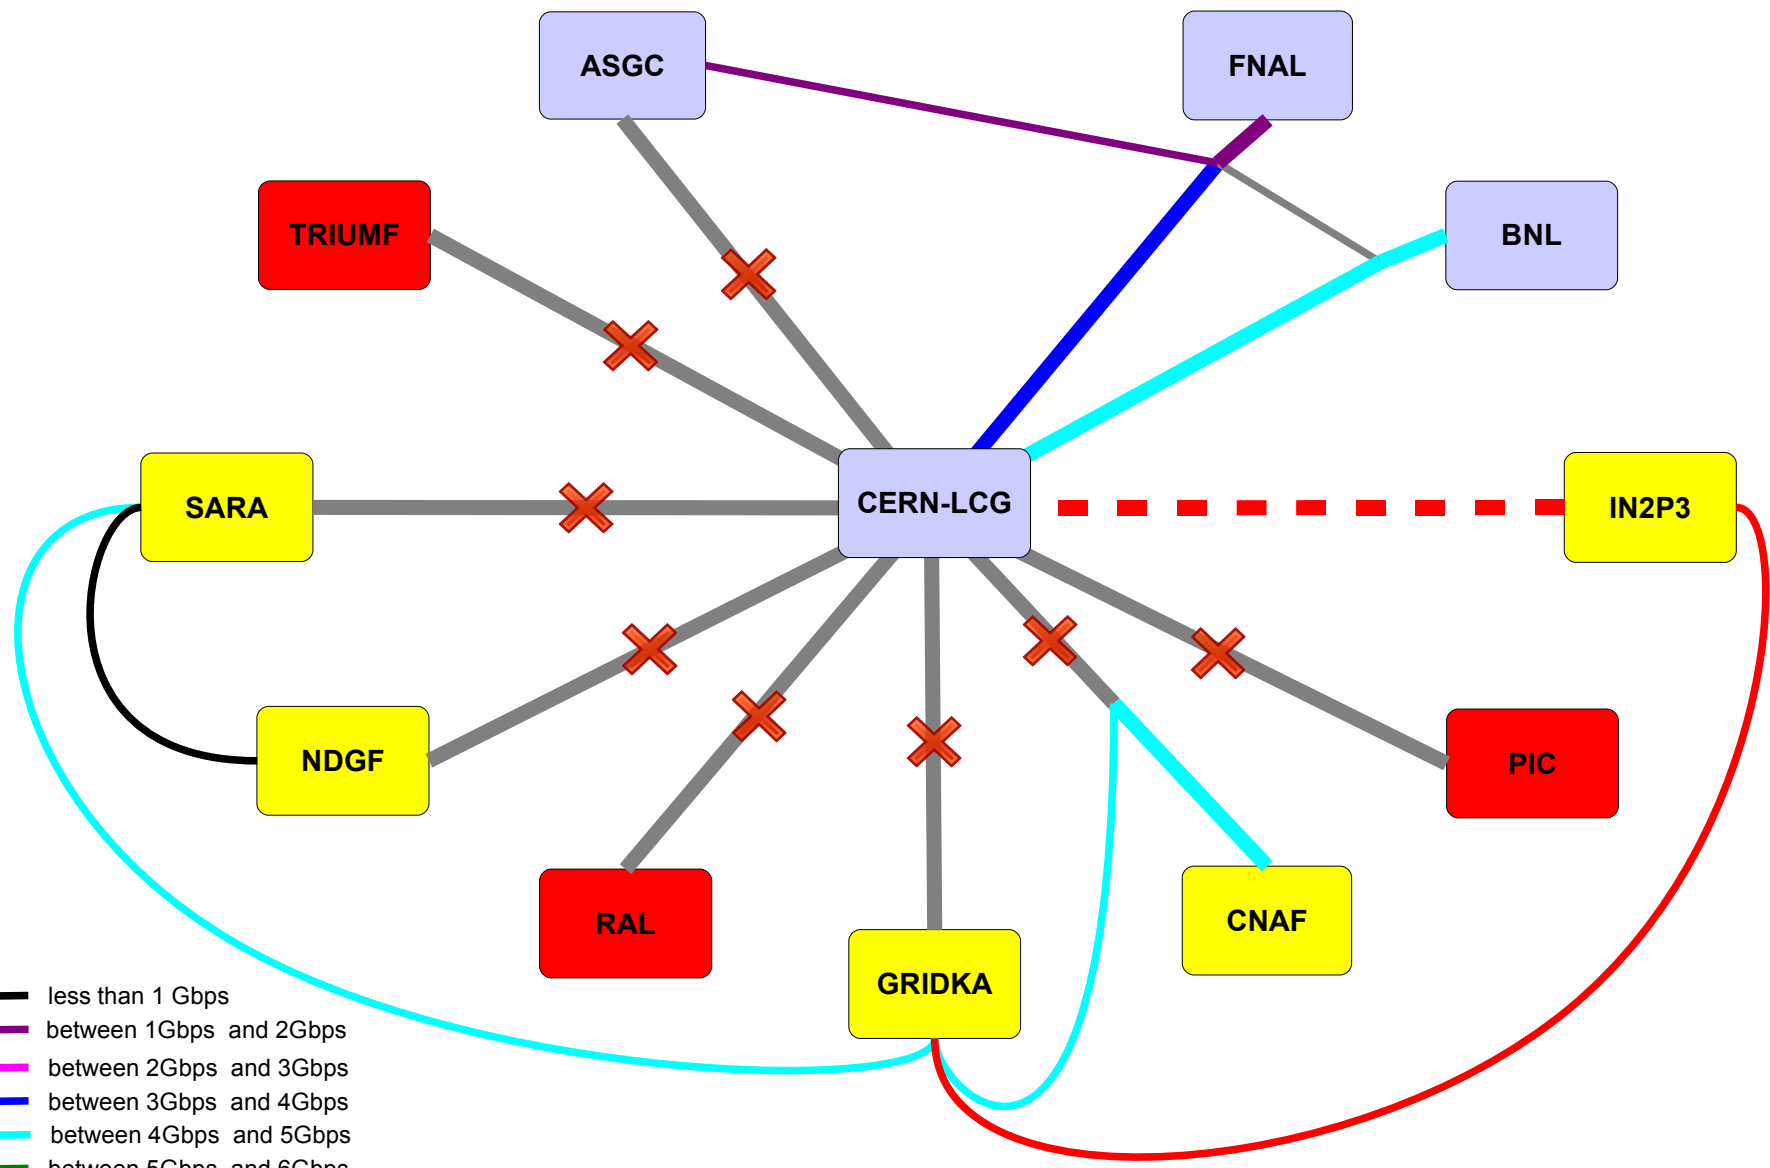
- less than 1 Gbps
- between 1Gbps and 2Gbps
- between 2Gbps and 3Gbps
- between 3Gbps and 4Gbps
- between 4Gbps and 5Gbps
- between 5Gbps and 6Gbps
- between 6Gbps and 7Gbps
- between 7Gbps and 8Gbps
- between 8Gbps and 9Gbps
- between 9Gbps and 10Gbps
- - - Overloaded



- less than 1 Gbps
- between 1Gbps and 2Gbps
- between 2Gbps and 3Gbps
- between 3Gbps and 4Gbps
- between 4Gbps and 5Gbps
- between 5Gbps and 6Gbps
- between 6Gbps and 7Gbps
- between 7Gbps and 8Gbps
- between 8Gbps and 9Gbps
- between 9Gbps and 10Gbps
- - - Overloaded

Sara should prefer NDGF for its backup

Links load – GN2 switch in Geneva down



- less than 1 Gbps
- between 1Gbps and 2Gbps
- between 2Gbps and 3Gbps
- between 3Gbps and 4Gbps
- between 4Gbps and 5Gbps
- between 5Gbps and 6Gbps
- between 6Gbps and 7Gbps
- between 7Gbps and 8Gbps
- between 8Gbps and 9Gbps
- between 9Gbps and 10Gbps
- Overloaded

IN2P3 link overloaded if it provides backup to all the T1s connected to GRIDKA

Backup via peering T1s is not very reliable because most of the Tier1s use the same carrier.

Also, a tier1 providing backup to others can have its data transfers impaired although its connectivity is fine.

Reliability can increase if GN2 protects the circuits.