Seamless Network-Wide Migrations

Laurent Vanbever

inl.info.ucl.ac.be/lvanbeve/

Network Automation 2011

Novembre 9, 2011

UCL Université catholique de Louvain

NETWORK automation CONFERENCE | EXHIBITION | 2011

Human factors are responsible for
50 to 80% of network device outages.

— Juniper Networks, 2008

Human factors are responsible for
50 to 80% of network device outages.

— Juniper Networks, 2008

Large (resp. medium) businesses lose an average of 3.6% (resp. 1%) in annual revenue due to network downtime.

— Infonetics Research, 2004 – 2006

seamless migrations

The progressive modification of the configuration of a running network without loosing any IP packets

Migrating the network can provide immediate benefits

Network migration can improve the

- manageability
- performance
- stability
- security

of the entire network

Migrating the network
can provide immediate benefits

Network migrations can improve the

- manageability
- performance
- stability
- security

introduce (remove) a hierarchy

support different features

improve route manipulation

of the entire network

Migrating the network
can provide immediate benefits

Network migrations can improve the

- manageability
- performance
- stability
- security

reduce control-plane load

reduce routing-table size

reduce convergence time

of the entire network

Migrating the network can provide immediate benefits

Network migrations can improve the

- manageability
- performance
- stability
- security

isolate parts of the network

reduce the churn

of the entire network

Migrating the network can provide immediate benefits

Network migrations can improve the

- manageability
- performance
- stability
- security

use stable code base

prevent potential attacks

of the entire network

Migrating the network is operationally complex

Reconfigure a running network
while respecting Service Level Agreement

Make highly distributed changes
on all the routers, in a coordinated manner

Face potential routing anomalies
as non-migrated routers interact with migrated ones

Seamless Network-Wide Migrations

1 IGP migrations

Ordering matters

2 BGP migrations

Ongoing work

3 How can we help ?

Manage complexity

Seamless Network-Wide Migrations

1 IGP migrations

Ordering matters

BGP migrations

Ongoing work

How can we help ?

Manage complexity

Problem | Replace the IGP configuration of a running network, router-by-router, without causing any anomalies

Details in: Laurent Vanbever, Stefano Vissicchio, Cristel Pelsser, Pierre Francois and Olivier Bonaventure. *Seamless Network-Wide IGP Migrations*. ACM SIGCOMM 2011, Aug. 2011.

Sub-problem 1	Replace the IGP configuration of a running network, router-by-router, without causing any anomalies
Current Practice	Run the two IGP configurations in parallel

Migrating the IGP usually requires running two routing planes

Abstract model of a router

Control-plane

Data-plane

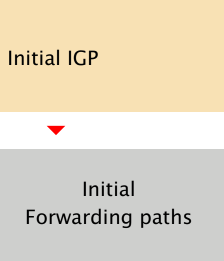

At first, the initial IGP dictates the forwarding paths being used

Migrating the IGP usually requires running two routing planes

Abstract model of a router

Control-plane

Data-plane

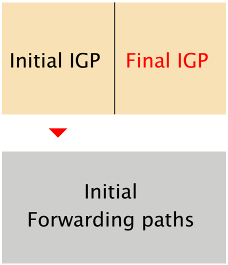

Then, the final IGP is introduced without changing the forwarding

Migrating the IGP usually requires running two routing planes

Abstract model of a router

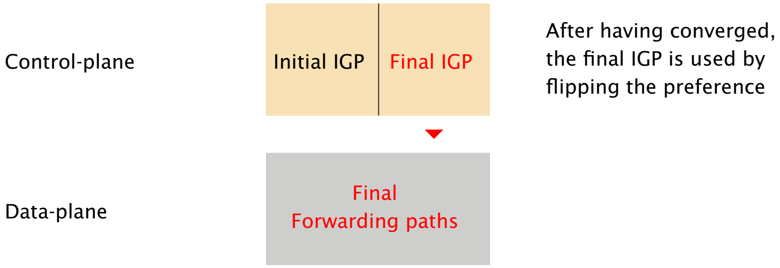

Migrating the IGP usually requires running two routing planes

Abstract model of a router

Control-plane

Initial IGP | Final IGP

After having converged, the final IGP is used by flipping the preference

Data-plane

Final
Forwarding paths

Migrating the IGP usually requires running two routing planes

Abstract model of a router

Control-plane

Final IGP

The initial IGP is removed as it is not used anymore

Data-plane

Final
Forwarding paths

Sub-problem 1

Replace the IGP configuration of a running network, router-by-router, without causing any anomalies

Sub-problem 2 Replace the IGP configuration of a running network, router-by-router, without causing any anomalies

Migrating the IGP can create *migration loops*

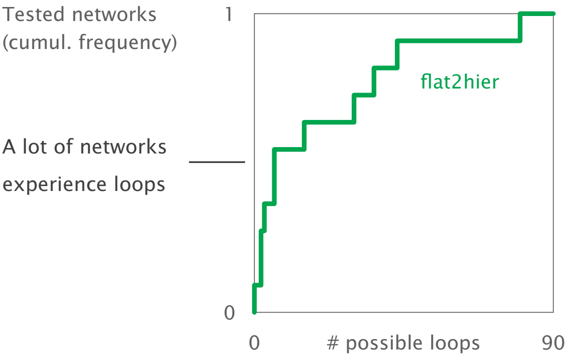

Up to 90 *migration loops* can arise during an IGP migration

Sub-problem 2

Replace the IGP configuration of a running network, router-by-router, without causing any anomalies

Sub-problem 2	Replace the IGP configuration of a running network, router-by-router, without causing any anomalies
Contributions	Seamless IGP migration is possible as long as the reconfiguration process follows a strict ordering

Contributions

Seamless IGP migration is possible as long as the reconfiguration process follows a strict ordering

which one ?

Reconfiguring the IGP might change the forwarding paths being used

In a flat IGP, routers forward traffic according to the shortest-path towards the destination.

In a flat IGP, R2 reaches R1 via R4

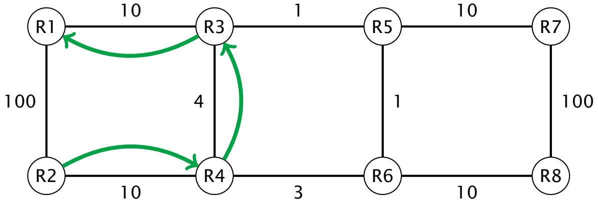

Reconfiguring the IGP might change the forwarding paths being used

In a hierarchical IGP, routers prefer paths contained within a single zone over the ones crossing several zones

In a hierarchical IGP, R2 reaches R1 directly

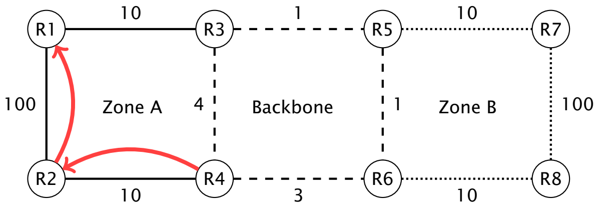

Whenever the forwarding paths change, forwarding loops can be created

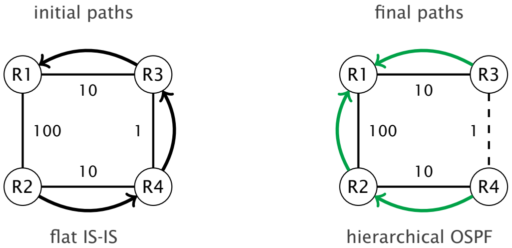

Forwarding paths towards R1

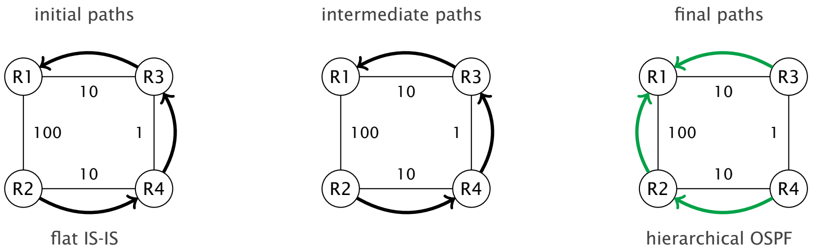

Forwarding paths towards R1

First, we migrate R3

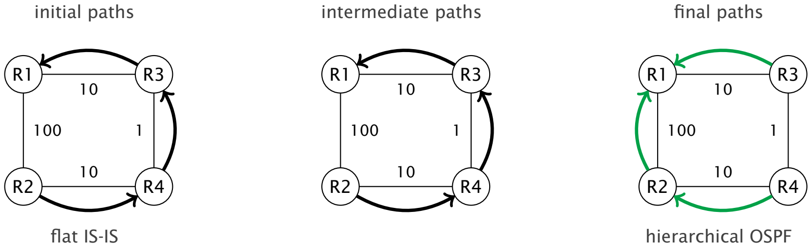

Forwarding paths towards R1

First, we migrate R3

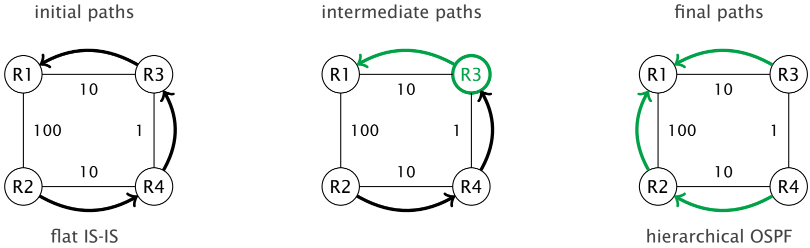

Forwarding paths towards R1

Then, we migrate R4

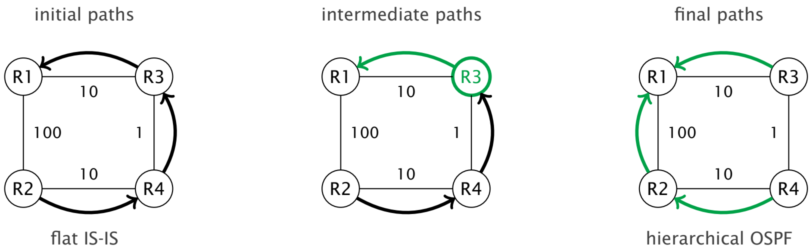

Forwarding paths towards R1

Then, we migrate R4

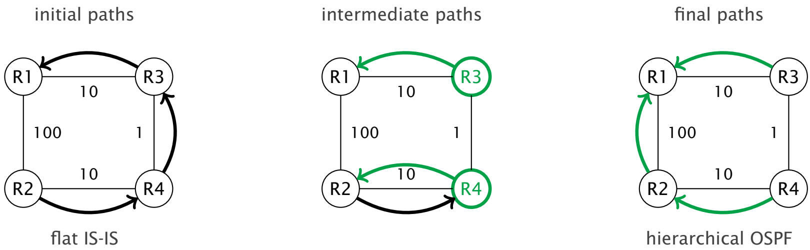

Forwarding paths towards R1

Whenever the forwarding paths change, forwarding loops can be created

A loop is created if R4 is migrated before R2

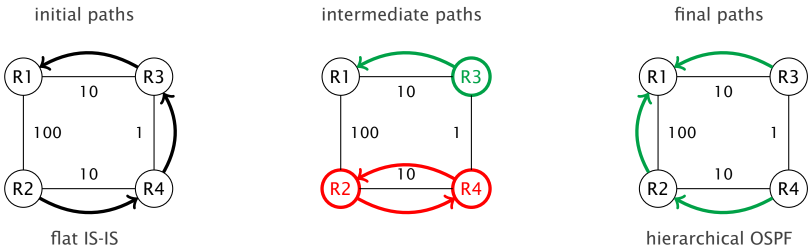

Forwarding paths towards R1

Migrations have to be performed following a precise ordering

No loop arises if R2 is migrated before R4

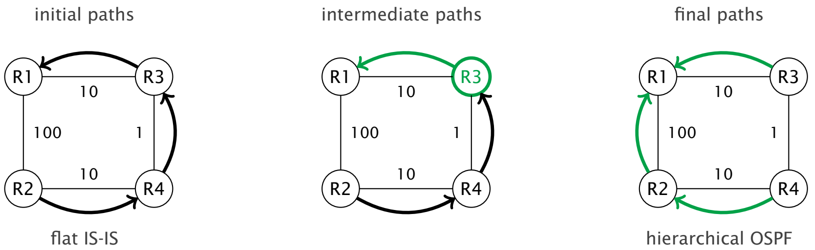

Forwarding paths towards R1

Migrations have to be performed following a precise ordering

No loop arises if R2 is migrated before R4

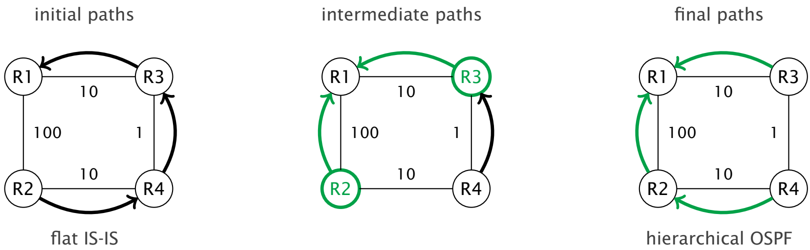

Forwarding paths towards R1

Migrations have to be performed following a precise ordering

No loop arises if R2 is migrated before R4

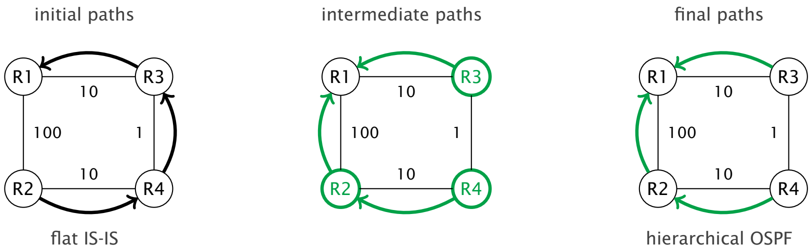

Forwarding paths towards R1

Although hard in theory, finding an ordering is possible in practice

Implementation of two ordering algorithms

1. correct, complete, but slow
2. correct, not-complete, but fast

Complete support of the

- introduction (removal) of an IGP hierarchy
- introduction (removal) of route summarization
- modifications of link weights

For all the tested networks,
we have been able to find an ordering

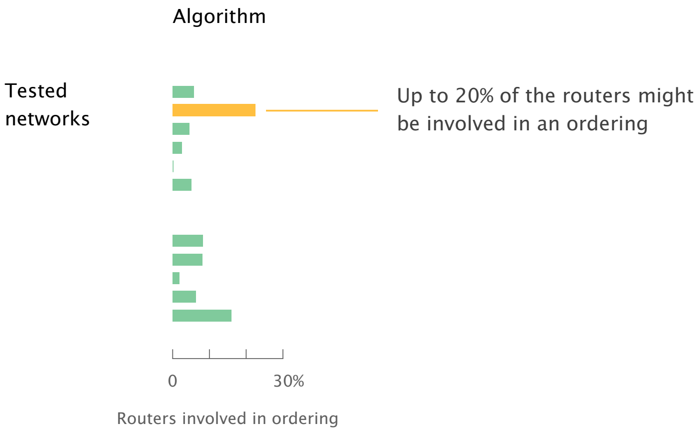

Our techniques can also deal with links and nodes failures during the migration

Failures can change the computed ordering
as they modify the underlying IGP topology

Solutions

- Precompute failover orderings
- Compute a new ordering when a failure is detected

We implemented a provisioning system which automates the process

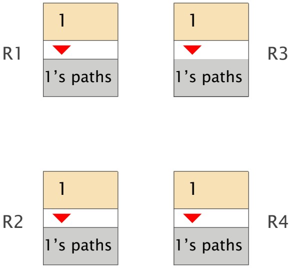

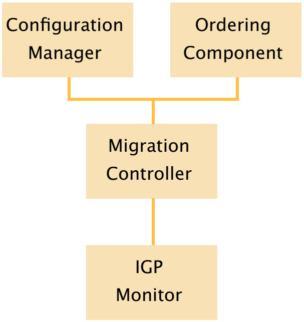

Network in which IGP 1 is replaced by IGP 2

First, the *Ordering Component* computes the ordering (if any)

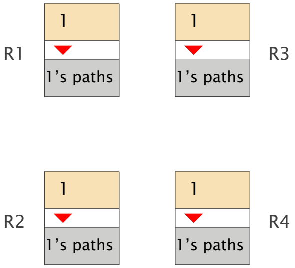

Network in which IGP 1 is replaced by IGP 2

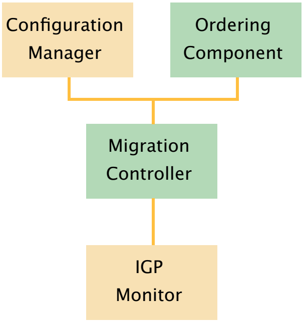

First, the *Ordering Component* computes the ordering (if any)

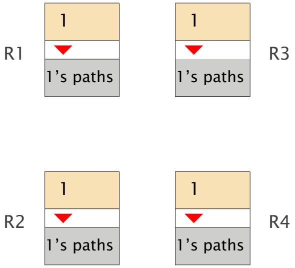

Network in which IGP 1 is replaced by IGP 2

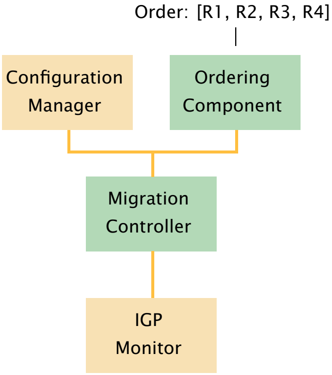

Second, the *IGP Monitor* builds a dynamic view of the IGP and assesses its stability

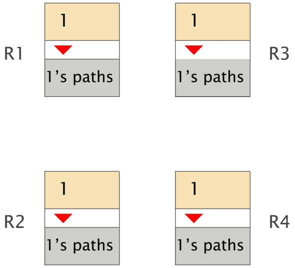

Network in which IGP 1 is replaced by IGP 2

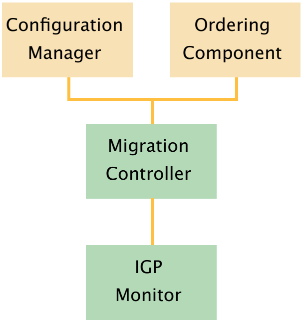

Third, the *Configuration Manager* introduces the, final configuration (not yet used) on all the routers

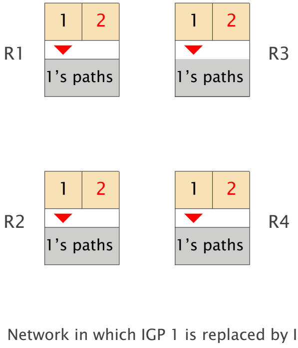

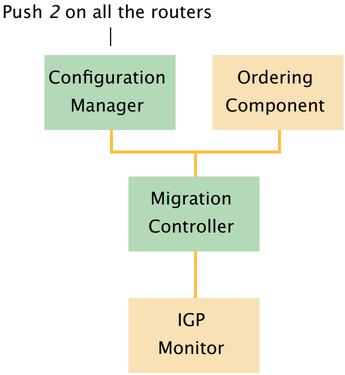

Fourth, the final IGP's completeness and stability are verified by the *IGP Monitor*

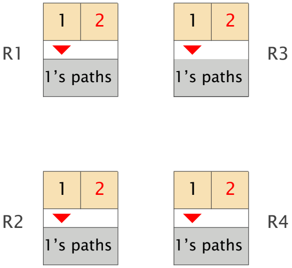

Network in which IGP 1 is replaced by IGP 2

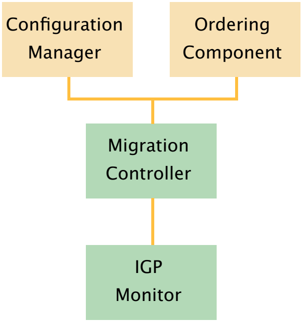

Fifth, the *Configuration Manager* reconfigures each router – according to the ordering — so that it uses the final IGP

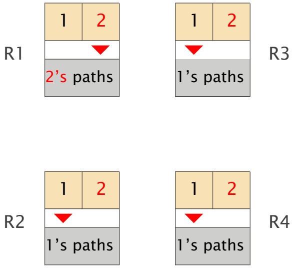

Network in which IGP 1 is replaced by IGP 2

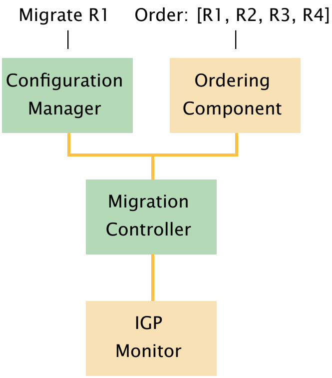

Fifth, the *Configuration Manager* reconfigures each router – according to the ordering — so that it uses the final IGP

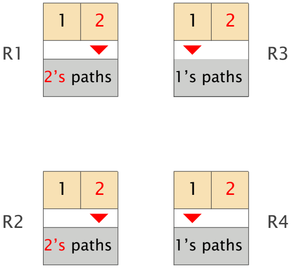

Network in which IGP 1 is replaced by IGP 2

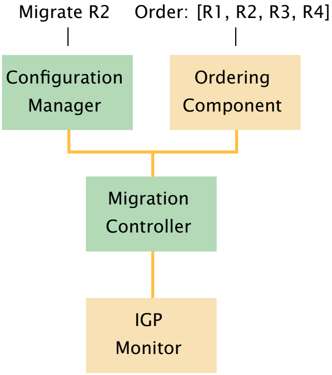

Fifth, the *Configuration Manager* reconfigures each router – according to the ordering — so that it uses the final IGP

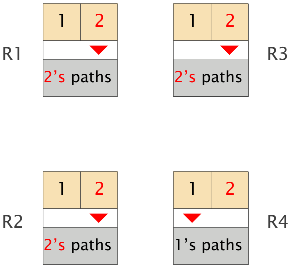

Network in which IGP 1 is replaced by IGP 2

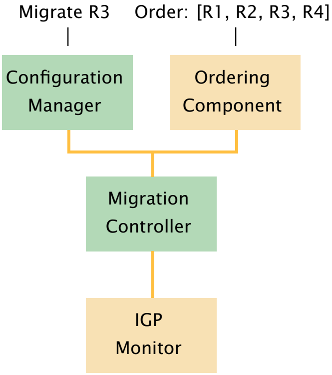

Fifth, the *Configuration Manager* reconfigures each router – according to the ordering — so that it uses the final IGP

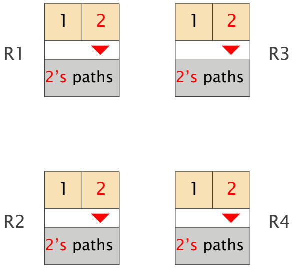

Network in which IGP 1 is replaced by IGP 2

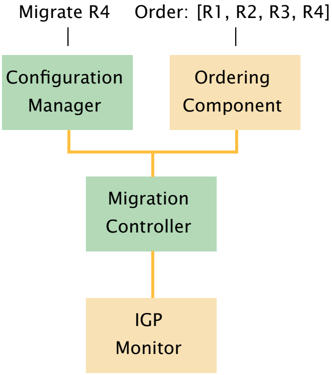

Sixth, the IGP migration is over. The *Configuration Manager* removes the initial IGP configuration from each router

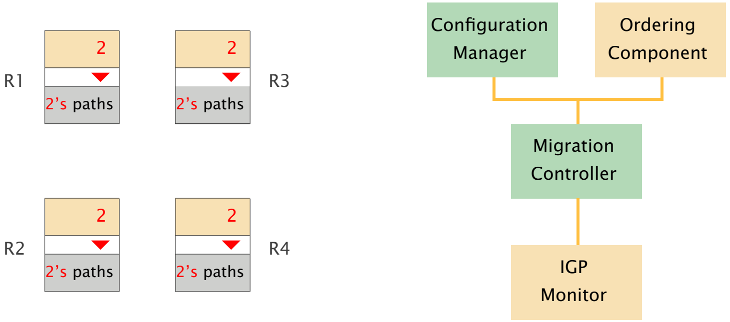

Network in which IGP 1 is replaced by IGP 2

Let’s reconfigure an existing network from a *flat* IGP ...

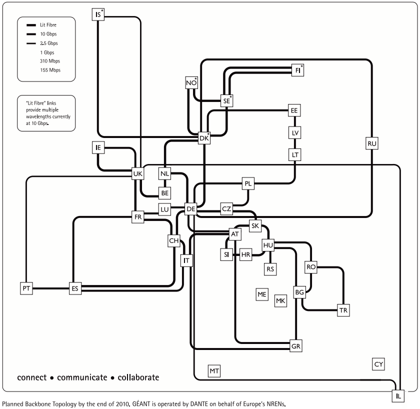

Planned Backbone Topology by the end of 2010. GÉANT is operated by DANTE on behalf of Europe's NRENs.

GEANT

European research network

36 routers

53 links

Let’s reconfigure an existing network from a *flat* IGP to a *hierarchical* IGP

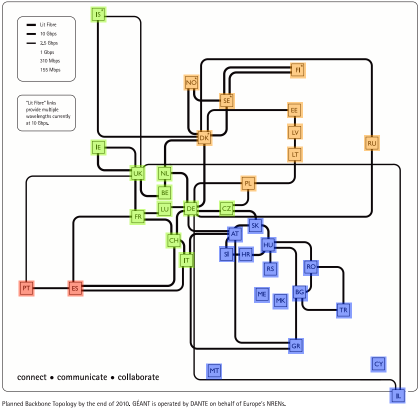

Planned Backbone Topology by the end of 2010. GÉANT is operated by DANTE on behalf of Europe’s NRENs.

GEANT

European research network

36 routers

53 links

Backbone zone

South-west zone

South-east zone

North-east zone

Lossless reconfiguration is possible, by following the precomputed ordering

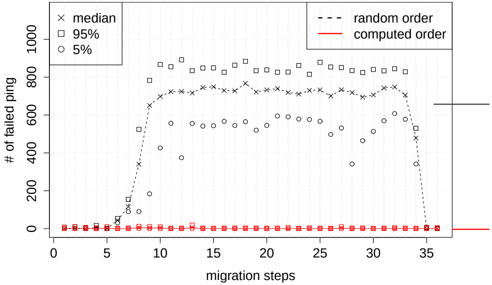

Traffic gets lost during more than *80%* of the process

No loss occurs with proper ordering

Average results (50 repetitions) computed on 700+ pings per step from every router to 5 problematic destinations

Seamless Network-Wide Migrations

IGP migrations

Ordering matters

2 BGP migrations

Ongoing work

How can we help ?

Manage complexity

BGP reconfiguration scenarios are numerous

	Scenarios
iBGP	Full-mesh to route-reflection
	Route-reflection to full-mesh
	Add, (re-)move sessions
eBGP	Add, (re-)move sessions
	Modify in(out)bound policies

BGP reconfiguration scenarios are numerous ... and problematic

	Scenarios	Problems
iBGP	Full-mesh to route-reflection	
	Route-reflection to full-mesh	Forwarding loops
	Add, (re-)move sessions	Routing instabilities
		Visibility issues
eBGP	Add, (re-)move sessions	Traffic shifts
	Modify in(out)bound policies	

BGP reconfiguration scenarios are numerous ... and problematic

Preliminary evaluations confirm that problems may arise even when following the best current practices

Problems

- Forwarding loops
- Routing instabilities
- Visibility issues
- Traffic shifts

As for the IGP, reconfiguring BGP can create problems

Initial configuration

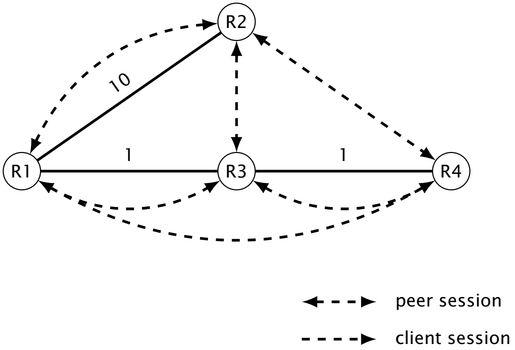

As for the IGP, reconfiguring BGP can create problems

Final configuration

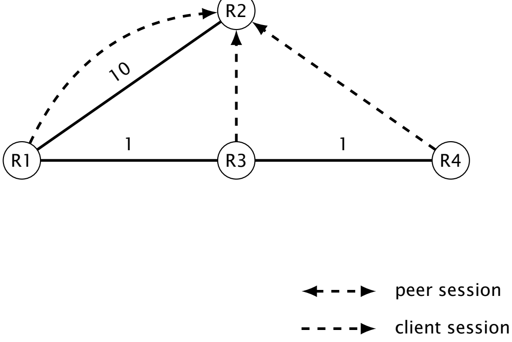

As for the IGP, reconfiguring BGP can create problems

In a full-mesh, every router knows 2 paths to reach the destination

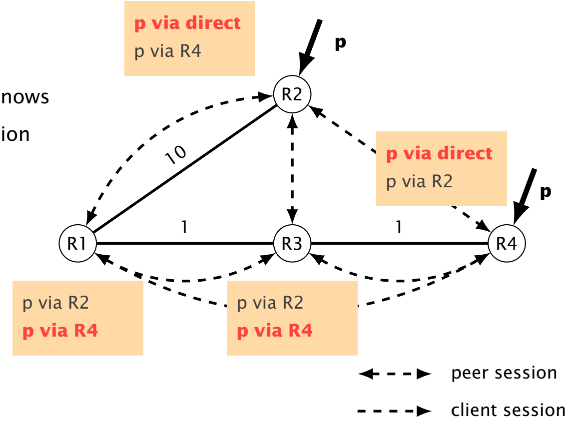

As for the IGP, reconfiguring BGP can create problems

Both R1 and R3 forwards traffic to R4

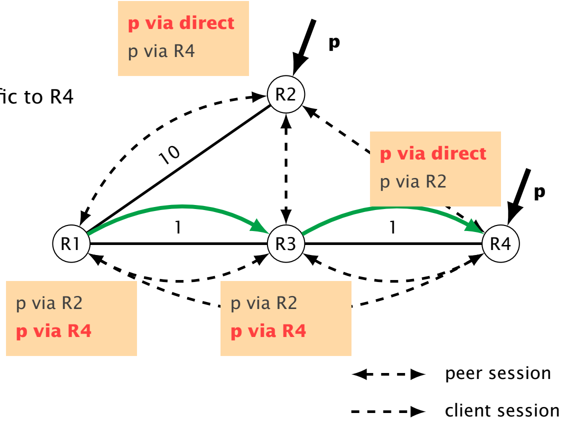

As for the IGP, reconfiguring BGP can create problems

R3 is reconfigured first by adding a client session towards R2

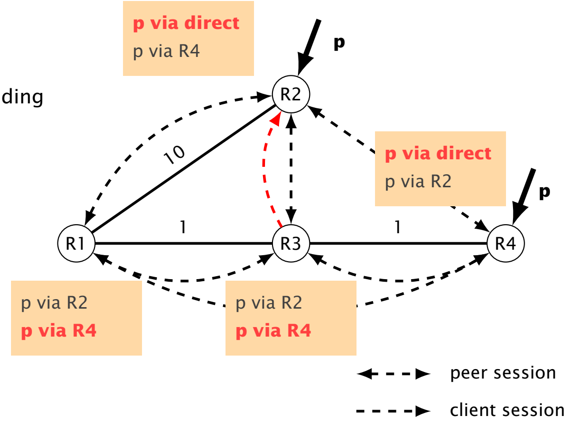

As for the IGP, reconfiguring BGP can create problems

R3 loses the visibility of R4's paths when the initial sessions are removed

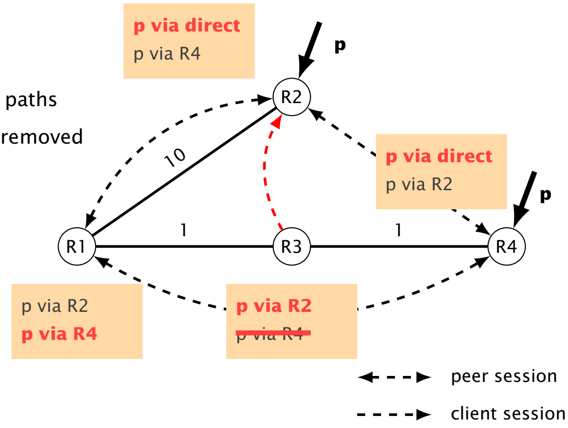

As for the IGP, reconfiguring BGP can create problems

A forwarding loop appears since

R3 and R4 use each other to reach P

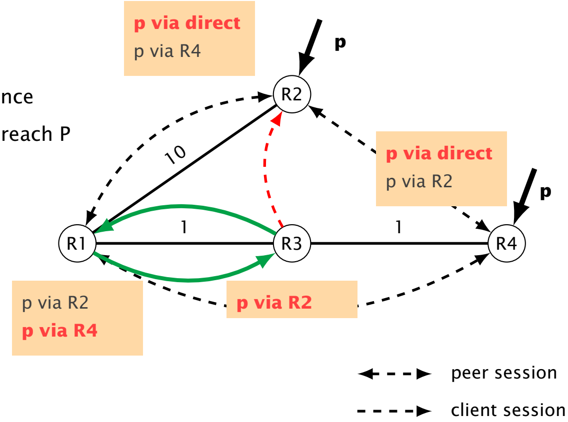

Adding iBGP sessions locally can worsen the visibility globally

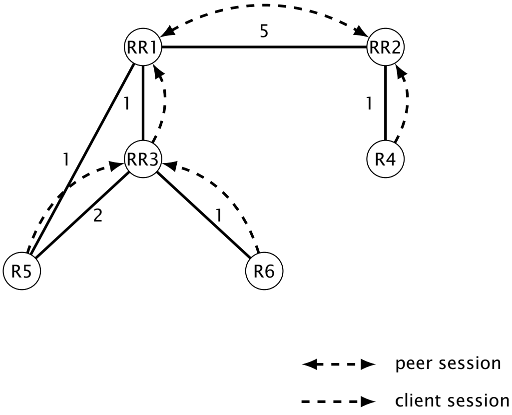

Adding iBGP sessions locally can worsen the visibility globally

The same prefix is learned on R4, R5 and R6

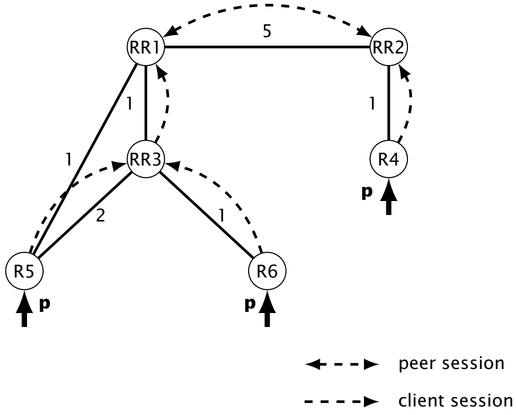

Adding iBGP sessions locally can worsen the visibility globally

Due to route-reflection, RR1 is not aware that R5 is an egress

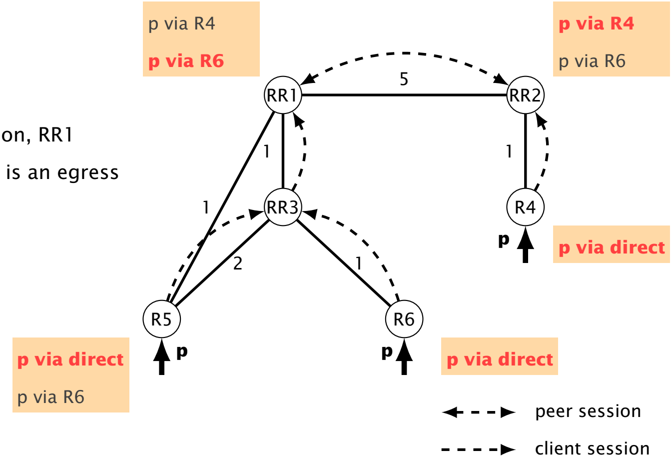

Adding iBGP sessions locally can worsen the visibility globally

Adding an iBGP session between R5 and RR1 solves the problem

... but worsen RR2's visibility

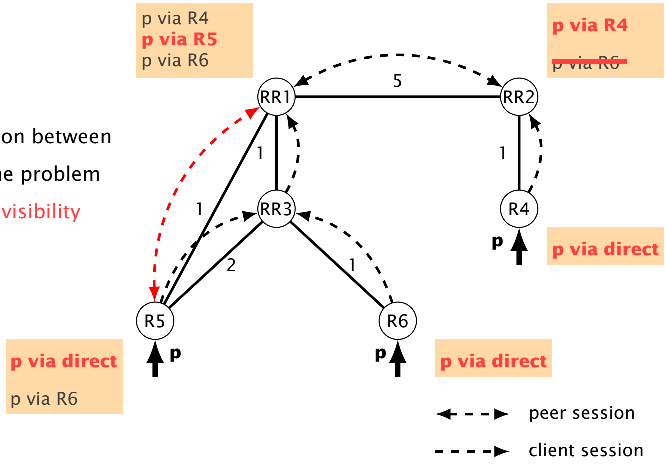

Adding iBGP sessions locally can worsen the visibility globally

RR1 does not propagate to RR2 paths learned on peer sessions

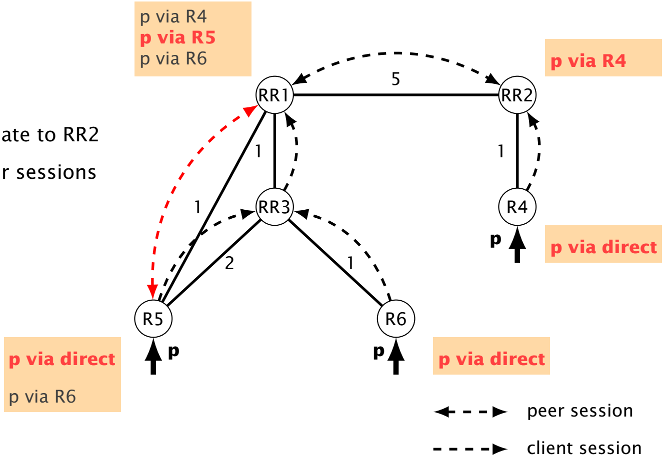

Adding iBGP sessions locally can worsen the visibility globally

Worse! If R4 looses its path, RR2 cannot reach P anymore

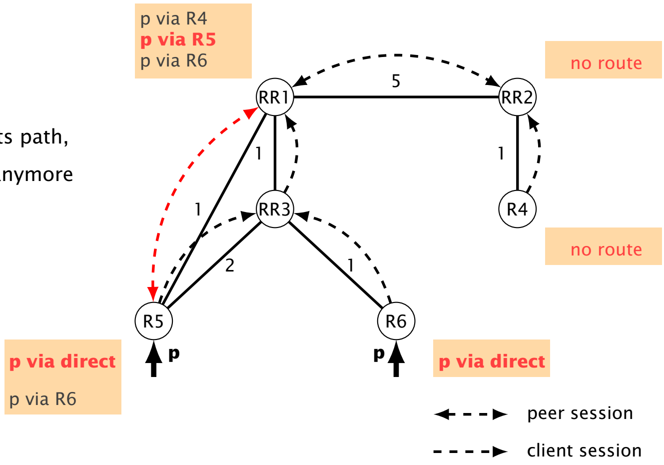

Reconfiguring BGP relying only on the protocol is quite challenging

Checking if adding (removing) an iBGP session will trigger

- forwarding loops
- routing instabilities
- visibility issues
- traffic shifts

is computationally intractable

Leveraging other technologies helps in achieving seamless BGP reconfiguration

Use the BGP/MPLS VPNs machinery

- Deploy the different configurations in different VRFs
- Migrate by switching the VRF used to forward traffic

Leveraging other technologies helps in achieving seamless BGP reconfiguration

Use the BGP/MPLS VPNs machinery

- Deploy the different configurations in different VRFs
- Migrate by switching the VRF used to forward traffic

Extend BGP graceful restart to the entire network

- Disconnect the RIB and the FIB before the migration
- Reconnect them together after the migration

Leveraging other technologies helps in achieving seamless BGP reconfiguration

Use the BGP/MPLS VPNs machinery

- Deploy the different configurations in different VRFs
- Migrate by switching the VRF used to forward traffic

Extend BGP graceful restart to the entire network

- Disconnect the RIB and the FIB before the migration
- Reconnect them together after the migration

A lot of research challenges still remain ...

Seamless Network-Wide Migrations

IGP migrations

Ordering matters

BGP migrations

Ongoing work

3 How can we help ?

Manage complexity

By using our techniques, we can help you achieve lossless migrations

Based on the initial and the final configurations, we can

IGP	Compute the operational ordering you should follow in order to not loose any packets
BGP	Evaluate the damages of the reconfiguration and work on ad-hoc solutions by modeling the network behavior

Please, come and talk to me if you are interested

Seamless Network-Wide Migrations

IGP migrations

Ordering matters

BGP migrations

Ongoing work

How can we help ?

Manage complexity

Don’t fear network reconfiguration, adapt the network to its environment

Seamless reconfigurations are possible
but they require careful planning

Our techniques provide theoretical guarantees
that every reconfiguration step is safe

Add flexibility to your network management
seamlessly move to the current best configuration

Seamless Network-wide Migrations

towards more agile networking

Laurent Vanbever

inl.info.ucl.ac.be/lvanbeve/

Network Automation 2011

Novembre 9, 2011