renesys

Internet Infrastructure: Virtual meets Reality

James Cowie, CTO

@jimcowie
@renesys

Netnod Autumn Meeting
Stockholm, Sweden
3 October 2013

Physical Fragility

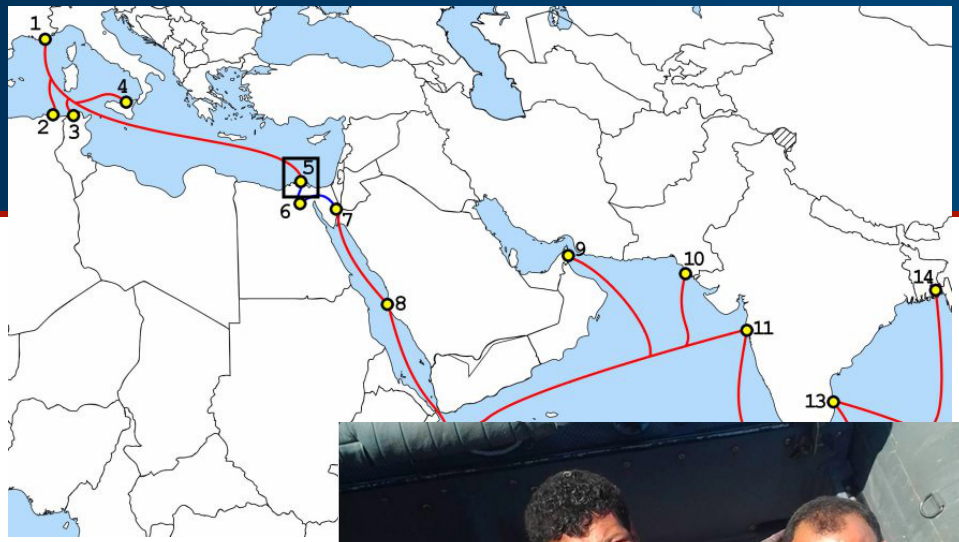

The Internet relies on locally fragile physical infrastructure:

- ✓ Submarine cables
- ✓ Terrestrial fiber networks
- ✓ Energy pipelines
- ✓ Power grids

renesys

The Internet “Breaks” Every Day

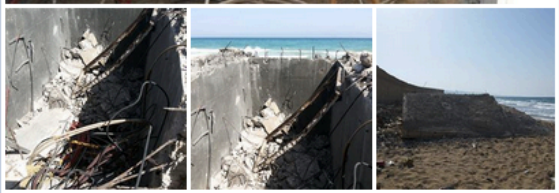

Like · Comment · Share 13

Like · Comment · Share 28

Omar Bendaher, Amore Ben Amer, حمد زوبي and 83 others like this. Top Comments

Sadam Swedan مشروع كابل السلفيوم هو مشروع غير مكتمل كان من المفترض ان يربط قبرص بليبيا عن طريق درنة وفي حالة اكتمل المشروع ستتضاعف سرعة الانترنت في ليبيا وكذلك ستتمكن من ليبيا من ادخال خدمة شبكة الالياف البصرية التي تصل سرعتها الي 10 ميغا في الثانية وان شاء الله ترجع الشركة المنفذة لاكمال المشروع
See Translation
1 · September 2 at 1:24pm

ثورة 17 فبراير ، درنه ، ليبيا replied · 2 Replies

محمد عباد الكوابل الالياف البصرية وليس الكهرباء وهي غرفة اللاتصالات الربط بي المنطقة الشرقية وهي اول غرفة رابط
See Translation
August 31 at 10:20am

ثورة 17 فبراير ، درنه ، ليبيا الكهرباء شورهم انارة الطريق الرئيسي... وليس لها علاقة بكابل الالياف البصرية... طريق الكورنيش كان طافي لا يسمح بالعمل حتى ساعات متاخرة وكذلك منطقة الغرفة مظلمة تسمح لأي مشبوه بالقيام بعمل تخريبي

Virtual Robustness

The Internet is a robust **virtual infrastructure** comprised of tens of thousands of communicating enterprises, each seeking to **maximize profit** according to local rules and business conditions

Credit: Tony Hisgett

renesys

How the Internet Survives and Flourishes

Designed for **simplicity**: rough consensus and running code, dumb core and smart edge, minimum viable interoperability.

renesys

How the Internet Survives and Flourishes

Designed for **simplicity**: rough consensus and running code, dumb core and smart edge, minimum viable interoperability.

Evolves toward **complexity**: more participants, more interconnection, more viable paths between arbitrary endpoints

renesys

Example: Crossover from Reality to Virtual

In this region, three geopolitical "Internet watersheds" meet

- Turkey
- Russia
- Iran

Europe's southern gas corridor clears the rights-of-way, Internet follows right behind

Credit: Al Jazeera

renesys

Oct 18th 2012: Iranian Internet Takes a Hit

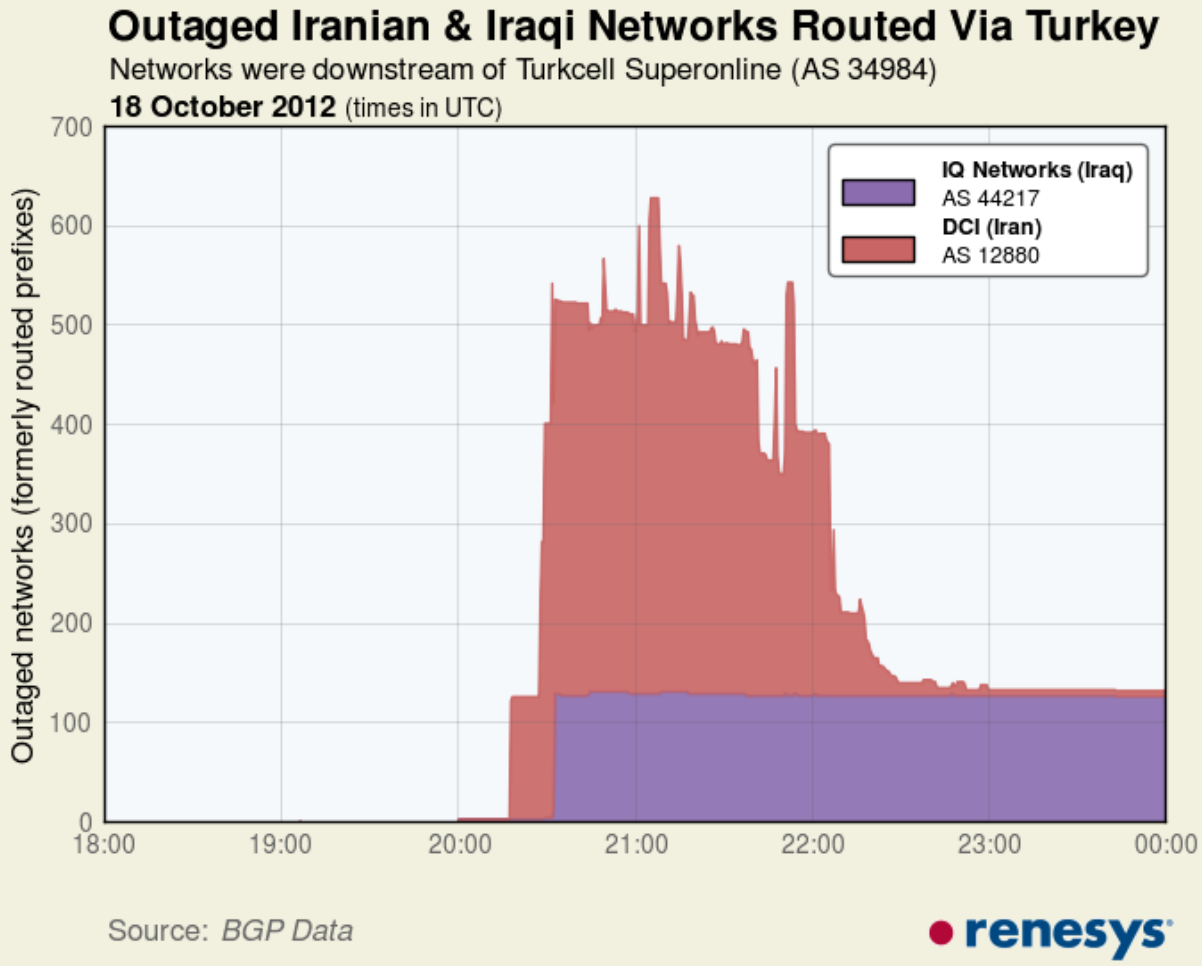

Iranian DCI loses Internet transit via Turkcell Superonline.

Iraq's IQ Networks also loses Internet transit via DCI.

Traffic re-routes.

Turkey gas pipeline blast ! 28 soldiers injured, Gazprom increases gas supplies to Turkey

19 October, 2012 | 14:47

Turkish Iranian Gas Pipeline sabotaged, Gazprom intervenes

28 Turkish soldiers were wounded as PKK terrorists blasted Iranian-Turkish natural gas pipeline between the eastern provinces of Ağrı and Erzurum. Monopolostic Gazprom supplied Turkey immediately with more gas.

Erzurum / NationalTurk – The gas flow through a pipeline carrying Iranian natural gas to Turkey was halted after the explosives planted by PKK terror organisation were detonated at around 8:30 p.m. as patrolling Turkish soldiers were nearby.

A fire broke out due to the explosion, but was subsequently extinguished after the flow of gas in the pipeline was cut off.

Active Measurement Confirms Outage

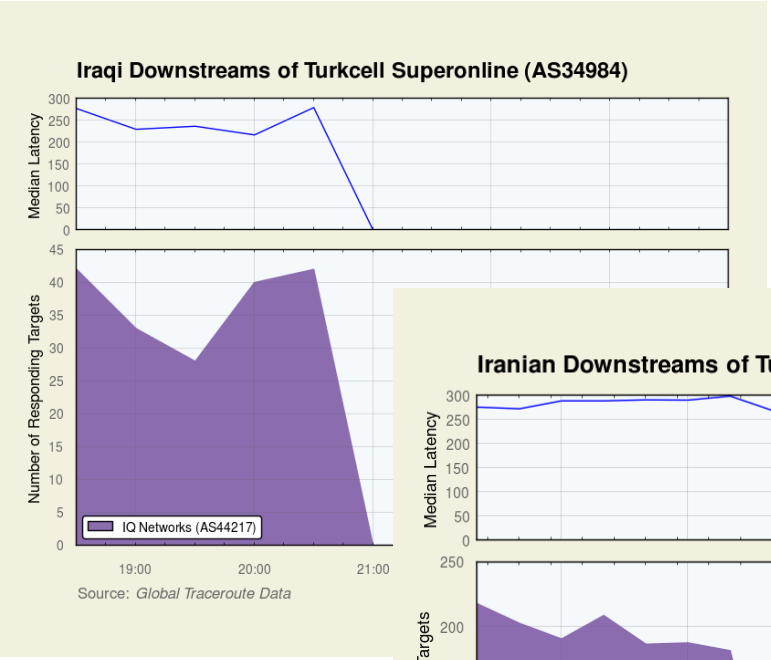

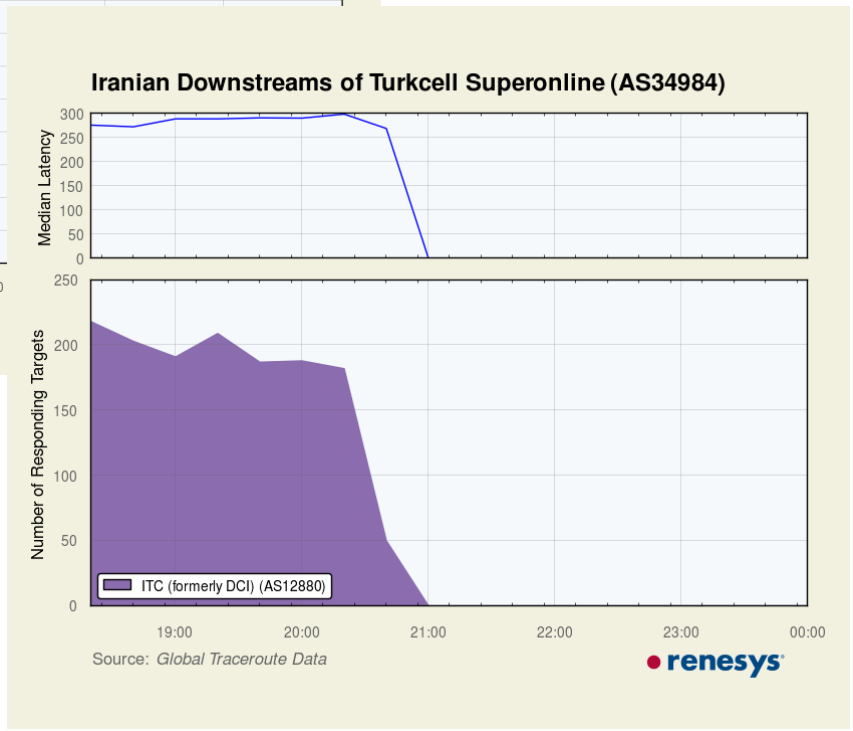

Iraqi, Iranian hosts within the affected networks downstream of Turkcell Superonline stop responding to traceroute via Turkish paths for several hours

renesys

Lasting Internet Impact: *None At All*

Colored bands indicate Iran's foreign Internet transit choices in October 2012

Traffic finds a restoration path, and the BGP-visible transit relationships are unaffected

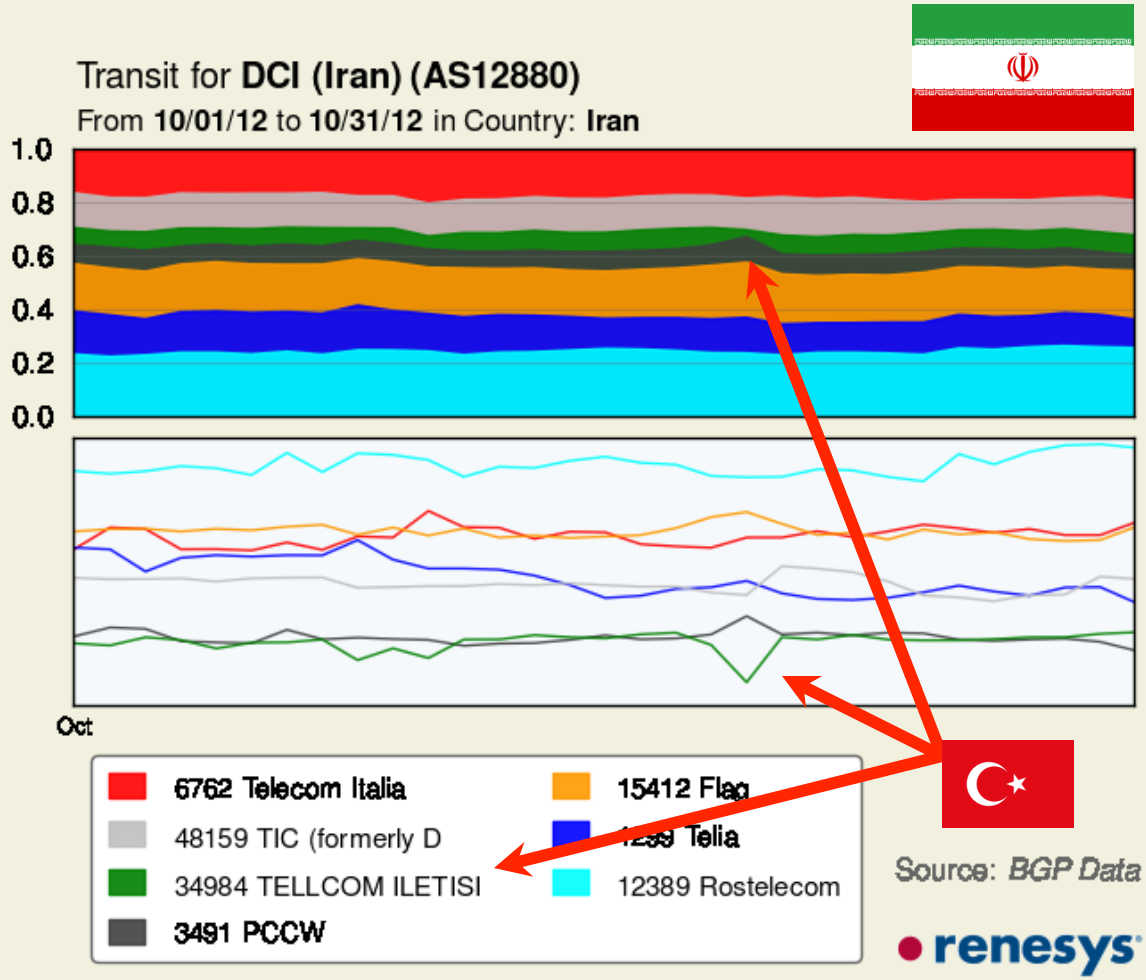

renesys

The Virtual Can Survive Physical Outages

This is what the Internet is good at: identifying damage and routing around it.

What damage can it *not* route around?

Syria, May 2013

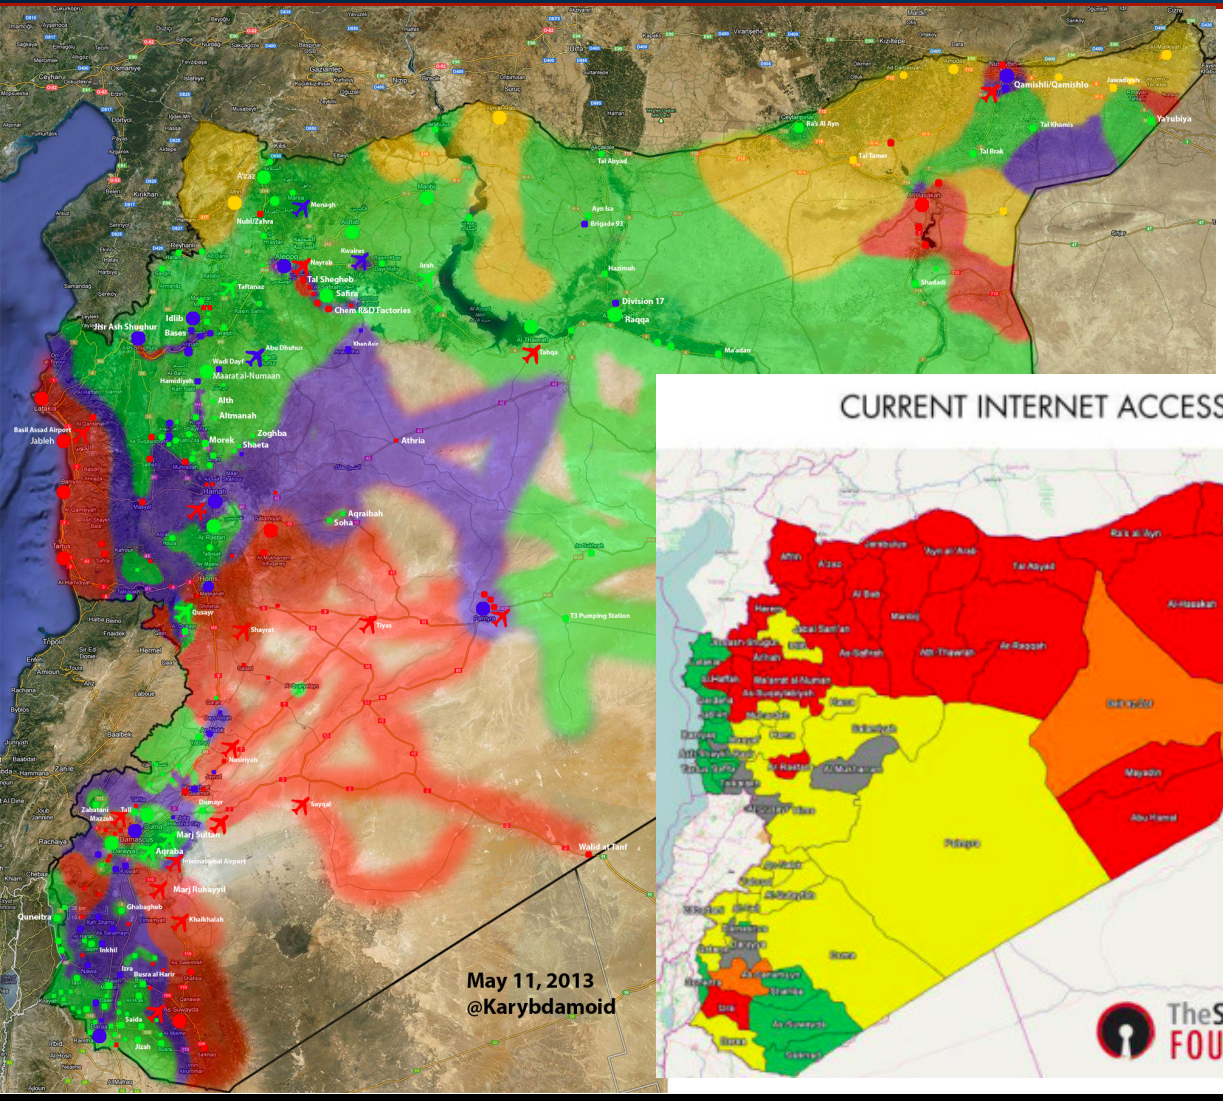

At left (11 May 2013):

- Govt control (red)
- Rebel control (green)
- Contested (purple)

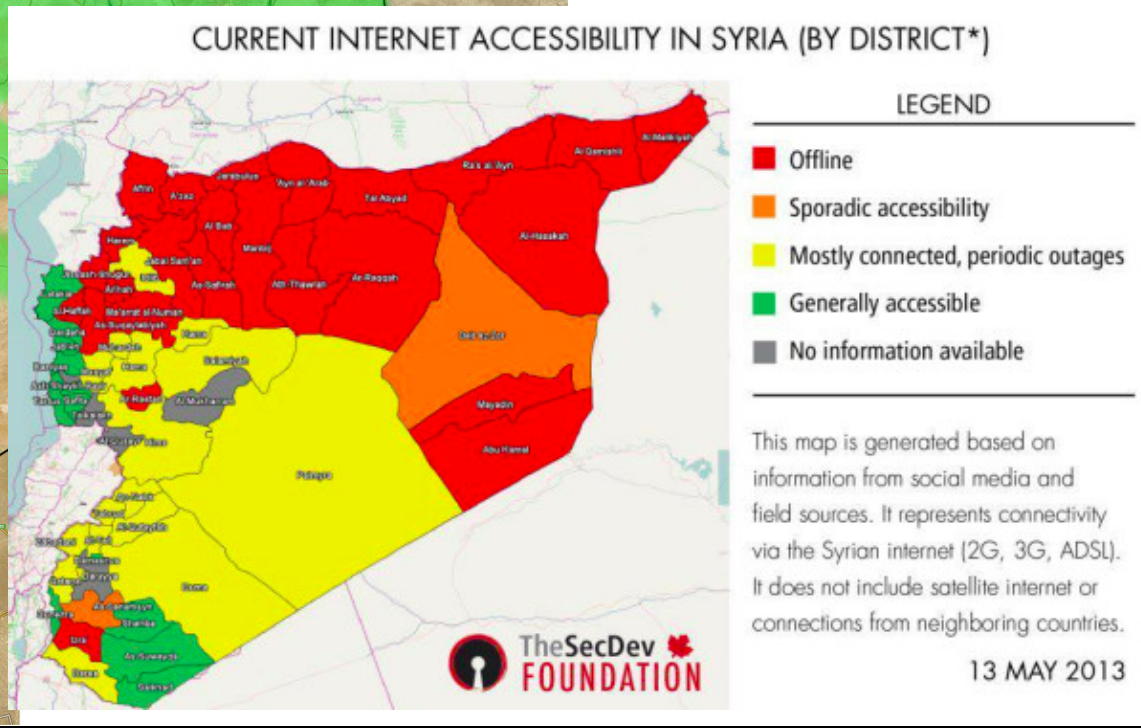

Nationwide outages still happen: Why?

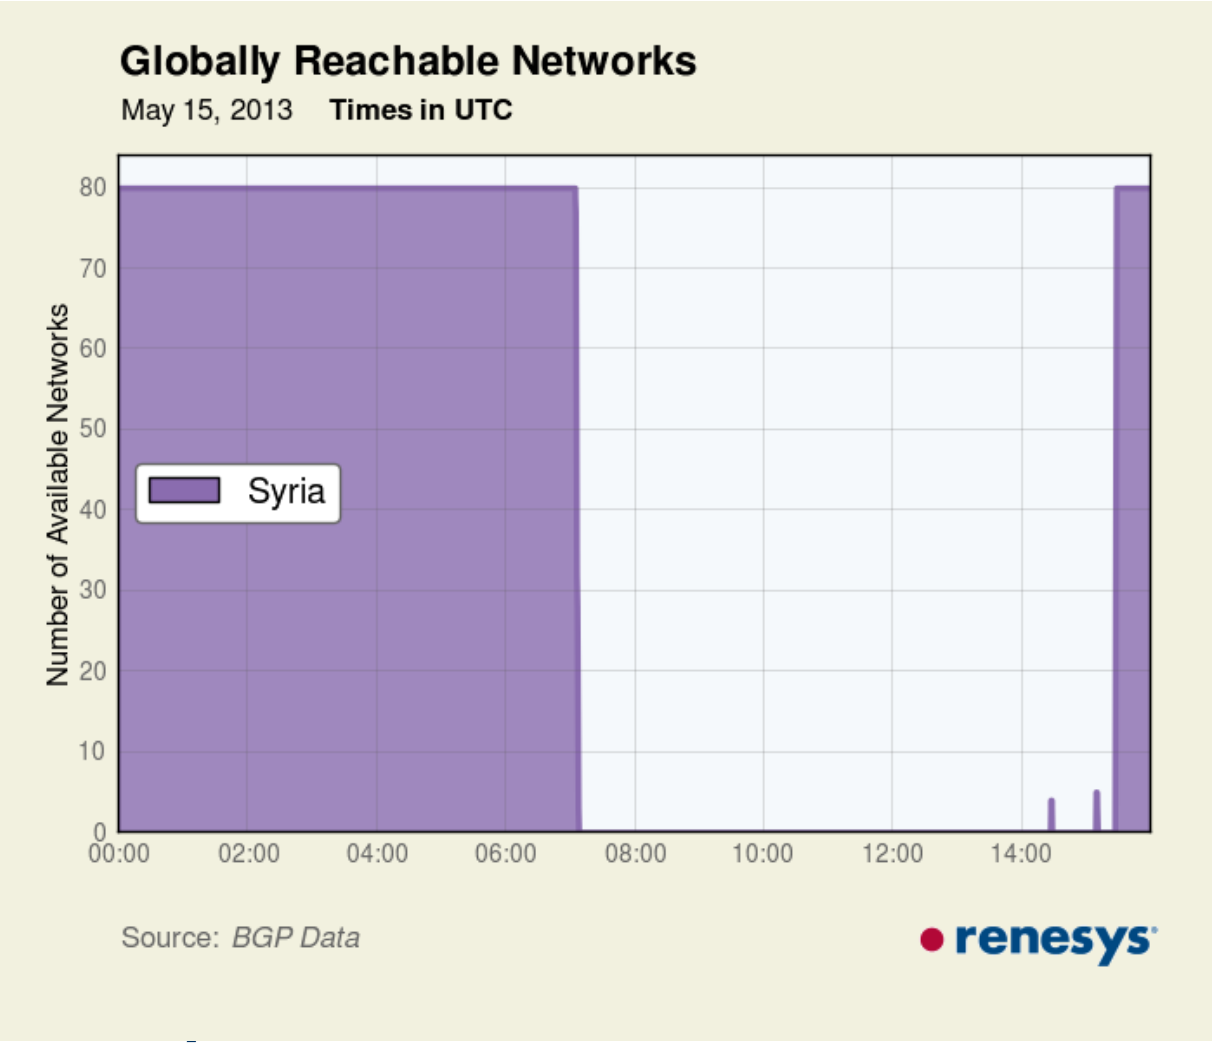

"Fiber damage 60km north of Damascus"

15 May 2013

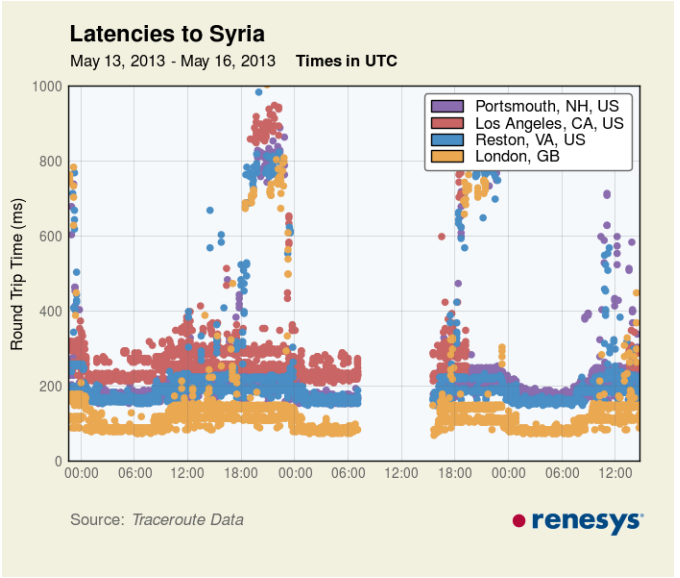

Nationwide outages still happen: Why?

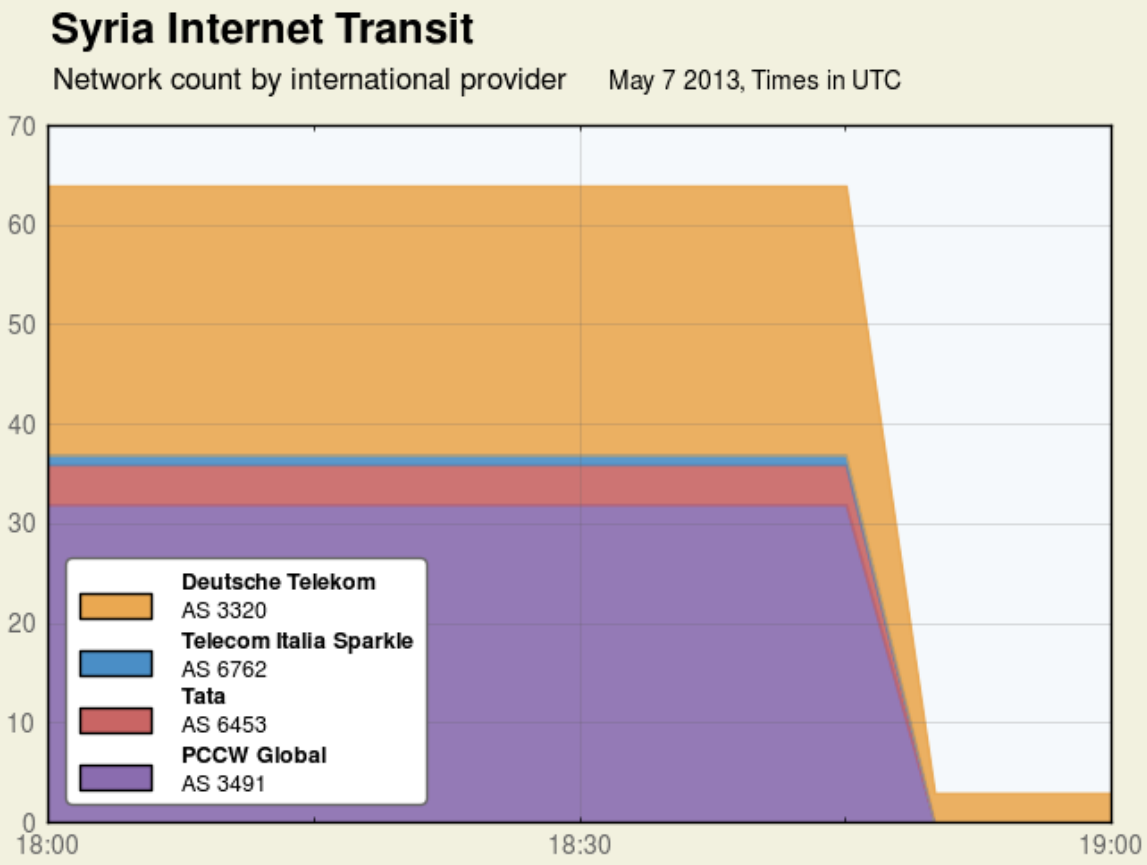

"Fault on Fiber Optic Cables"

7 May 2013

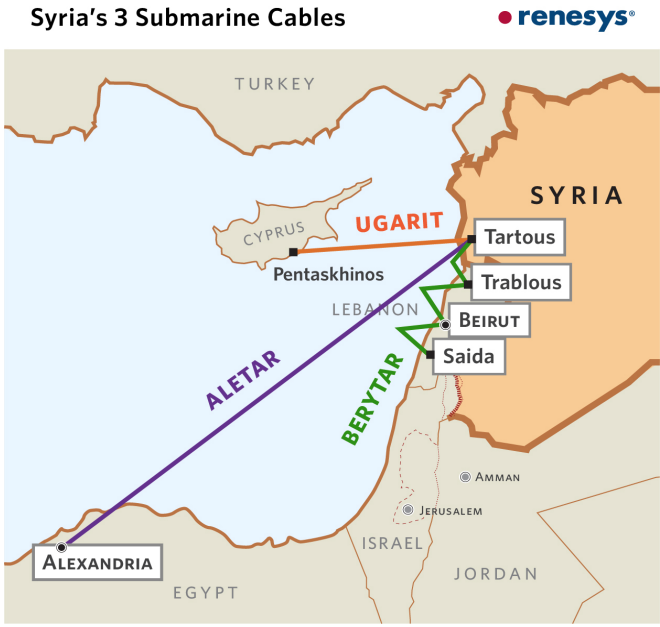

Nationwide outages still happen: Why?

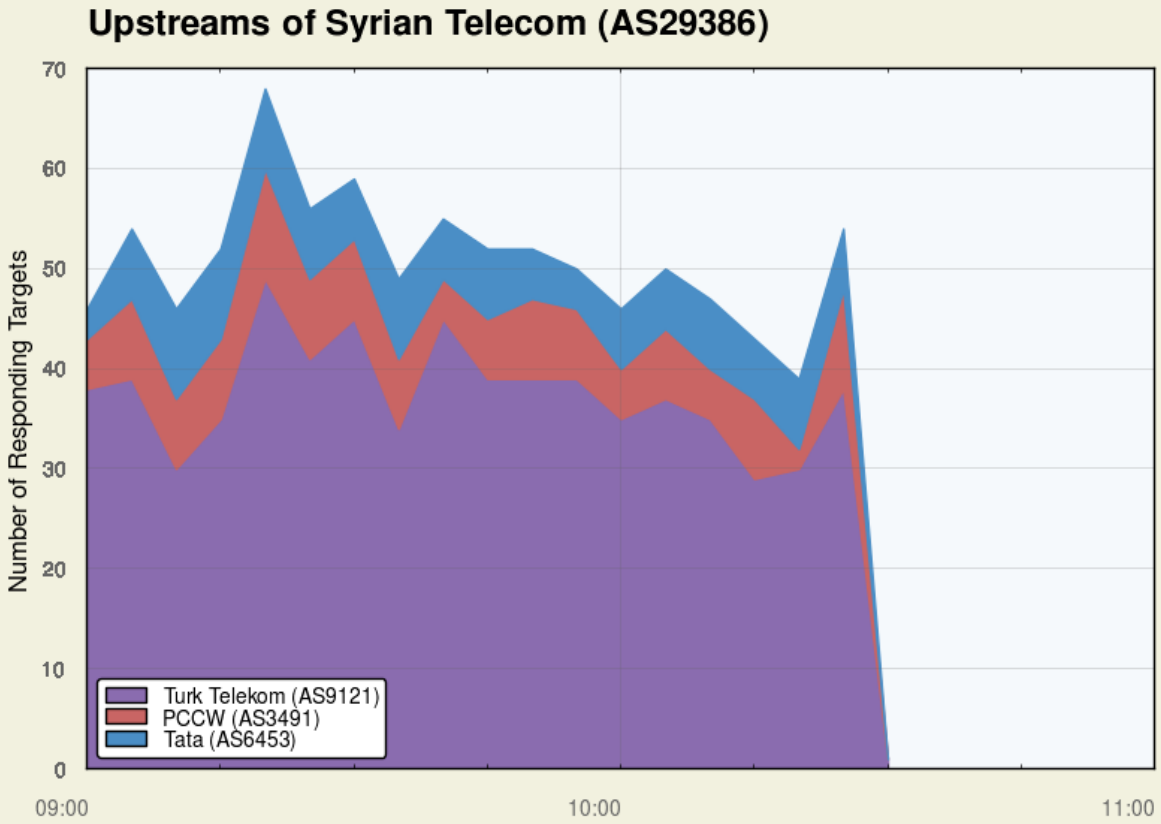

"Terrorists"

29 Nov 2012

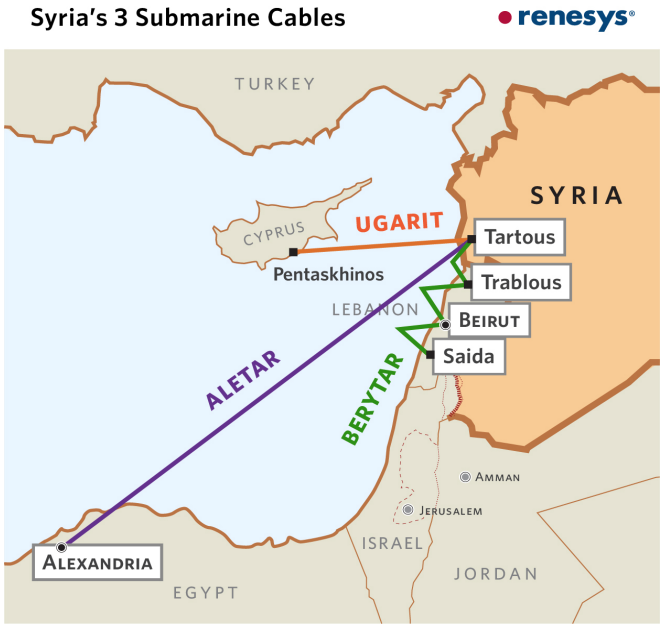

Nationwide outages still happen: Why?

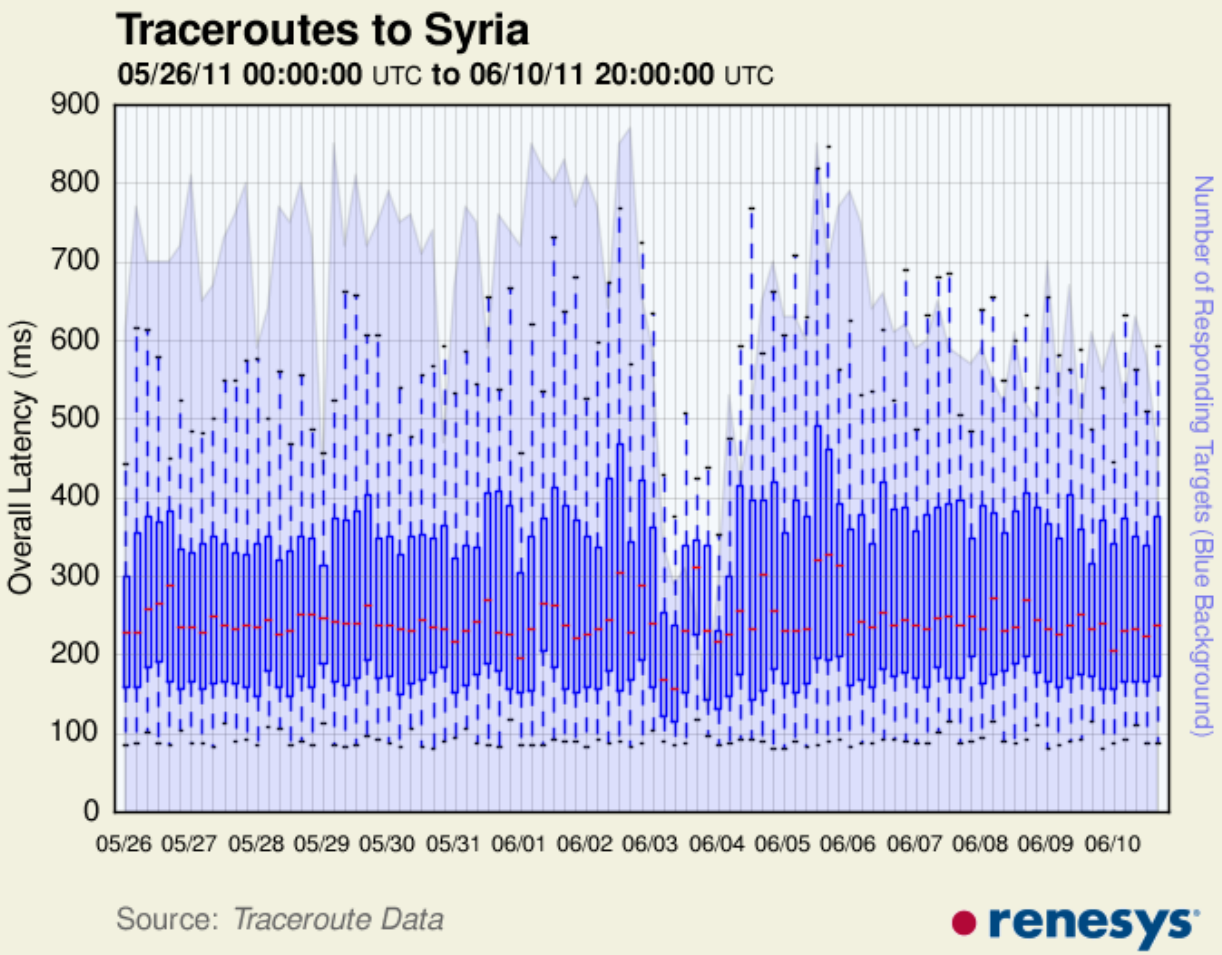

"The government has shut the Internet down"

3 June 2011

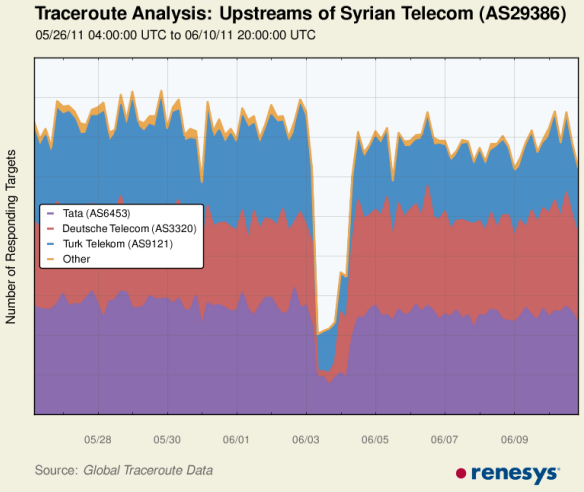

Nationwide outages still happen: Why?

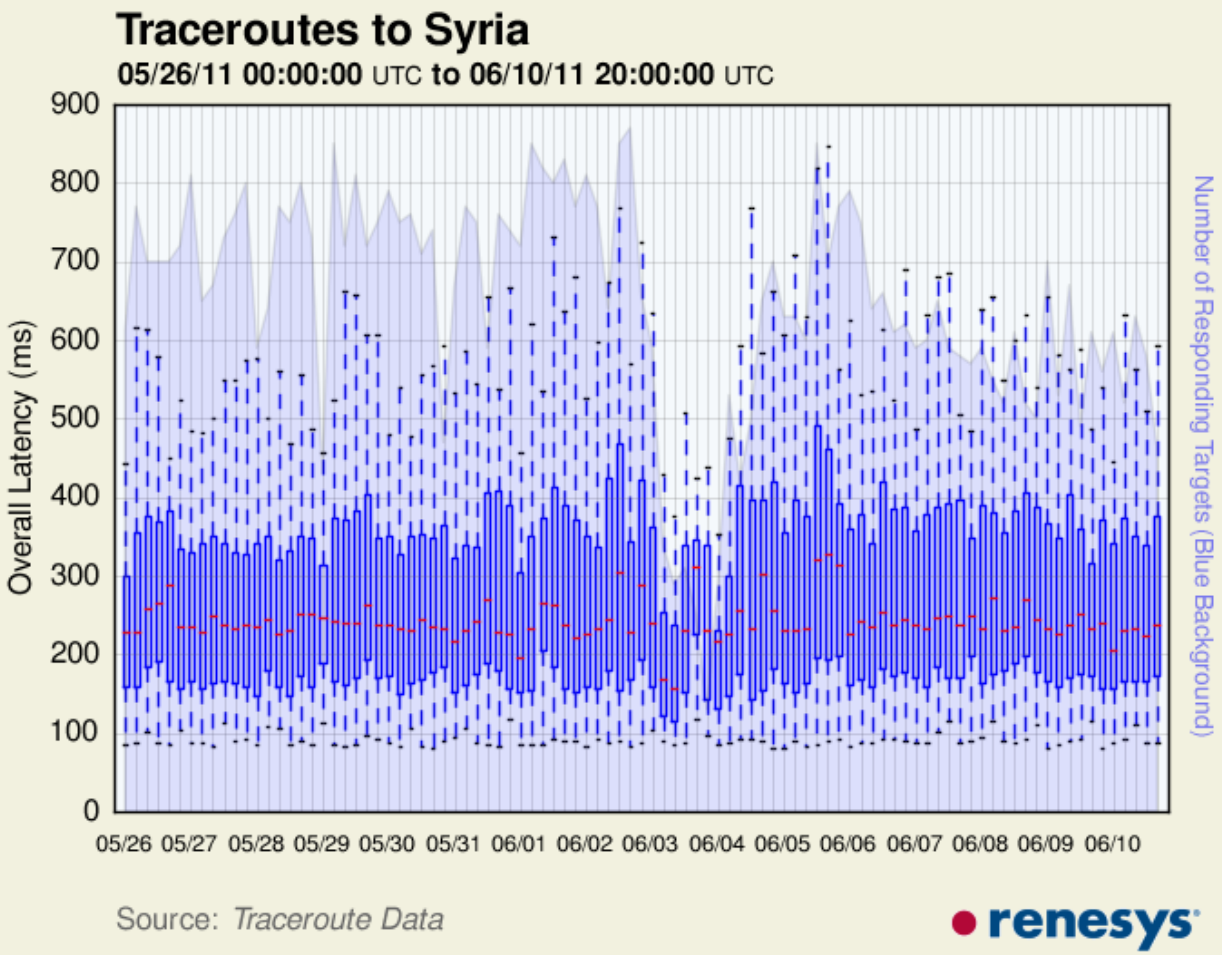

~~"The government has shut the Internet down"~~

"Internet is broken"

3 June 2011

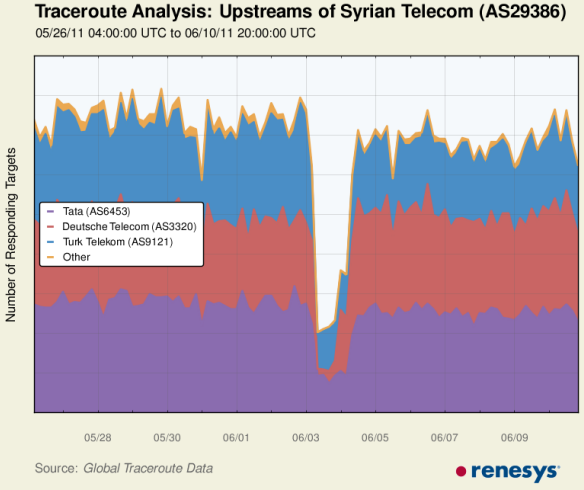

This leads us to ask the obvious question

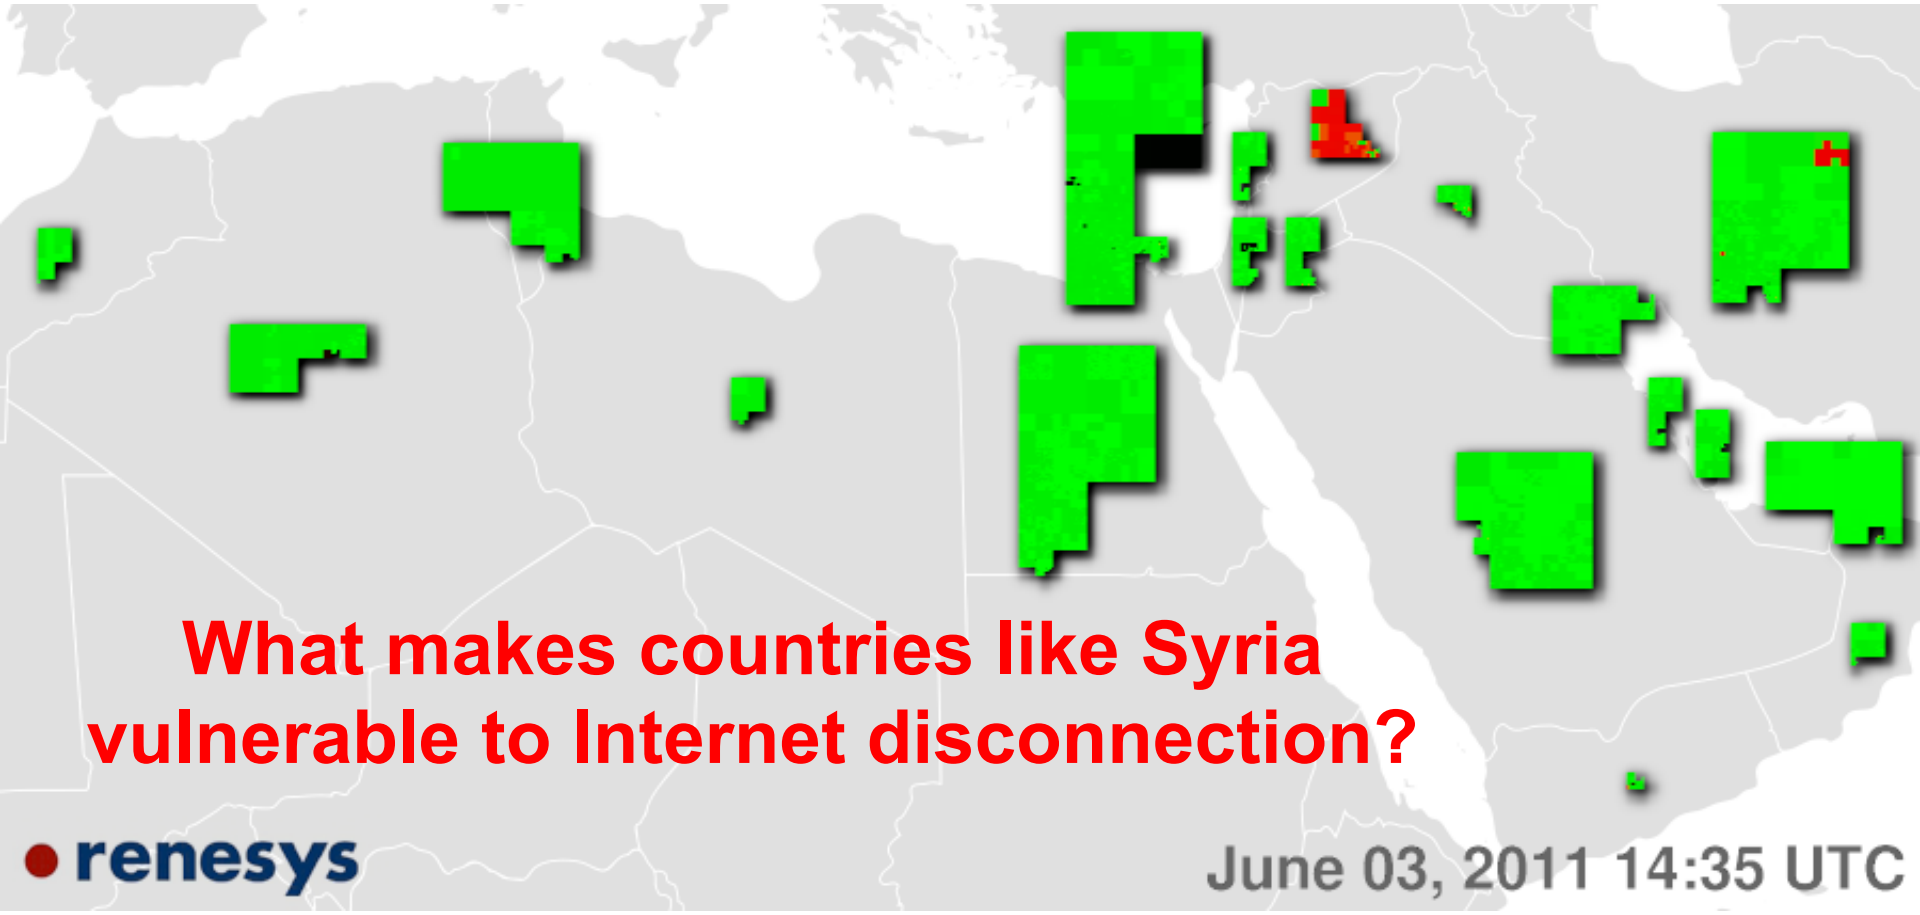

Hypothesis: Provider Diversity

How many distinct institutions in your country have direct BGP transit relationships with international Internet providers?

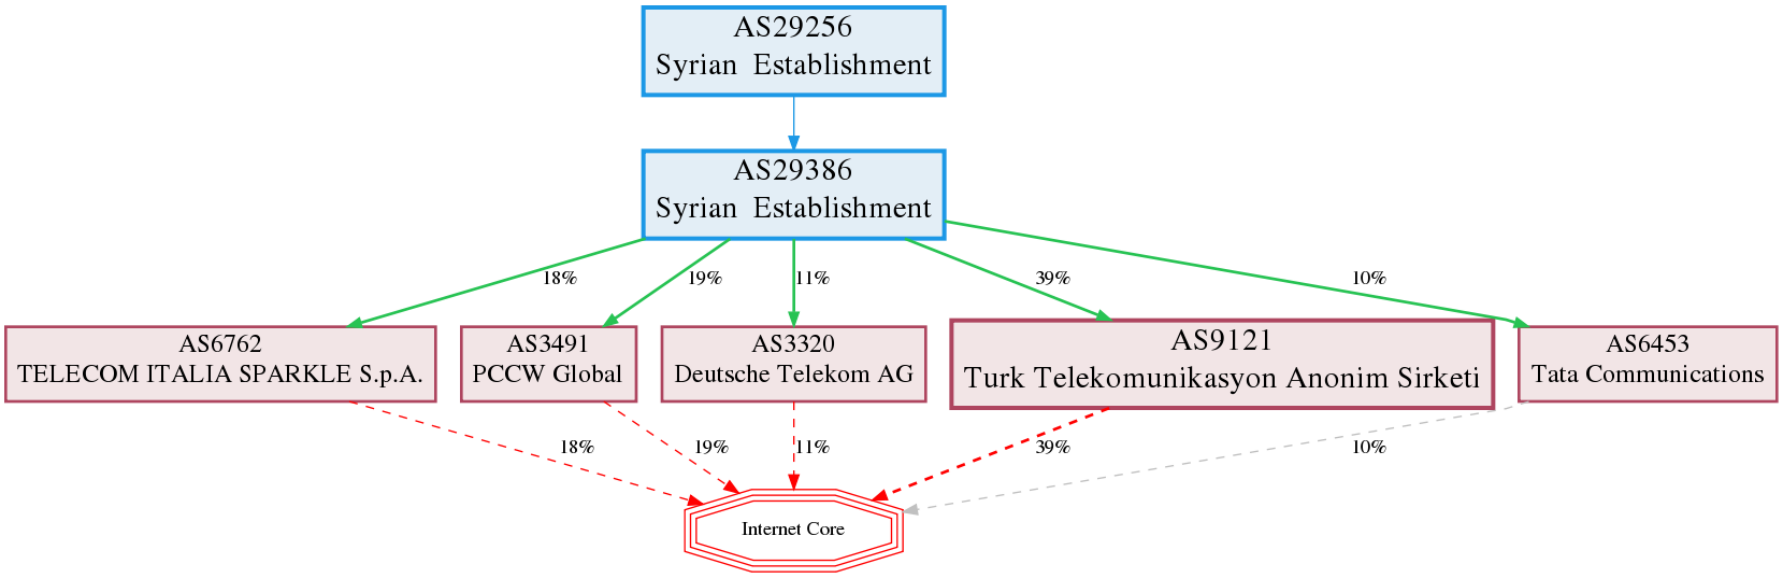

renesys

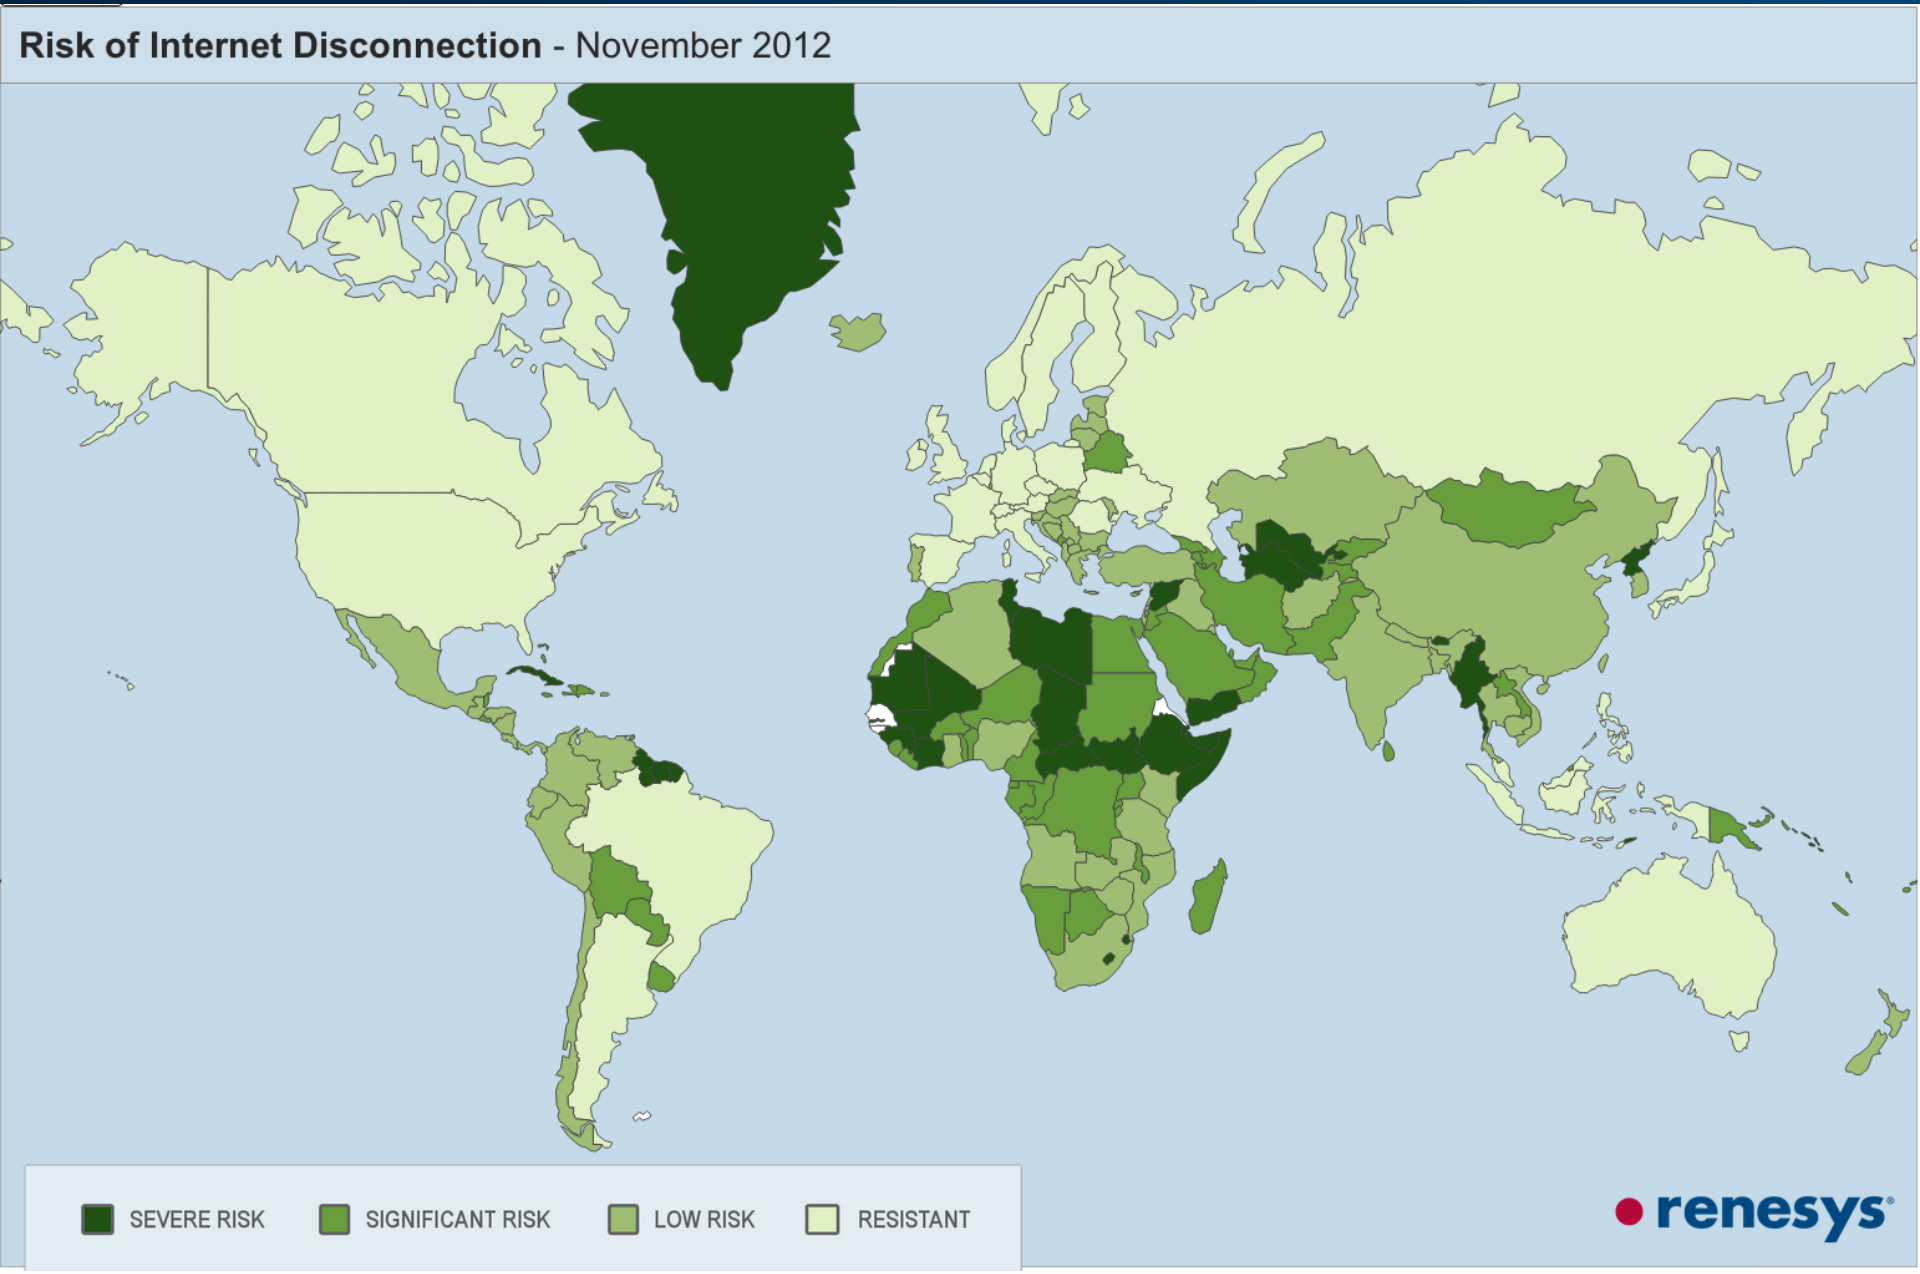
Risk of Internet Disconnection - November 2012
SEVERE RISK
SIGNIFICANT RISK
LOW RISK
RESISTANT
renesys

renesys

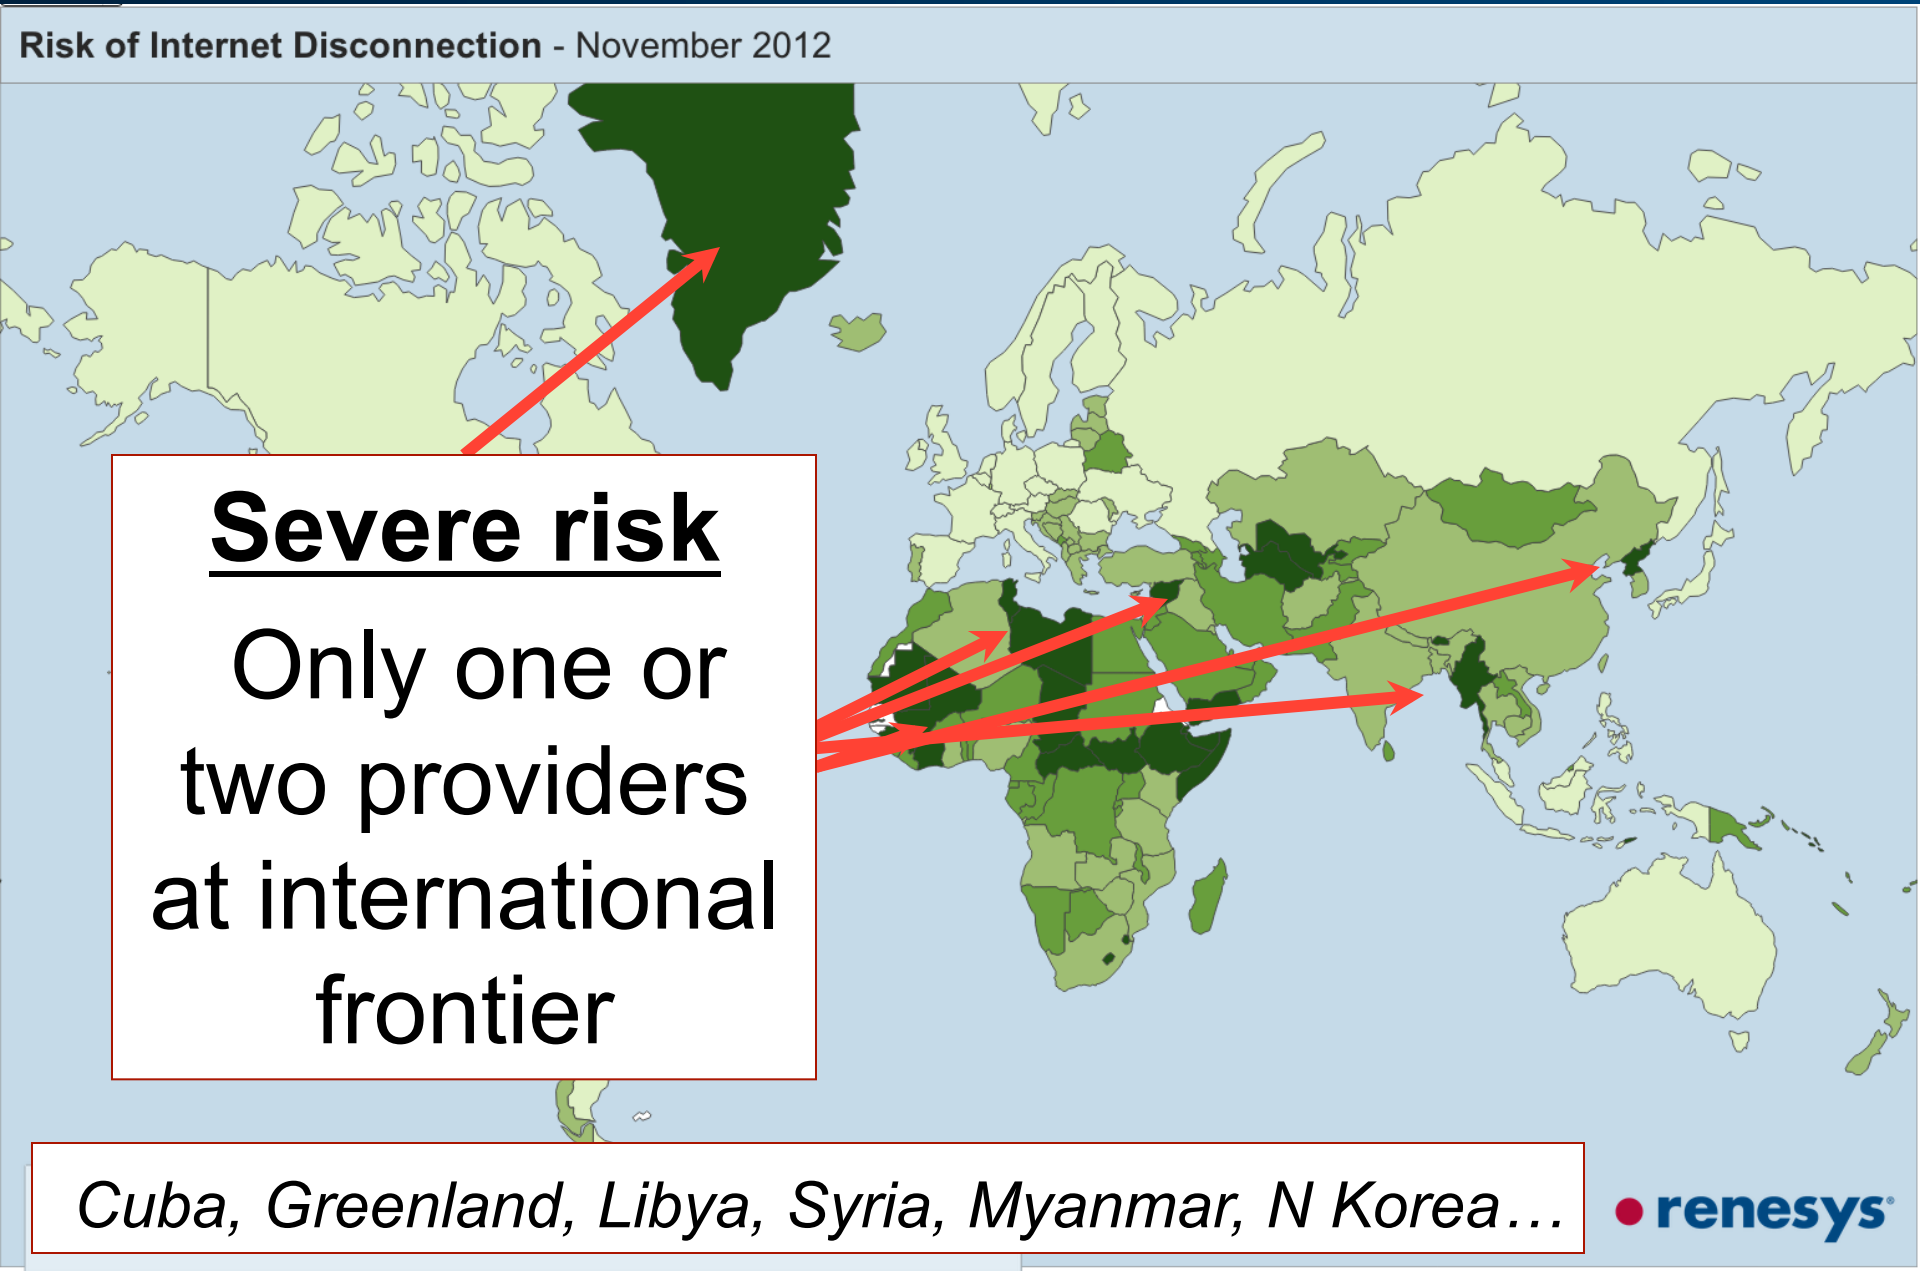
Risk of Internet Disconnection - November 2012
Severe risk
Only one or two providers at international frontier
Cuba, Greenland, Libya, Syria, Myanmar, N Korea…
renesys

renesys

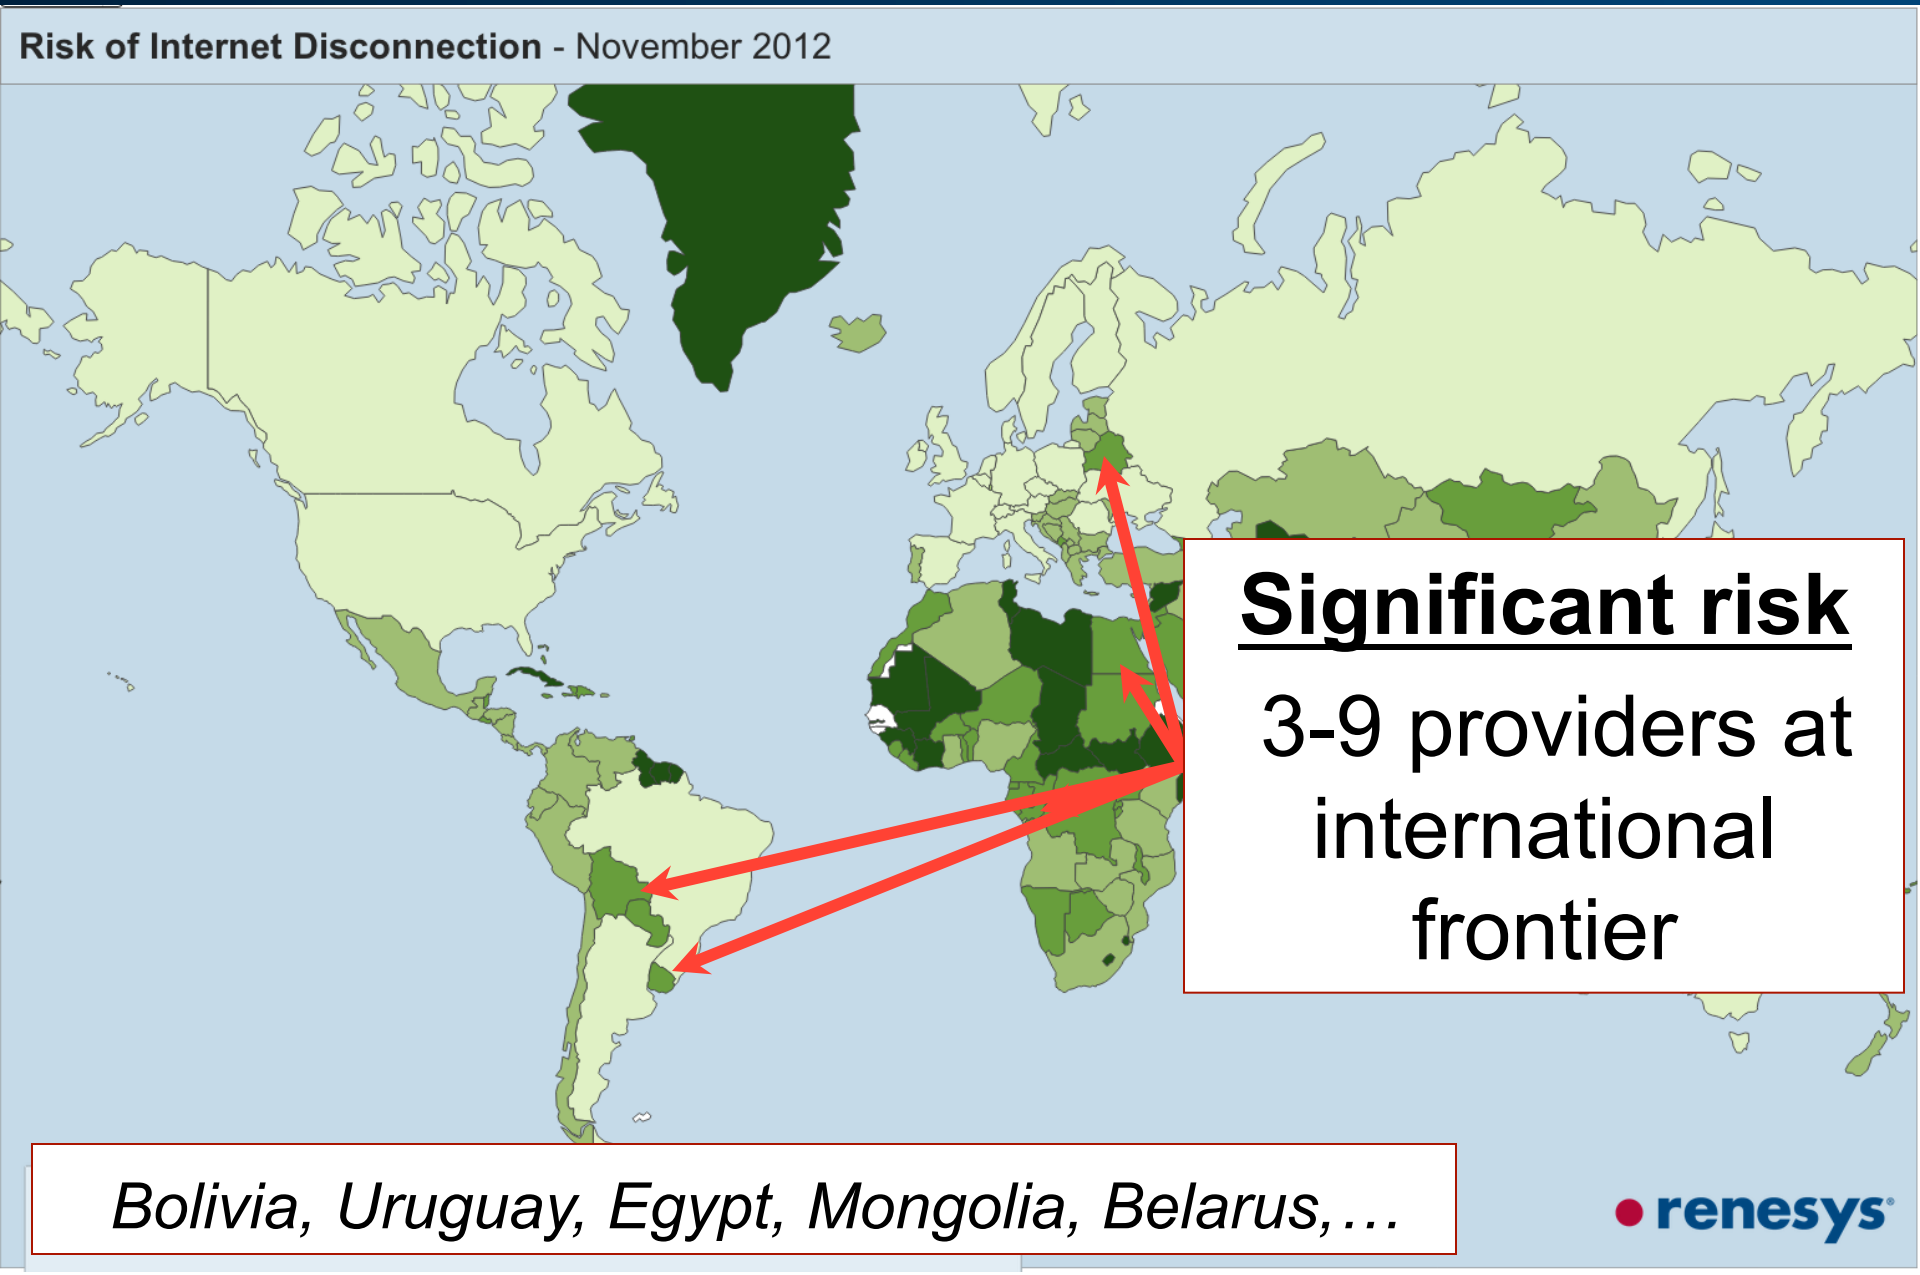
Risk of Internet Disconnection - November 2012
Significant risk
3-9 providers at international frontier
Bolivia, Uruguay, Egypt, Mongolia, Belarus,…
renesys

renesys


Risk of Internet Disconnection - November 2012
Low risk
10-39 providers at international frontier
Mexico, Venezuela, Iceland, China, Afghanistan
renesys

renesys

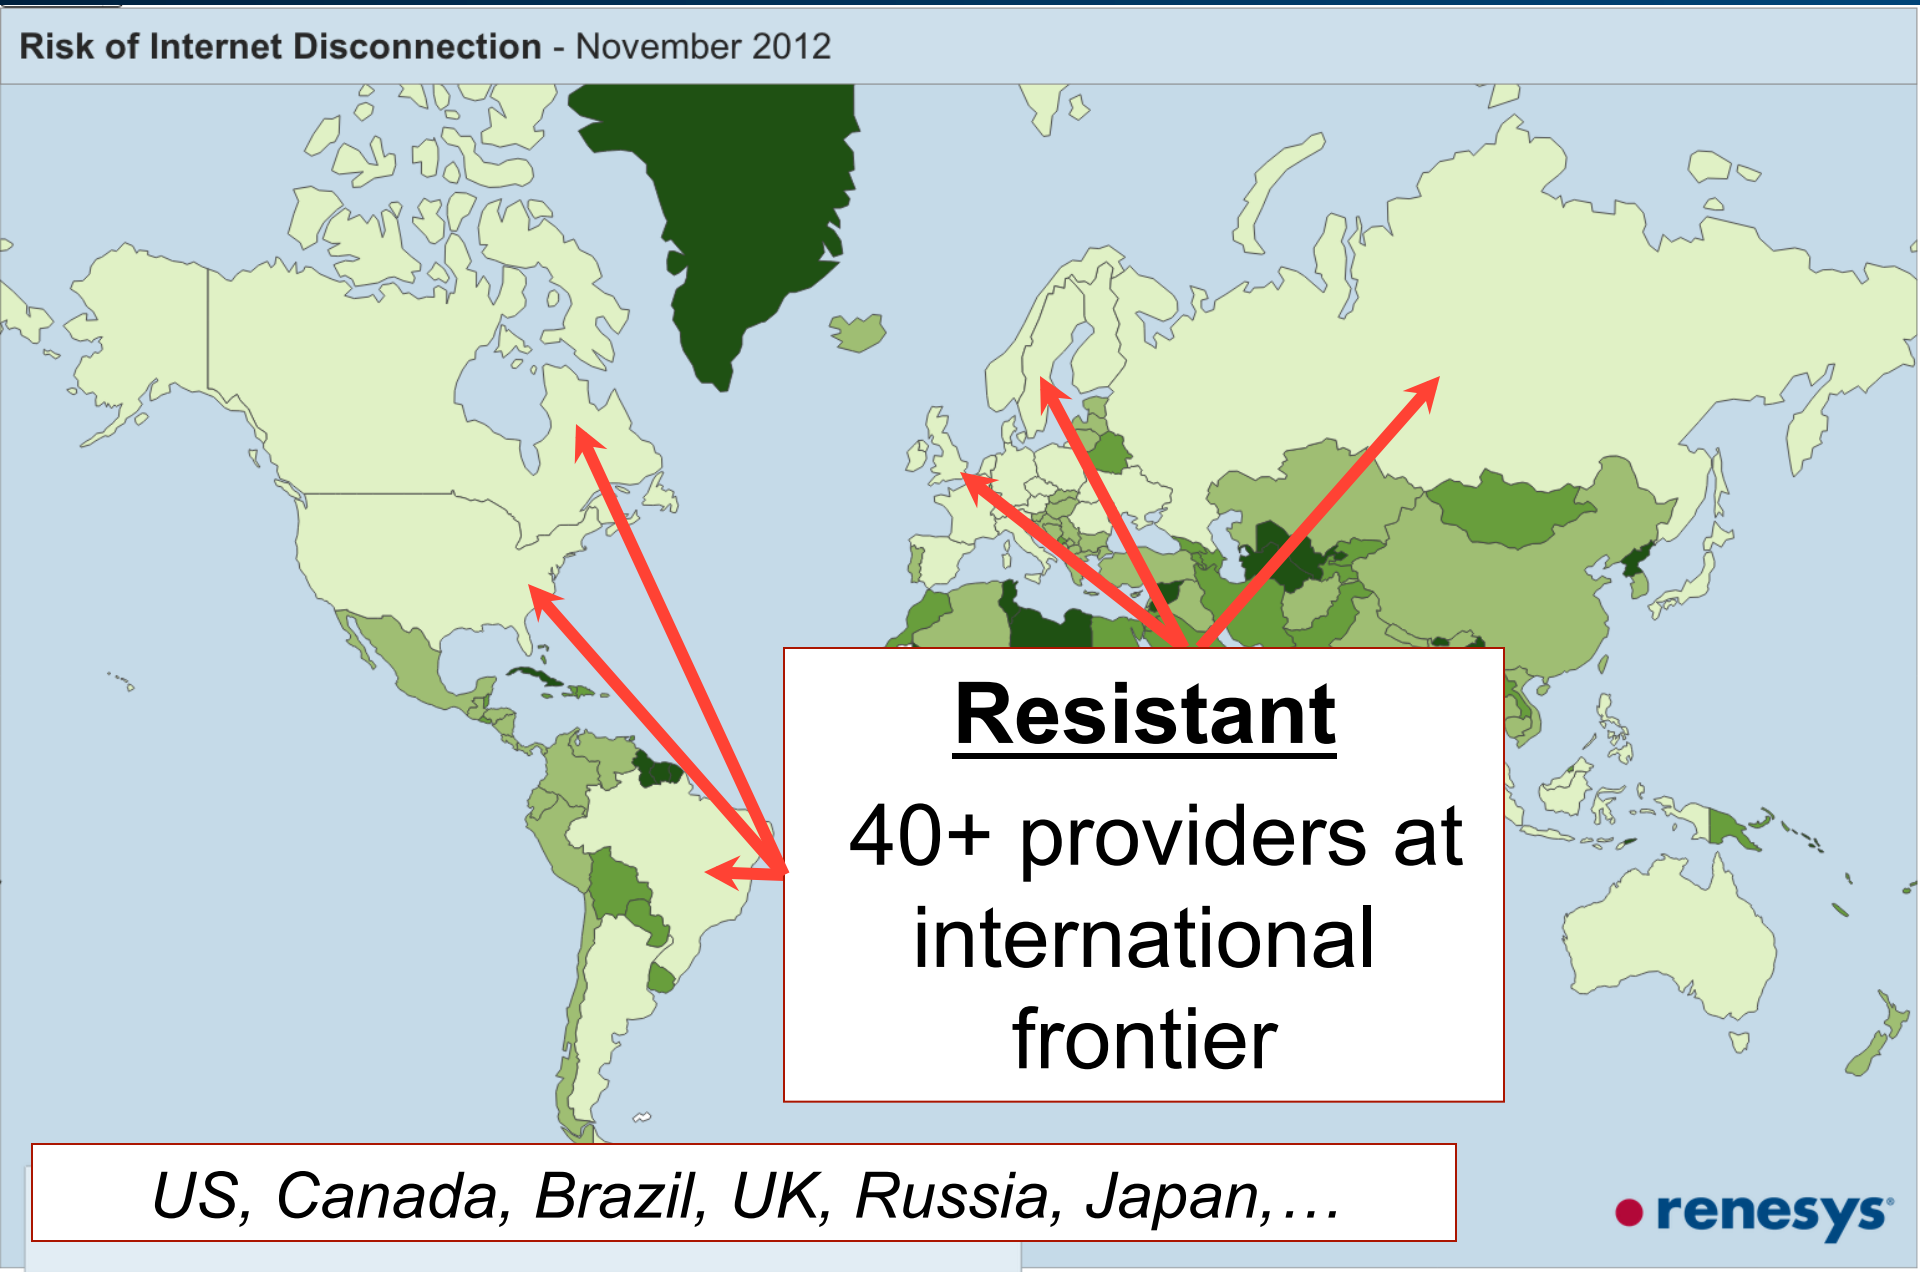

renesys

Let's do an experiment

- Examine a year's worth of Internet routing traffic, more than **40 billion** BGP protocol messages
- Filter and group these into **428,000** distinct outage events affecting groups of networks in **229 countries**

- In all, **17 countries** were affected by a significant (90%+) national outage at least once since January 2012.

renesys

Of the 17 countries affected ….

12 were at "severe risk" (1-2 at the frontier)

- **Syria, North Korea, Mali, Uzbekistan, Nauru, Palau, Suriname, Guyana, Gambia, Cook Islands, Marshall Islands, Comoros (19% of 61 countries; many small islands)**

3 were at "significant risk" (3-10 at the frontier):

- **Brunei, Macao, Equatorial Guinea (4% of 72 countries)**

2 were considered "low risk" (10-39 at the frontier):

- **Bangladesh** (multiple Sea-Me-We-4 cuts)
- **Lebanon** (Maintenance on IMEWE, July 2013)
- 3% of 60 low-risk countries

renesys

IMEWE Cable Maintenance, 2 July 2012

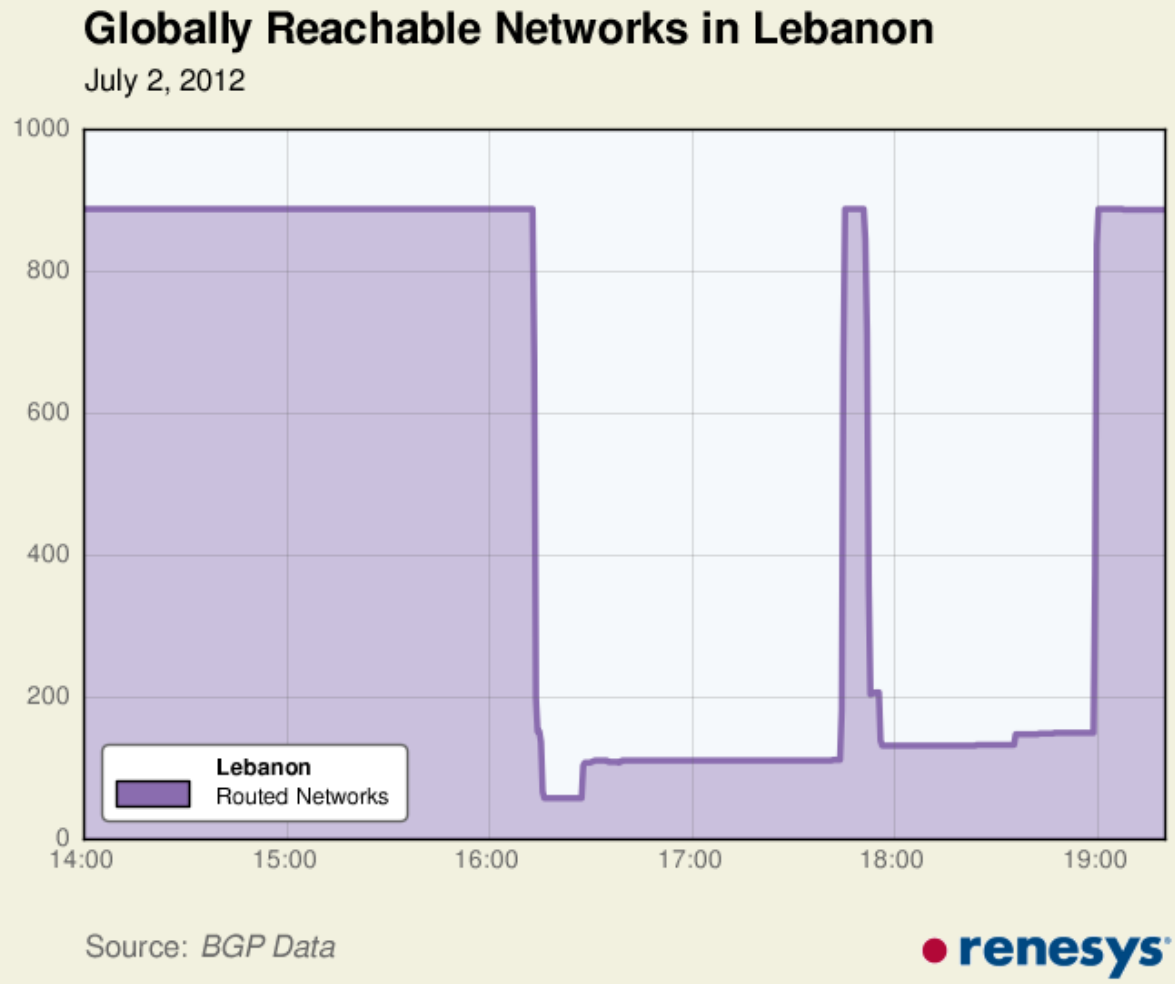

IMEWE Cable Maintenance, 2 July 2012

Lebanon: IMEWE Cable Concentrates Risk

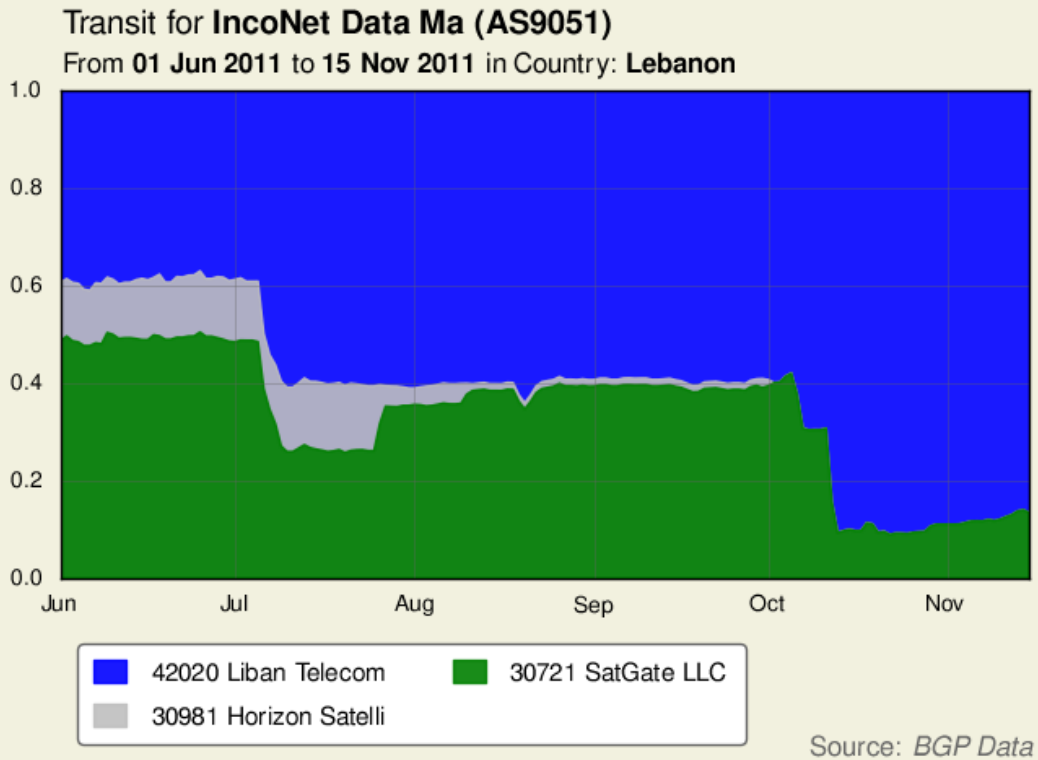

AS42020 %pct Lebanon On-Net

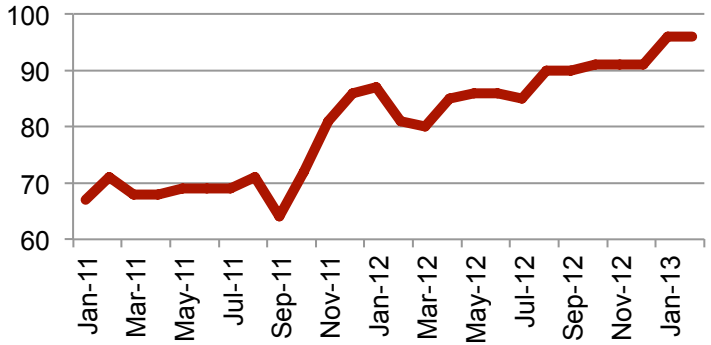

ASNs with Cross-Border Connectivity

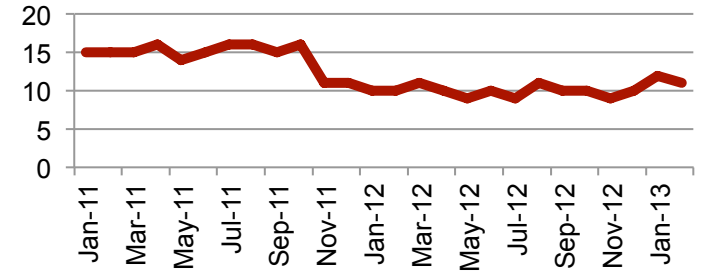

renesys

“Resistant To Disconnection”

In the last year, there were **no countrywide outages** affecting countries with **40+** companies at the international frontier.

Not a one.

renesys

Fine, then!

How *does* a country go about gaining 40+ providers at its national frontier?

renesys

Diversification. But It Takes Time.

- Government has a role to play in encouraging competition and diversification, particularly in low-diversity markets.
- Over time, a self-sustaining Internet market that is large and competitive enough should require minimal regulation.

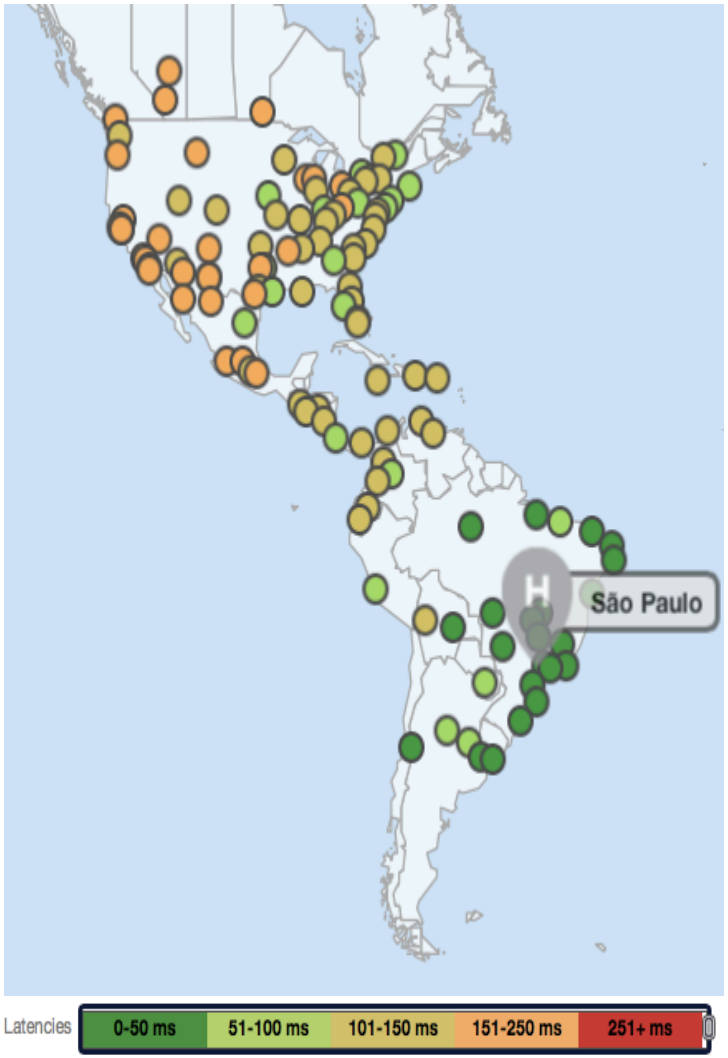

renesys

Physical Diversification Is Not Sufficient

Cuba activated a first submarine fiberoptic link to Venezuela in January 2013..

..and a second segment to Jamaica in May 2013.

But they still have a single company controlling the international frontier, AS11960.

renesys

Cuba's First Steps: Physical Diversity

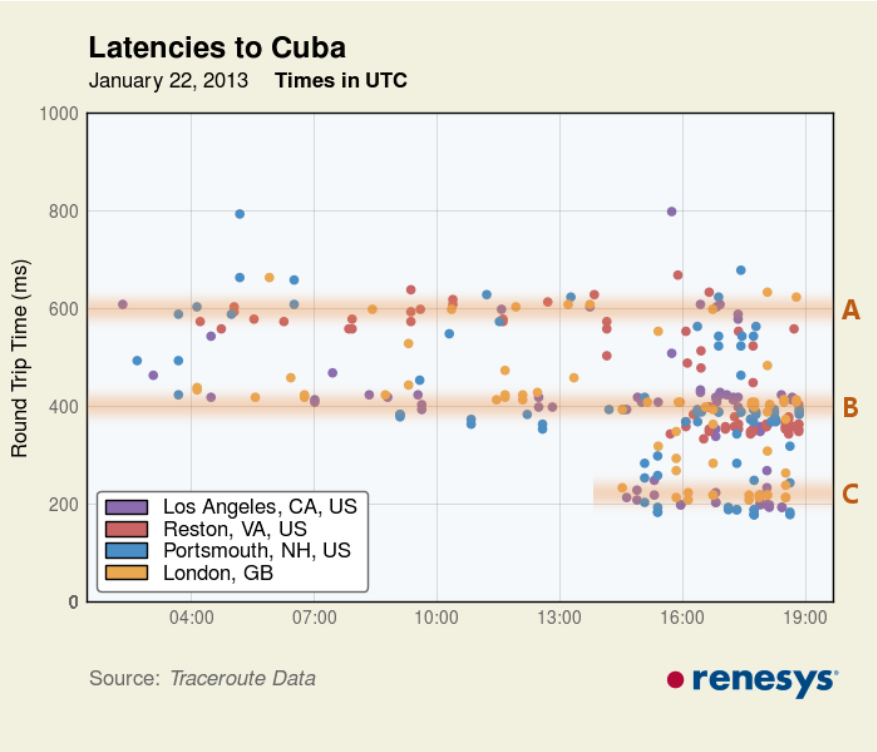

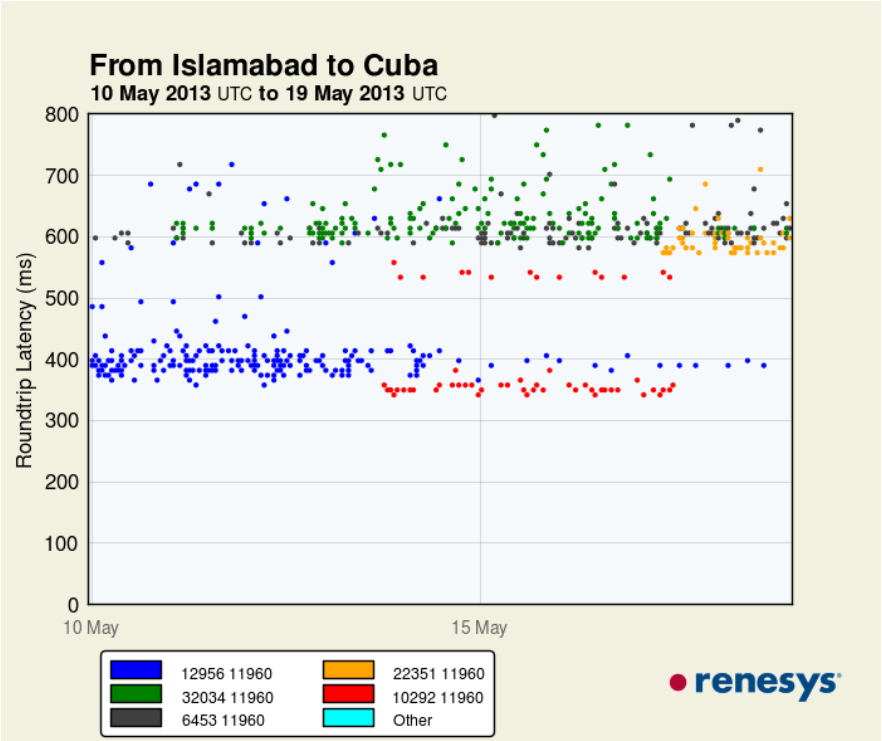

Jan 2013: satellite (a) gives way to faster submarine connectivity to VZ (b,c)

May 2013: Telefonica transit through VZ is itself supplemented by faster C&W transit through Jamaica

renesys

Brazil: On The Fast Track

Brazil adds “2 Mexicos” each year

- 340% increase in ASN count since 2010
- Already exceeds the rest of LATAM combined
- Up to **218 ASNs** at international frontier (Feb 2013)

Municipal IXPs (PTTs) facilitate significant domestic Internet growth

renesys

Costa Rica: Growing Even Faster

Five years later, 2008 telecom law overhaul yielding results

- Was a requirement of CAFTA-DR agreement
- Created telecom regulator: Sutel
- Ended ICE monopoly

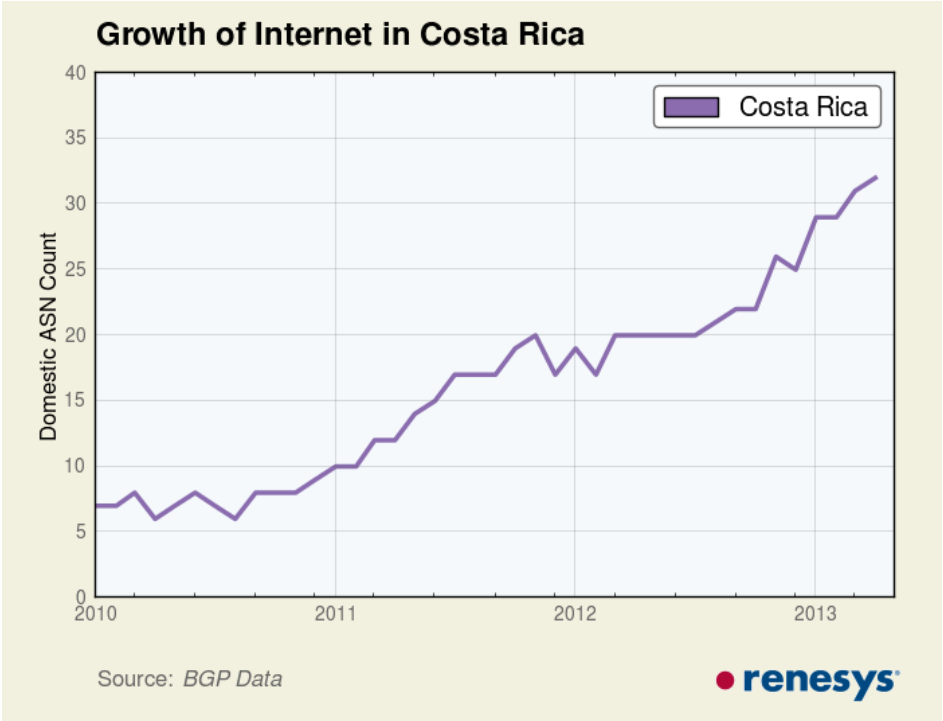

- Went from **6 ASNs** at frontier (Jan2011) to **19** (Feb2013)

…. By our metrics, now at “low risk of disconnection”

renesys

Uruguay isn't growing at all

- Number of routed ASNs unchanged in four years
- State-owned Antel faces no fixed-line competition
- A fine provider, but the market is stagnant

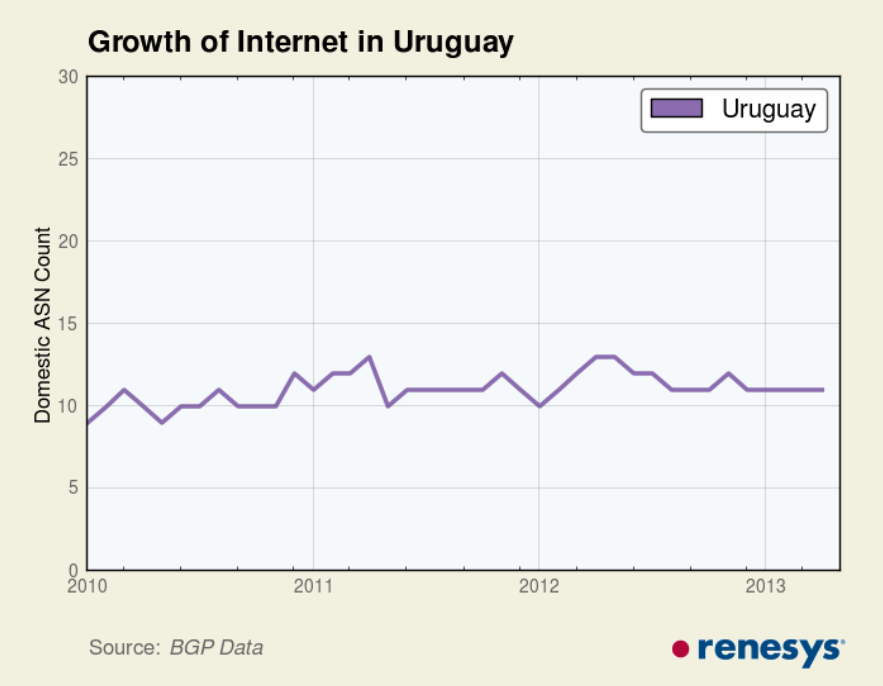

- Number of ASNs at the international frontier actually *dropped* since 2011, from 7 to just 5.
- "It's great now .. what if something changes?"

renesys

Conclusions

~~Conclusions~~

Aphorisms

renesys

Aphorism #1

The human vulnerabilities of the Internet (temptations to meddle, monitor, censor, control, regulate) are now a greater danger than its physical weaknesses.

renesys

Aphorism #1, Updated

Great dangers can create great opportunities,
especially for those in smaller/emerging markets.

- Wake up call for countries who do “trombone routing” of their domestic traffic to Euro/USA
- Powerful stimulus for domestic peering and local exchanges
- *This is going to be a good thing for the Internet as a whole*

renesys

Aphorism #2

We aren't smart enough to rebuild the Internet from the ground up "the way it should be."

If we were, you'd all be using IPv6 right now.

renesys

Aphorism #3

The best way to address the problems of “Internet as critical infrastructure” is simply to **build more of it.**

Our job is to seek out single points of failure (or control) and help the Internet diversify around them.

Beyond some point of self-organized complexity, there is no more “kill switch.”

renesys

Thank you!

renesys

Internet Infrastructure: Virtual meets Reality

James Cowie, CTO

@jimcowie
@renesys

Netnod Autumn Meeting
Stockholm, Sweden
3 October 2013