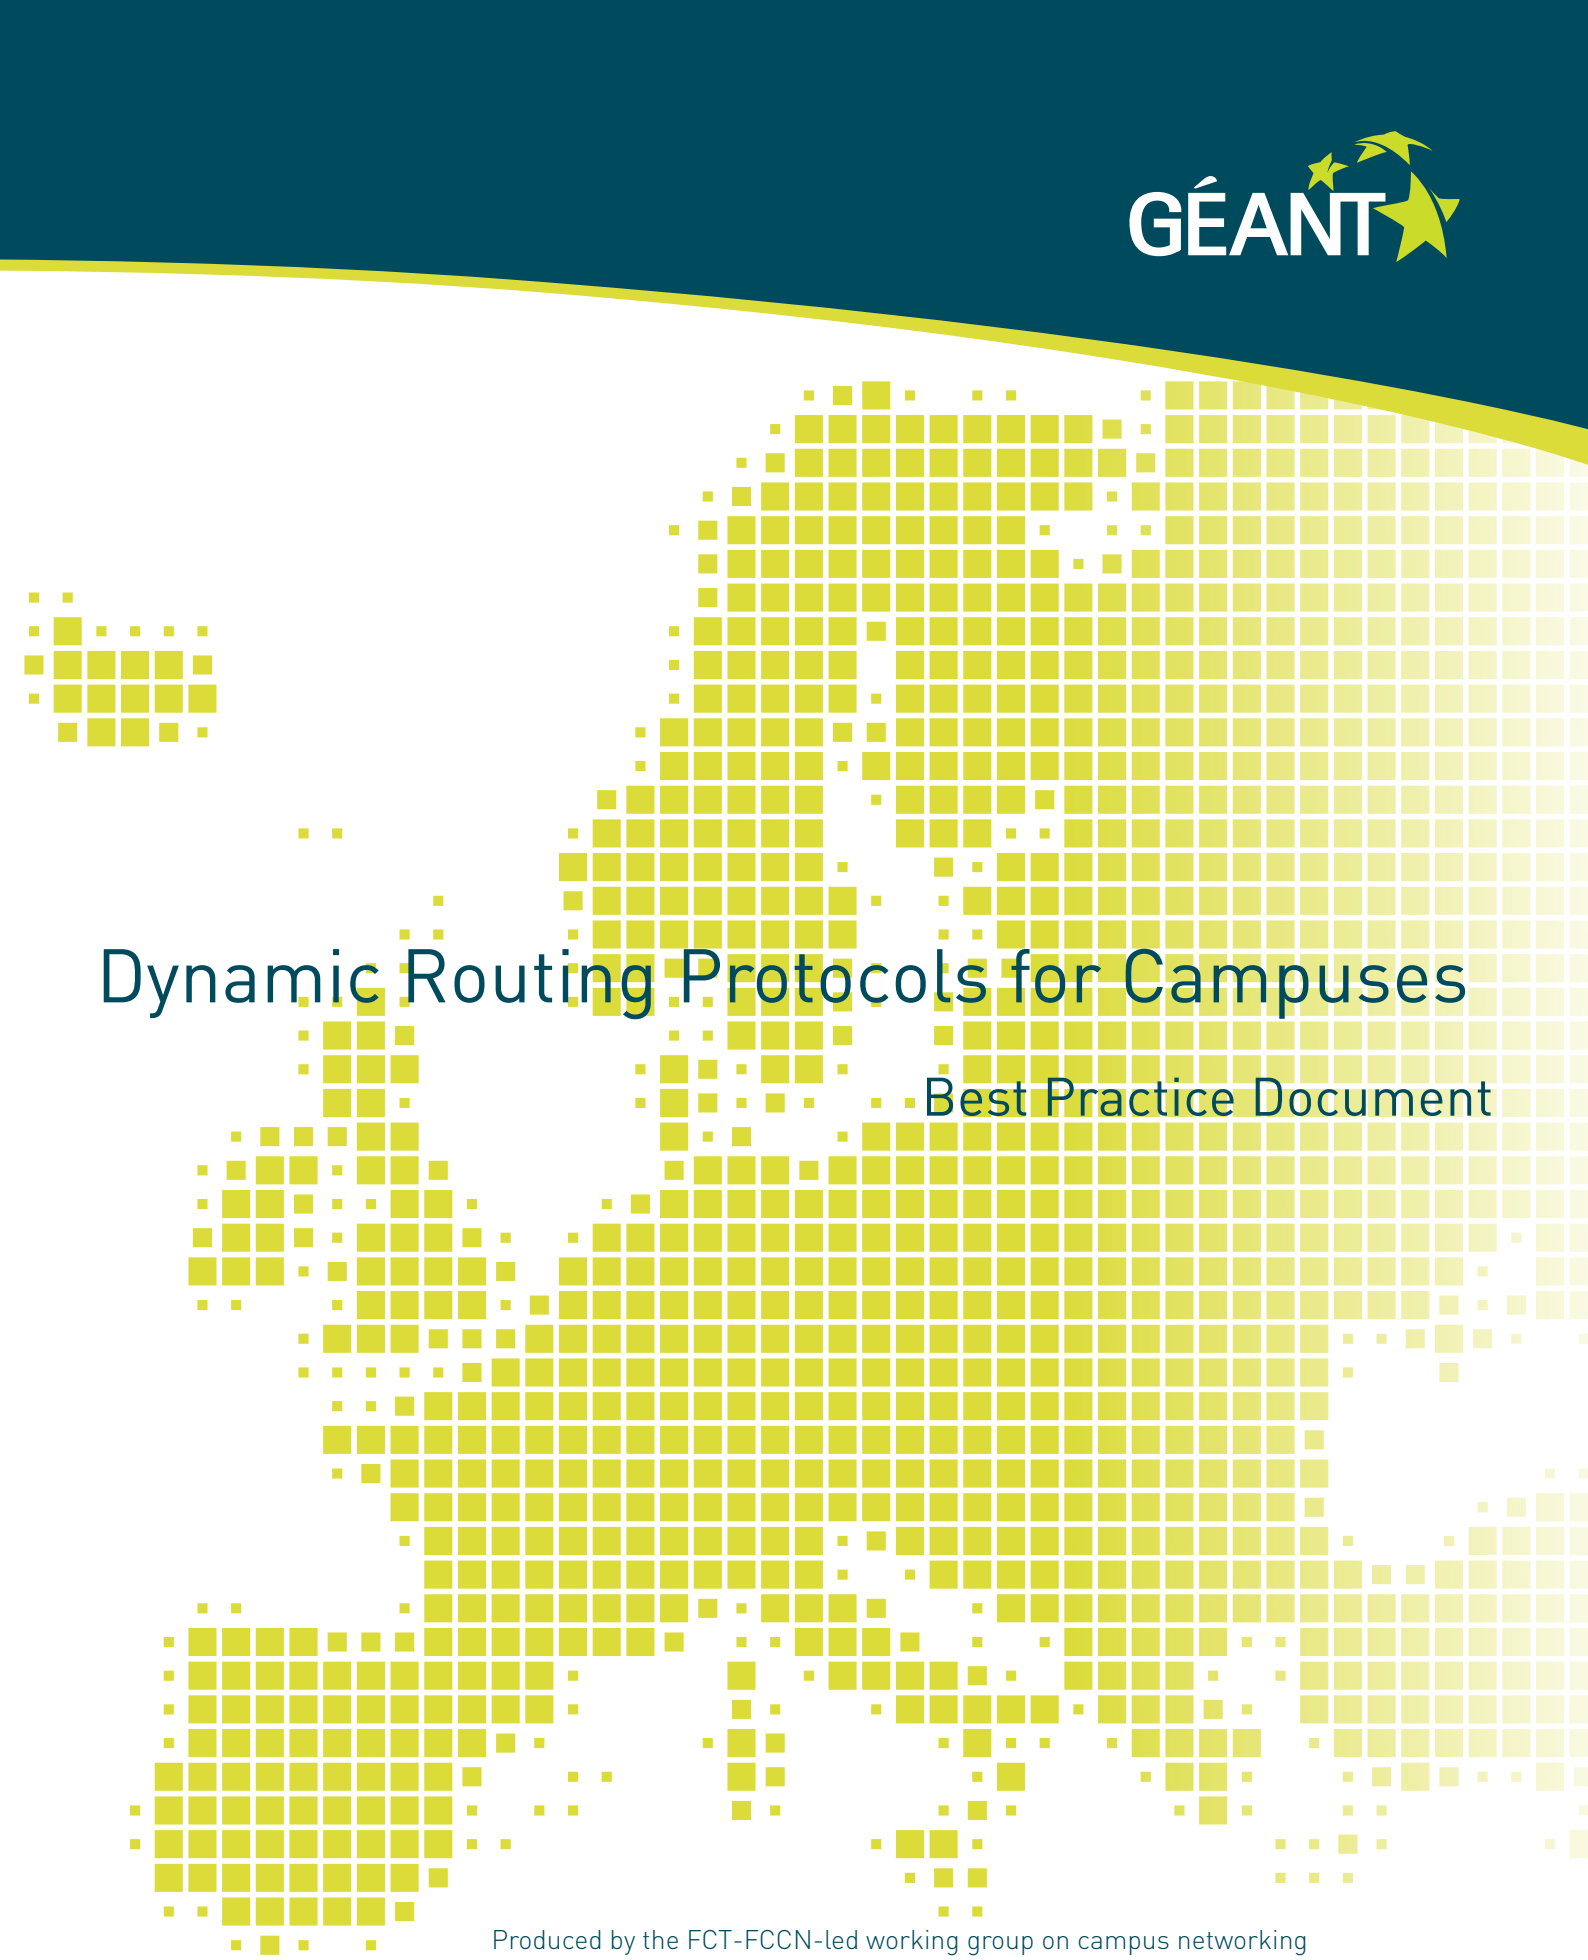

Dynamic Routing Protocols for Campuses

Best Practice Document

Produced by the FCT-FCCN-led working group on campus networking

Author: C. Friaças (FCCN); Contributors: P. Ribeiro (IPLNet); P. Costa (University of Minho); R. Ribeiro (Instituto Universitário de Lisboa); I. Lugo (Instituto Superior Técnico, Lisbon)

October 2014

Document No: GN3-DBPC-301
Version / date: Version 2.4; October 2014
Original language : English
Original title: “Dynamic routing protocols for campuses”
Original version / date: Version 1.0; August 2014
Contact: cfriacas@fccn.pt

FCT-FCCN is responsible for the contents of this document. The document was developed by a FCT-FCCN-led working group on campus networking.

The research leading to these results has received funding from the European Community’s Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 605243 (GN3plus).

Table of Contents

Table of Tables

Executive Summary

Numerous documents about dynamic routing protocols from several different sources can nowadays be easily found. This document tries to provide a campus-oriented perspective, based on the experience of several campus networks' managers.

The main audience for this work is fellow campus networks' managers, looking to build new IP network infrastructures or enhance their existing setup.

IPv6 aspects are covered throughout this document, given that higher education institutes often pursue innovation, and the new Internet Protocol is a technology ready to be used, instead of only advocating its usage by third parties.

The content of this document is designed to be light enough to be used as a quick reference in terms of dynamic routing protocols' usage, by any (even unexperienced) network administrator. Sample configurations for Cisco IOS and Open Source software Quagga – which can run over simple servers – are the main focus of this work.

1 Introduction

When you consider the deployment of a dynamic routing protocol within your campus the issue you need to solve first is sorting out if you really need it. Depending on your network's size, its topology and any local constraints (such as protocol availability in routers' operating systems), static routing could be the most adequate solution. Regarding size, a very small network, with only one or two IP routers wouldn't greatly benefit from a change to dynamic routing.

If you feel dynamic routing is the way to go, the next step is to choose which protocols will be used, and this document aims to support you in reaching a decision.

There are several options, which may imply other subsequent choices. Going for an open routing protocol may enable you to choose generic hardware (i.e. server-based routing), instead of proprietary hardware. In the case you will need to manage an external routing protocol, you can also use it for internal purposes. You can also choose a protocol that will enable routing both for IPv4 and IPv6.

During the following chapters, we will focus on practical details regarding several routing protocols. One concept that you need to keep in mind is the difference between an internal routing protocol and an external routing protocol. The first type is used with the same autonomous system (AS) and the second type to exchange routes with a different autonomous system. An autonomous system is a set of networks managed by the same organisation, which share the same routing policy. It's also perfectly normal to find two or more autonomous systems within the same organisation, if networks belong to different departments, or if they are managed by different teams or have different locations.

2 Choosing an internal routing protocol

The IGP (Interior Gateway Protocol) deployed within your network is something that you will always need to remember, when performing any change. You will need to understand how these protocols work in order to prevent any unwanted behaviour or any performance issues. There is a variety of options, and this document will only focus on the most popular ones.

2.1 RIP

RIP (Routing Information Protocol) is one of the oldest protocols based on distance-vector. Scalability and convergence time are its main weaknesses. If your network is so simple that you can manage to run it with RIP, then you might prefer doing it with static routing. RIP's original version (version 1) doesn't transport network mask information, making it unusable in classless networking, which is now the *de facto* standard. It's advisable to use RIP version 2 with authentication, in order to avoid any issues related with rogue equipment that might have RIP enabled. The interfaces serving user networks should be declared as passive, to avoid unuseful traffic and unwanted new adjacencies.

Nevertheless,Table 2.1 and Table 2.2 are configuration examples, for Cisco and Quagga respectively:

```
router rip
    version 2
    passive-interface default
    no passive-interface <intf>
    network x.y.w.z/<mask>                    # the networks to be announced need to be added
    …
    network x.y.w.z/<mask>
ipv6 unicast-routing
interface <intf>
    ip address x.y.w.z/<mask>
    ip rip send version 2
    ip rip receive version 2
    ipv6 enable
```

```
    ipv6 address x:y:w:z/<mask>
    ipv6 rip <id> enable
```

Table 2.1: Cisco RIP basic configuration example

```
(/etc/quagga/ripd.conf)
router rip
    network x.y.w.z/mask
    network <interface>
(/etc/quagga/ripngd.conf)
router ripng
    network x:y:w:z/mask
    network <interface>
```

Table 2.2: Quagga RIP basic configuration example

As you might notice from the above examples, you'll need to run separate processes in order to use RIP on a dual-stack network.

On Quagga (Table 2.2), you should also include the "link-detect" feature in every active interface, so that routing protocols may react quicker to link outages.

2.2 EIGRP

EIGRP stands for Enhanced Interior Gateway Routing Protocol. It was originally a proprietary protocol (from Cisco Systems), which was transformed in an open standard during 2013. Like RIP, EIGRP is also a distance- vector protocol which also uses some link-state techniques in an algorithm called DUAL (Diffuse Update Algorithm). While running a dual-stack network, if EIGRP is your choice, only one process will be running. The same advices, regarding security and stability, given previously regarding RIP should also apply. Declare all interfaces that shouldn't form an adjacency as passive and use authentication to avoid unconfigured equipment to being integrated in your routing plane.

Please see below a basic EIGRP configuration example. At this time, open-source implementations of EIGRP are still not available.

```
ipv6 unicast-routing
router eigrp 100
    no auto-summary
    network x.y.w.z <wildcard mask>
    …
    network x.y.w.z <wildcard mask>
```

```
ipv6 router eigrp 100
    router-id x.y.w.z
interface <interface>
    ip address x.y.w.z/<mask>
    ipv6 address x:y:w:z/<mask>
    ip authentication mode eigrp 100 md5
    ipv6 enable
    ipv6 eigrp 100
```

Table 2.3: Cisco EIGRP basic configuration example

2.3 OSPF and OSPFv3

Open Shortest Path First (and its version 3 for IPv6) is a link-state routing protocol, adequate to campus networks with several IP routing nodes. The same advices given before regarding RIP and EIGRP also apply to OSPF. Declare all interfaces that shouldn't form adjacencies as passive and use authentication to avoid any wrongful interference with your routing domain.

In Cisco IOS equipment, if available, use the command "ip ospf <instance> area <area> at interface level, this should be the preferred way of associating interfaces with OSPF and, in this case, the "network" command in the "router ospf <instance>" context shouldn't be used.

When redistribution is used, some care should be taken to limit unwanted routes -- route-maps can do this task, with the help of access-lists or prefix-lists.

Sample configurations are given in Table 2.4 and Table 2.5.

```
ipv6 unicast-routing
router ospf 100
    router-id x.y.w.z
    redistribute static metric 1000 metric-type 1 subnets route-map ospf-static-redist
    network x.y.w.z <wildcard mask> area 0
    …
    network x.y.w.z <wildcard mask> area 0
    passive-interface default
    no passive-interface <intf>
ipv6 router ospf 100
    router-id x.y.w.z
    redistribute static metric 1000 metric-type 1 subnets route-map ospf-static-redist
interface <interface>
```

```
    ip address x.y.w.z/<mask>
    ipv6 address x:y:w:z/<mask>
    ipv6 enable
    ipv6 ospf 100 area 0
route-map ospf-static-redist permit 10
   match tag 100
```

Table 2.4: Cisco OSPF and OSPFv3 basic configuration example

```
(/etc/quagga/ospfd.conf)
router ospf
     network x.y.w.z/nn area 0
(/etc/quagga/ospf6d.conf)
router ospf6
     router-id a.a.a.a
     interface eth0 area 0
```

Table 2.5: Quagga OSPF and OSPFv3 basic configuration example

There are two curious details regarding OSPFv3 configuration: The router ID is a 32-bit address, which is often confused with an IPv4 address, because designers chose to allow its input using the same syntax. The other odd design choice was the need to link the OSPF process with each interface to be used, instead of the need to describe only network/mask (which is the original IPv4 OSPF method). It's not mandatory to add a loopback, or specify a router ID, but it's recommended to create one (which will be reachable if any interface is available), and configure it as the router ID.

2.4 IS-IS

IS-IS (Intermediate System to Intermediate System) runs over CLNS (Connectionless-mode Network Service).

```
ipv6 unicast-routing

router isis 100

    net nn.oooo.xxxx.yyyy.wwwww.zz

interface <interface>
    ip address x.y.w.z/<mask>
    ipv6 address x:y:w:z/<mask>
    ip router isis 100
```

```
    ipv6 router isis 100
```

Table 2.6: Cisco IS-IS basic configuration example

The most awkward detail within IS-IS are CLNS addresses (nn.oooo.xxxx.yyyy.wwwww.zz). This CLNS address is a router identifier, which means that each router must have defined a unique CLNS address.

Running IS-IS on Quagga isn't a simple choice. The protocol is not supported on the main distribution packages, but alpha/beta source-code is freely available. The official Quagga website [4] has an explicit reference indicating IS-IS for IPv6 has known unsolved issues.

2.5 iBGP

iBGP is the internal version of Border Gateway Protocol. The most visible detail that distinguishes iBGP from eBGP (external BGP) is using different AS numbers in order to express neighborhoods. Some behaviours of the protocol are different depending on the internal/external flavor. (e.g. In iBGP local-preference attributes are preserved in the AS, whereas in eBGP relations, the routes' receiver always overwrites the local-preference). Table 2.7 and Table 2.8 are simple examples – but security shouldn't be forgotten, hence the use of a password for each neighbor is a strong suggestion.

One aspect to be considered (if iBGP is your choice) is scalability, and the possible usage of route reflectors instead of configuring full mesh (i.e. all routers have one session to every other router also running iBGP).

```
ipv6 unicast-routing
router bgp 100
    router-id x.y.w.z
    neighbor <group> peer-group
    neighbor <group> version 4
    neighbor a.a.a.a remote-as 100
    neighbor a.a.a.a password <password>
    neighbor b.b.b.b remote-as 100
    neighbor b.b.b.b password <password>
    neighbor c.c.c.c remote-as 100
    neighbor c.c.c.c password <password>
    neighbor a:a:a:a::a remote-as 100
    neighbor a:a:a:a::a password <password>
    neighbor b:b:b:b::b remote-as 100
    neighbor b:b:b:b::b password <password>
    neighbor c:c:c:c::c remote-as 100
    neighbor c:c:c:c::c password <password>
```

```
!
address-family ipv4 unicast
neighbor <group> activate
no auto-summary
no synchronization
network x.y.w.z mask <mask>
neighbor a.a.a.a activate
neighbor b.b.b.b activate
neighbor c.c.c.c activate
exit-address-family
!
address-family ipv6 unicast
network x:y:w:z/<mask>
neighbor a:a:a:a::a activate
neighbor b.b.b.b::b activate
neighbor c.c.c.c::c activate
exit-address-family
!
```

Table 2.7: Cisco iBGP basic configuration example

(/etc/quagga/bgpd.conf)

```
router bgp 100
    router-id x.y.w.z
    neighbor <group> peer-group
    neighbor <group> version 4
    neighbor a.a.a.a remote-as 100
    neighbor a.a.a.a password <password>
    neighbor b.b.b.b remote-as 100
    neighbor b.b.b.b password <password>
    neighbor c.c.c.c remote-as 100
    neighbor c.c.c.c password <password>
    neighbor a:a:a:a::a remote-as 100
    neighbor a:a:a:a::a password <password>
    neighbor b:b:b:b::b remote-as 100
    neighbor b:b:b:b::b password <password>
    neighbor c:c:c:c::c remote-as 100
```

```
neighbor c:c:c:c::c password <password>
no auto-summary
no synchronization
network x.y.w.z mask <mask>
neighbor a.a.a.a activate
neighbor b.b.b.b activate
neighbor c.c.c.c activate
exit-address-family
!
address-family ipv6
network x:y:w:z/<mask>
neighbor a:a:a:a::a activate
neighbor b.b.b.b::b activate
neighbor c.c.c.c::c activate
exit-address-family
!
```

Table 2.8: Quagga iBGP basic configuration example

The system's kernel where Quagga is running should be compiled with TCP MD5 support signature (RFC2385), if authentication is to be used.

3 External routing protocol

3.1 eBGP

eBGP is the external mode of Border Gateway Protocol, in order to enable routing information exchange between different autonomous systems (i.e. networks). eBGP (external BGP) is configured by using other networks' autonomous system ID through the command remote-as. In eBGP it's especially important to control which routes to accept and which routes to be sent to your peers. One should guarantee that your own routes are not announced to you from the outside world, and that you are not exporting any private routes towards other networks. To avoid unusually long AS paths the "bgp maxas-limit" can be used. IPv6 routes can be transported over IPv4 sessions and the inverse, but that introduces several dependencies, risks and limitations. There should be separate sessions for IPv4 and IPv6 routes, each established over the same IP protocol version.

Table 3.1 and Table 3.2 give eGBP basic configuration examples for Cisco and Quagga respectively.

```
ipv6 unicast-routing
router bgp 100
    bgp router-id x.y.w.z
    bgp log-neighbor-changes
    bgp maxas-limit 32
    neighbor <group> peer-group
    neighbor <group> version 4
    neighbor a.a.a.a remote-as 200
    neighbor a.a.a.a password <password>
    neighbor a.a.a.a maximum-prefix 500
    neighbor a.a.a.a prefix-list MY_NETWORKS out
    neighbor a.a.a.a prefix-list DROP_NETWORKS in
    neighbor a:a:a:a::a remote-as 200
    neighbor a:a:a:a::a password <password>
    neighbor a:a:a:a::a maximum-prefix 500
    !
```

```
    address-family ipv4 unicast
    neighbor <group> activate
    no auto-summary
    no synchronization
    network x.y.w.z mask <mask>
    neighbor a.a.a.a activate
    exit-address-family
    !
    address-family ipv6 unicast
    network x:y:w:z/<mask>
    neighbor a:a:a:a::a activate
    exit-address-family
!
```

Table 3.1: Cisco eBGP basic configuration example

```
(/etc/quagga/bgpd.conf)
router bgp 100
    router-id x.y.w.z
    no auto-summary
    no synchronization
    neighbor a.a.a.a remote-as 200
    neighbor a.a.a.a password <password>
    neighbor a.a.a.a update-source eth0
    neighbor a.a.a.a next-hop-self
    neighbor a.a.a.a maximum-prefix 500
    neighbor a.a.a.a prefix-list MY_NETWORKS out
    neighbor a.a.a.a prefix-list DROP_NETWORKS in
    neighbor a.a.a.a filter-list 301 out
    neighbor a.a.a.a filter-list 401 in
    neighbor a:a:a:a::a remote-as 200
    neighbor a:a:a:a::a password <password>
    no neighbor a:a:a:a::a activate
    network x.y.w.z mask <mask>
!
    address-family ipv6
```

```
    network x:y:w:z/<mask>
    neighbor a:a:a:a::a activate
    exit-address-family
!
```

Table 3.2: Quagga eBGP basic configuration example

When the choice is running BGP, it's important to define if a public or private ASN will be used. A private ASN is the cheapest option, provided that your upstream provider can "mask" that private ID when there is the need to propagate one of your routes to the Internet (i.e. your upstream's upstreams). A public ASN can be an option if your organisation is already a RIR (RIPE NCC in the case of Europe) Local Internet Registry, or if finding a sponsoring LIR that will charge your organisation a sensible cost associated to this numbering resource.

4 Conclusion

Experience has shown that tools are available in order to evolve from static routing to dynamic routing in the context of R&E campus networks. While for external routing purposes, BGP is the only viable option, each campus network manager must choose which IGP is the best fit for its context. The authors also feel that choosing RIP is not a good option, and that using static routing would be a better solution.

This document provides several basic configuration samples, either for a vendor-based solution (Cisco Systems) and a free open-source «build your own router» solution (Quagga). When your choice is using a Cisco device, you will also have to check if the protocols that you wish to run are supported in your current operating system distribution/version. If this is not the case, you will need to contact your local reseller, in order to evaluate options on how to add protocol related features.

From reading this document, the reader should be able to understand how simple it is to move to dynamic routing, after that decision is taken. It's also important to understand the security implications, and the need for cautious filtering. While using BGP (iBGP or eBGP), neighborhoods have to be declared, which is not the case in other protocols. This should encourage a network manager to proceed cautiously when determining which interfaces run dynamic routing protocols.

References

[1] Recommended Resilient Campus Network Design,
Tomas Podermanski, Vladimir Zahorik, March 2010 (CBPD114, Czech Republic)
http://www.terena.org/activities/campus-bp/pdf/gn3-na3-t4-cbpd114.pdf

[2] Recommendations for a Redundant Campus Network,
Gunnar Bøe, Vidar Faltinsen, Einar Lillebrygfjeld, December 2011 (UFS114, Norway)
http://www.terena.org/activities/campus-bp/pdf/gn3-na3-t4-ufs114.pdf

[3] Edge Device for a Campus Network,
Jani Myyry, Kaisa Haapala, Janne Oksanen, Janne Niemi, May 2011 (Finland)
http://www.terena.org/activities/campus-bp/pdf/gn3-na3-t4-edge-device.pdf

[4] http://www.nongnu.org/quagga

Glossary

ACL	Access Control List
ARP	Address Resolution Protocol
AS	Autonomous System
ASN	Autonomous System Number
BGP	Border Gateway Protocol
CLNS	Connectionless-mode Network Service
DUAL	Diffuse Update Algorithm
eBGP	Exterior Border Gateway Protocol
EIGRP	Enhanced Interior Gateway Routing Protocol
ID	Identification
iBGP	Interior Border Gateway Protocol
IGP	Interior Gateway Protocol
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IS-IS	Intermediate System To Intermediate System
OSPF	Open Shortest Path First
OSPFv3	Open Shortest Path First version 3
R&E	Research & Education
RIP	Routing Information Protocol
RIPE NCC	Réseaux IP Européens Network Coordination Centre
RIR	Regional Internet Registry

Complete BPDs are available at http://services.geant.net/cbp/Pages/Home.aspx

campus-bp-announcements@terena.org