

# GCC Peering Status & Peering curiosities

Bernd Spiess

[bernd.spiess@uae-ix.net](mailto:bernd.spiess@uae-ix.net)

+43 676 848267 401

**UAE-IX**   
powered by DE-CIX

## How to create a “perfect” network for whole GCC

Target:

- Keep traffic local
- Short RTT in region

How could such a (theoretical) network look like?

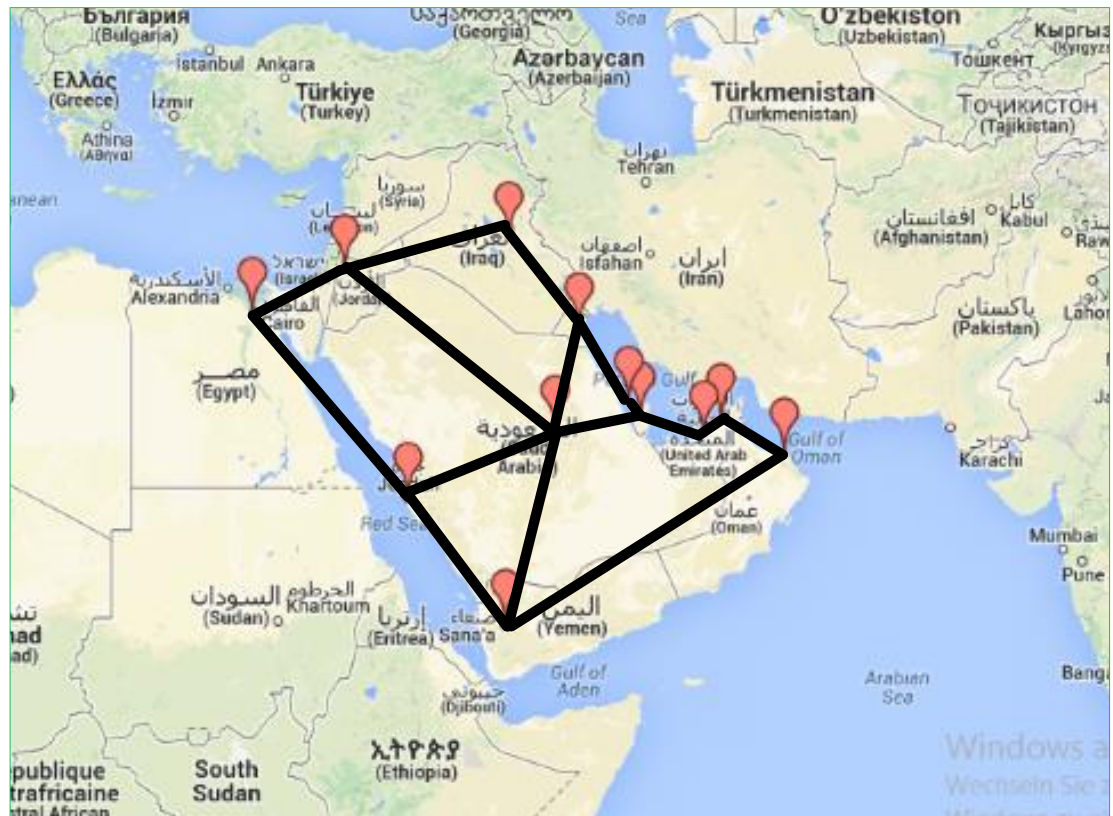


## How to create a “perfect” network for whole GCC

Solution:

- Invent a “GCC supercarrier”
- Draw as many lines as possible...

Now lets bring in some interconnections



## How to create a “perfect” network for whole GCC

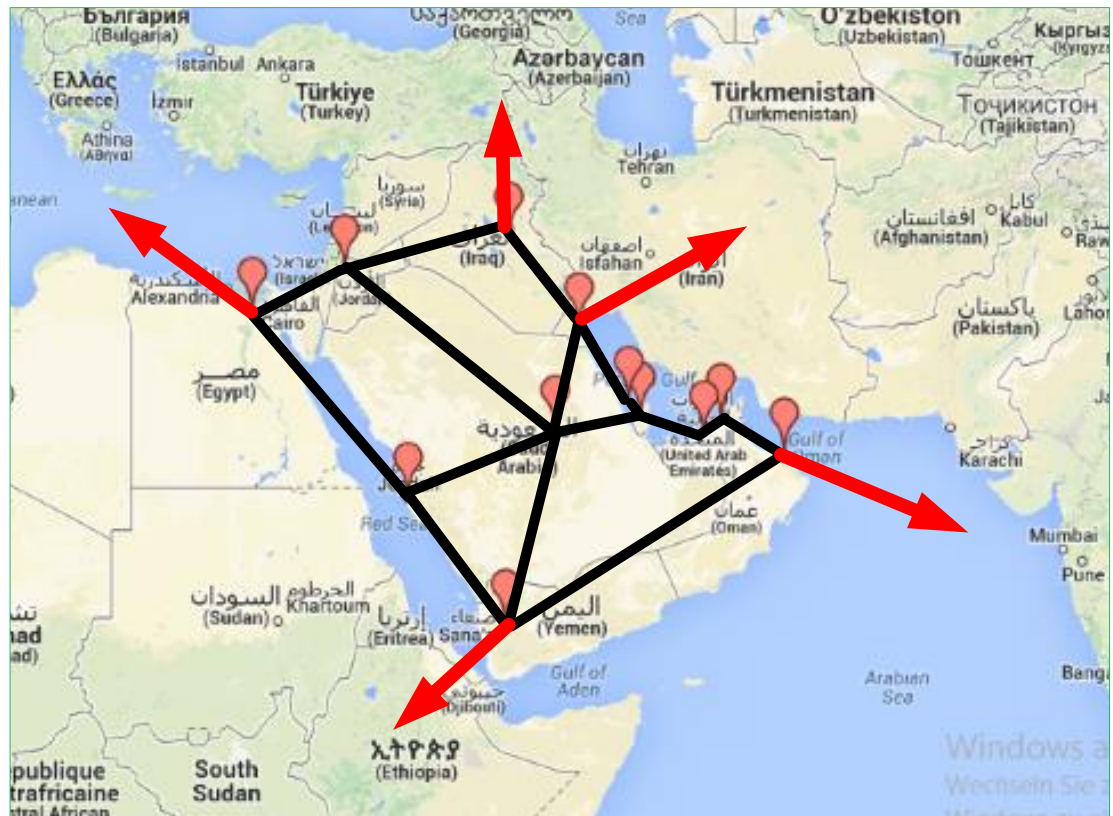
Interconnection?

Solution:

- Follow the logic...

Now lets bring in some content –

Where should the central content hub for GCC be located?

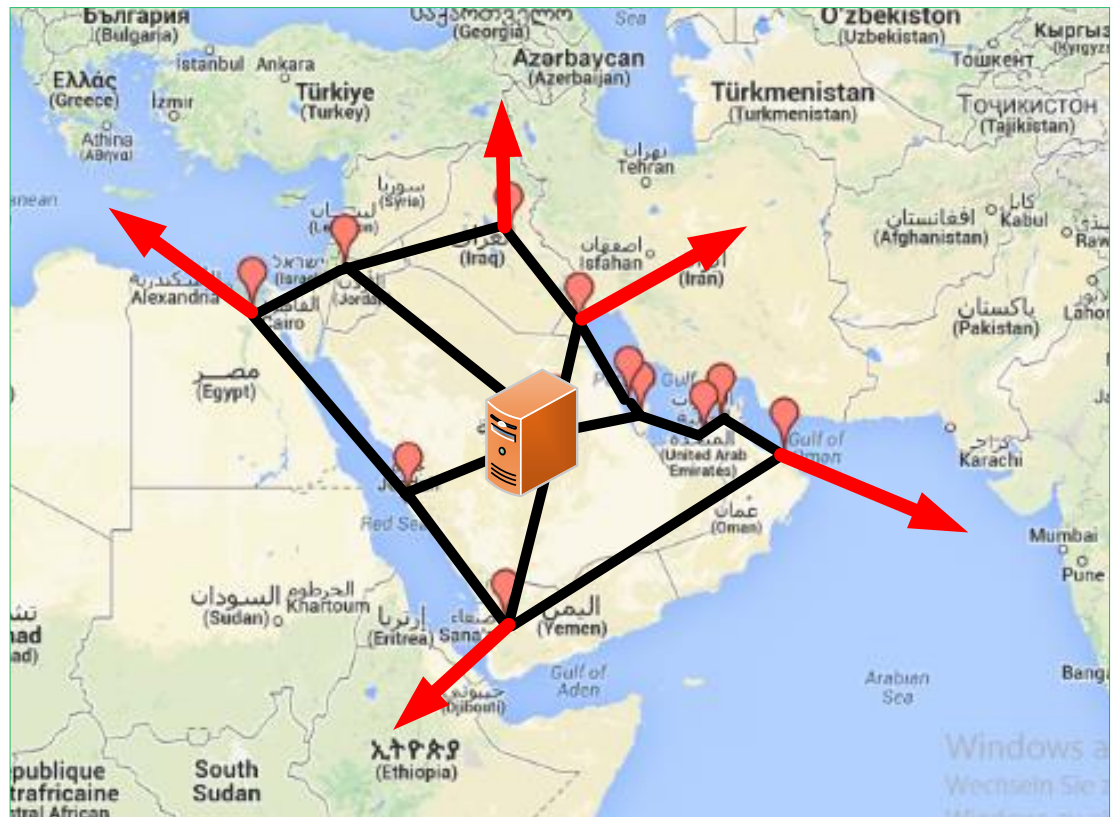


## How to create a “perfect” network for whole GCC

Solution:

- of course ... it would be in the middle... (in this theoretic “supercarrier model”)

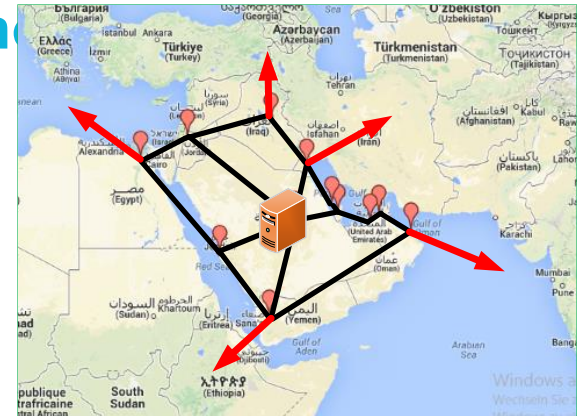
Now bringing up some questions...



## How to create a “perfect” network for who

### Questions

- Would this network be fast + low RTT?
- What is the name of this “Supercarrier”?
- Political possible?



=> conclusion: technical nice – would work nowhere

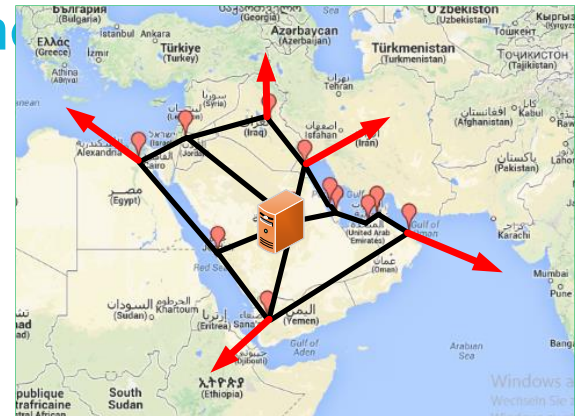
## How to create a “perfect” network for who

So Idea #1 was crap...

A quick Idea #2 to achieve low latency for the whole region:

Every ISP in GCC choose the same Tier1 IP-Transit Provider and this one does run his network as shown in example #1...

... better now ... ?



## How to create a “perfect” network for whole GCC

So lets have a look where the real challenge is located...

... and keep the idea with full mesh and low RTT in mind ...




## All needed peering in GCC is already done?

1) Saudi Arabia

= 4 players

= 6 privat interconnects for  
full mesh peering



## All needed peering in GCC is already done?

2) UAE

2 players

= 1 privat interconnect for  
full mesh peering



## All needed peering in GCC is already done?

Skipping pictures for:

- Kuwait (4-7 players)
- Yemen (1 player)
- Oman (2 player)
- Qatar (2 player)
- Bahrain (~5 player)

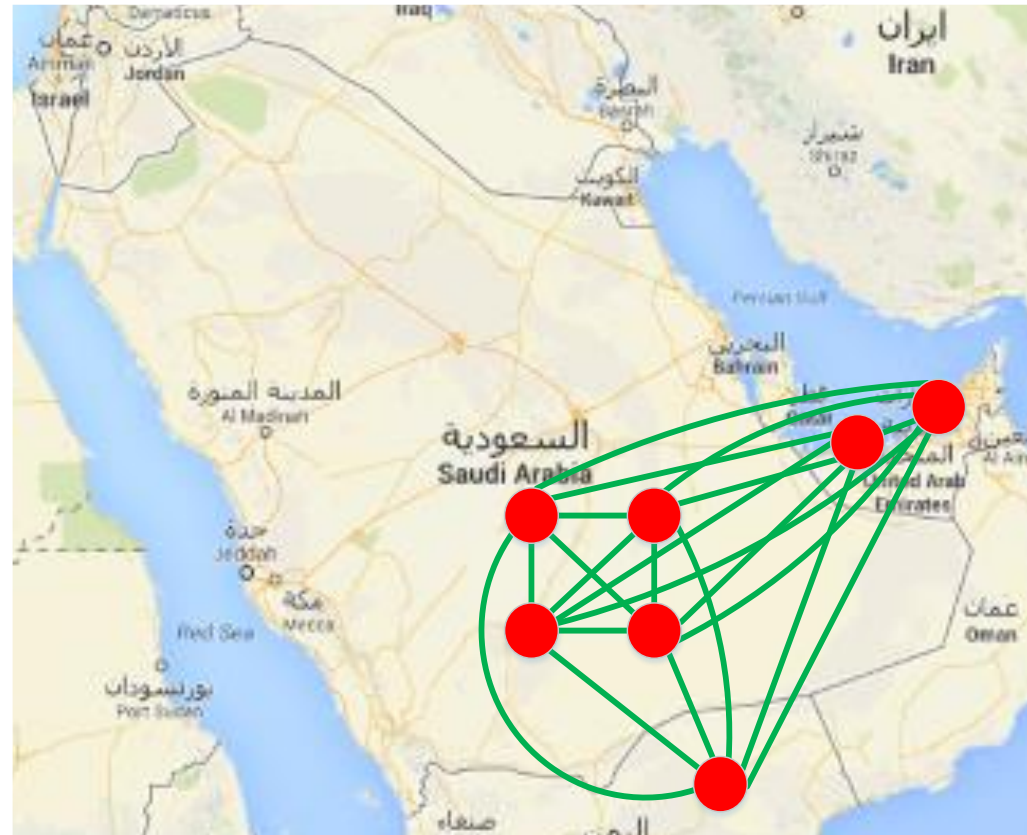
## All needed peering in GCC is already done?

Now start to bring this networks together...  
(reminder: target: low RTT)

First: full mesh:

Saudia Arabia + UAE + Yemen

Means: 21 direct links for a full mesh

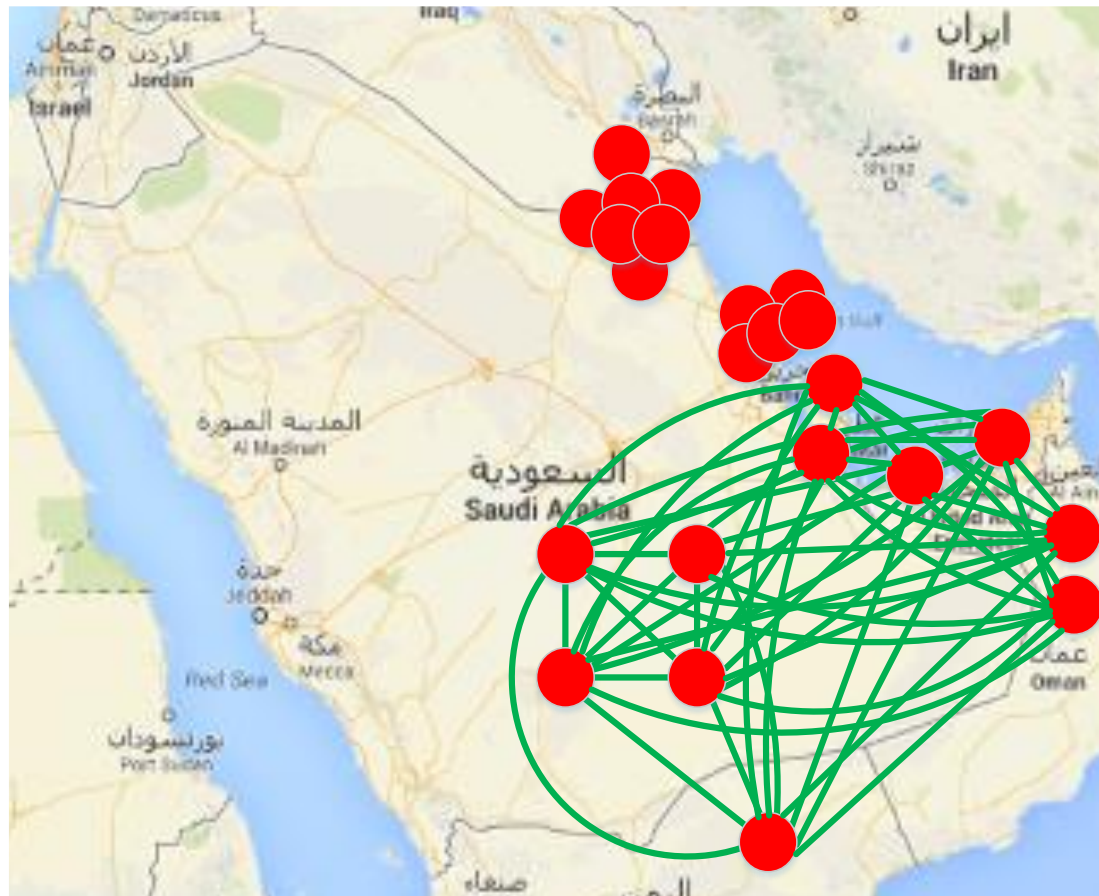


## All needed peering in GCC is already done?

Second: full mesh:  
add Oman + Qatar

**Skip lines for Kuwait and  
Bahrain**

Means: 23 players =  
253 interconnects



## All needed peering in GCC

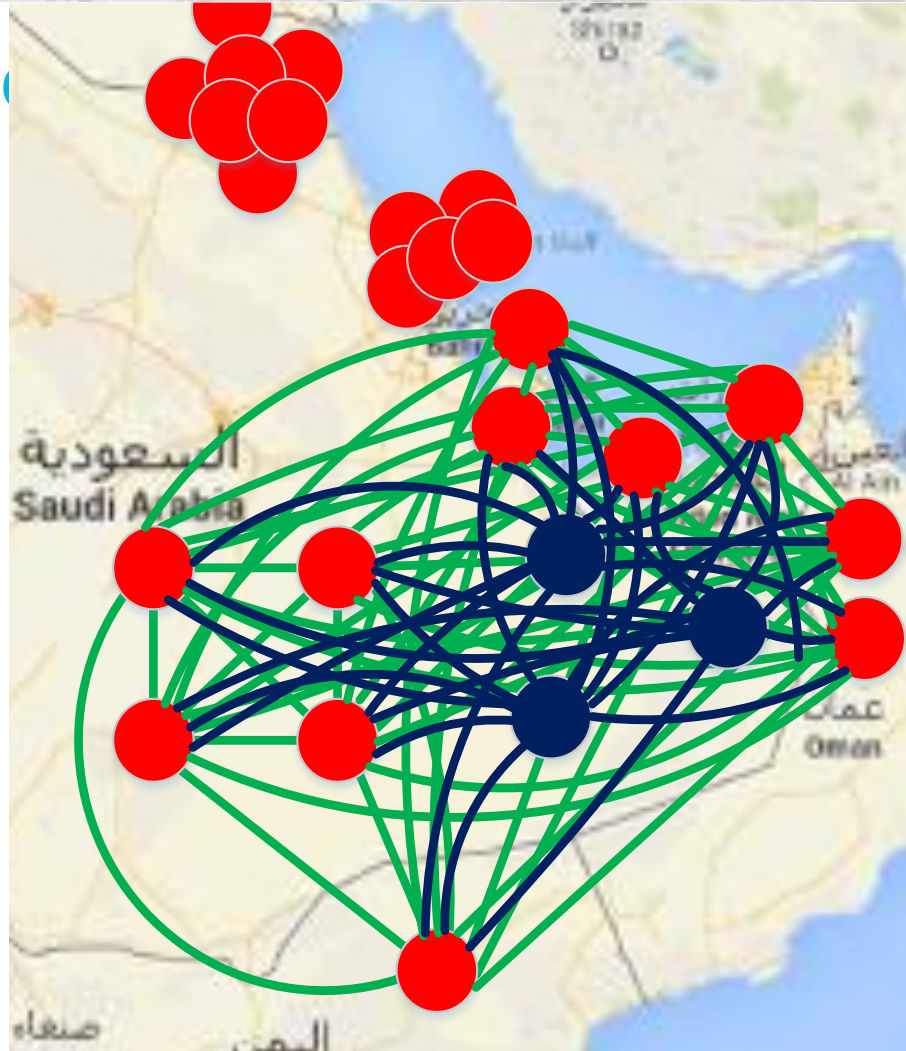
Third: full mesh:

Add 3 content players

(still without Bahrain/Kuwait...)

... add those lines in your  
Head...)

...325 interconnects...



## All needed peering in GCC is already done?

4<sup>th</sup>: full mesh:

Adding gaming, more content, more players, international players

=> No picture of that is possible

means:

- 10 players = 45 private (10 Gbit ?) interconnections
- 50 players = 1.225 interconnections
- 100 players = 4.950 interconnections

Number of players in GCC is > 100 !

## All needed peering in GCC is already done?

Conclusio:

- Does not scale
- Nobody would pay
- => full Mesh is not possible



# So what is the outcome?

Internal full mesh not possible

=

The existing internal private peerings

+ the Tier1 Upstreams has to do the low RTT job “inner GCC”

But: how does this look like today?

## Test 1:

Try: Tracing **from Dubai to Dubai**

Side A: p... lookingglass

Side B: ip of i... in dubai

**Expected RTT: 1-2 msec**

## comparings: dubai to dubai



	Abort / Remo
traceroute ip 84.233.221.50	
Tracing the route to Gi0-3.dxb-003-access-3.interoute.net (84.233.221.50)	
1 ge5-0-1.var01.dub01.pccwbtn.net (63.218.176.66)	0 msec 0 msec 0 msec
2 pos4-6.cr03.ldn01.pccwbtn.net (63.218.176.38)	136 msec 136 msec 136 n
3 TenGE11-2.br02.ldn01.pccwbtn.net (63.218.12.146)	136 msec 136 msec
4 xe-11-1-1.lon21.ip4.tinet.net (77.67.94.153)	136 msec 136 msec 136 msec
5 xe-11-3-0.par72.ip4.tinet.net (141.136.111.246)	144 msec
xe-2-2-2.par72.ip4.tinet.net (141.136.111.250)	148 msec 144 msec
6 interoute-gw.ip4.tinet.net (77.67.75.238)	148 msec 144 msec 144 msec
7 ae1-0.mrs-001-score-1-re0.interoute.net (217.118.118.74)	156 msec 156 r
8 Gi0-1.mrs-boi-access-2.interoute.net (217.118.118.86)	160 msec 160 msec
Gi0-3.mrs-boi-access-2.interoute.net (217.118.118.82)	160 msec
9 so-1-0-0-0.dxb-003-access-1-re1.interoute.net (84.233.221.41) [MPLS: La	
10 ge-0-0-0-0.dxb-003-access-2-re1.interoute.net (84.233.221.30)	260 msec
11 Gi0-3.dxb-003-access-3.interoute.net (84.233.221.50)	264 msec * 260 m
Query Complete	


## Test 2:

Try: from Jeddah to Jeddah (~ 0 km)

Side A: t... lookingglass

Side B: L... – router ip in jeddah

Expected RTT: 1 msec



## comparings:

## jeddah to jeddah

- 1 195.219.167.42 [MPLS: Label 2414 Exp 0] 52 msec
- if-2-0-0.core1.WYN-Marseille.as6453.net (80.231.165.77) [MPLS: Label 2414 Exp 0] 48 msec
- 195.219.167.42 [MPLS: Label 2414 Exp 0] 48 msec
- 2 if-1-2-1-29.tcore2.WYN-Marseille.as6453.net (80.231.200.21) [MPLS: Label 616320 Exp 0] 48 msec
- 3 if-9-2.tcore2.L78-London.as6453.net (80.231.200.14) [MPLS: Label 682656 Exp 0] 48 msec
- 4 if-2-2.tcore1.L78-London.as6453.net (80.231.131.2) [MPLS: Label 616320 Exp 0] 48 msec
- if-1-2.tcore1.L78-London.as6453.net (80.231.130.121) [MPLS: Label 616320 Exp 0] 48 msec
- if-2-2.tcore1.L78-London.as6453.net (80.231.131.2) [MPLS: Label 616320 Exp 0] 48 msec
- 5 if-17-2.tcore1.LDN-London.as6453.net (80.231.130.130) 72 msec 72 msec
- 6 195.219.83.102 72 msec 72 msec 72 msec
- 7 vl-3503-ve-117.csw1.London1.Level3.net (4.69.166.137) [AS 3356] 80 msec
- vl-3502-ve-116.csw1.London1.Level3.net (4.69.166.133) [AS 3356] 80 msec
- vl-3504-ve-118.csw1.London1.Level3.net (4.69.166.141) [AS 3356] 72 msec
- 8 ae-56-111.ebr1.London1.Level3.net (4.69.153.113) [AS 3356] 72 msec
- ae-58-113.ebr1.London1.Level3.net (4.69.153.121) [AS 3356] 76 msec
- ae-57-112.ebr1.London1.Level3.net (4.69.153.117) [AS 3356] 80 msec
- 9 ae-46-46.ebr1.Paris1.Level3.net (4.69.143.106) [AS 3356] 68 msec 68 msec
- ae-47-47.ebr1.Paris1.Level3.net (4.69.143.110) [AS 3356] 72 msec
- 10 ae-81-81.csw3.Paris1.Level3.net (4.69.161.86) [AS 3356] 68 msec
- ae-61-61.csw1.Paris1.Level3.net (4.69.161.78) [AS 3356] 68 msec
- ae-91-91.csw4.Paris1.Level3.net (4.69.161.90) [AS 3356] 72 msec
- 11 ae-93-93.ebr3.Paris1.Level3.net (4.69.161.121) [AS 3356] 68 msec 68 msec
- ae-73-73.ebr3.Paris1.Level3.net (4.69.161.113) [AS 3356] 68 msec
- 12 \*
- ae-1-7.bar1.Marseille1.Level3.net (4.69.143.237) [AS 3356] 80 msec 84 msec
- 13 ae-5-6.bar1.Jeddah1.Level3.net (4.69.143.157) [AS 3356] 128 msec 128 msec



## Test 3:

Try: from F1... Kuwait to Dubai

Side A: F1... Kuwait (T.../H.../E...)

Side B: P... Dubai

**Expected RTT: 50-100 msec**

## Comparings: Dubai to Kuwait

tracertoute ip 62.215.0.98

Tracing the route to tec-sw1-g216.fasttelco.net

1 ge5-0-1.var01.dub01.pccwbtn.net (63.218.17

2 pos4-6.cr03.lhn01.pccwbtn.net (63.218.176.3

3 63.218.222.33 212 msec 208 msec 208 msec

4 63-218-223-46.static.pccwglobal.net (63.218.

5 ae-7.r23.nycmny01.us.bb.gin.ntt.net (129.250

ae-1.r23.nycmny01.us.bb.gin.ntt.net (129.250.4

6 ae-9.r20.asbnva02.us.bb.gin.ntt.net (129.250

7 \*

ae-0.r21.asbnva02.us.bb.gin.ntt.net (129.250.4.5) [MPLS: Label 550724 Exp 0] 212 msec 210 msec

8 ae-4.r20.snjsca04.us.bb.gin.ntt.net (129.250.4.103) [MPLS: Label 573344 Exp 0] 284 msec 292 msec 284 msec

9 ae-2.r20.sngpsi05.sg.bb.gin.ntt.net (129.250.3.49) [MPLS: Label 740369 Exp 0] 288 msec 300 msec 280 msec

10 ae-4.r21.sngpsi02.sg.bb.gin.ntt.net (129.250.3.183) 264 msec 288 msec 272 msec

11 116.51.17.246 368 msec

116.51.27.14 364 msec 376 msec

12 89.211.1.165 372 msec 368 msec 384 msec

Query Complete



## Test 4:

Try: from Dubai to Dubai

Side A: P... Dubai

Side B: F2... Dubai

**Expected RTT: 1-2 msec**



## Comparings: Dubai to Dubai



tracertoe ip 62.216.134.193

Tracing the route to lo0.cjr01.dxb001.flagtel.com (62.216.134.193):

- 1 ge5-0-1.var01.dub01.pccwbtn.net (63.218.176.66) 4 ms
- 2 pos4-6.cr03.ldn01.pccwbtn.net (63.218.176.38) 136 msec
- 3 TenGE11-2.br02.ldn01.pccwbtn.net (63.218.12.146) 136 msec
- 4 peer2.ldn1.flagtel.com (195.66.236.146) 136 msec 148 ms
- 5 ge-0-1-0-0.pjr01.ldn001.flagtel.com (62.216.129.137) 136 msec
- 6 62.216.128.113 136 msec
- 7 xe-11-1-0-0.pjr03.dxb001.flagtel.com (85.95.26.189) 224 msec
- 8 lo0.cjr01.dxb001.flagtel.com (62.216.134.193) 224 msec

Query Complete

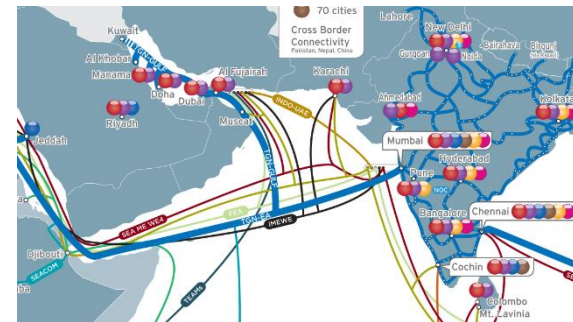
## Test 5:

Try: from P... Dubai to T... Jeddah

Side A: P... Dubai

Side B: T... Jeddah

Expected RTT: 50-100 msec



## Comparings: Dubai to Jeddah



tracertoute ip 80.231.165.78

Tracing the route to if-6-0.core1.JSD-Jeddah.as6453.

1 ge5-0-1.var01.dub01.pccwbtn.net (63.218.176.66) (

2 pos4-6.cr03.ldn01.pccwbtn.net (63.218.176.38) 136

3 tenge13-3.br02.ldn01.pccwbtn.net (63.218.12.246)

4 208.178.62.197 192 msec 136 msec 136 msec

5 tata-1.ar7.lon3.gblx.net (64.215.195.238) 136 msec

6 if-15-2.tcore2.L78-London.as6453.net (80.231.131.

7 if-9-2.tcore2.WYN-Marseille.as6453.net (80.231.200.

8 if-6-0.core1.JSD-Jeddah.as6453.net (80.231.165.78)

Query Complete

## Test 6:

Try: from Dubai to Jeddah

Side A: P... Dubai

Side B: L... Jeddah

Expected Runtime: 50-100 msec



## Comparings: Dubai to Jeddah

traceroute ip 4.69.148.222

```
Tracing the route to ae-1-11.bar2.Jeddah1.Level3.net
1 ge5-0-1.var01.dub01.pccwbtn.net (63.218.176.66)
2 pos4-6.cr03.lhn01.pccwbtn.net (63.218.176.38) 15
3 tenge13-3.br02.lhn01.pccwbtn.net (63.218.12.246)
4 208.178.62.197 172 msec 272 msec 216 msec
5 4.68.110.157 180 msec 136 msec 136 msec
6 vl-3503-ve-117.csw1.London1.Level3.net (4.69.166.166)
vl-3504-ve-118.csw1.London1.Level3.net (4.69.166.166)
7 ae-59-114.ebr1.London1.Level3.net (4.69.153.121)
ae-56-111.ebr1.London1.Level3.net (4.69.153.113)
8 ae-58-113.ebr1.London1.Level3.net (4.69.153.121)
ae-47-47.ebr1.Paris1.Level3.net (4.69.143.114) 15
ae-45-45.ebr1.Paris1.Level3.net (4.69.143.102) 144
9 ae-61-61.csw1.Paris1.Level3.net (4.69.161.78) 144
```

Query Complete



## Test 7:

Try: from P... Dubai to Oman

Side A: P... Dubai

Side B: O... Oman

**Expected RTT: 30-50 msec**

## Comparings: Dubai to Oman

tracertool ip 82.178.72.20

Tracing the route to 82.178.72.20

```
1 ge5-0-1.var01.dub01.pccwbtn.net (63.218.176.66) 0 msec 0
2 pos4-6.cr03.lhn01.pccwbtn.net (63.218.176.38) 136 msec 136
3 omantel.pos2-2.cr02.nyc02.pccwbtn.net (63.218.109.118) 352 msec 352
4 82.178.33.101 496 msec 284 msec 284 msec
5 82.178.33.105 344 msec 344 msec 344 msec
6 82.178.159.198 348 msec 344 msec 348 msec
7 212.72.10.129 372 msec 360 msec 360 msec
8 ***
9 ***
10 82.178.127.25 344 msec 344 msec 348 msec
11 82.178.72.20 348 msec 344 msec 348 msec
12 ***
13 ***
14 ***
15 ***
```

Query Complete



## Conclusio:

- ! All used example networks are great carriers / operators
- This way of interconnecting GCC is great for Europe –  
GCC is financing Europes low transit costs story!



## So what is the outcome?

Example 2:

Very high RTT to the DNS ROOT Servers

## A view at route-servers.net...

- i-root
- F-root
- K-root
- L-root



## A view at route-servers...

RTT seen from Dubai

Good: K F I L all low

What about the other 9?

Btw: where is IPv6 ??

You have it all running?

Probe ID: 2243

Name Server	Response time	Last update
K (IPv4)	15.8 ms	2013-08-24 09:31:46 UTC
F (IPv4)	20.4 ms	2013-08-24 09:31:51 UTC
I (IPv4)	26.2 ms	2013-08-24 09:31:39 UTC
L (IPv4)	26.3 ms	2013-08-24 09:31:54 UTC
A (IPv4)	189.1 ms	2013-08-24 09:32:11 UTC
D (IPv4)	224.9 ms	2013-08-24 09:32:19 UTC
J (IPv4)	274.8 ms	2013-08-24 09:32:45 UTC
M (IPv4)	281.4 ms	2013-08-24 09:31:52 UTC
E (IPv4)	286.7 ms	2013-08-24 09:32:30 UTC
B (IPv4)	289.2 ms	2013-08-24 09:32:12 UTC
C (IPv4)	292.0 ms	2013-08-24 09:32:15 UTC
H (IPv4)	314.4 ms	2013-08-24 09:32:42 UTC
G (IPv4)	319.4 ms	2013-08-24 09:32:26 UTC

## So what is the outcome?

Example 3:


Very high RTT to major content (CDN)

## round trip time to content

This is the situation in GCC

(taken November 2013)

seen from Riad



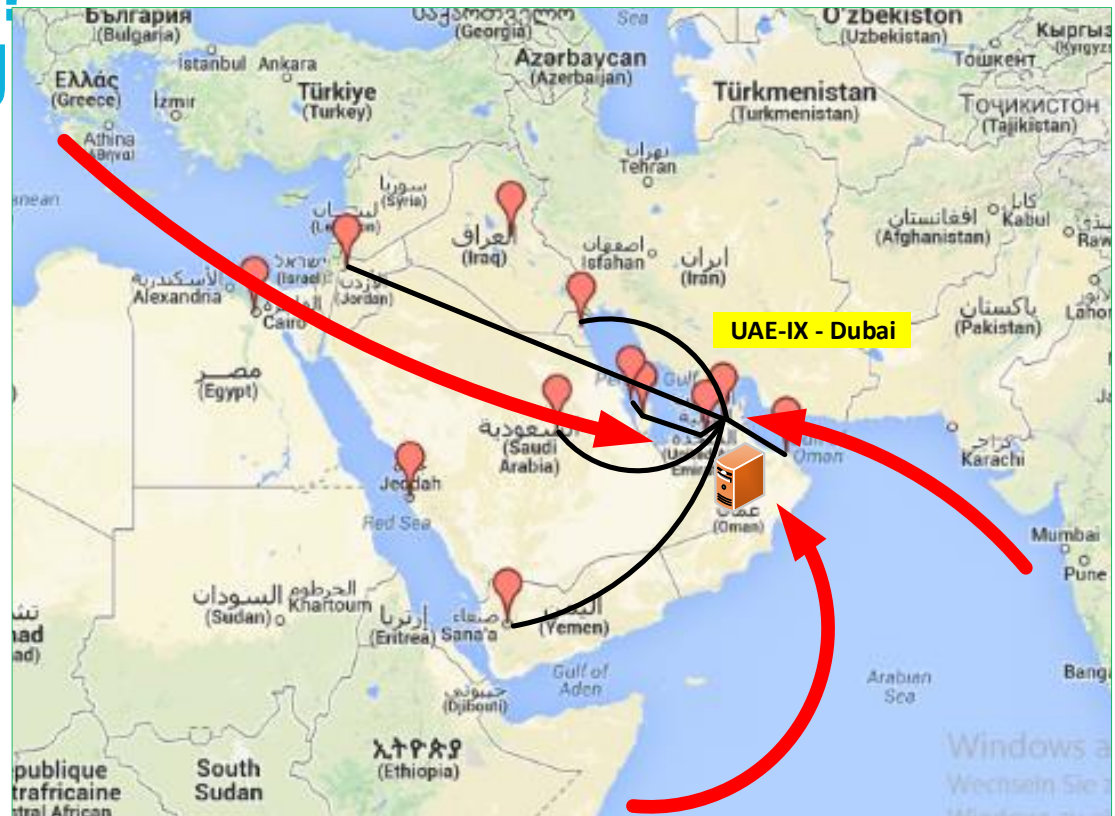
## So why am I standing here? What is the solution? (Att: Marketing)

Regional Peering does solve the RTT problem...

International content comes into the region

Regional content comes back to the region.

International player are connecting on their cost!



# Showing some peering curiosities

## Showing some peering curiosities:

### Example 1: Missing Prefixes at routeservers

#### DUBAI

Neighbor	AS	PfxRcd
...		
185.1.8.14	42	44
...		

#### FRANKFURT

Neighbor	AS	PfxRcd
...		
80.81.194.42	42	50
...		



## Showing some peering curiosities:

Example 2: Having RS configured, but missing all prefixes

### DUBAI

Neighbor	AS	PfxRcd
...		
185.1.8.xx	xxxxx	0
...		

AS Info | Graph v4 | Graph v6 | Prefixes v

Company Website:

Country of Origin:

Prefixes Originated (all): 1,364  
Prefixes Originated (v4): 1,313  
Prefixes Originated (v6): 51

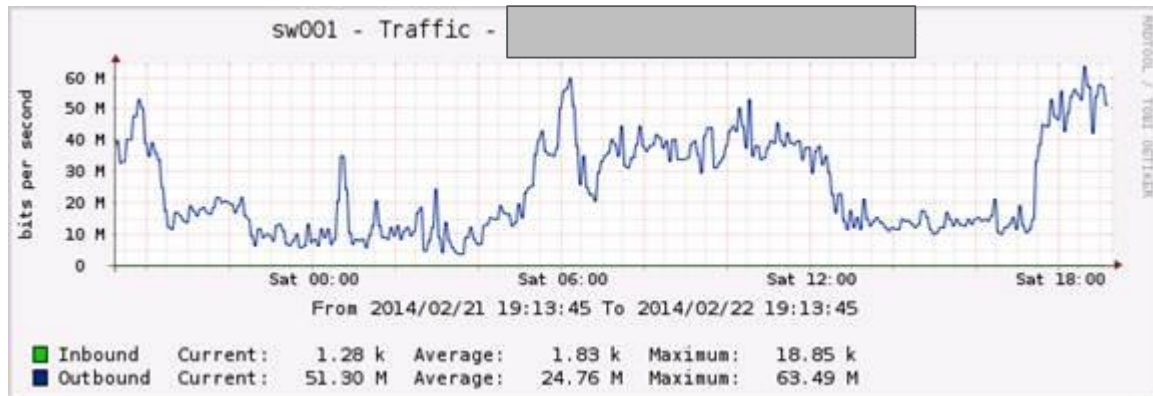
Prefixes Announced (all): 1,369  
Prefixes Announced (v4): 1,313  
Prefixes Announced (v6): 56

## Showing some peering curiosities:

Example 3: make no use of learnt prefixes

### DUBAI

Neighbor	AS	PfxRcd
...		
185.1.8.yy	yyyy	zz
...		



### ATTENTION:

This leads to heavy asymmetric RTT picture

## Showing some peering curiosities:

### Example 4: missing RRDB entries drops prefixes

#### DUBAI - ROUTEVIEW SERVER

```
BGP routing table entry for 192.xx.yy.0/24
 8xxx 2yyyy, (received & used)
 185.1.8.xx from 185.1.8.xx
   Origin IGP, metric 0, localpref 100, valid, external, best
```

#### DUBAI - ROUTESERVER

```
% Network not in table
```


**THINK ON REGISTERING ALL PREFIXES IN RRDB's -  
ROUTESERVER KICK's UNREGISTERED PREFIXES!**

## Showing some peering curiosities:

Example 5: 10 Gbit peering port – but backhaul to small

10 Gbit Peering port

1.2 Gbit = 2 x STM-4



## Showing some peering curiosities:

### Example 6: multiple different sourced locations

Seen szenario in Frankfurt:

ROOT Server Prefix:

Announced from Carrier A with 2 AS path length from location FRANKFURT

Announced from Carrier B with 2 AS path length from location PARIS

Announced from Carrier C with 2 AS path length from location USA

Question to the audience: Which prefix source does win?

[www.uae-ix.net](http://www.uae-ix.net)

Bernd Spiess  
**bernd.spiess@uae-ix.net**

+43 676 848267 401

sales@uae-ix.net