

IP Over Anything (anything over IP)

By: Blaine Christian

Before we Begin

- All the technologies about to be mentioned have a valid and useful set of features.
- Everybody has a different recipe and most recipes work given enough perspiration.

Frame

- Did not increase interface capability fast enough (at least on my platform). ATM drove it out with interfaces at OC12 and eventually OC48 speed.
- Extra gear in the core... Do we really need more layers of complexity or can we bundle them into one device?
- Things are different today.

Frame

- Still using it though!
 - Interesting method for sharing POS links.
 - Inexpensive Stat-MUX gain to offset those expensive router interfaces.
 - What can I do with policy routing/MPLS etc... here?
- In the core it is not clear that cost savings on interfaces helps that much (statmux towards customers makes more sense).
- Link State prioritization issues abound.

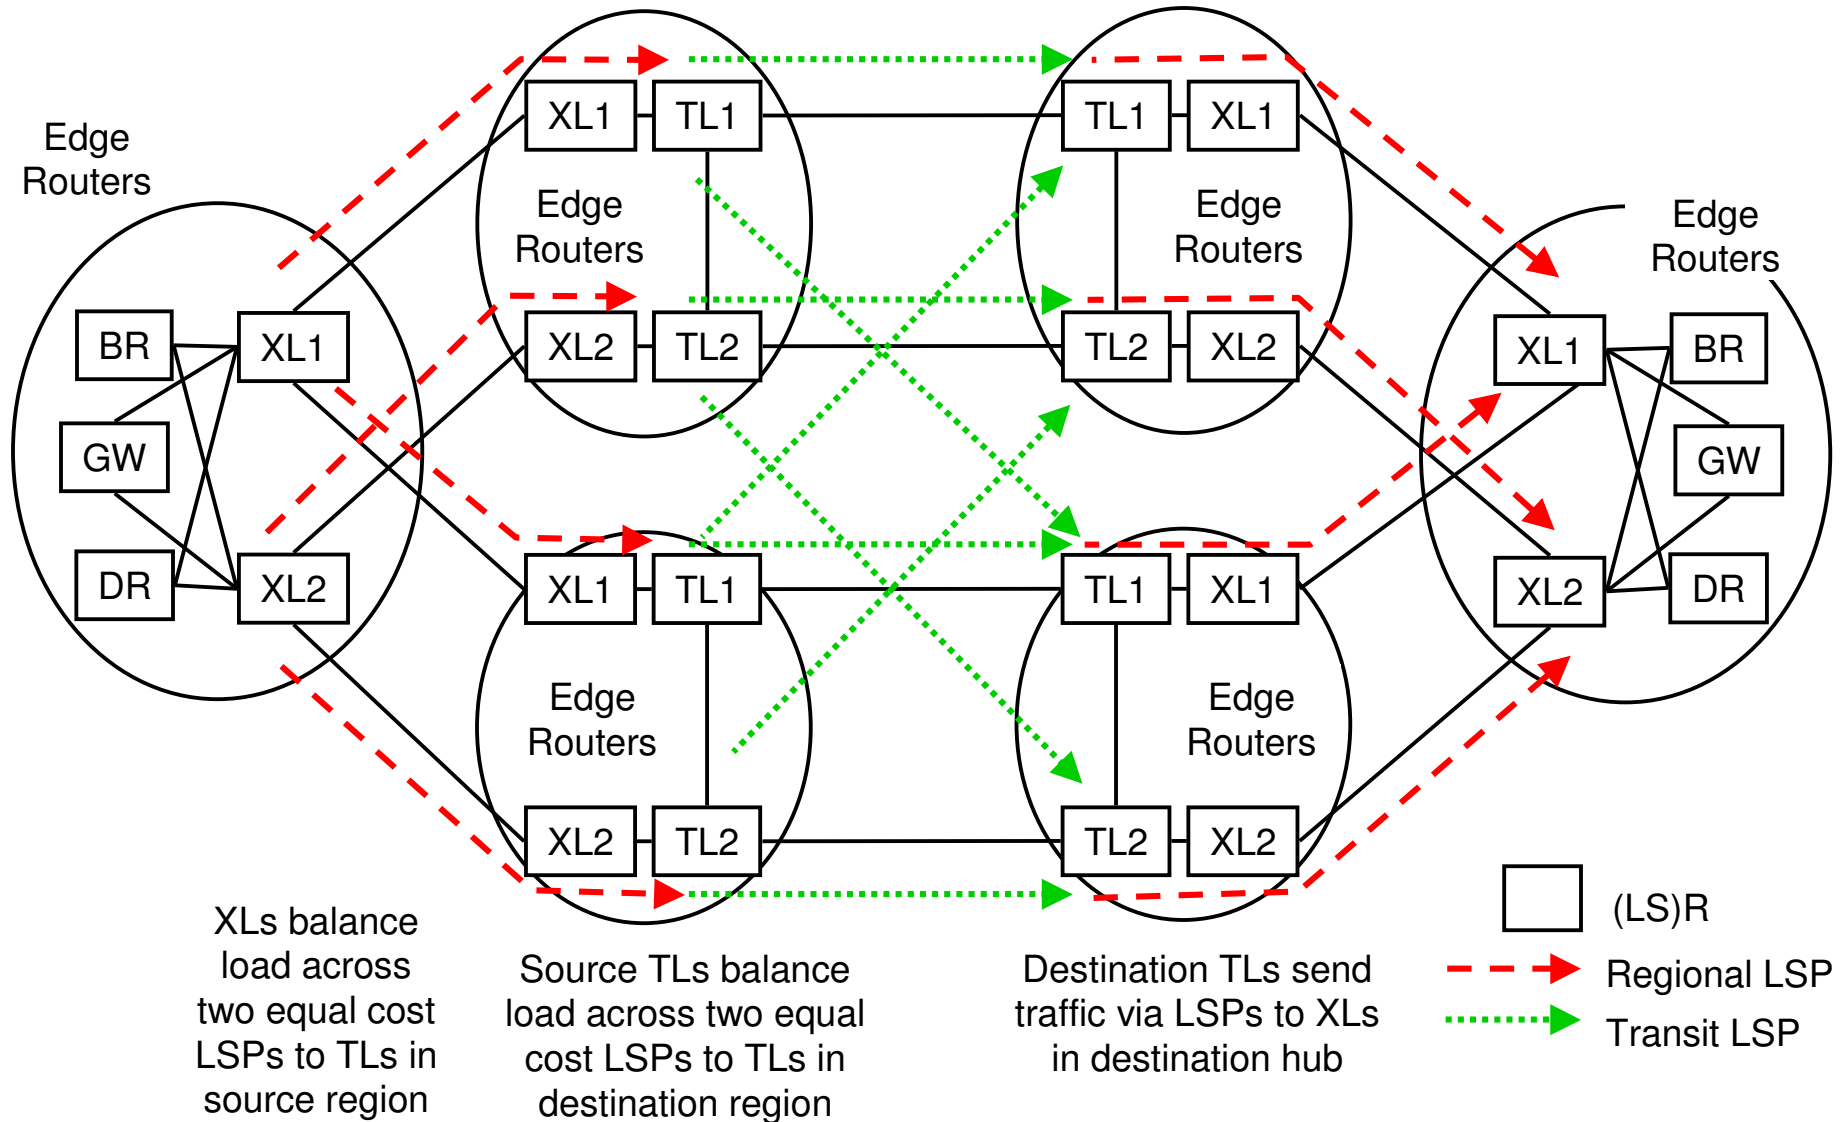
ATM

- SAR became an issue (OC12 was the limit for a long time).
- ECMP helps to overcome size based limitations.
- Once again, extra gear in your network. Can you find a mechanism that allows you to use a single device instead of logging into a dozen to determine the fate of your traffic?
- Once again, things are different today, high speed (OC48) SAR is available.

ATM (Cont.)

- For a time a two profile UBR/VBR network was all the rage. It worked, sort of, but became more of a headache than anything else so we pushed every thing into VBR.
- Frame to ATM mapping adds some life to edge switches.
- Link State prioritization again!

MPLS Architecture



IP/MPLS

- Latest greatest.
- A bit of herd instinct going on but for some good reasons.
 - Simplified troubleshooting
 - Manage your flows directly on routers instead of through intermediate devices.
 - Interesting failure protection mechanisms.
- A means to an end

IP/MPLS

- For RSVP it is good to “pretest” your LSPs before re-carving bandwidth.
- On a large scale flows are quite stable so don't change them often.
- An easy way to attach a latency penalty to a circuit and reserve a specific flow across it!
- Mind numbingly easy to manage
- Link state prioritization built in!
- Link Specific Pre-Emption (sound good?)

Which Do you Prefer?

- 8 Parallel uplinks?
- 2 Parallel uplinks?
- 1 uplink?

Which do you prefer part deux

- ATM drove massive parallelism
- Redundancy excessive and costly
- ECMP not effective for tunneled/chunky traffic
- One link obviously not the answer
- Two is “just” right for IP forwarding and the reduction of fate sharing (hmm MPLS has to “fail over” though, as do the chunky/tunneled flows).
- Fate sharing becomes the 64k dollar question (ask executive how much redundancy they want and see them squirm).

IGP

- Good for general indication of traffic direction.
- Some folks use it for much more.
- A bit slow on old devices.
 - MPLS and other L2 methods can help you avoid this (think insertion of shortcuts into your IGP with long lived timers)

ECMP

- Increased the lifespan of ATM dramatically.
- Has some implementation problems (Chunky flows).
 - Potential solution use a range of source addresses towards a single destination.
 - Create Multiple tunnels (not quite so optimal)
 - IPsec makes “looking deeper” into packets a difficult proposition
 - IPv6 Flow labels can also help
- Use it! It can reduce your fate sharing... It is better to have 50% of your customers traffic impacted than 100%.

ECMP

- Watch those hash functions... You can occasionally land yourself in corner cases.
 - A key that let's you mix up the hash can help here.

BGP

- Not a good tool for traffic management.
 - Yet people still persist (and occasionally there is no other way).
- MEDs are broken. Oscillation is annoying and potentially damaging to your ability to converge.
- Hierarchy = Good
 - Does every edge device need to know exactly what is happening on every single other edge device or can we reduce some of that information load? Too much loss of information creates sub-optimal routing though so beware!

Conclusion

- IP Hierarchy is good
- Massive Parallelism is great as long as you have perfect hashing and tiny flows
- MPLS and “native” backbone I2 mechanisms are a boon to good TE.