

# **Firewalling beyond 10Gbps**

## ***Securing the LCG***

**Terena conference - 24 May 2007**

**Edoardo.Martelli@cern.ch**

The security related aspects of this work have been co-funded by the European Commission ISSeG project, <http://www.isseg.eu/>

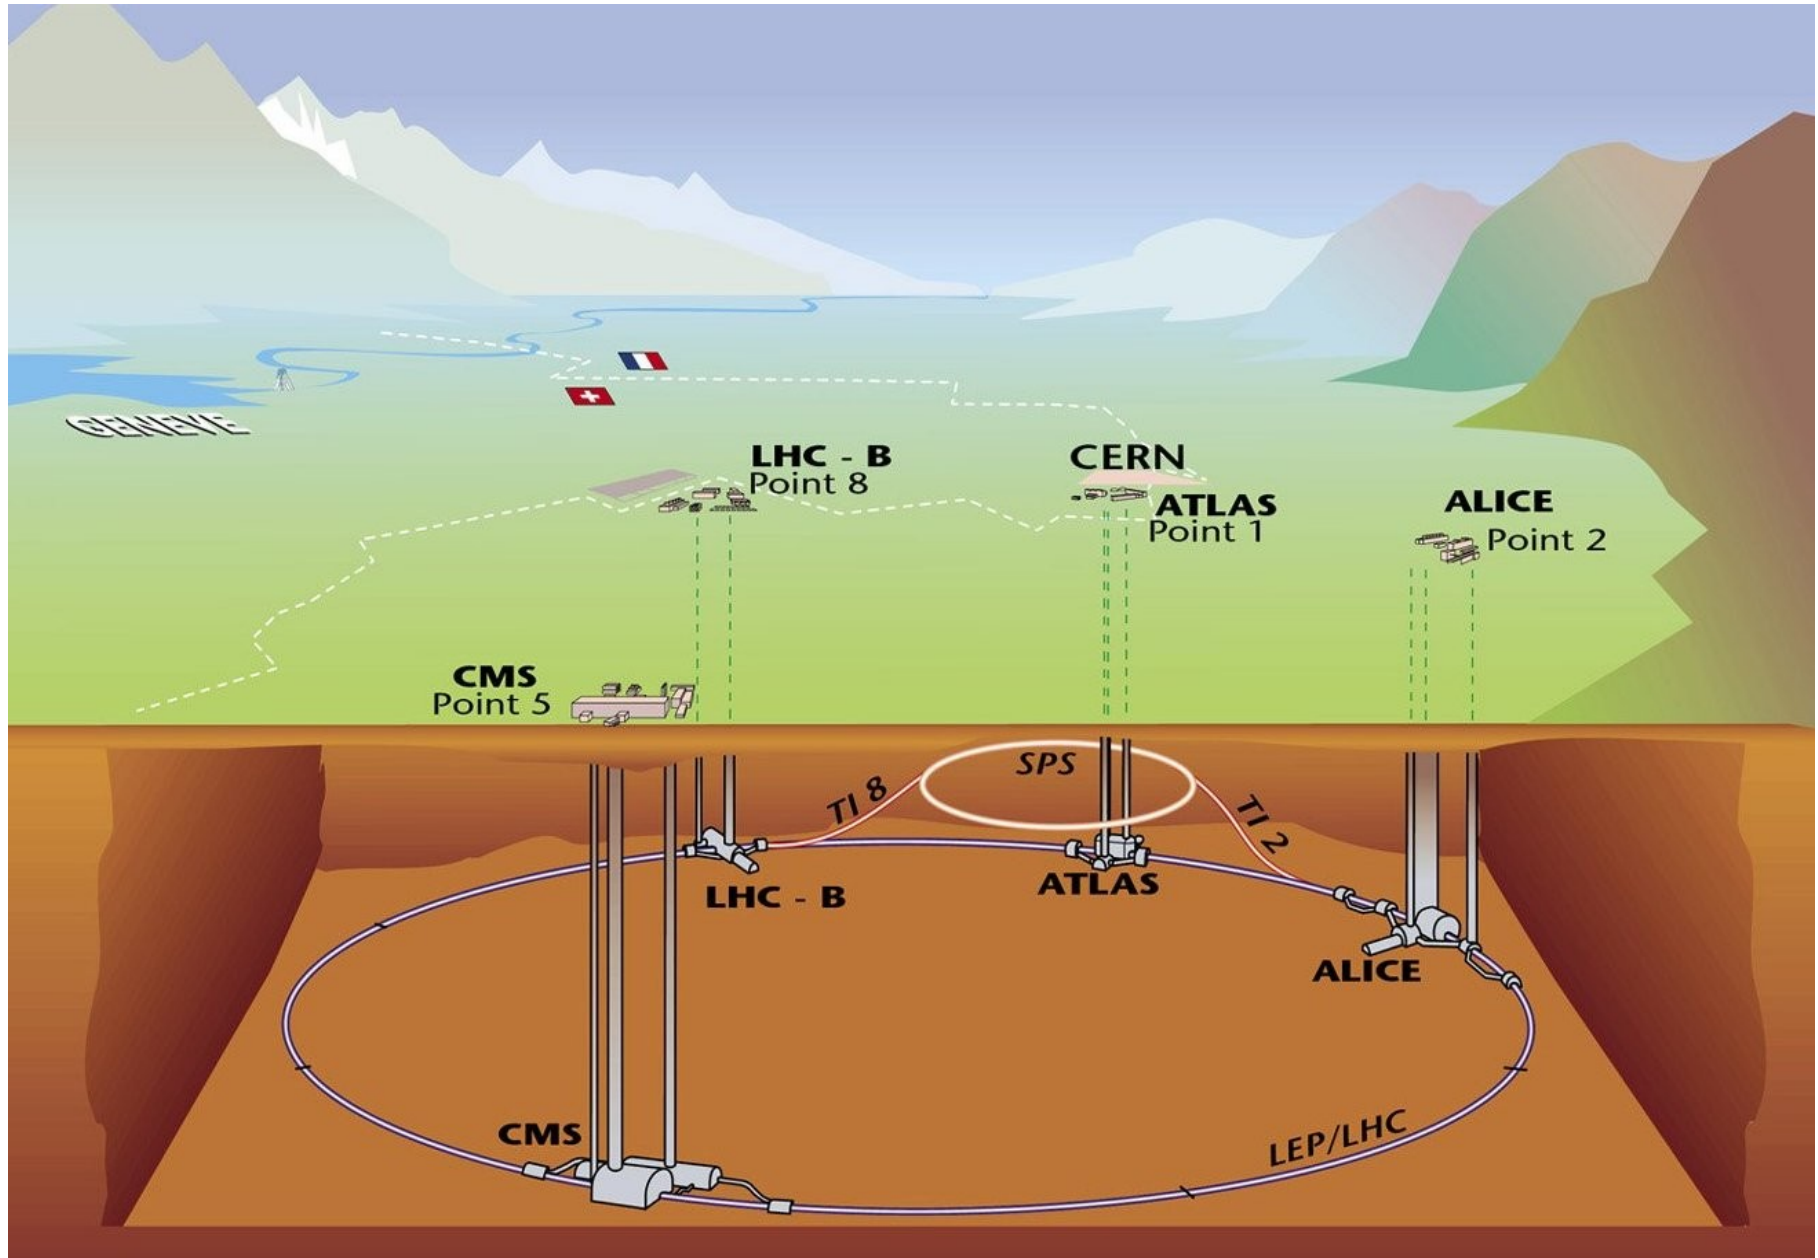
# Content

- **Network upgrade for the LCG**
- **Requirements for the CERN main firewall**
- **Hardware architecture**
- **The management framework**
- **Implementation experiences**
- **Conclusion**

# Content

- **Network upgrade for the LCG**
- Requirements for the CERN main firewall
- Hardware architecture
- The management framework
- Implementation experiences
- Conclusion

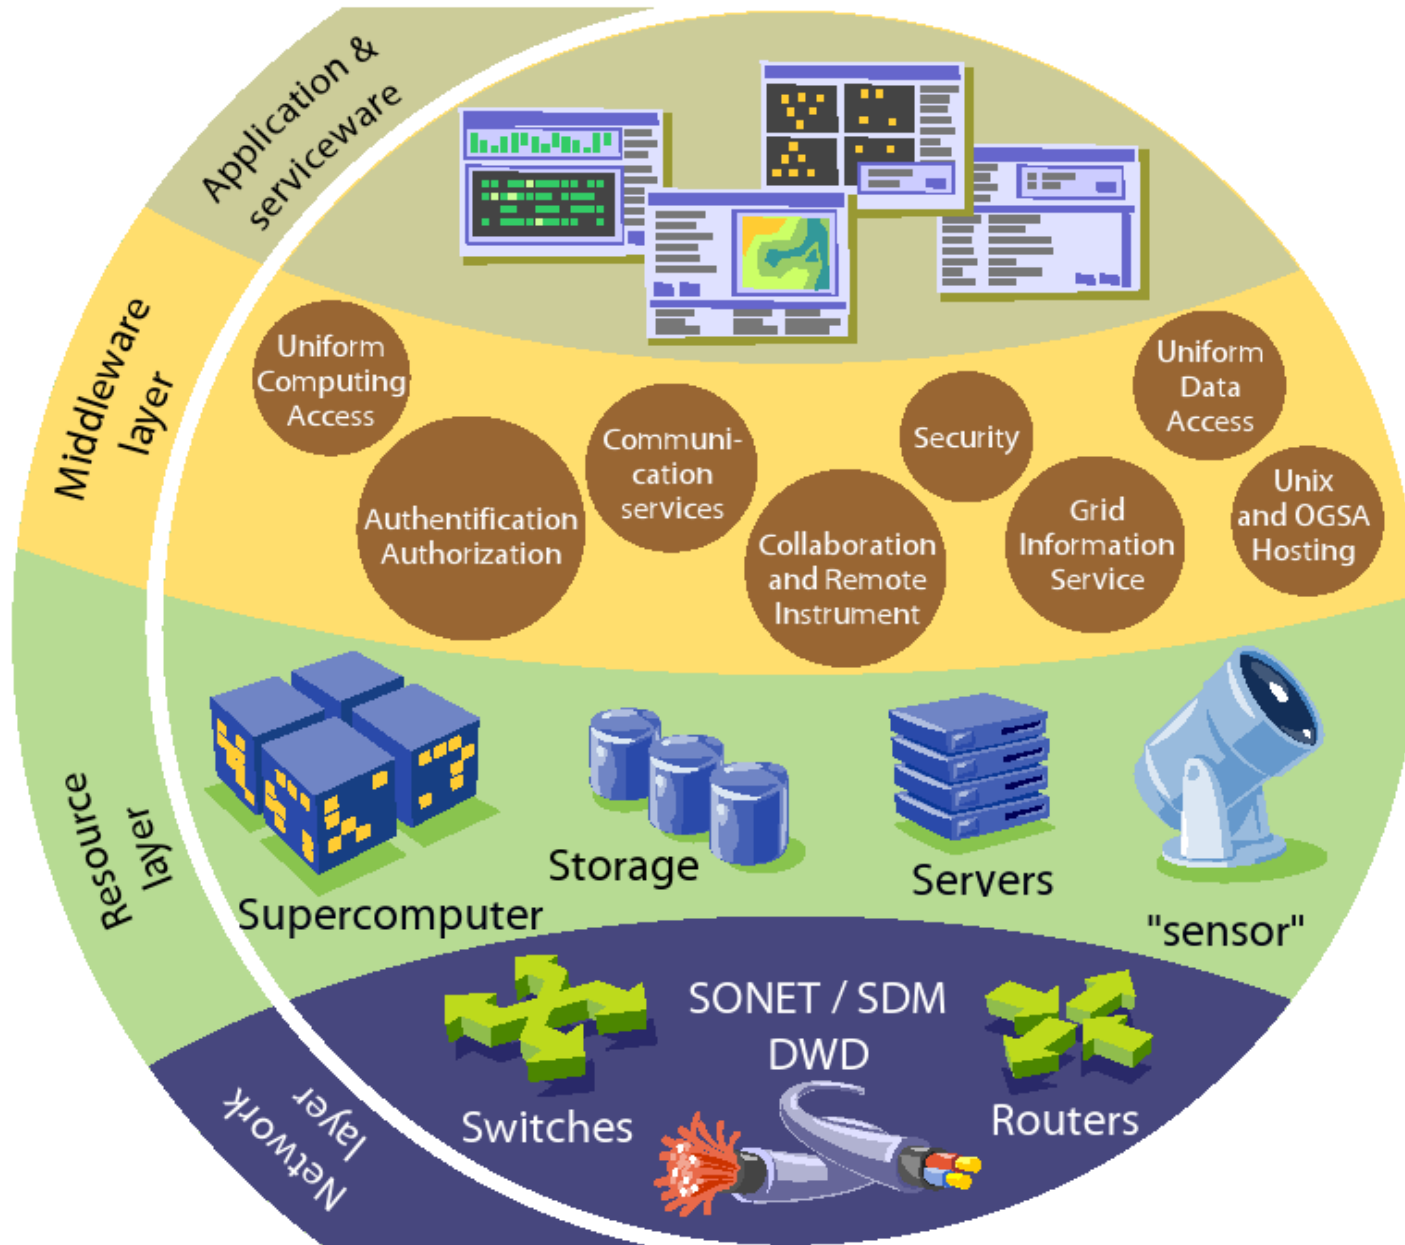
# LHC: the Large Hadron Collider



# LHC data

- **40 million collisions per second, of which only few hundreds will be kept**
- **3-4 MB of data for each collision**
- **The four LHC experiments will generate over 10 Peta bytes of data per year [Peta =  $10^{15}$ ]**
- **The data analysis will require more than 100,000 of today's best CPUs.**

# LCG: LHC Computing Grid



# LCG: the Tier model

## **Tier-0 – the accelerator centre**

- **CERN**
- Data acquisition and pre-processing
- Long-term data storage
- Distribution of data to Tier-1 centres

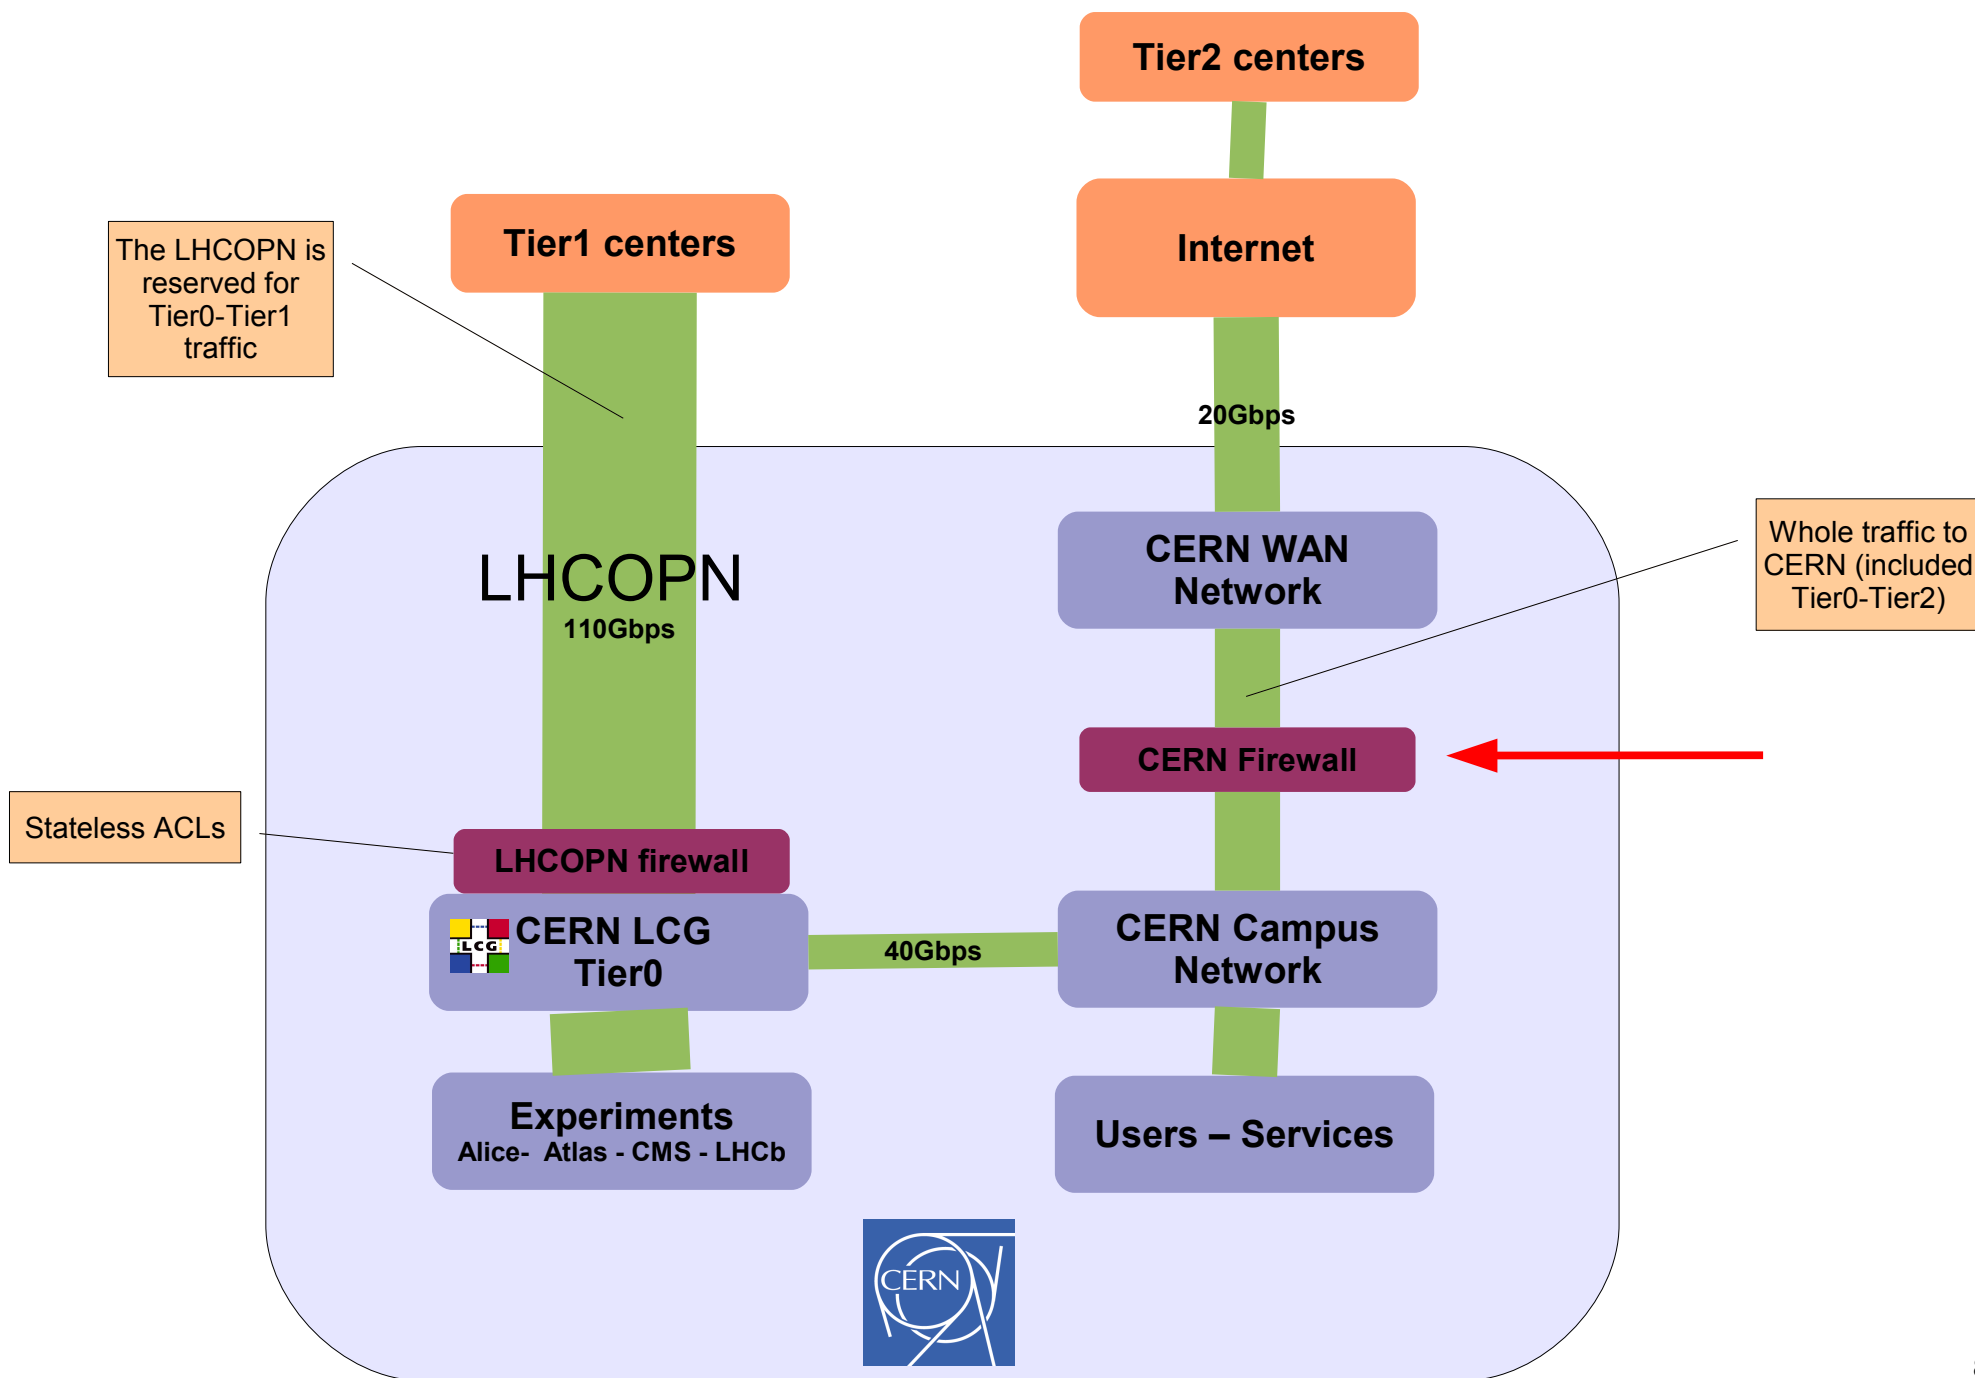
## **Tier-1 - data processing and distribution**

- **11 computer centres directly connected to CERN**
- Managed Mass Storage
- Grid-enabled data service
- Heavy data analysis

## **Tier-2**

- **more than one hundred centres around the world**
- Simulation
- End-user analysis
- batch and interactive

# The CERN network





# Content

- Network upgrade for the LCG
- Requirements for the CERN main firewall**
- Hardware architecture
- The management framework
- Implementation experiences
- Conclusion

# Challenges

***Allow easy sharing of data and services with any location on the Internet, but still guarantee protection and security***

***High throughput for the LCG and the hosted scientific community***

***Finite budget***

# Hardware requirements

## **Bandwidth, Reliability, Flexibility:**

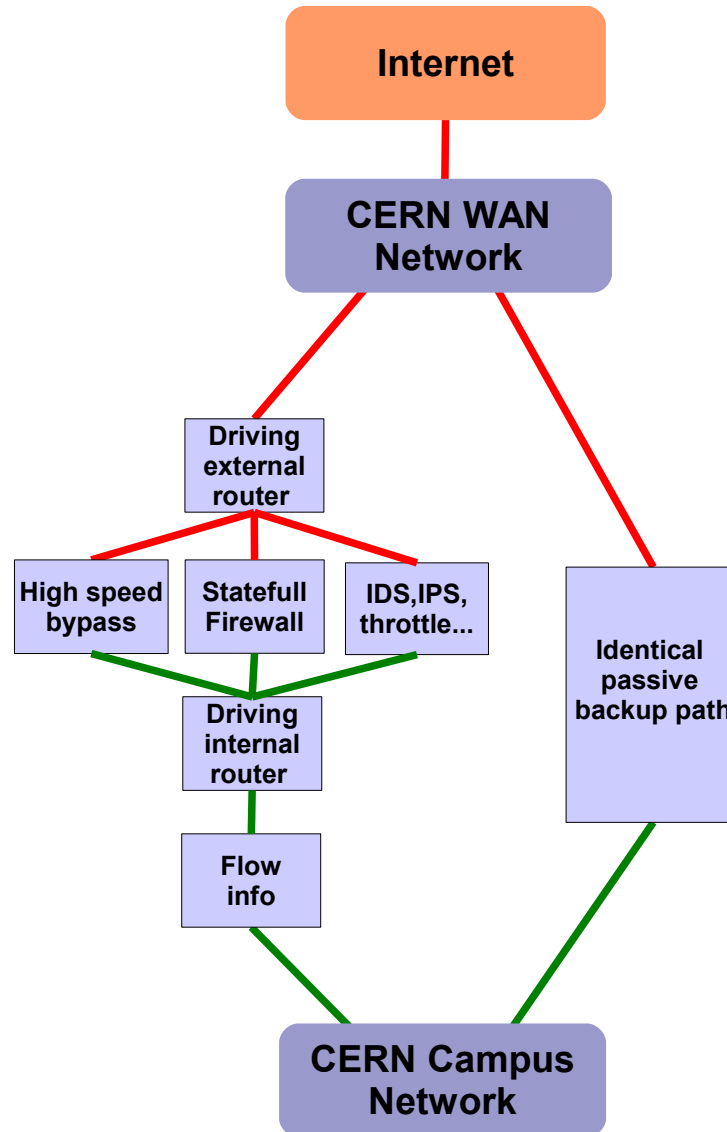
- at least 2Gbps of stateful inspection for generic traffic
- at least 20Gbps for high speed data transfer
- fully redundant system
- flexibility to add security features (IDS, tapping, throttling....)

# Security requirements

## **Fine granularity, Full monitoring:**

- possibility to filter down to the TCP/UDP port of every single host
- inspection of all the generic traffic
- full information about flows
- offload of well defined and trusted traffic

# Architecture characteristics



**Keys:**

- external link
- internal link

# Management framework requirements



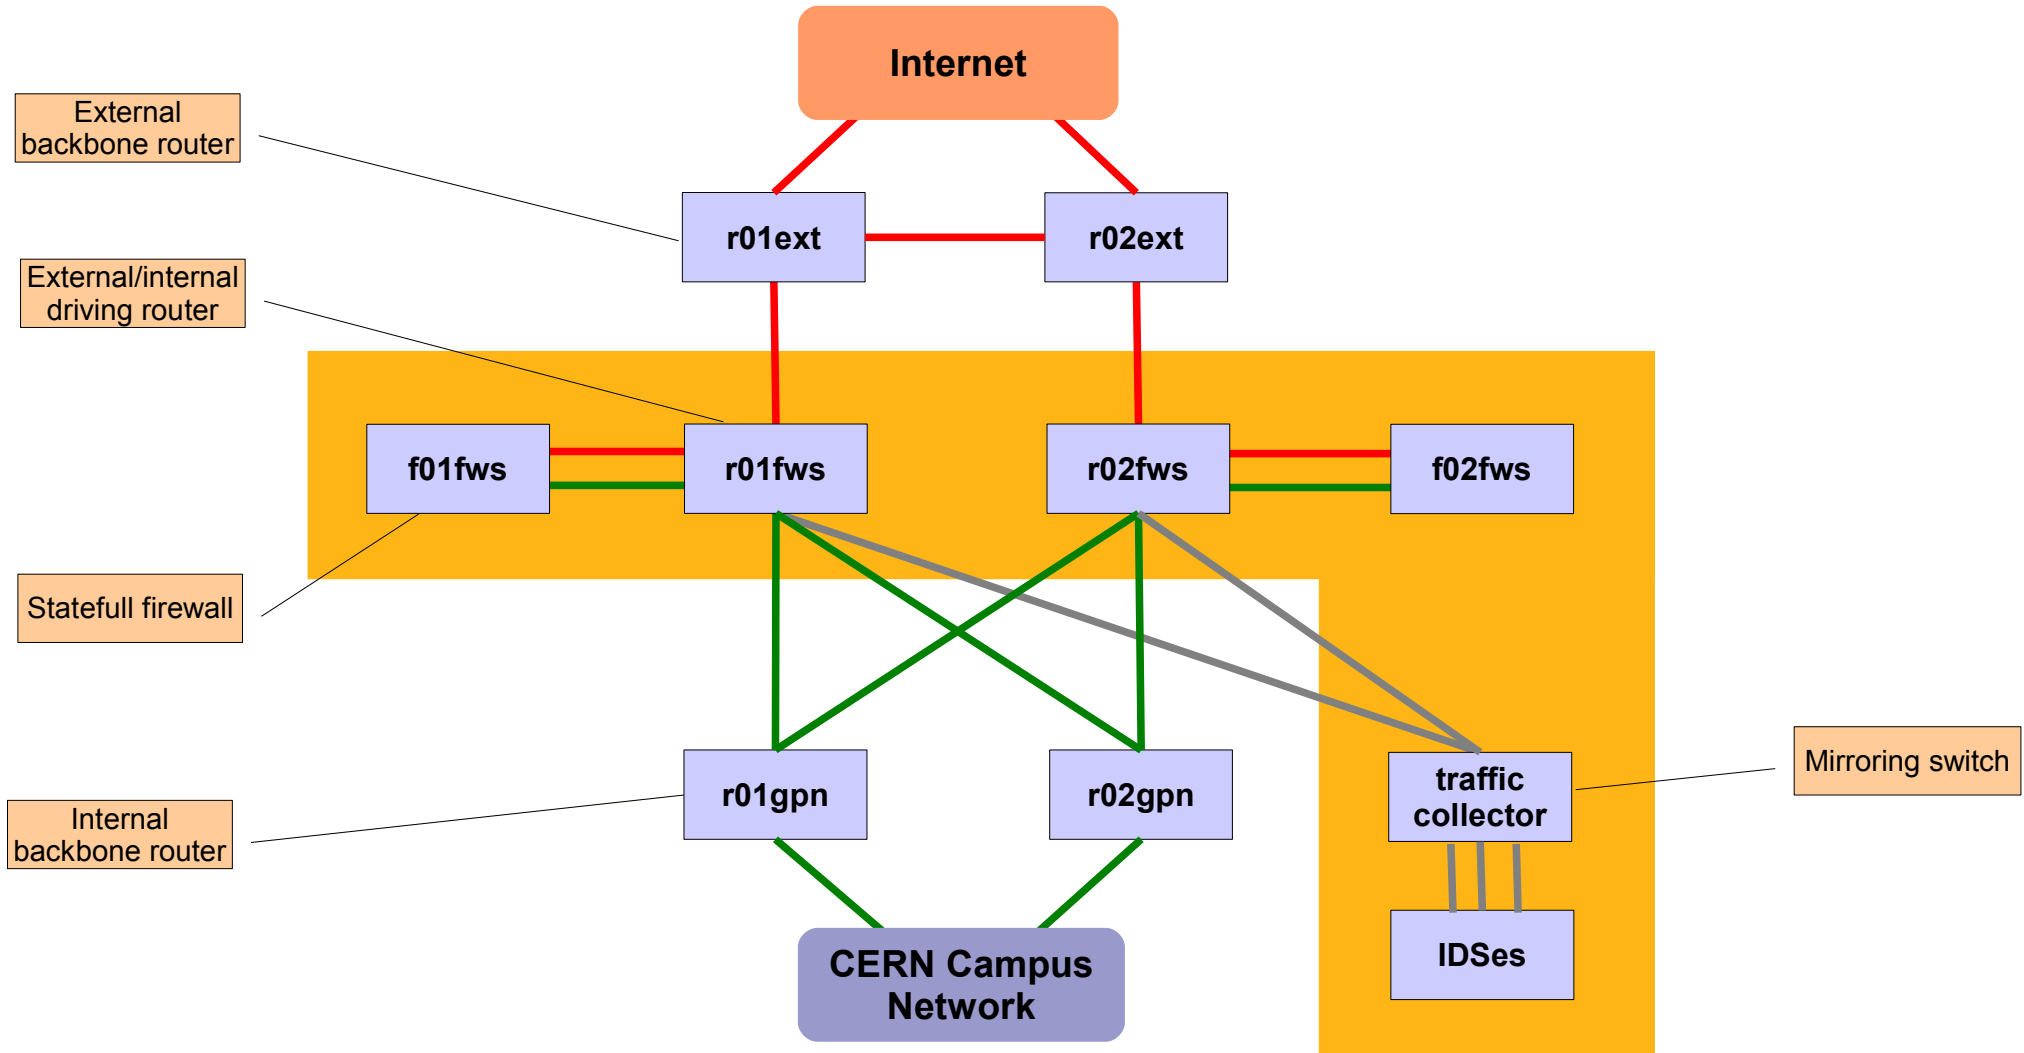
## **Integrated, Automated, Flexible:**

- fully integrated with the existing Network Management System and the Network Database
- hardware independent
- architecture independent
- automatic updates

# Content

- Network upgrade for the LCG
- Requirements for the CERN main firewall
- **Hardware architecture**
- The management framework
- Implementation experiences
- Conclusion

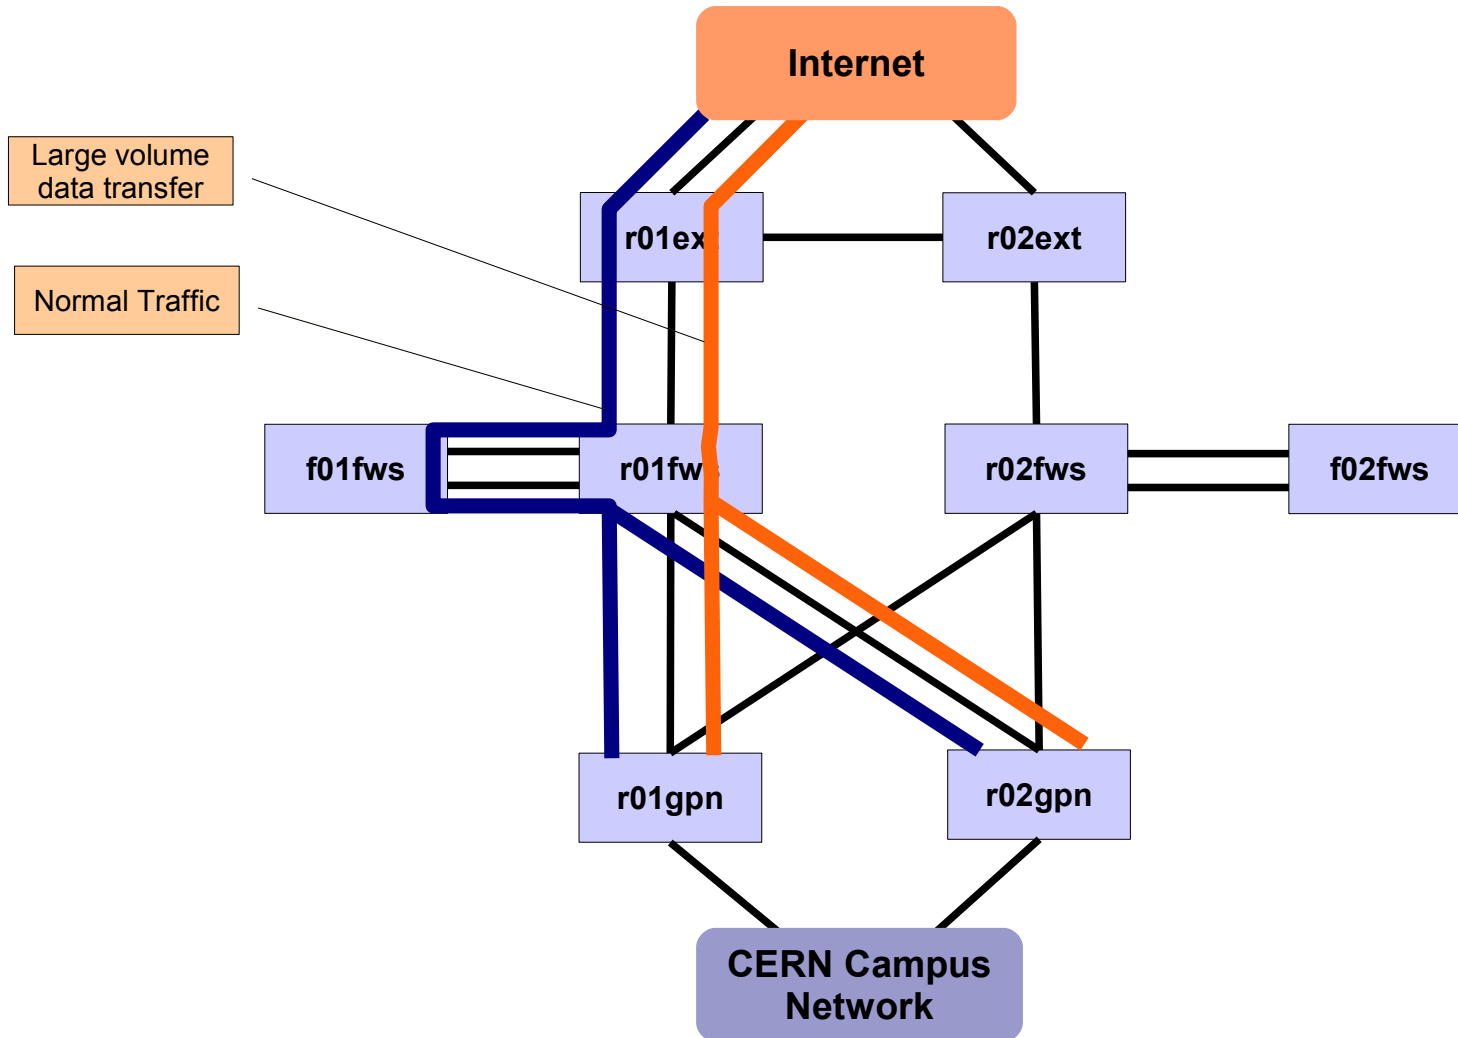
# Architecture



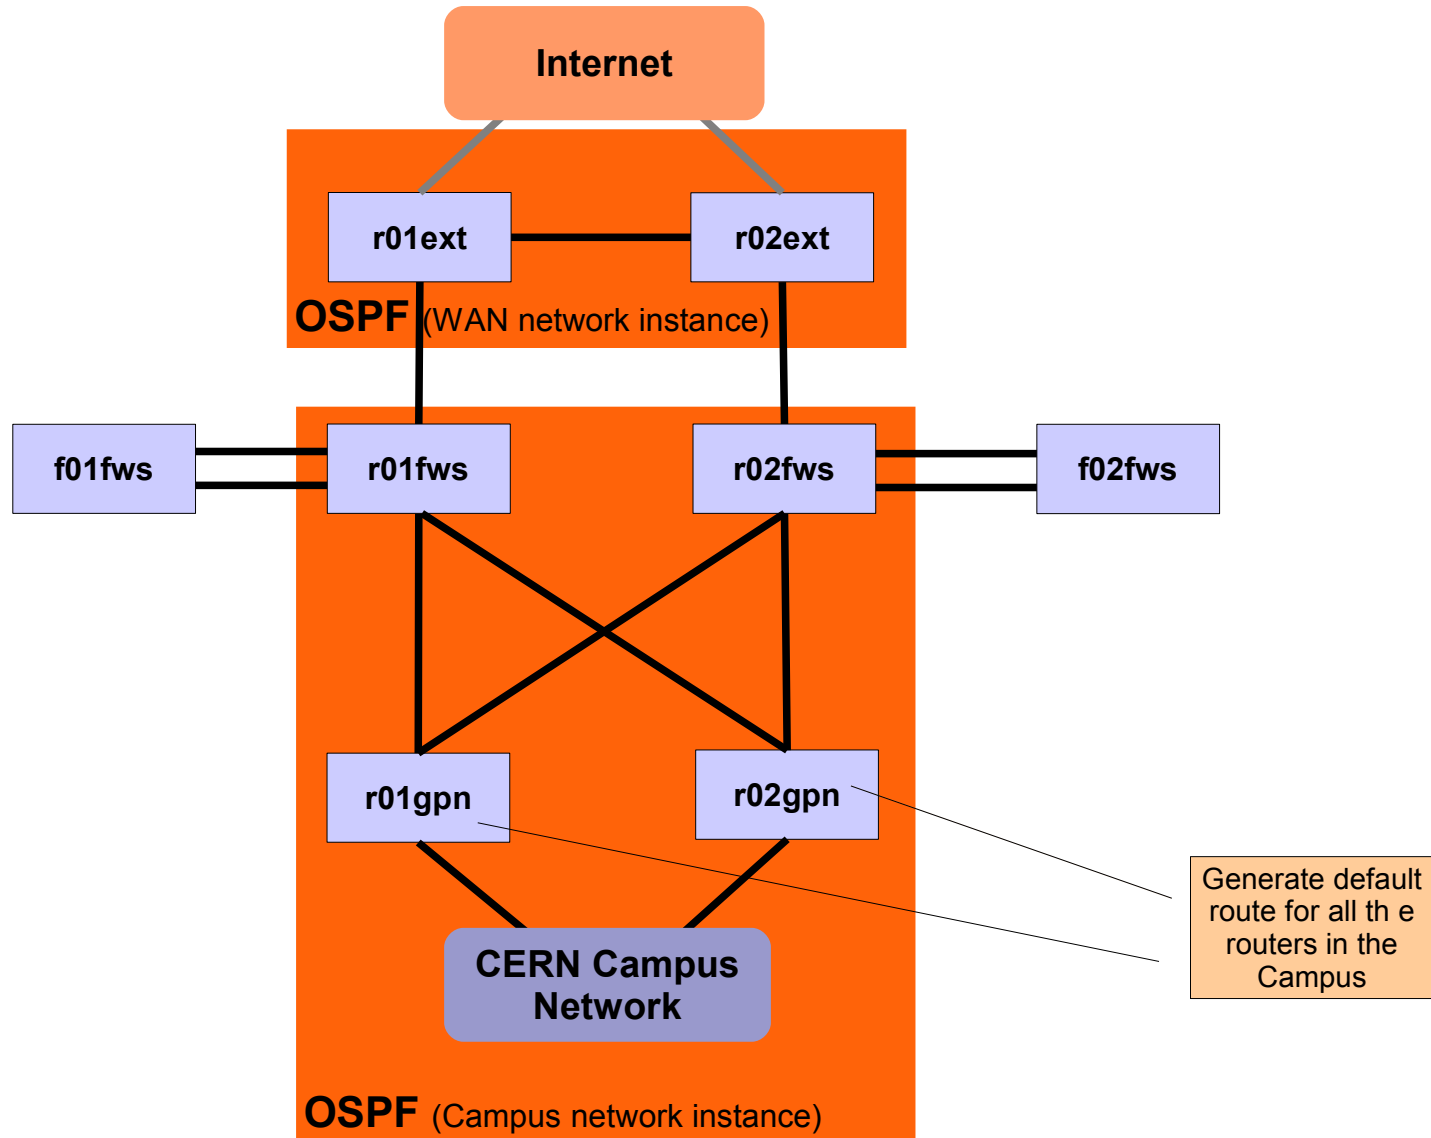
- Keys:**
- external link
  - internal link
  - mirrored traffic link



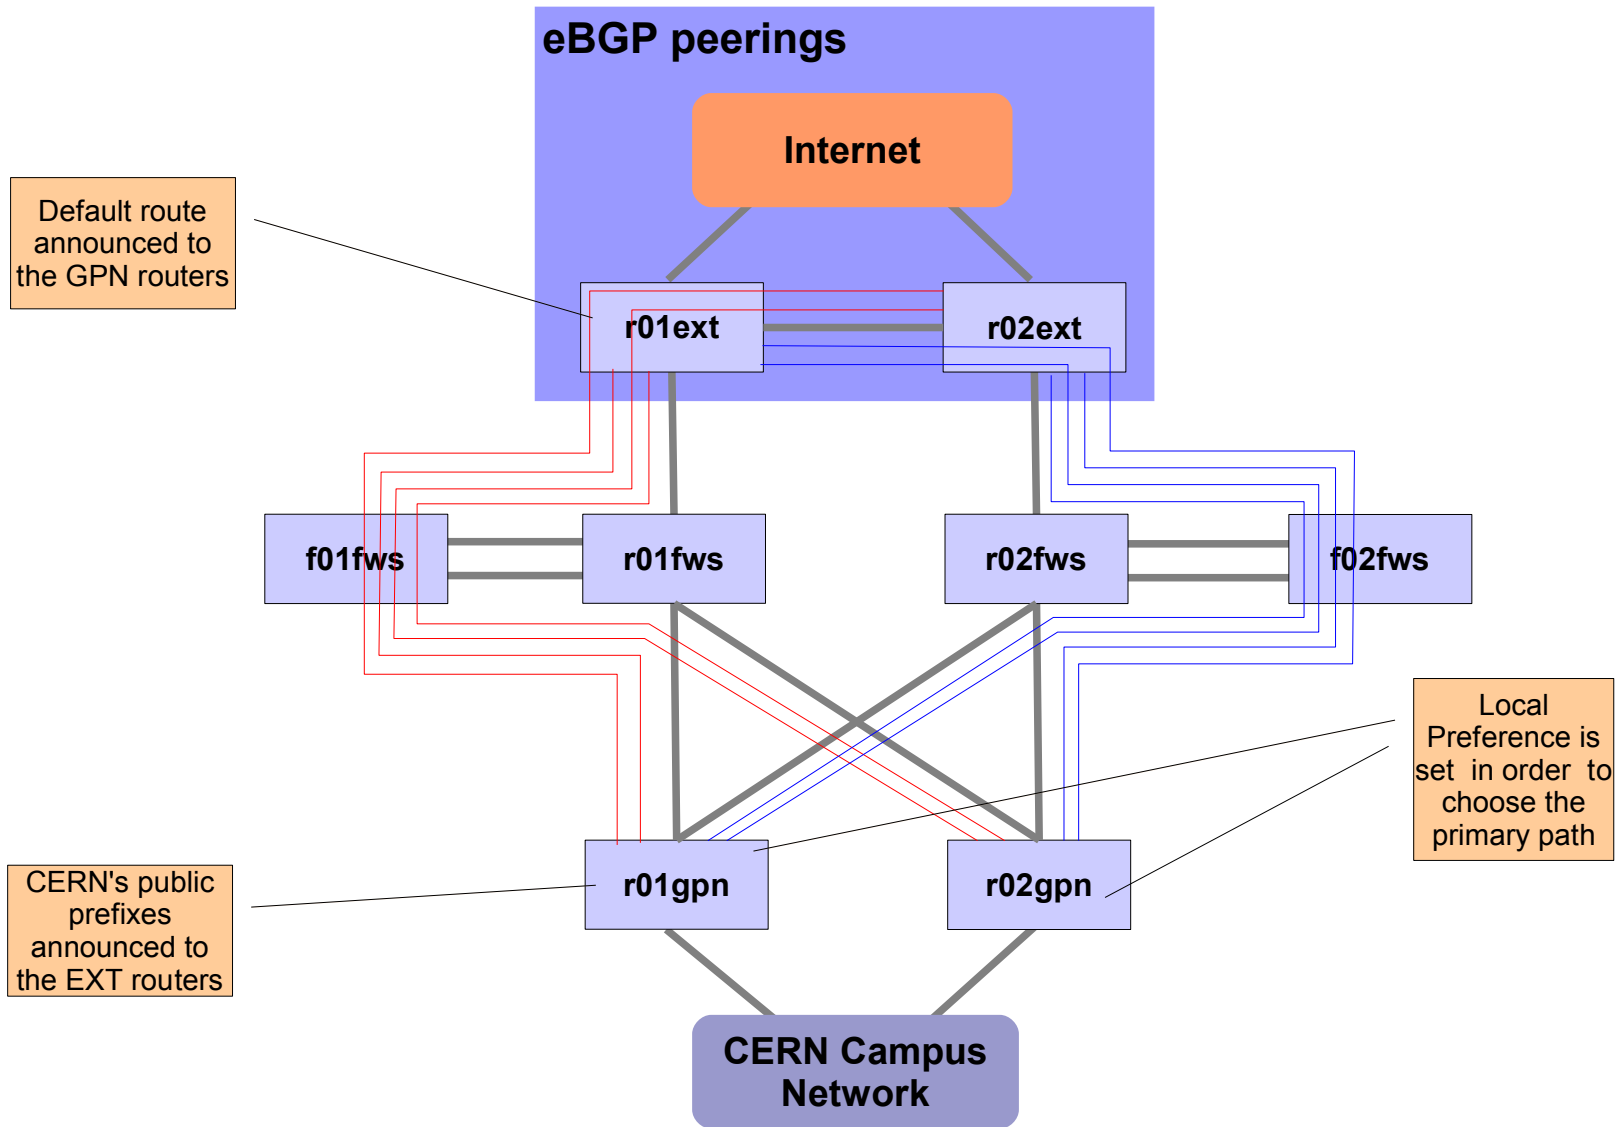
# Traffic flows



# Routing: OSPF



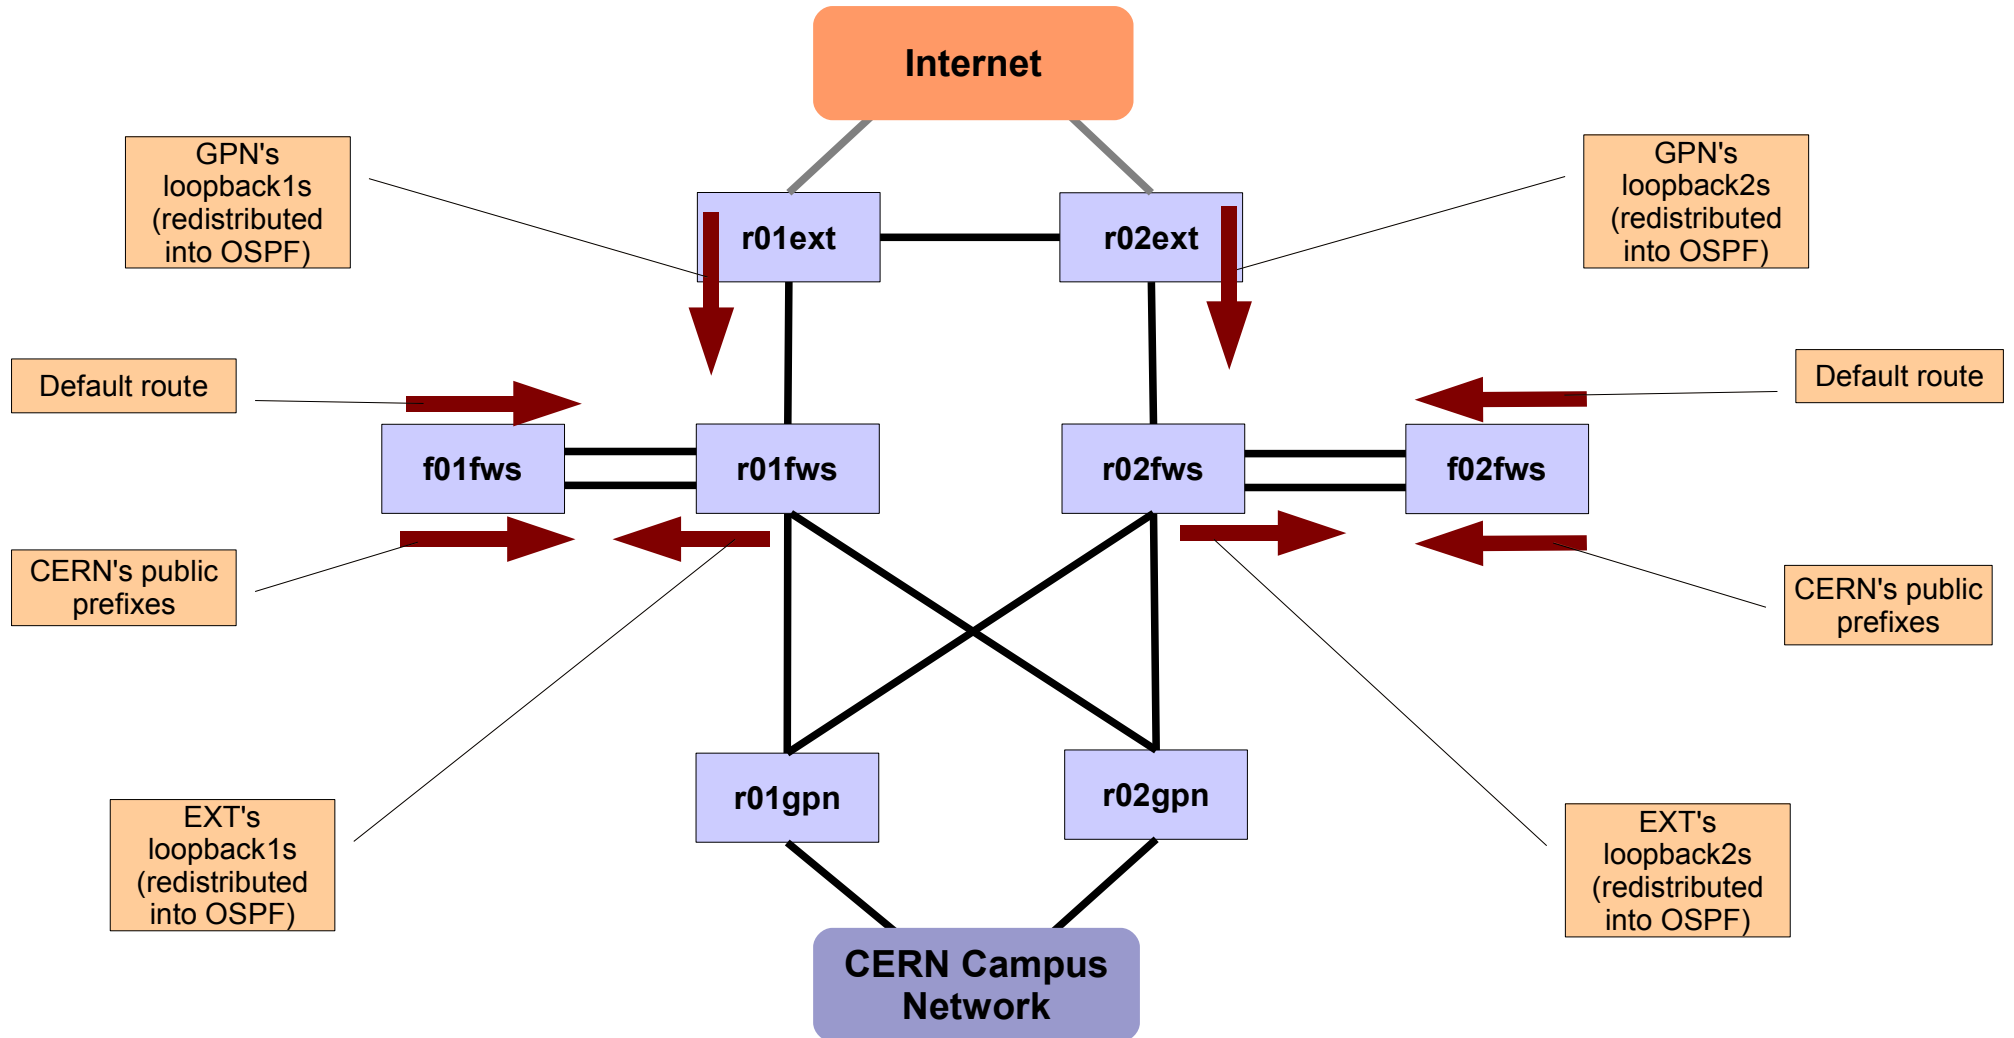
# Routing: BGP



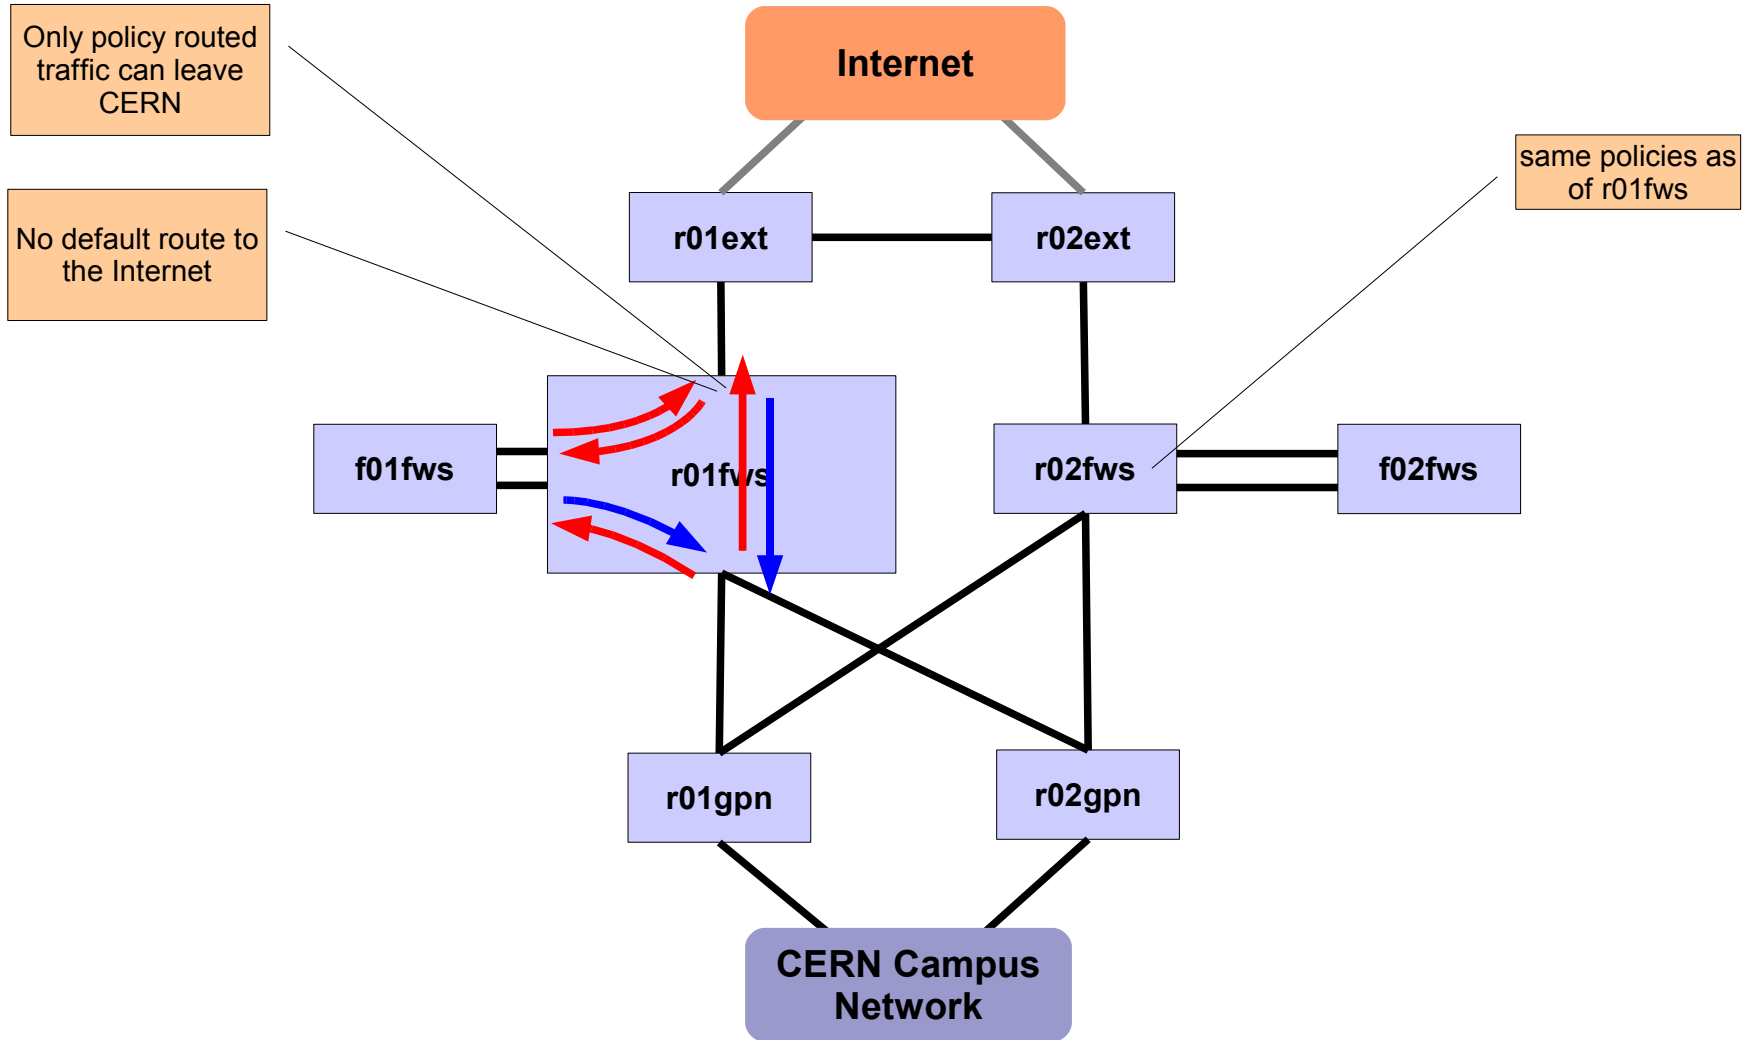
**Keys:**

- iBGP peerings between Loopback1s
- iBGP peerings between Loopback2s



# Routing: Static routes



# Routing: Policy Based



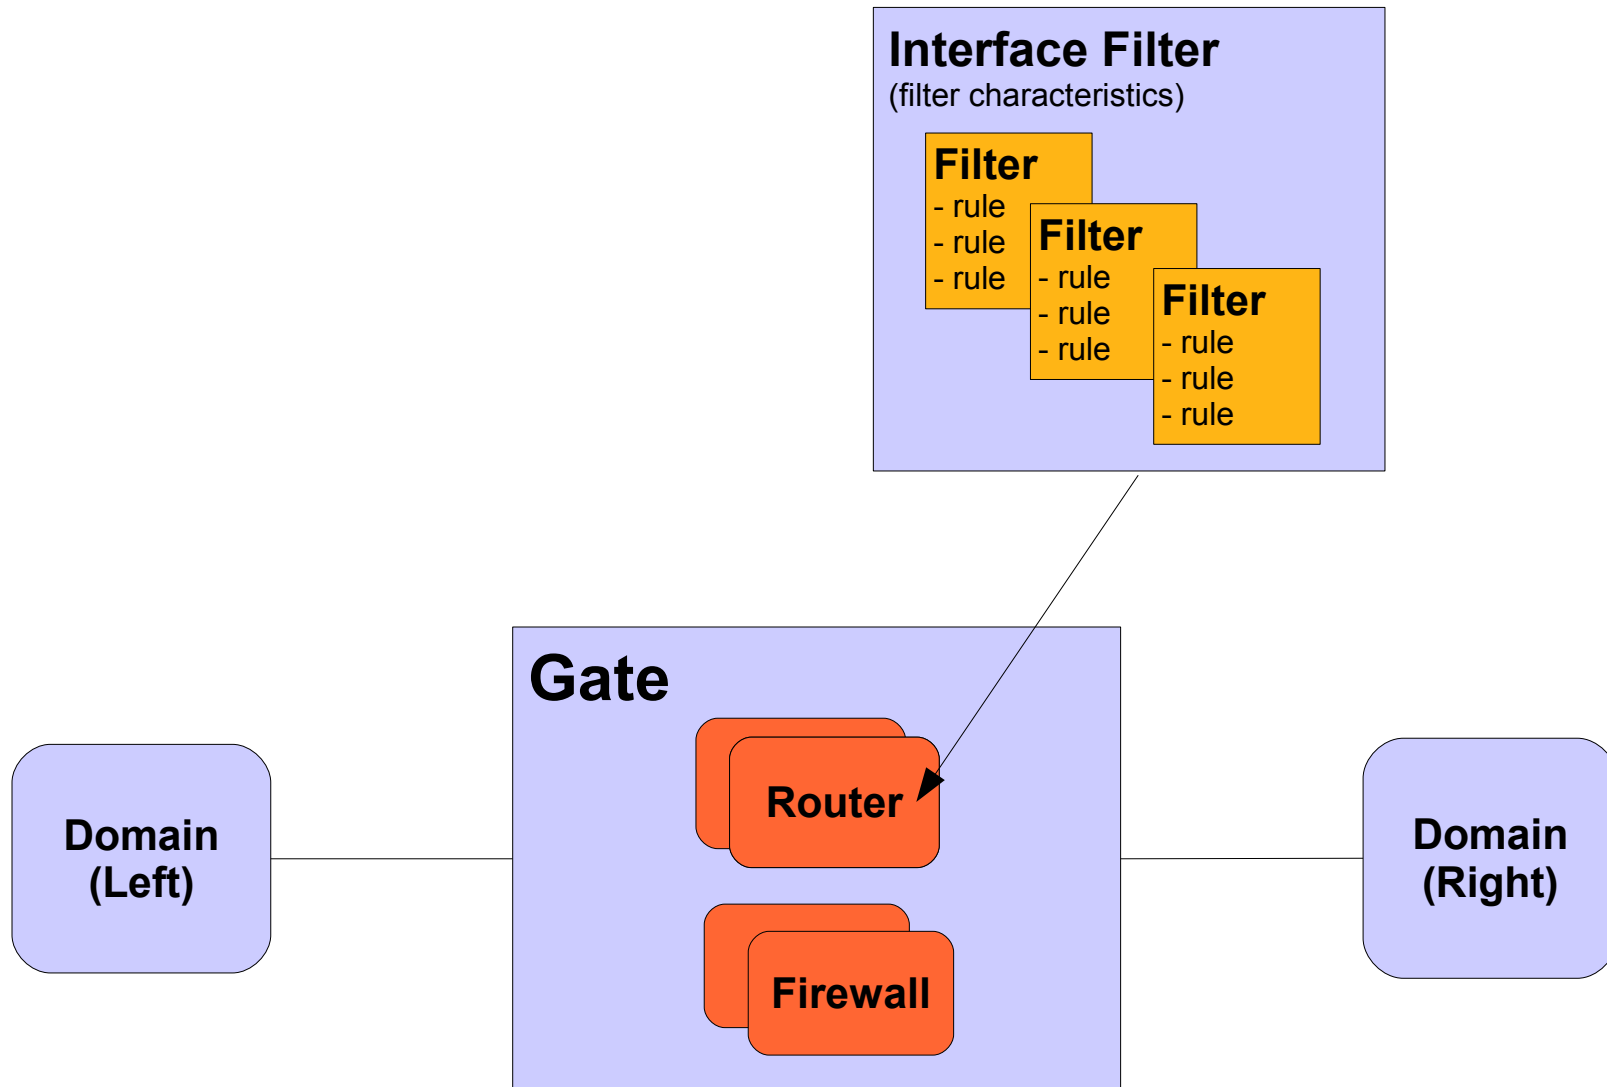
**Keys:**

-  Policy Routed traffic
-  Routed traffic

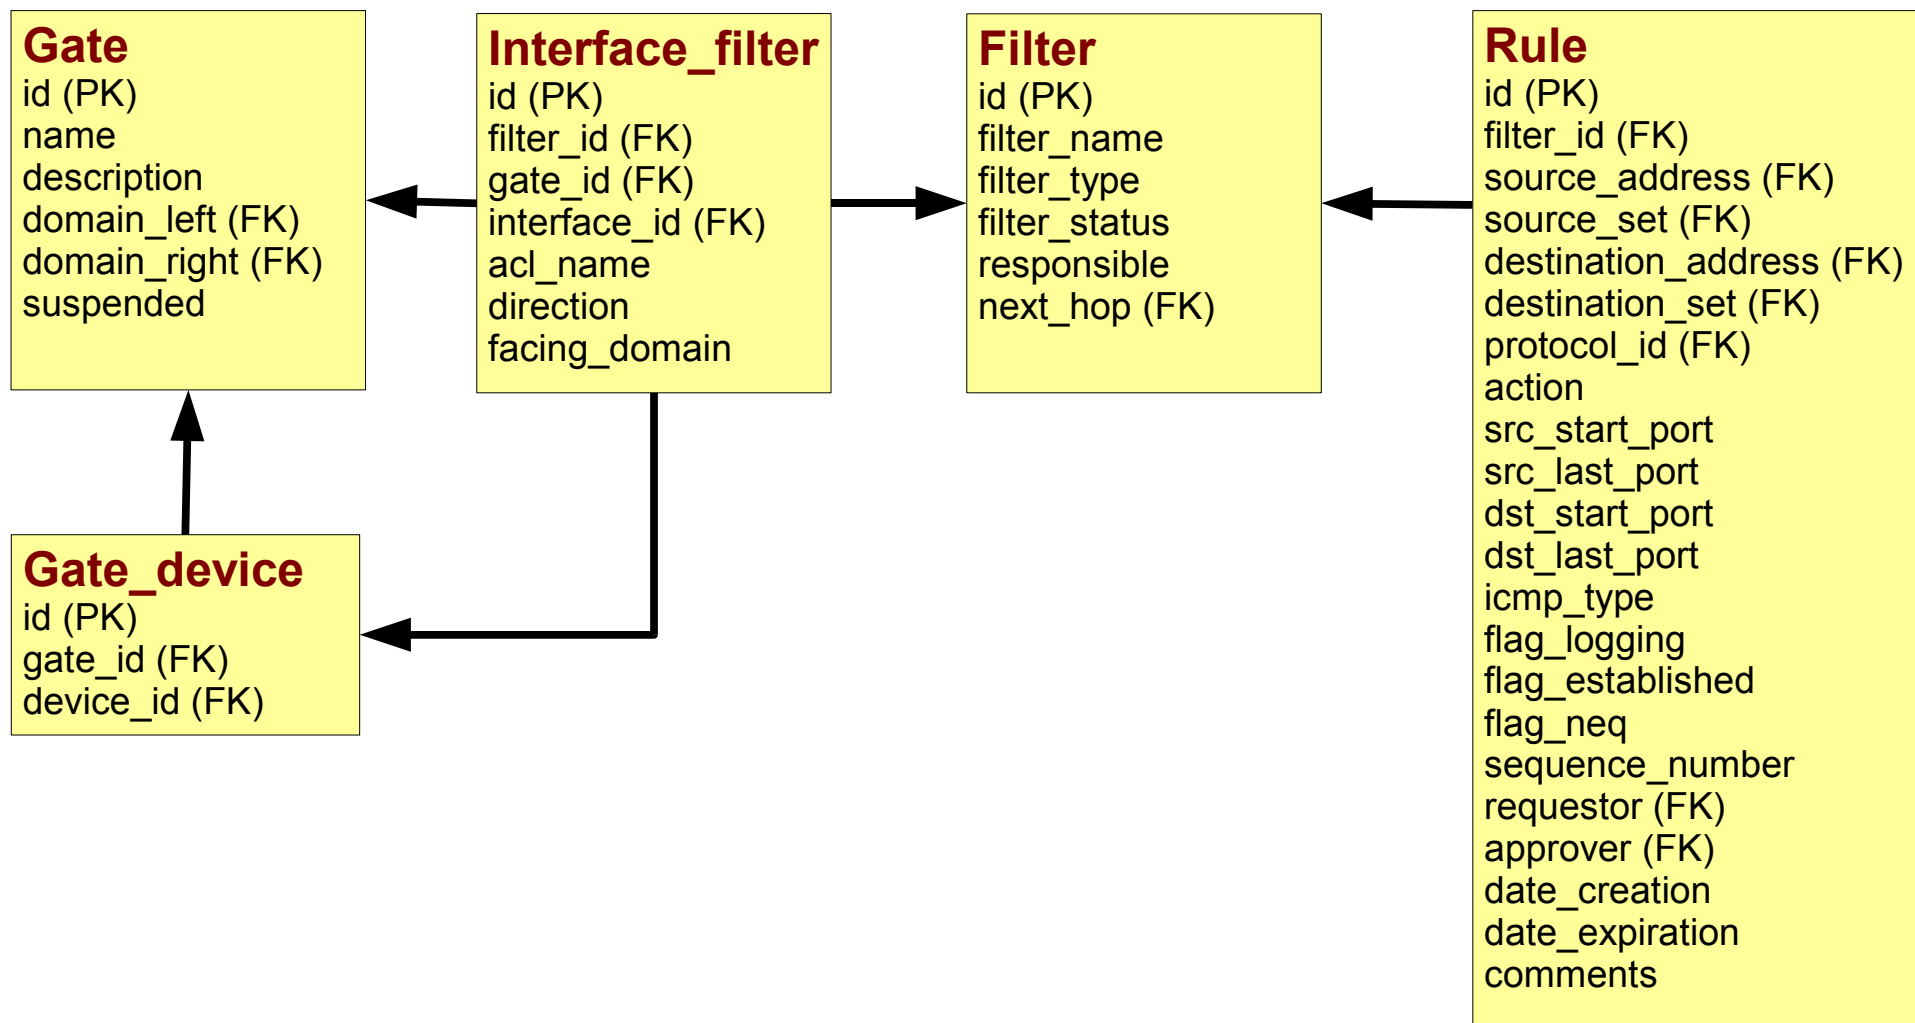
# Content

- Network upgrade for the LCG
- Requirements for the CERN main firewall
- Hardware architecture
- **The management framework**
- Implementation experiences
- Conclusion

# The Gate model

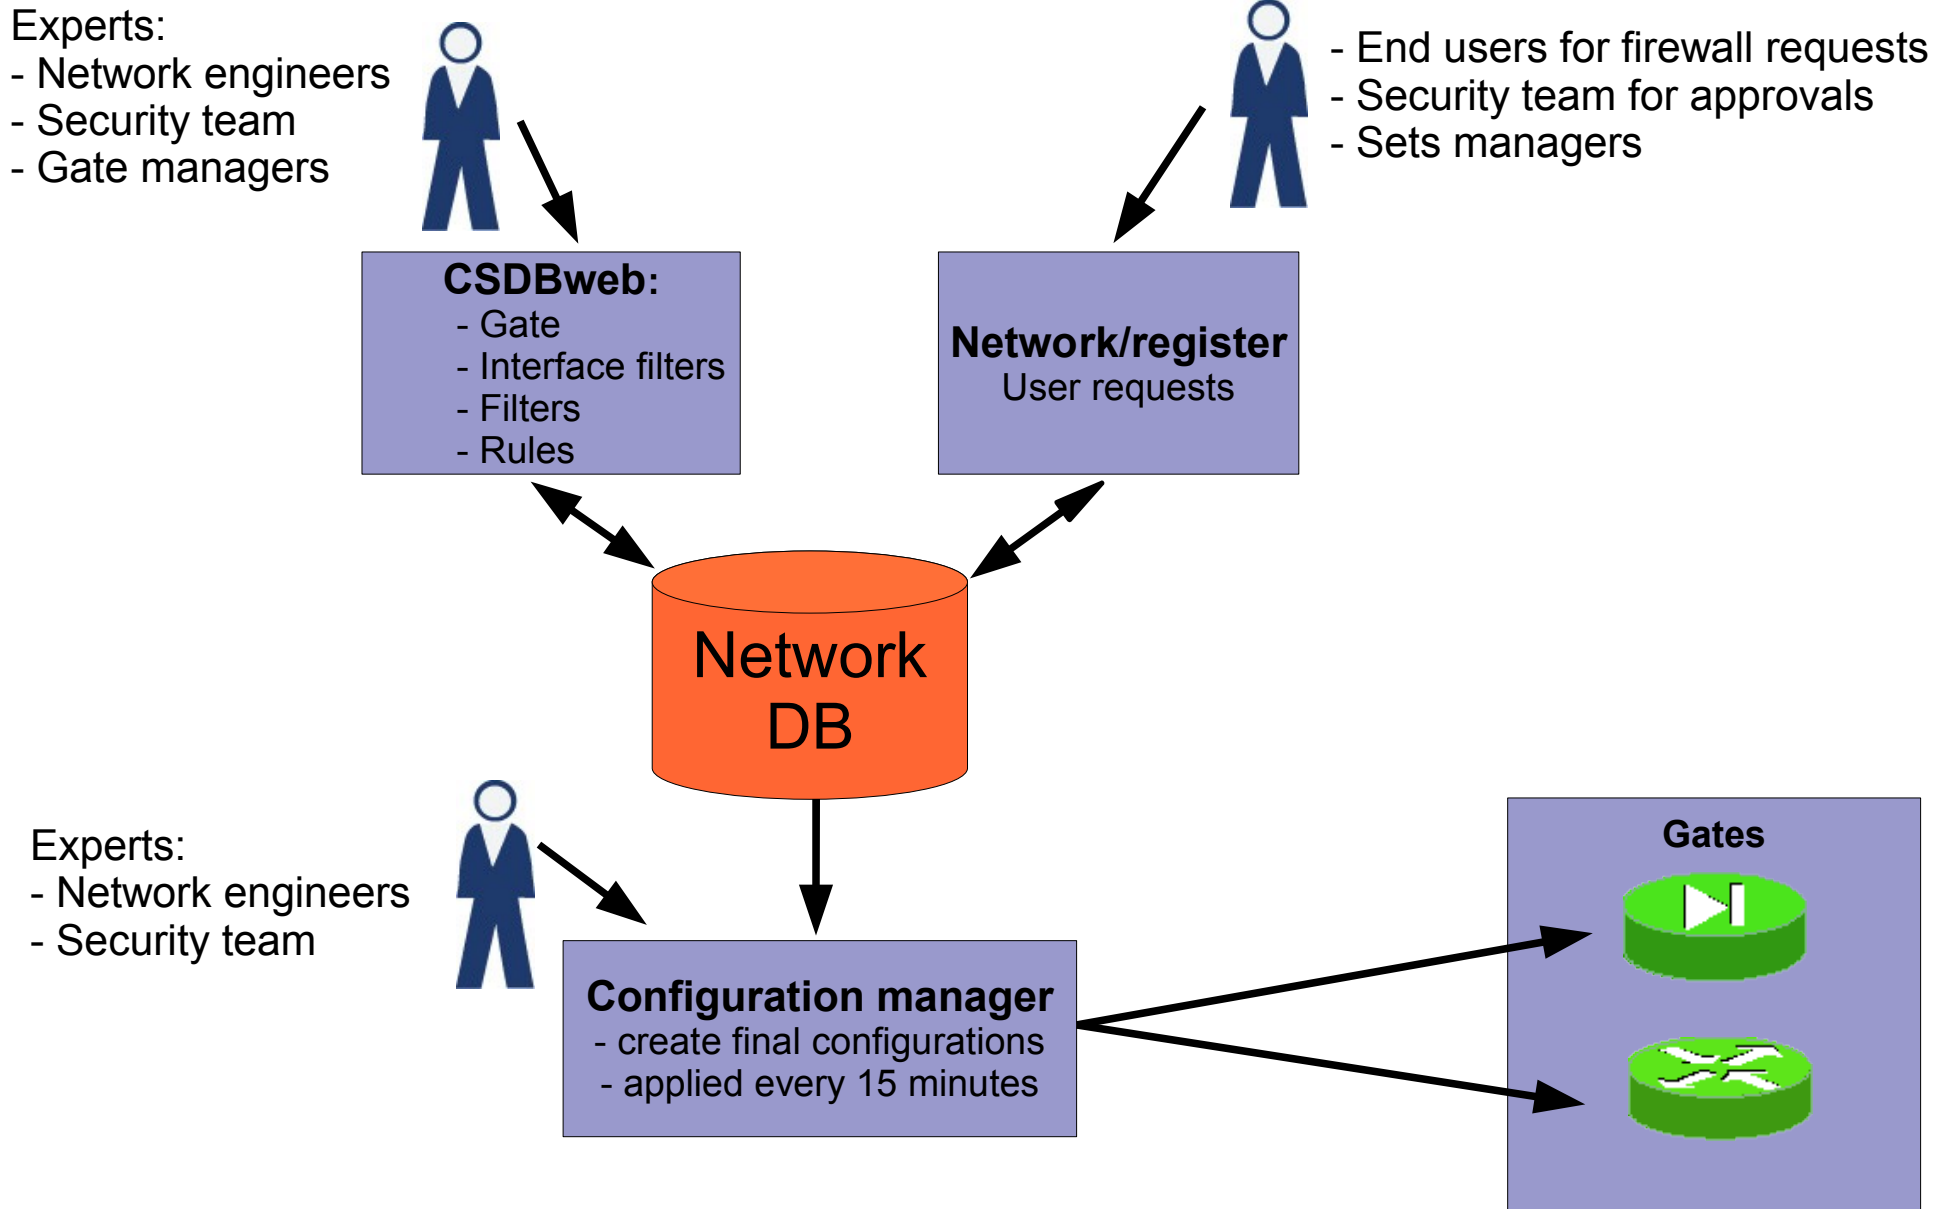


# The Database schema





# Gates framework's components



# CSDBweb Gates

**Web interface to define the Gate, its  
Interface filters, Filters and Rules.**

Developed in Java

For expert user

# Gate



**Gate**

**Gate Name:**

Domain Left:

Domain Right:

Record 1 of 2

**Devices:**

Sel	Name
<input type="checkbox"/>	<a href="#">GPN-E-FCIFM-1</a>
<input type="checkbox"/>	<a href="#">GPN-E-RCI</a>

**Interface Filters:**

Sel	Name	Type	Interface
<input type="checkbox"/>	<a href="#">IFW1_GPN_ACL_IN</a>	ACL	GPN-E-FCIFM-1-FE1
<input type="checkbox"/>	<a href="#">IFW1_GPN_ACL_OUT</a>	ACL	GPN-E-FCIFM-1-FE1
<input type="checkbox"/>	<a href="#">IFW1_INTERNET_ACL_IN</a>	ACL	GPN-E-FCIFM-1-FE1
<input type="checkbox"/>	<a href="#">IFW1_INTERNET_ACL_OUT</a>	ACL	GPN-E-FCIFM-1-FE1
<input type="checkbox"/>	<a href="#">IPR1_FW_GPN_ACL_IN</a>	ACL	GPN-E-RCI76-1-FE1
<input type="checkbox"/>	<a href="#">IPR1_FW_GPN_ACL_OUT</a>	ACL	GPN-E-RCI76-1-FE1
<input type="checkbox"/>	<a href="#">IPR1_FW_GPN_PBR_ANY</a>	PBR	GPN-E-RCI76-1-FE1
<input type="checkbox"/>	<a href="#">IPR1_FW_INTERNET_ACL_IN</a>	ACL	GPN-E-RCI76-1-FE1
<input type="checkbox"/>	<a href="#">IPR1_FW_INTERNET_ACL_OUT</a>	ACL	GPN-E-RCI76-1-FE1
<input type="checkbox"/>	<a href="#">IPR1_GPN1_ACL_IN</a>	ACL	GPN-E-RCI76-1-PG5
<input type="checkbox"/>	<a href="#">IPR1_GPN1_ACL_OUT</a>	ACL	GPN-E-RCI76-1-PG5
<input type="checkbox"/>	<a href="#">IPR1_GPN1_PBR_FW</a>	PBR	GPN-E-RCI76-1-PG5
<input type="checkbox"/>	<a href="#">IPR1_GPN1_PBR_HTAR</a>	PBR	GPN-E-RCI76-1-PG5
<input type="checkbox"/>	<a href="#">IPR1_GPN1_PBR_NULL</a>	PBR	GPN-E-RCI76-1-PG5
<input type="checkbox"/>	<a href="#">IPR1_GPN2_ACL_IN</a>	ACL	GPN-E-RCI76-1-PG6
<input type="checkbox"/>	<a href="#">IPR1_GPN2_ACL_OUT</a>	ACL	GPN-E-RCI76-1-PG6
<input type="checkbox"/>	<a href="#">IPR1_GPN2_PBR_FW</a>	PBR	GPN-E-RCI76-1-PG6
<input type="checkbox"/>	<a href="#">IPR1_GPN2_PBR_HTAR</a>	PBR	GPN-E-RCI76-1-PG6
<input type="checkbox"/>	<a href="#">IPR1_GPN2_PBR_NULL</a>	PBR	GPN-E-RCI76-1-PG6
<input type="checkbox"/>	<a href="#">IPR1_INTERNET_ACL_IN</a>	ACL	GPN-E-RCI76-1-PE1
<input type="checkbox"/>	<a href="#">IPR1_INTERNET_ACL_OUT</a>	ACL	GPN-E-RCI76-1-PE1
<input type="checkbox"/>	<a href="#">IPR1_INTERNET_PBR_FW</a>	PBR	GPN-E-RCI76-1-PE1
<input type="checkbox"/>	<a href="#">IPR1_INTERNET_PBR_NULL</a>	PBR	GPN-E-RCI76-1-PE1

Description:

# Interface Filter

## Interface Filter

I-face Filter Name:   Record 1 of 1

Gate:  Interface:

Direction:  Facing Domain:

Type:  ACL Name:

Status:  Default Action:

**Responsible:**

Description:

## Filters

Sel	Filter	Type	Status	Responsible
<input type="checkbox"/>	<a href="#">CNIC-TN</a>	NORMAL	ACTIVE	TECHNICAL-NETWORK ADMINISTRATOR

# Filter



**Filter**

Filter Name:   Record 1 of 1

Type:  Status:

Responsible:

Description:

**Rules**

Sel	SeqNo	Action	Protocol	Left Address	Ports	Right Address	Ports
<input type="checkbox"/>	<a href="#">20</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC LXPLUS]	eq 22
<input type="checkbox"/>	<a href="#">40</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR ALICE DISKSERVER]	eq 2811
<input type="checkbox"/>	<a href="#">45</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR ATLAS DISKSERVER]	eq 2811
<input type="checkbox"/>	<a href="#">50</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR CMS DISKSERVER]	eq 2811
<input type="checkbox"/>	<a href="#">55</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR ITDC DISKSERVER]	eq 2811
<input type="checkbox"/>	<a href="#">60</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR LHCB DISKSERVER]	eq 2811
<input type="checkbox"/>	<a href="#">65</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR PUBLIC DISKSERVER]	eq 2811
<input type="checkbox"/>	<a href="#">70</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR ALICE DISKSERVER]	range 20000-21000
<input type="checkbox"/>	<a href="#">75</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR ATLAS DISKSERVER]	range 20000-21000
<input type="checkbox"/>	<a href="#">80</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR CMS DISKSERVER]	range 20000-21000
<input type="checkbox"/>	<a href="#">85</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR ITDC DISKSERVER]	range 20000-21000
<input type="checkbox"/>	<a href="#">90</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR LHCB DISKSERVER]	range 20000-21000
<input type="checkbox"/>	<a href="#">95</a>	PERMIT	TCP	0.0.0.0/255.255.255.255 [Any]		.../... [IT CC CASTOR PUBLIC DISKSERVER]	range 20000-21000

# Rule



**Traffic Rule**

Filter Name:

Action:  Protocol:

Rule Seq. No:  App Type:

ICMP Type:  Code:  Connection established:  Logging On:

Rule Commented:

**Left:**

Type:  Name:

Address:  .  .  .  Mask:  .  .  .

Ports:  -  NEQ

**Right:**

Type:  Name:

Address:  .  .  .  Mask:  .  .  .

Ports:  NEQ

**Metric:**

Create Date:  Effective Start Date:

Expiration Date:  Reconfirmation Date:

Justification:

# <http://Network/Register>

**Web form to request firewall openings.**

Developed in Java

For end users

# Network/register for End-Users

End Users can request firewall openings for their devices:

## Network Connection Request Form v 9.0

<a href="#">Main Menu</a>	<a href="#">Update Information</a>	<a href="#">New Connection</a>	<a href="#">New Terminal Connection</a>	<a href="#">Move System</a>	<a href="#">Disconnect/Delete</a>
<a href="#">Display Information</a>	<a href="#">ServiceChange</a>	<a href="#">Register Portable</a>	<a href="#">New Portable Outlet</a>	<a href="#">Disconnect Portable Outlet</a>	<a href="#">Last Operation</a>

**emartell** logged in

[Logout](#)

**Visitor Requests**

[Procedure](#)

[Submit](#)

[Sign](#)

**Blocked Systems**

[By IP](#)

[By Hardware](#)

**Register**

[About](#)

[Problems?](#)

[SOAP access](#)

[MIKE](#)

[Set Mgmt](#)

[Admin Requests](#)

[FAQ](#)

[News Subscribe](#)

[HELP!!!](#)

**Topology**

[By Building](#)

[By StarPoint](#)

**Apropos...**

[Portables](#)

[DHCP](#)

### Network Connection Request Forms - Update Information

The following information about a device which is already connected to the CERN Computer network corresponds to what we have in our databases at the moment. Please modify this information if necessary.

However, for modifications to the [CERN Network Domain](#) or [Medium](#), please go back and select the appropriate option.

Mandatory fields are marked with (\*). Please do not forget to submit your request by selecting the **'Send Request'** button at the end of this page. [HELP](#) is available by selecting the links on this page.

For any questions or comments, please contact [NETOPS](#).

The fields have been filled with the information we have in our databases. Please change them as desired.

#### Update PCITCSEM

◆ <b>Device Name:</b>	PCITCSEM [Last Operation]	◆ <b>Rename To:</b>	<input type="text" value="PCITCSEM"/>
◆ <b>Location:</b>	0031 R-0024	( <b>Zone:</b> <input type="text"/> )	
◆ <b>Outlet:</b>	0024/04 (This plug connects a fan-out) (To change location, use "Move System")		
◆ <b>Manufacturer: (*)</b>	<input type="text" value="IBM"/>		
◆ <b>Model/Type: (*)</b>	<input type="text" value="THINKPAD T42"/>		
◆ <b>Generic Type:</b>	COMPUTER		
◆ <b>Operating System: (*)</b>	<input type="text" value="WINDOWS XP + LINUX"/>		
◆ <b>Op. Syst. version: (*)</b>	<input type="text" value="DEBIAN"/>		

#### Central Firewall Configuration (connections from outside CERN)

This device has a default firewall configuration which allows connections from CERN to the Internet. Requests for additional firewall access **are not** normally accepted. If you have a justified request endorsed by your Group Leader or Experiment technical Coordinator then please select the box below.

[Make Firewall Request >>>](#)

(Please note that firewall configurations can be overridden for devices which are members of a SET.)



# Network/register for End-Users

## Standards request might be approved more quickly

### Firewall Authorisation Requests - New Request

You are requesting direct Internet access for your machine in the main CERN Firewall.

Please be aware that machines directly exposed to the Internet will be continually attacked and create a risk for the rest of the site. To avoid this you should access your machine from off-site using an intermediate gateway system as described at <http://cern.ch/security/Internet>.

Note that direct off-site SSH access is not normally accepted. You should reach your system via LXPLUS which has additional intrusion checks.

If the methods described above cannot be used and you still wish to request direct Internet access then complete the following form:

Mandatory fields are marked with (\*). Please do not forget to submit your request by selecting the **'Send Request'** button at the end of this page. [HELP is available](#) by selecting the links on this page.

For any questions or comments related to this firewall request form please contact the CERN Computer Security Team · [computer.security@cern.ch](mailto:computer.security@cern.ch).

#### Request Information

◆ <b>Interface name: (*)</b>	PCITCSEM
◆ <b>Service: (*)</b>	Other
If Other, then please specify:	
- Port number:	
- Protocol:	
- Application:	
Hint: Give the name of the application	
◆ <b>Expiry date:</b>	
Use this date when firewall access is needed	and on this date.
◆ <b>Describe the professional requirement for which direct Internet access is required. (*)</b>	

Other

- SSH server on port 22/tcp
- Web server on port 80/tcp
- Secure web server on port 443/tcp
- Dual Web server (HTTP and HTTPS) on ports 80/tcp and 443/tcp
- Alternate Web on tcp/8080
- Alternate Secure Web on tcp/8443
- MySQL server on port 3306/tcp
- GridFTP on tcp/2811
- Globus TCP Port Range on tcp/20000-25000
- Globus UDP Port Range on udp/20000-25000
- Oracle TNSLSNR on tcp/1521
- CVS Server on tcp/2401
- Other

# Network/register for End-Users

Any request can be submitted:

◆ **Interface name: (\*)**

◆ **Service: (\*)**

If Other, then please specify:

- Port number:
- Protocol:
- Application:

Hint: Give the name of the application listening on the port.

◆ **Expiry date:**  /  /  (DD/MM/YYYY)

Use this date when firewall access is needed for a known duration. The firewall authorisation will be automatically removed on this date.

◆ **Describe the professional requirement for which direct Internet access is required. (\*)**

◆ **Explain why the recommended methods for accessing CERN from the Internet, as described at <http://cern.ch/security/Internet>, cannot be used.** If appropriate, explain why CERN's central services (e.g. [web](#), [J2EE](#), [CVS](#), ...) cannot be used. (\*)

◆ **Explain how you ensure that the device for which firewall access is requested will be kept pro-actively secured for security updates.** (This includes the operating system and all installed applications). (\*)

◆ Enter the name of the Group Leader or the Experiment Technical Coordinator who has endorsed this request. If this request is related to an experiment then please also enter the experiment name. (\*)

- ◆ **Name: (\*)**
- ◆ **Department:**
- ◆ **Experiment Name (if relevant):**
- ◆ **First Name: (\*)**
- ◆ **Group:**

# Approval procedure

1 – User fills the request form



2 - Request sent to Security Team

3 - Request denied - User notified

3 - Request approved – User notified



4 - Rule inserted in the Network Database



5 - Rule configured in the Gate



# Network/register for End-Users

Users can check their current firewall openings

## Interface(s) Information

[>>Network Service HELP<<](#) [>>Network Interface Card\(s\) HELP<<](#)

<b>Interface Name</b> PCITC...CERN.CH	<b>IP Address</b> 137.138.100.0	<b>Service Name</b> S31-S-AN5	<b>External TCP/IP Connectivity</b> OUTGOING
<b>Subnet Mask:</b> 255.255.0.0 <b>Default Gateway:</b> 137.130.100.1		<b>Name Servers:</b> 137.138.100.1, 137.138.100.2 <b>Time Servers:</b> 137.138.100.10, 137.130.100.10	
<b>IP Aliases:</b> NONE			
<b>Interface belongs to set(s):</b> TESTECS			
<b>Bound Interface Card(s):</b> NONE			
<b>Outlet</b> 0020/04	<b>CERN Network Domain</b> GPN	<b>Medium</b> FASTETHERNET	

## Central Firewall Configuration (connections from outside CERN)

Application	Port or Port Range	Protocol	Justification	Expiration date
GLOBUS-TCP	20000 - 25000	TCP	TEST	12-DEC-2009
SSH	22	TCP	SADC	
OTHER	4443	UDP	TESTDES	
HTTP	80	TCP	CASD	12-DEC-2010
GRIDFTP	2811	TCP	SADC	12-DEC-2030

(Please note that firewall configurations can be overridden for devices which are members of a SET.)

# Network/register for Set Managers

Set-managers can control their sets and have firewall openings automatically applied:

## Sets sub-menu

[Display Set](#) [New Set](#) [Update or Delete Set](#)

## Update Set

[>>Sets HELP<<](#)

Set Name: (\*)

Domain: (\*)

Responsible: (\*)

Department:  Group:

Description: (\*)

Project Url:

**INFO: You DO NOT have privileges to change set IT CC CASTOR ATLAS DISKSERVER**

## Set Type

This is an [Inter-Domain Set \(Domain filters\)](#)

## Contents(\*) [\(Click here for a list of all computers in set\)](#)

<ul style="list-style-type: none"><li>LXFS...01.CERN.CH</li><li>LXFS...02.CERN.CH</li><li>LXFS...03.CERN.CH</li><li>LXFS...04.CERN.CH</li><li>LXFS...05.CERN.CH</li><li>LXFS...06.CERN.CH</li><li>LXFS...07.CERN.CH</li></ul>	<p>Search for Device, Service</p> <p>Device Name: <input type="text"/></p> <p>IP-Service: <input type="text"/></p> <p>Inter-Domain Set: <input type="text"/></p>
---	--

### Traffic Rule

Filter Name:  [Reset](#) [Show Gates](#)

Action:  Protocol:

Rule Seq. No:  App Type:

ICMP Type:  Code:  Connection Established:  Logging On:

Rule Commented:

**Left:**

Type:  Name:  [Find](#)

Address:  Mask:

Ports:  -  NEQ

**Right:**

Type:  Name:  [Find](#)

Address:  Mask:

Ports:  -  NEQ

**Metric:**

**Requestor:**  **Approver:**

**Creator:**

Create Date:  Effective Start Date:

Expiration Date:  Reconfirmation Date:

Justification:

Comments:

Sets are also managed using SOAP interfaces.

# Configuration manager

**cfmgr-gates creates and applies the configuration**

Developed in Perl

For expert users

# cfmgr-gates

Every 15 minutes:

- extracts all the Gate information and rules from the Network Database
- builds the configurations for all the devices
- optimizes access lists
- checks consistency
- estimates hardware resource utilization
- if everything is OK, **automatically downloads every configuration that needs to be updated**

# Content

- Network upgrade for the LCG
- Requirements for the CERN main firewall
- Hardware architecture
- The management framework
- **Implementation experiences**
- Conclusion



# ACL's memory depletion

***Powerful software framework with almost no limitation in term of number of rules, but the policies have to be implemented in real hardware***

```
cfmgr gate check> aclresources
```

```
*** Routers ACL resource consumption ***
```

```
Router    Used resources
```

```
F01FWS    ACL entries 3831 (5%)
```

```
F02FWS    ACL entries 3831 (5%)
```

```
R01EXT    Masks 1008 (24%), entries 3775 (11%), LOUs 18 (14%)
```

```
R02EXT    Masks 1008 (24%), entries 3775 (11%), LOUs 18 (14%)
```

# ACL's memory depletion (2)

***Caution: depending on the hardware, some ACL operators can use a huge amount of ACLs memory:***

***30 lines like this:***

```
permit tcp src range 2000-20000 dst range 2000-10000
```

***can use all the available memory***

# PBR and CPU utilization

***Caution: some policy operators are executed in the Route Processor's CPU and not in the line cards' network processor (and such behavior is not always documented).***

# Other BGP configurations

Several possibilities:

## **eBGP:**

- public AS in the WAN network, private AS in the Campus.
- BGP confederation.

*pro: eBGP requires less peerings*

*cons: too many changes in the live WAN network*

## **iBGP:**

*pro: easily implemented in the live network*

*cons: many more peerings, the full-mesh must be preserved*

# Content

- Network upgrade for the LCG
- Requirements for the CERN main firewall
- Hardware architecture
- The management framework
- Implementation experiences
- **Conclusion**

# Conclusions

## **Stateful firewall**

- ready for increasing traffic load
- easily expandable when the market will be ready

## **Overall bandwidth**

- matches the CERN Internet connectivity
- can scale beyond 80Gbps

## **Management framework**

- fully use of the Network database
- manages any gate

## **Automatic Gate updates**

- in case of any change in the database

***A complex system that makes everybody's life simpler***

# Questions?