geographical aggregation

(to support multihoming)

in IPv6

Iljitsch van Beijnum

multihoming problem

- Today: announce own block over BGP
- BGP doesn't scale all that great
- But it's not just the protocol: routing gurus haven't been able to come up with something that is orders of magnitude better
- In IPv6: no lack of address space anymore, increase in multihoming?

IETF multi6 wg

- Little progress for a long time
- I sent in an independent submission: "Provider-Internal Aggregation based on Geography to Support Multihoming in IPv6"
- Now two design teams working on multi-address solutions

my goals

- Enable multihoming in IPv6 as soon as possible. So:
  - No new code
  - No cooperation between networks
  - Intermediate term scalability, ultimately be replaced by long term solution
  - No support for very exceptional cases

how it works

- Distribute full global routing table over the different routers in a network rather than give each router a full copy
- Could be done without geography, but then "scenic routing"
- So in every router "local" more specifics + global aggregates
- This means addresses must be assigned geographically

provider internal

- Each (ISP) network generates aggregates for internal use
- Aggregates are ***NOT*** announced to other networks

interconnection

- Every network announces all customer routes to all its peers everywhere
- Every network gets to aggregate in a way that fits its topology
- Interconnection doesn't have to be in target area!
- Bad interconnection means bad aggregation, but still reachable

example

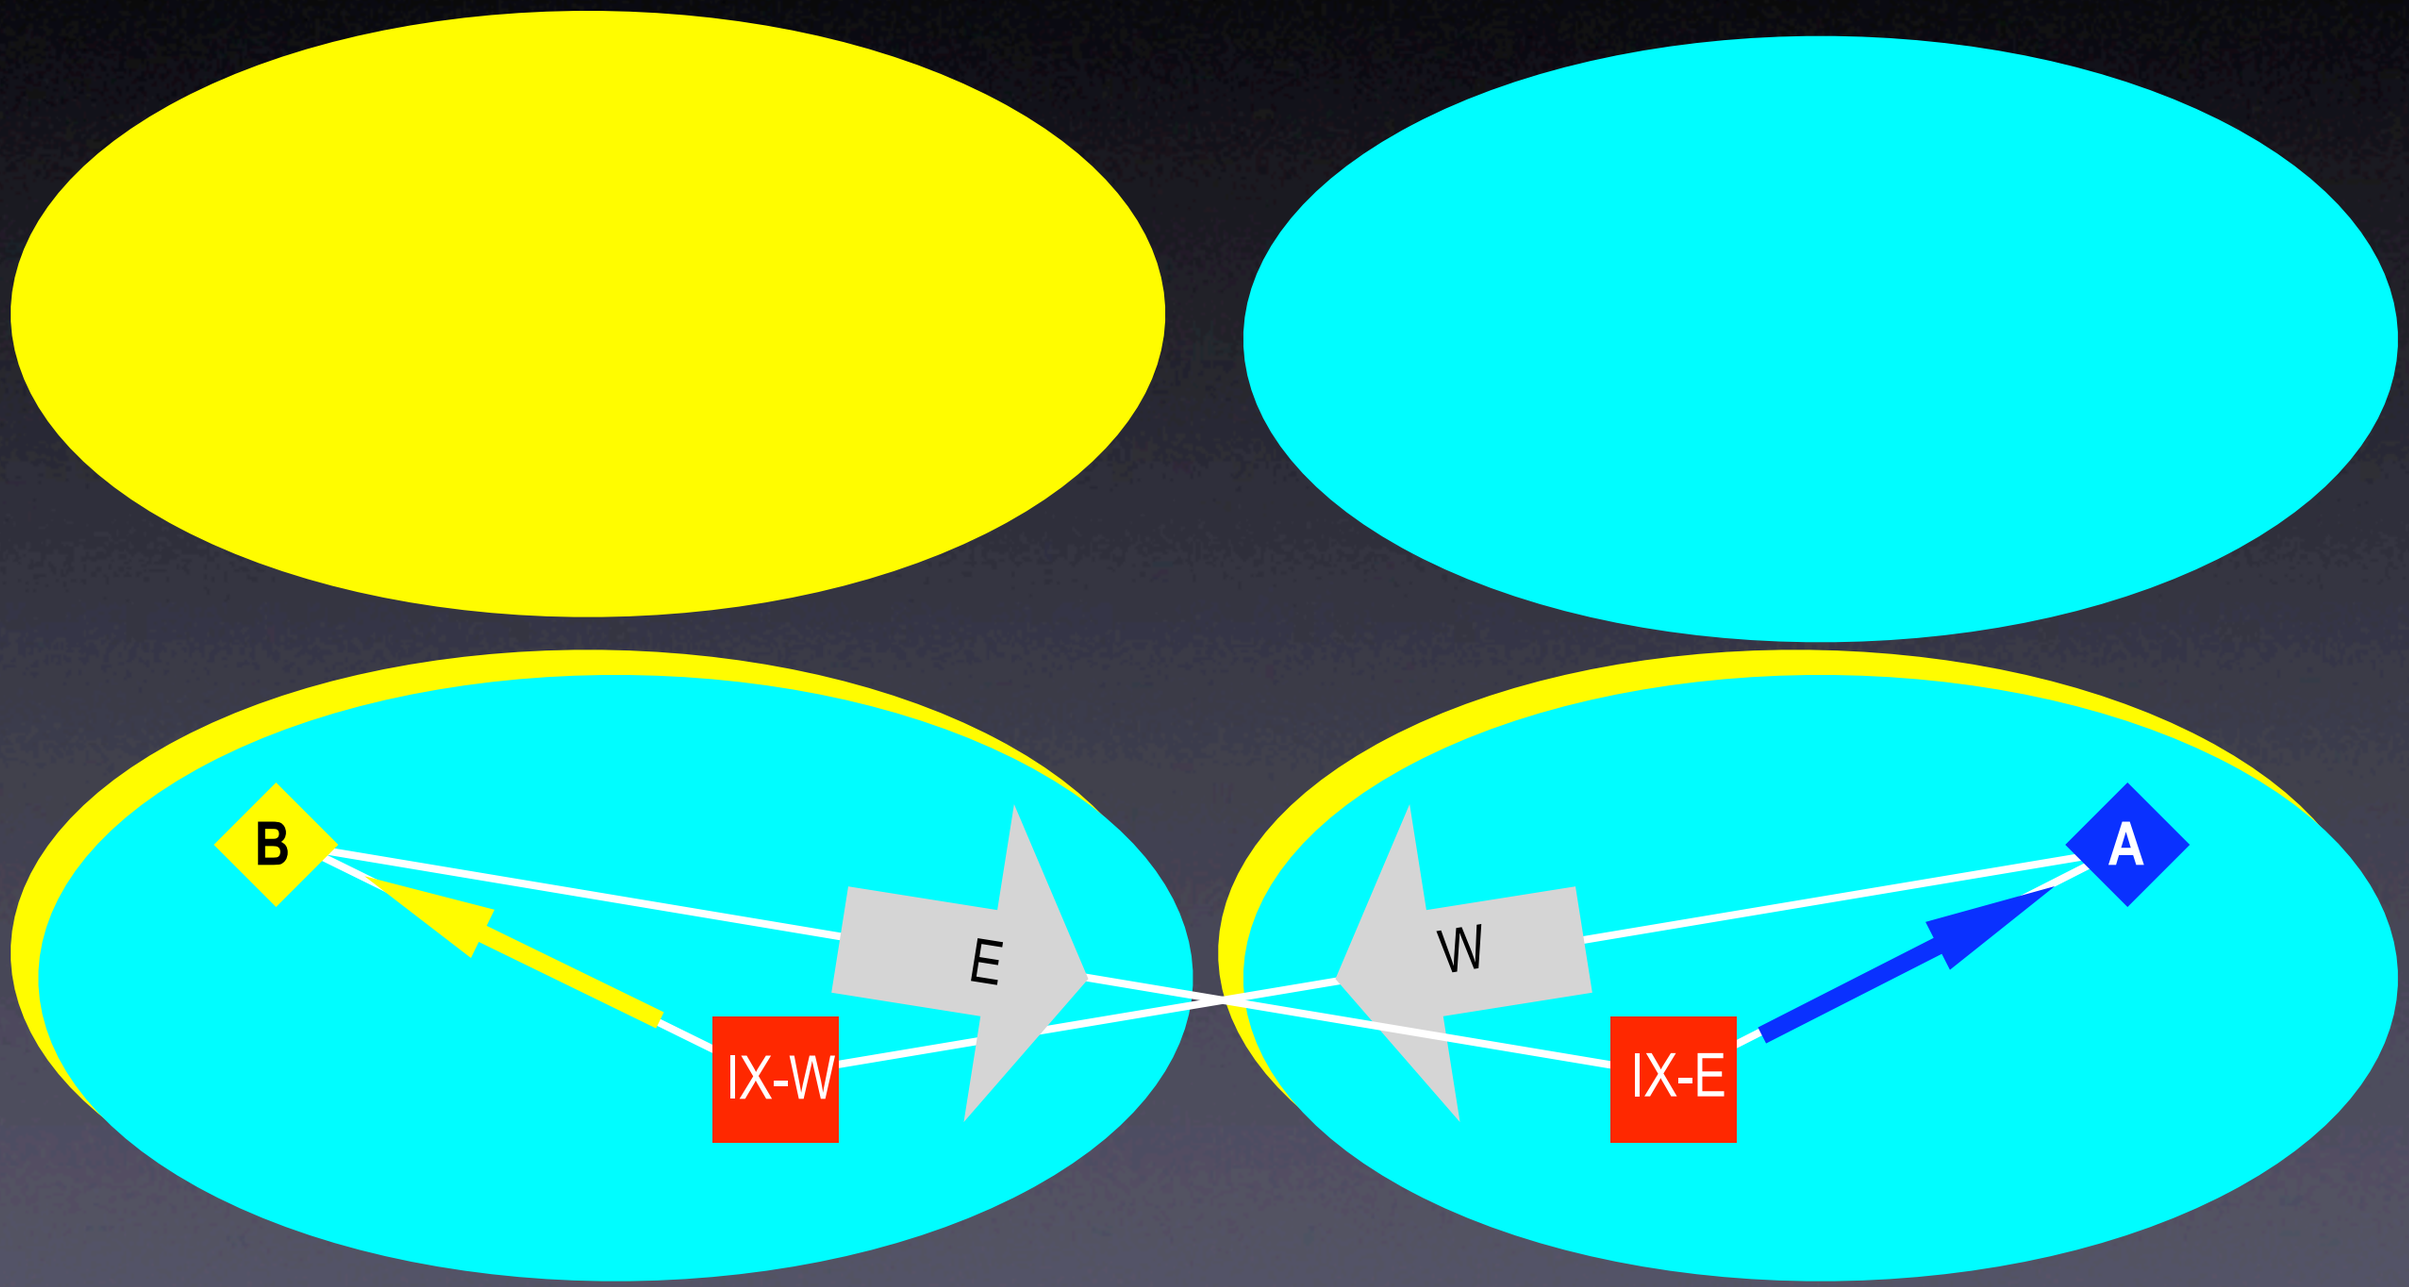

very few downsides

- Only need geographical address assignment to start multihoming *immediately*
- Implement aggregation in each network independently when deemed desirable
- When we get locator/identifier separation geo addresses can be identifiers, automatically clean up routing table

why still useful?

- We can start multihoming *now*
- Possibility of aggregating allows flexible memory/processing vs optimum path tradeoff
- Future developments...?
  - more dynamic routing?
  - several addresses with different routing?

routing vs switching

- Routing: complex software, lots of memory and per-packet processing (but smart)
- Switching: power-hungry CAMs (but fast)
- Optical switching: 15 ms circuit switching time (but cheap, simple and future-proof)
- Cost: 100 : 10 : 1
- In the future dynamic optical paths? Introduces dynamics into routing!

multiaddressing

- Multiple addresses with different routing, let host figure it out?

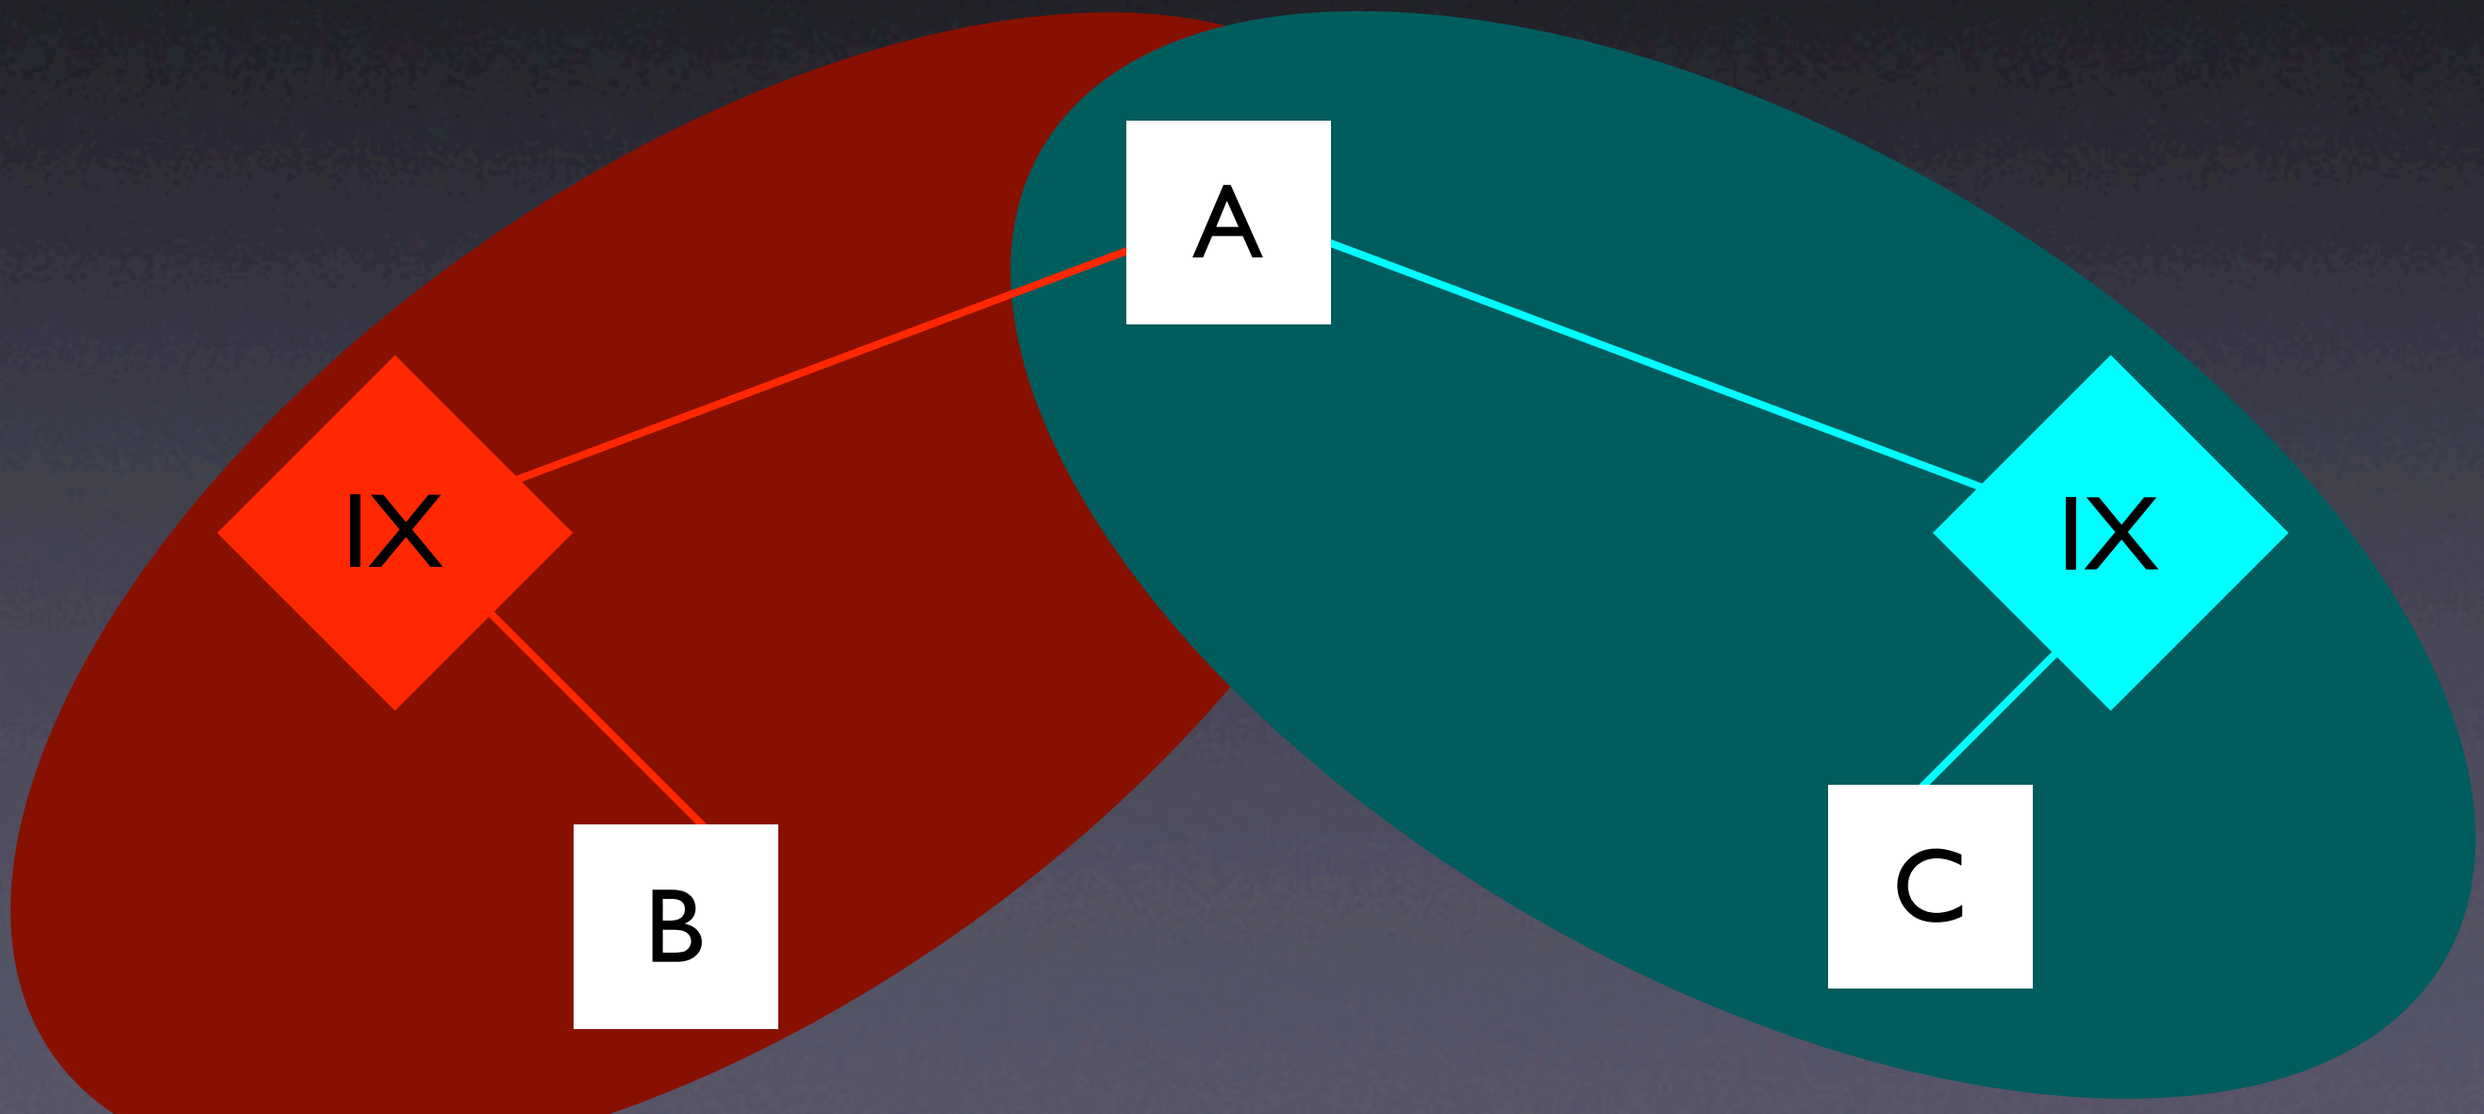

questions?

- URLs:
  - http://www.ietf.org/html.charters/multi6-charter.html
  - http://www.ietf.org/internet-drafts/draft-van-beijnum-multi6-isp-int-aggr-01.txt
  - http://www.bgpexpert.com/presentations/