

Impact of the Tohoku Quake & Tsunami as Seen by a Japanese ISP

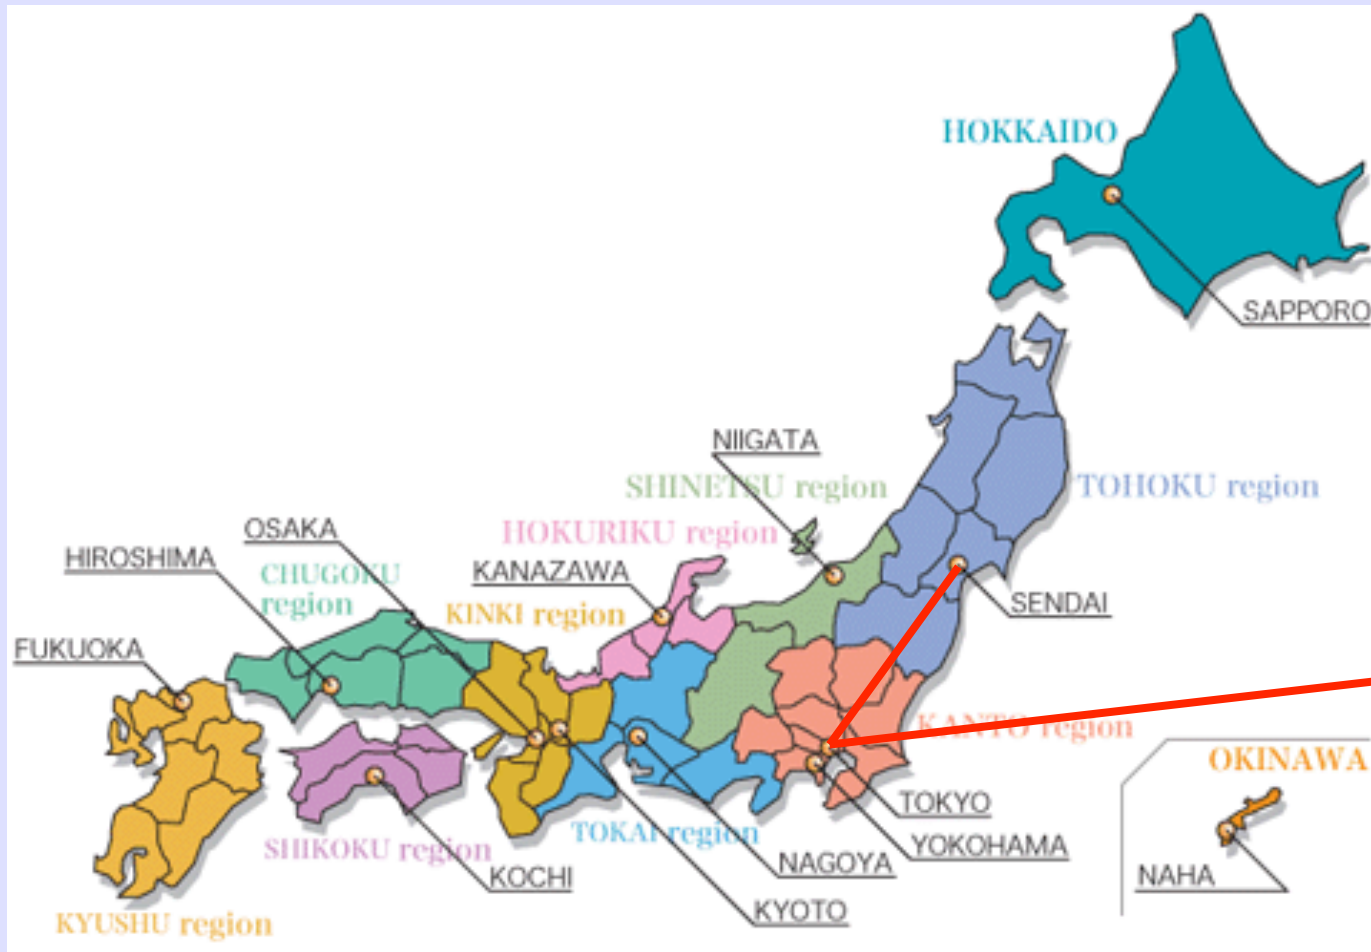
Kenjiro Cho
Cristel Pelsser
Randy Bush
Youngjoon Won

Appeared in the *Special Workshop on the Internet
and Disasters*, CoNEXT 2011

Agenda

- Real world events
- Routing observations
 - OSPF
 - BGP
- Traffic observations

Japan



March 11th	
14:46	The earthquake of Magnitude 9.0 about 130km east of Sendai city
14:48	Sendai Datacenter switched to in-house power generator
14:48	The two links to Sendai are lost
21:50	One link to the US fails
March 12th	
01:13	Two more links between Tokyo and the US fail
06:16	One of two links to Sendai is recovered
11:30 (~)	External power supply of Sendai Datacenter is restored
20:41	Recovery of one of three Japan- US links
20:51	Recovery of a second Japan- US link (link was disabled at the beginning of the month)
March 13th	
04:36	Recovery of a third Japan- US link
21:20	Recovery of the second link to Sendai
March 14th	Monday - Back in Business

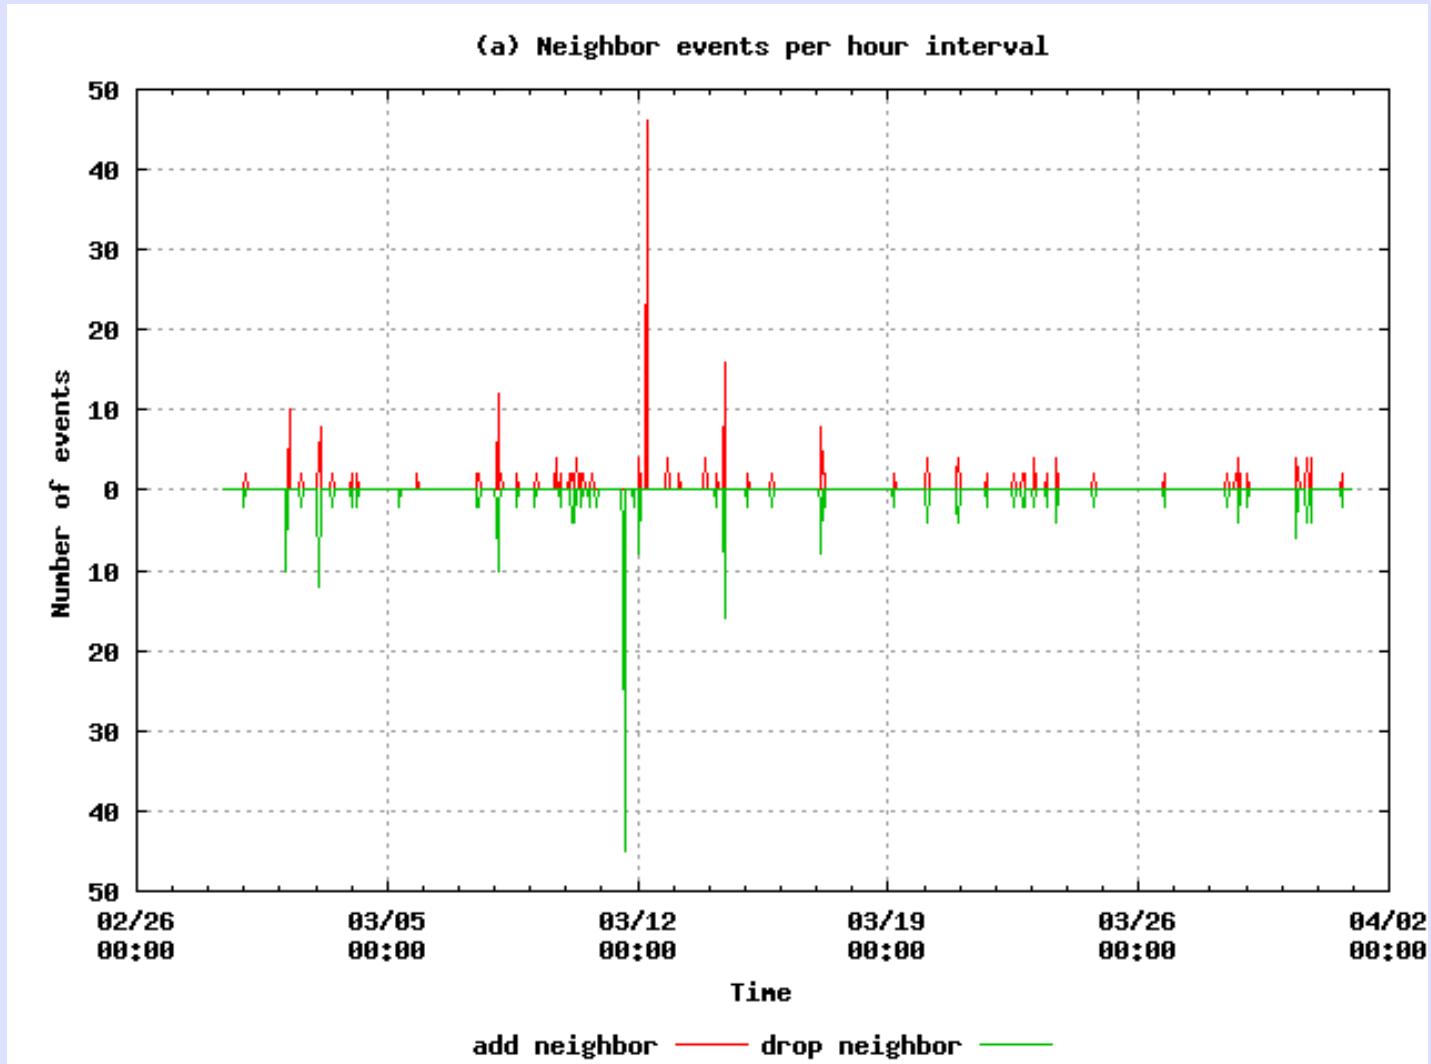
Routing Viewpoints

- Internal behavior: OSPF
- An external view: BGP in a neighboring ISP (Provider)

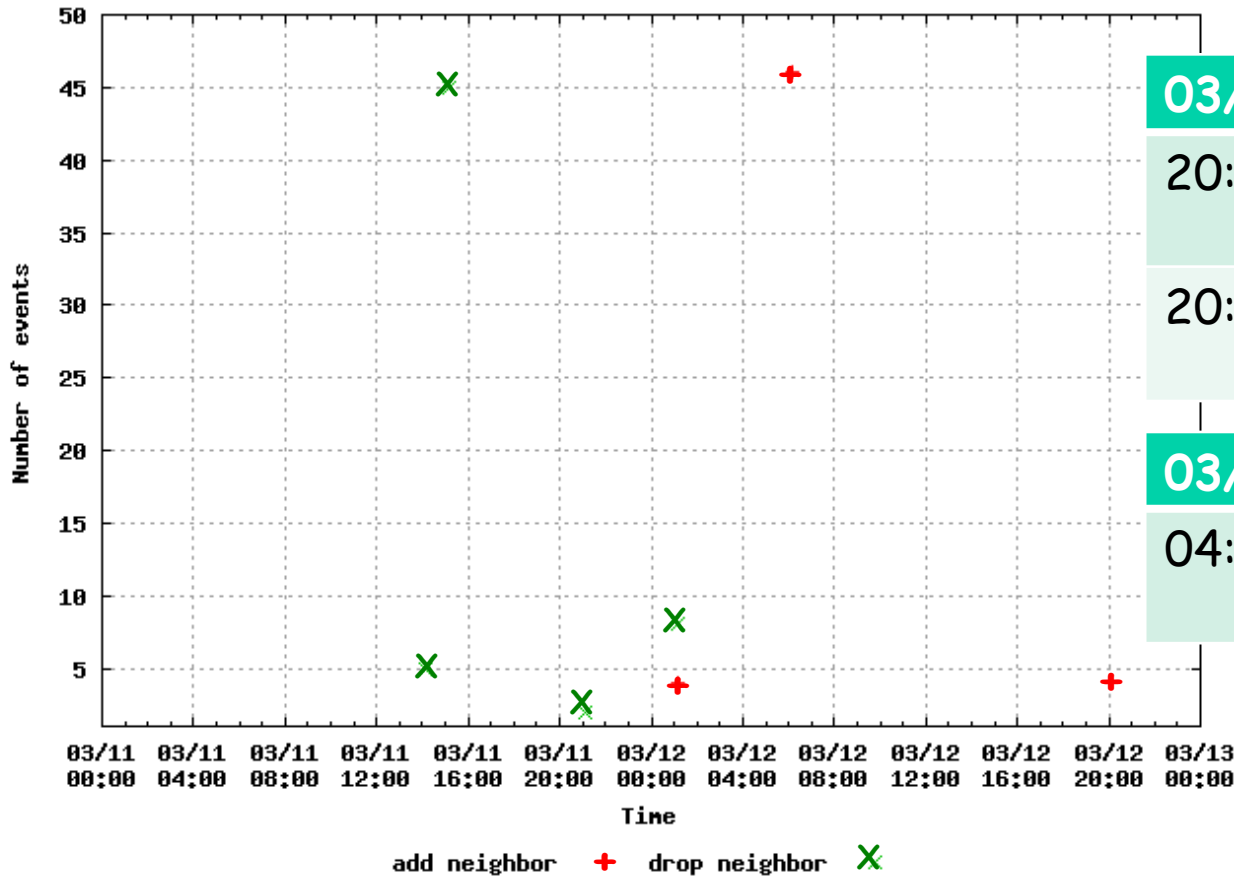
OSPF Behavior

- Route Explorer - Packet Design
- Trace for the backbone (area 0)
 - ~ 1525 links (Feb 28th 00:00:00)
 - ~ 325 nodes (Feb 28th 00:00:00)
- Count number of events occurring within an hour

OSPF Analysis



(a) Neighbor events per hour interval




03/12	
20:41	Recovery of one Japan-US link
20:51	Recovery of second Japan-US link
03/13	
04:36	Recovery of third Japan-US link

03/11		03/12	
14:48	Two links to Sendai lost	1:13	2 more links between Tokyo and the US flap and fail (drop - add - drop events)
15:34	OSPF timeouts for Tohoku		
21:50	One link to the US fails	6:16	One link to Sendai is recovered

Internal Behavior: OSPF

- Connectivity to Sendai lost for 15 hours 28 min
- Out of a dozen or so trans-Pacific links, three links fail
- OSPF churn is very low compared to the number of refresh LSAs

Prefix Events in OSPF



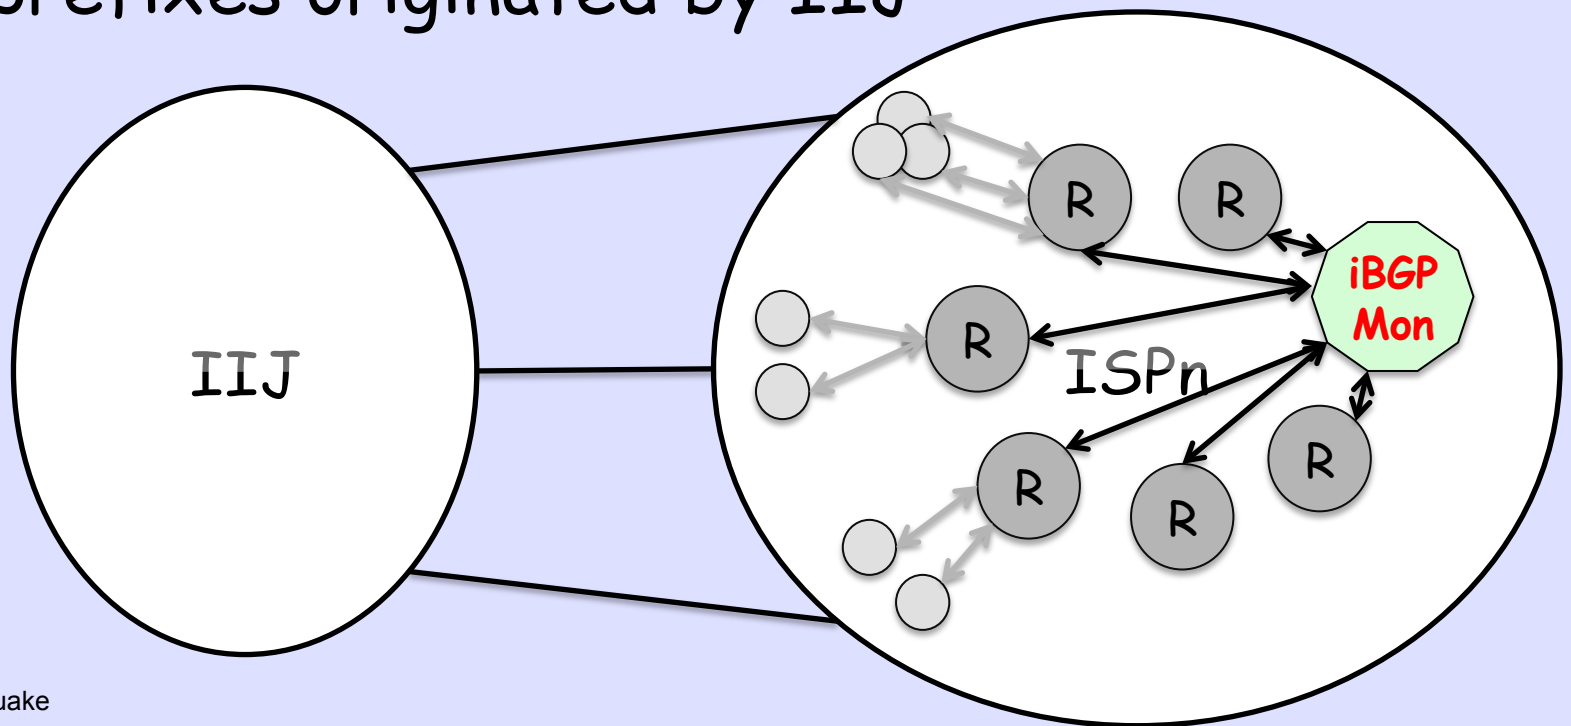
Around 100 prefixes disappear

External View: BGP

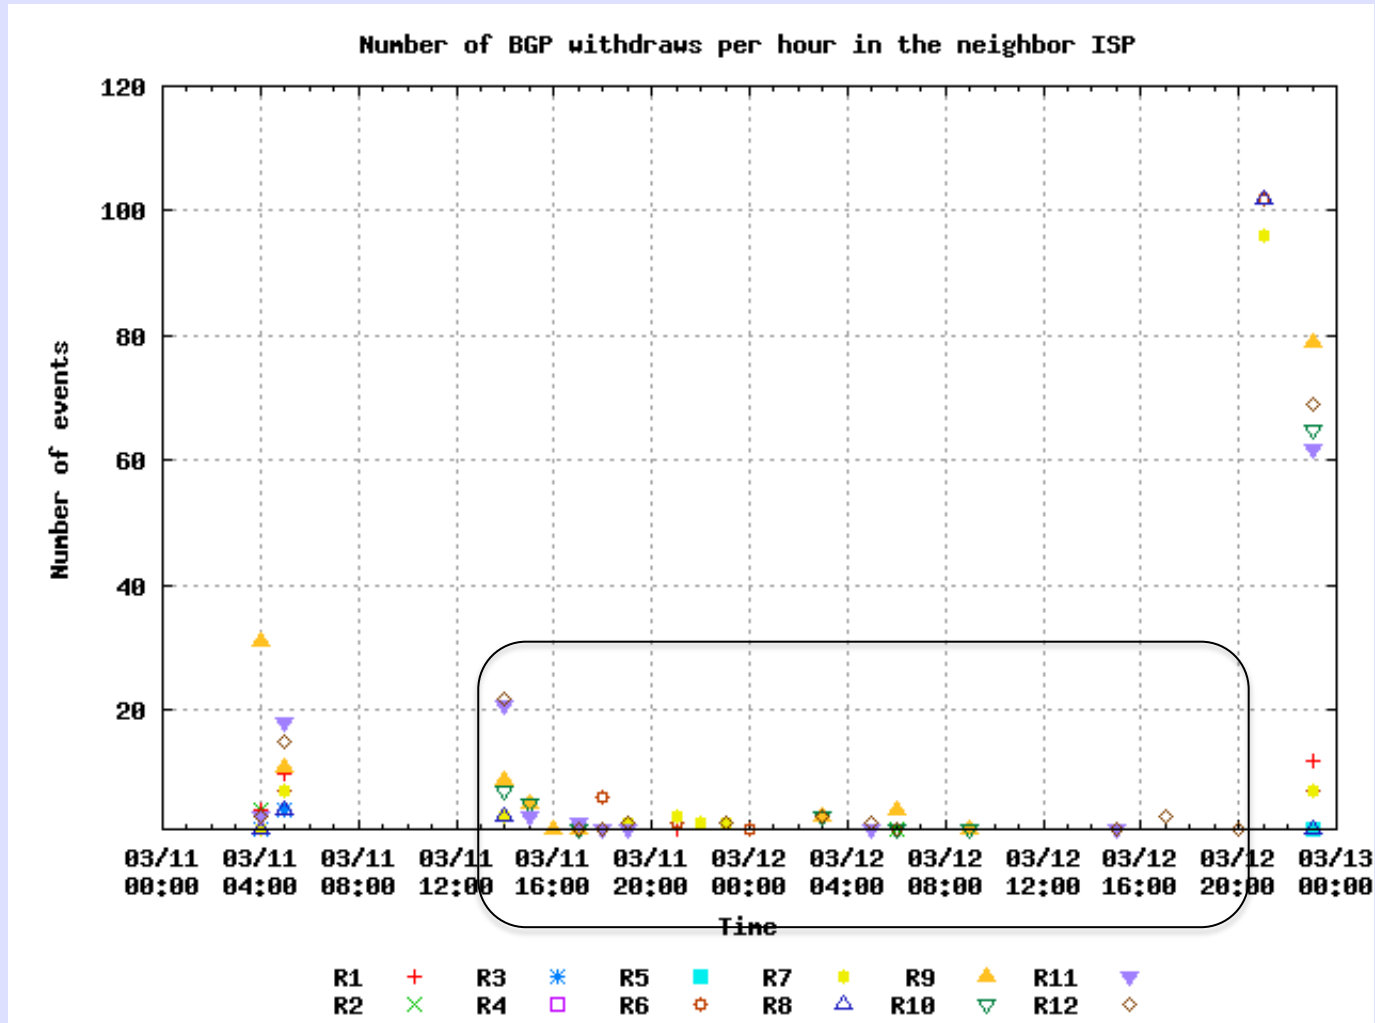
Analysis

iBGP data inside a neighboring ISP (ISPn)

What does the iBGP monitor in ISPn see for prefixes originated by IIJ

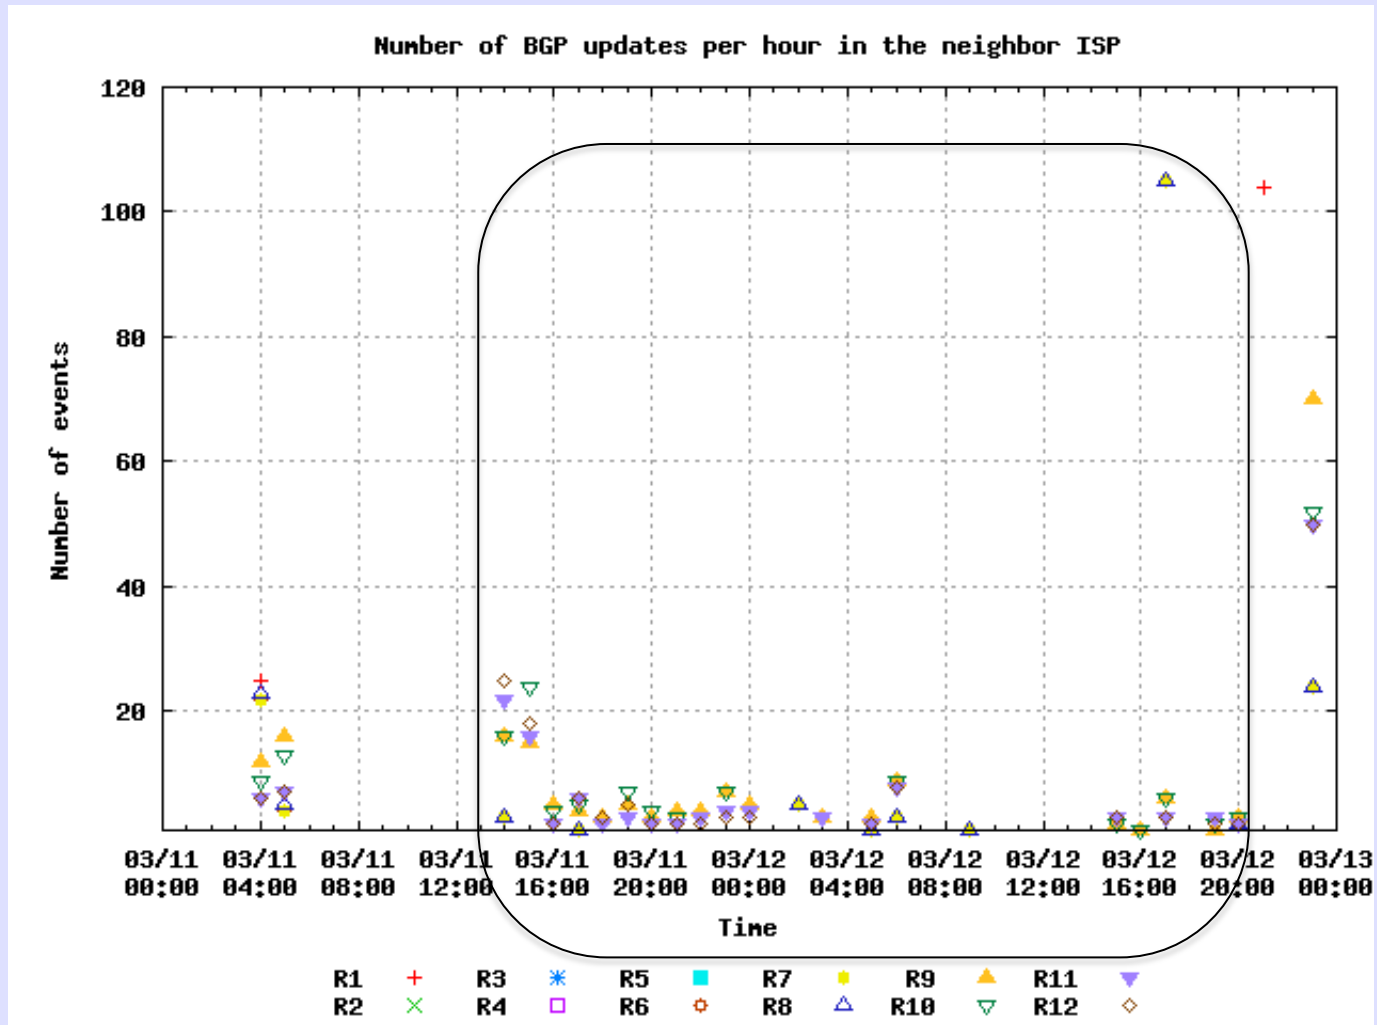


Withdraws Seen by Peer

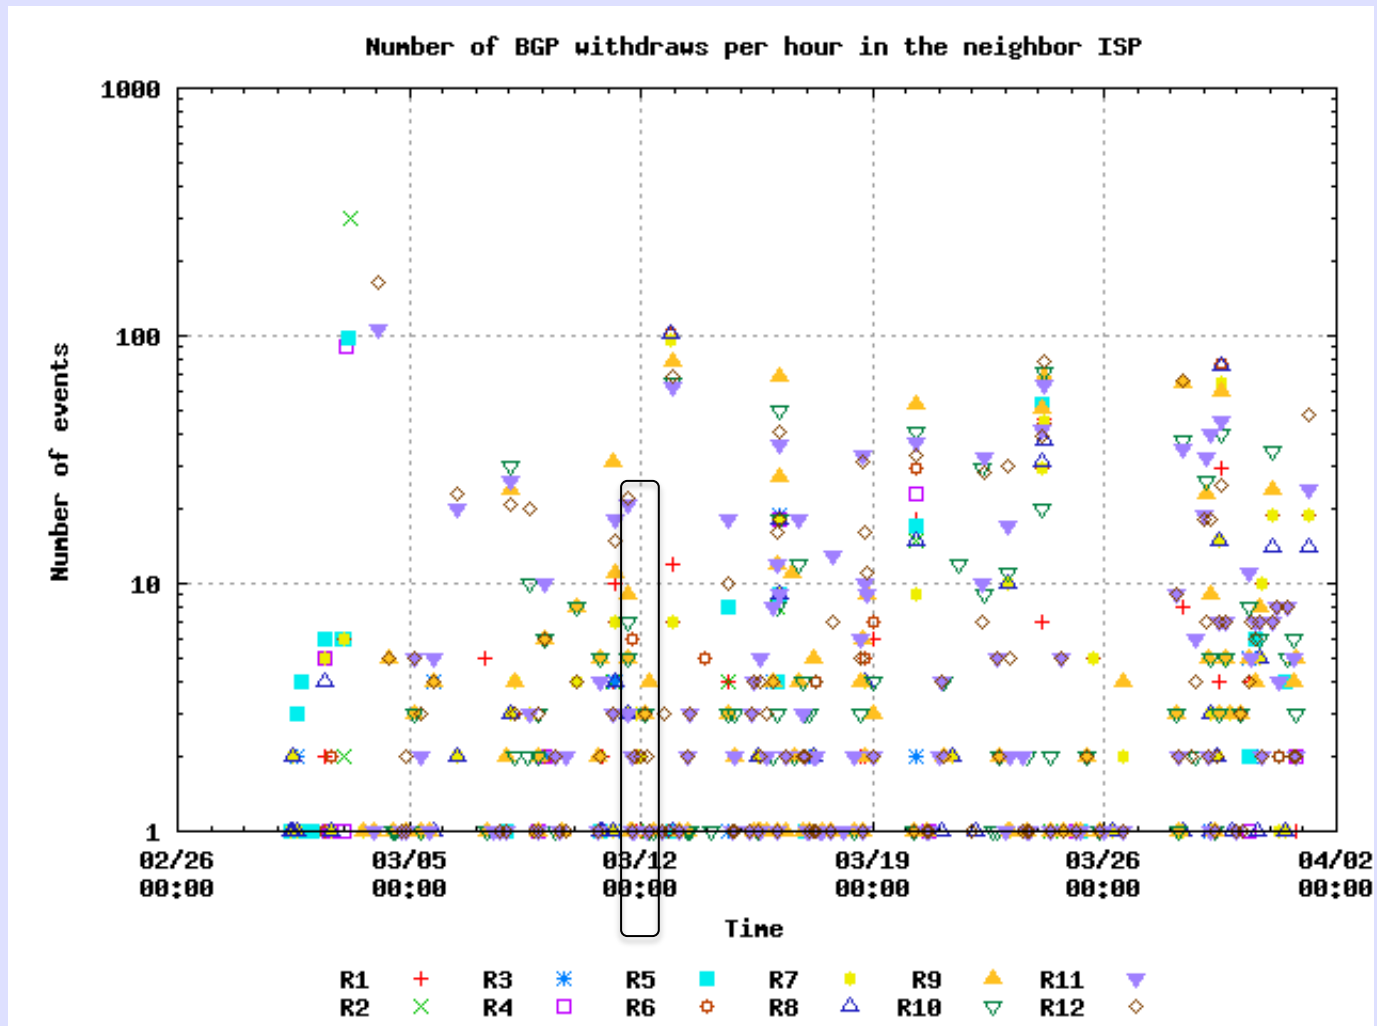


1412 prefixes advertised by IIJ

Updates Seen by Peer





But are Drowned in Normal BGP





Traffic

- Broadband traffic
- 3 trans-Pacific links under the microscope
- JPNAP

Broadband Traffic




2 Trans-Pacific Links




Traffic reroutes to another existing link

Another Trans-Pacific

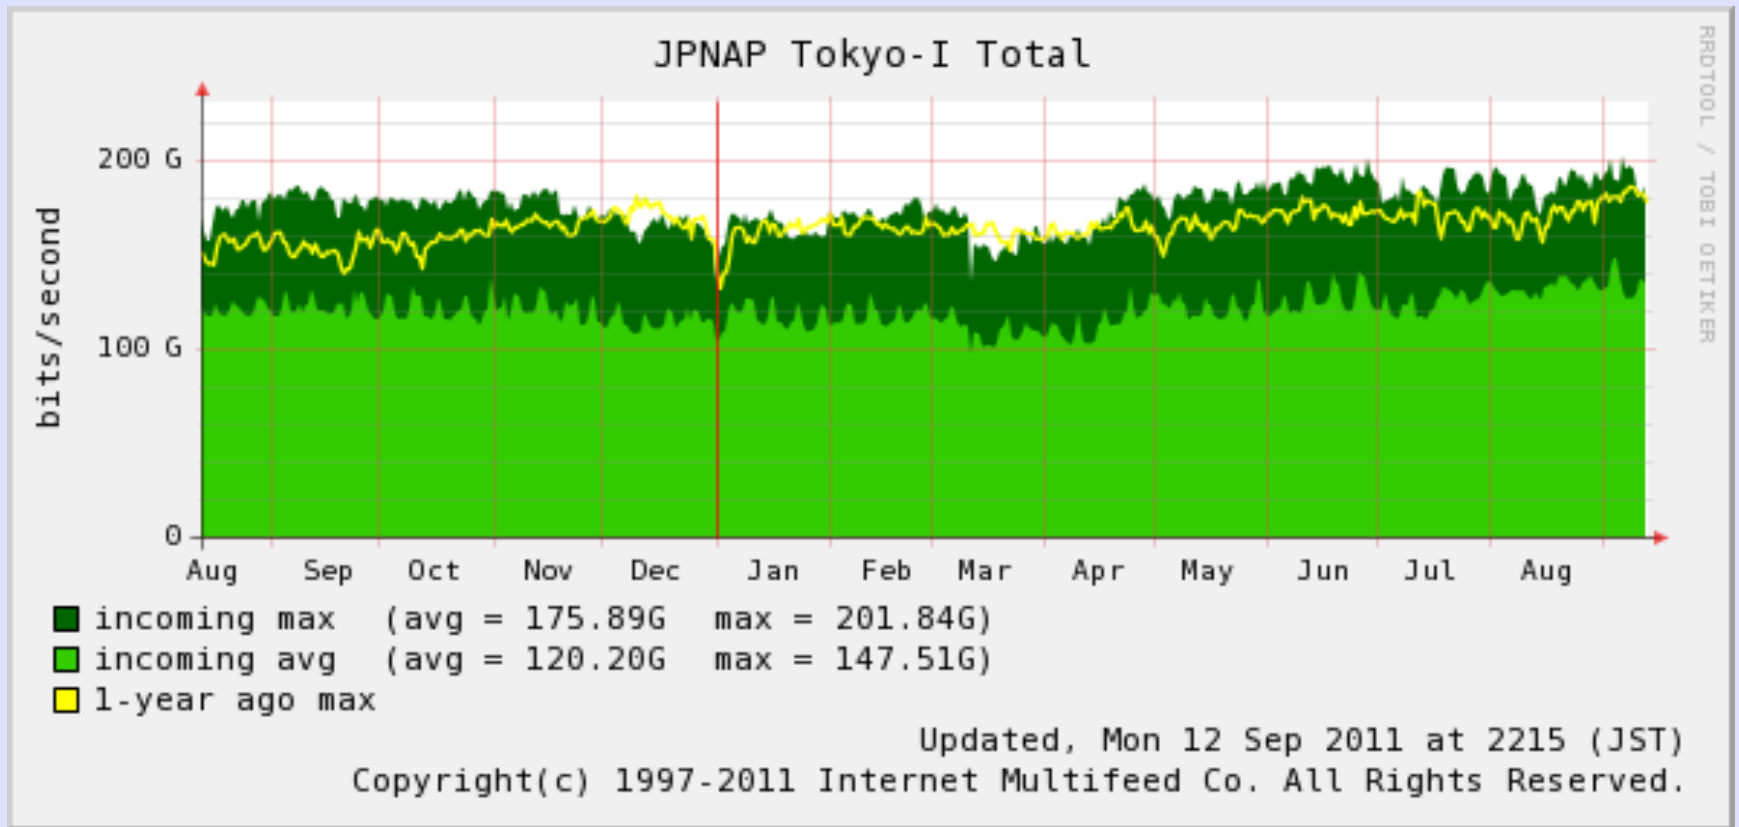


Typical un-Cut Link Stayed Up
and Had No Congestion

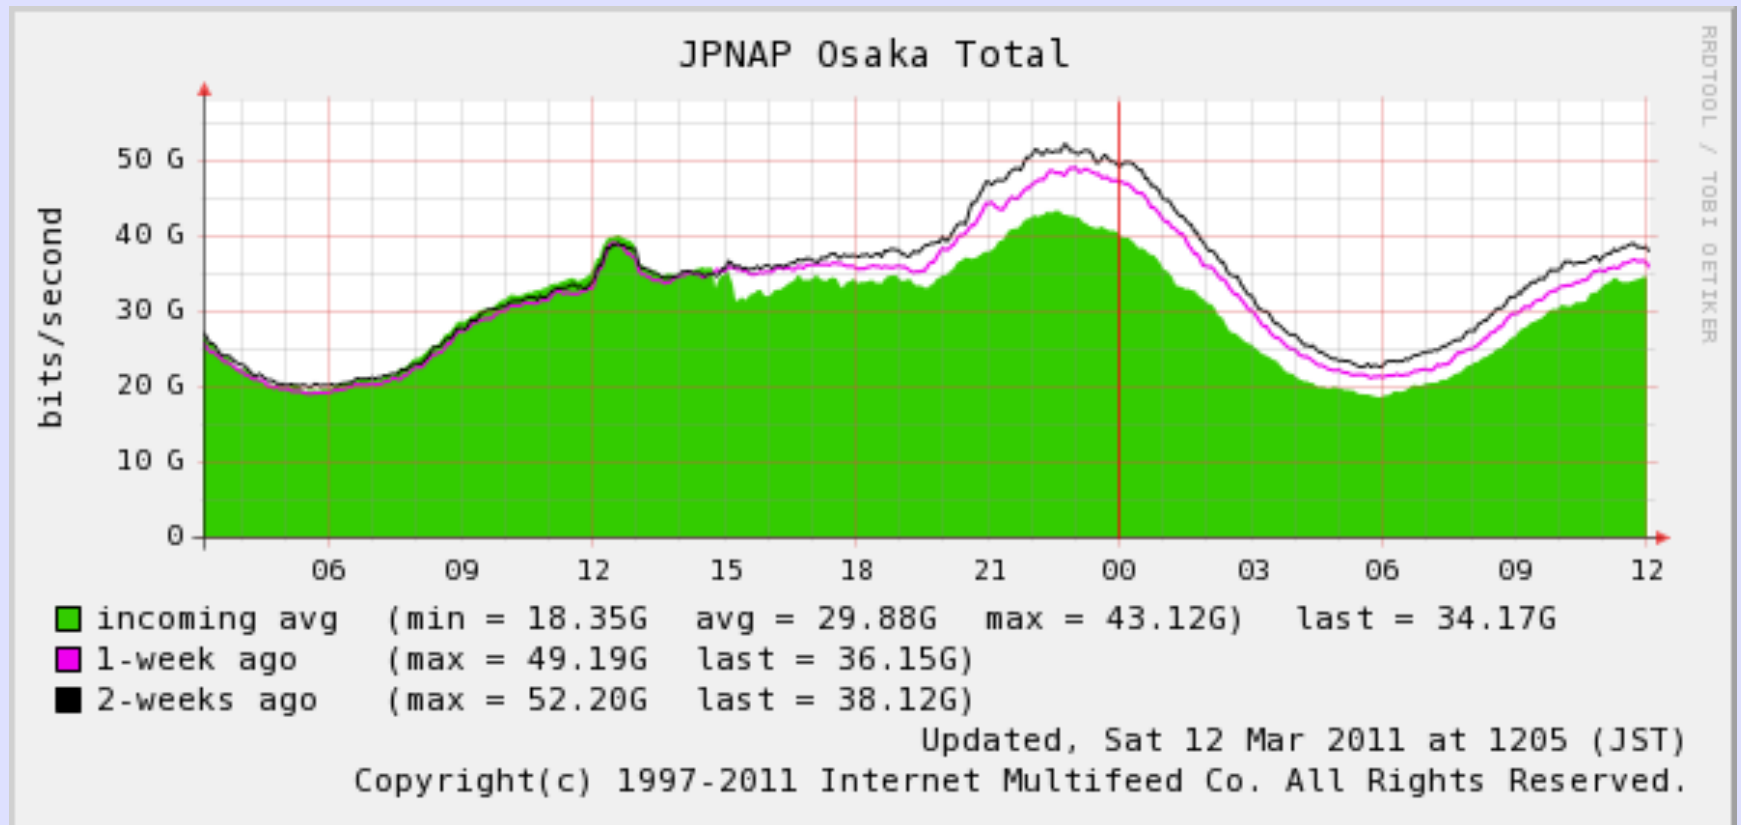
JPNAP: Tokyo 1



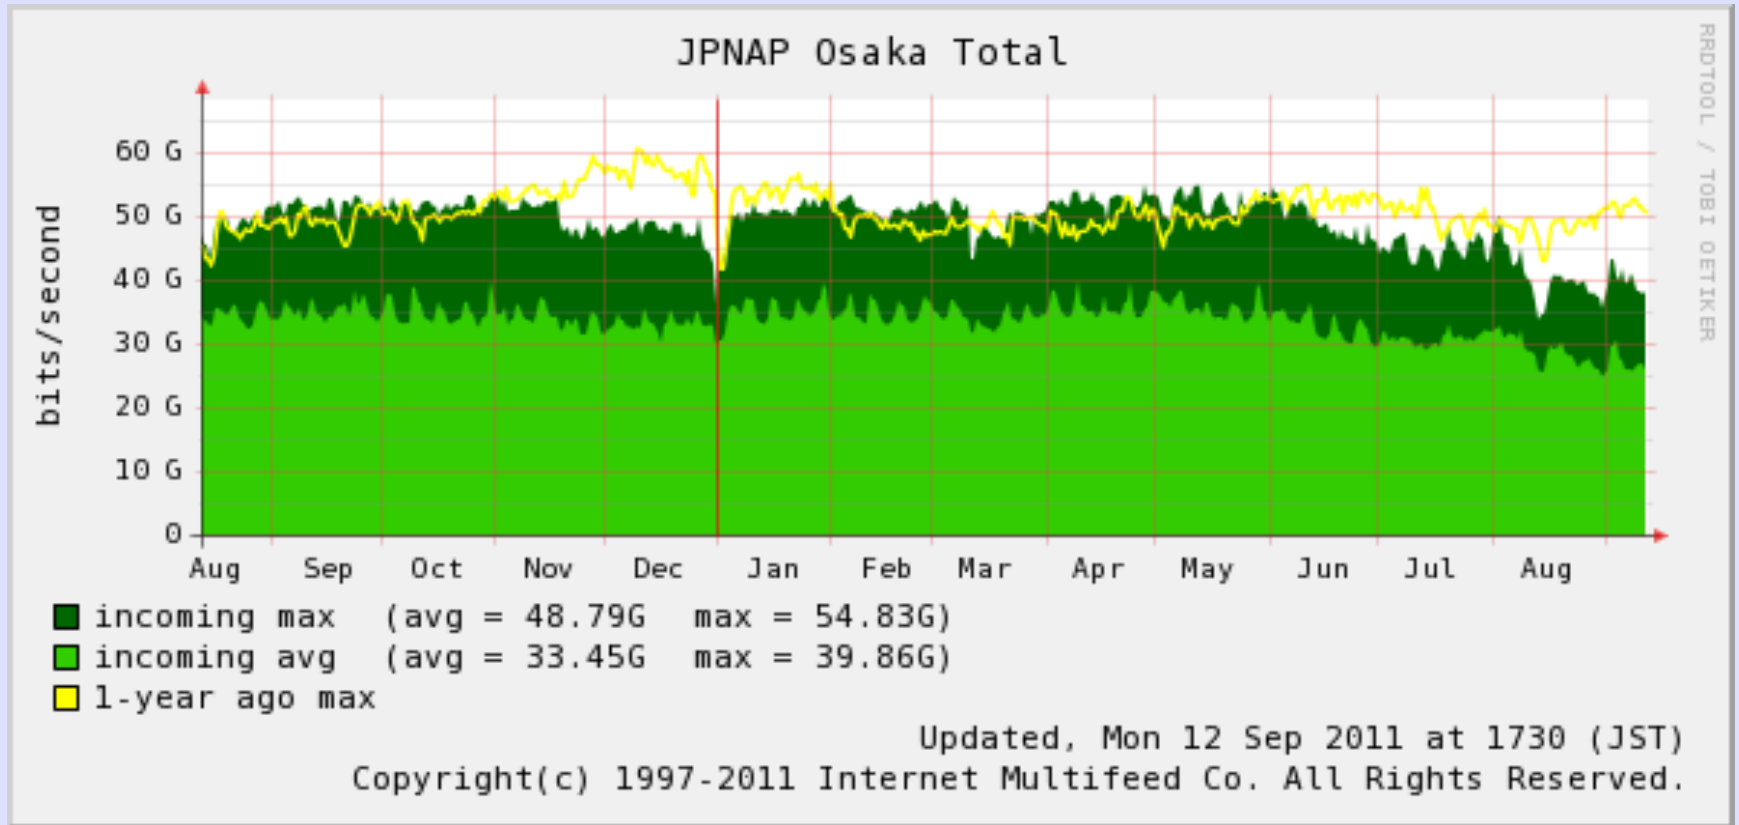
JPNAP: Tokyo 1



JPNAP: Osaka



JPNAP: Osaka



What Else Did We See?

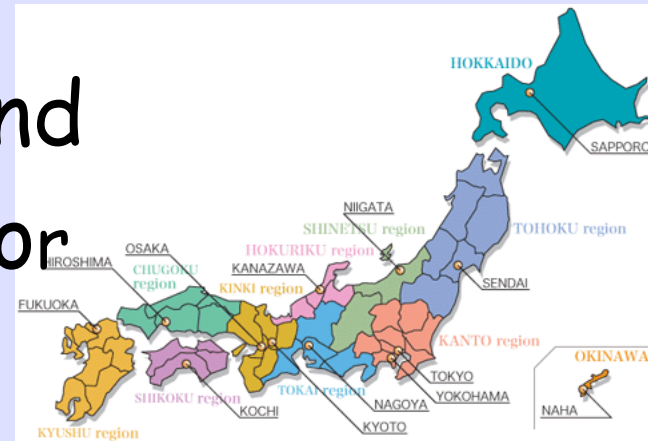
OSPF

One link to Sapporo failed because it shared fate with one of the links to Sendai, no customer effect

BGP

One neighbor router sent peaks of updates every day at the same time

It is useful to find abnormal behavior



What Happened

- Sendai was disconnected for more than 15 hours
- No effect on non-Sendai customers
- Significant trans-Pacific links impacted by the quake and aftershocks
- Almost nothing to see as IGP and BGP healed the wounds

A Simple ISP

- No MPLS, real IPv6, no secret sauce
- Just IP routing and prudent operations
- Aside from down-times due to physical isolation of Sendai, no impact on non-Sendai traffic
- Routing spikes to work around cuts
- Boring, as it should be

Routing Works

Over-Provisioning Works

The Internet Works