AtlanticWave:
It's finally happening

Dan Magorian, MAX
Director of Engineering and Operations
magorian@maxgigapop.net

Quilt Peering Workshop, St Louis MO Oct 12, 2006

Here's what we told everyone 18 months ago about what A-Wave would be, still true today:

- A-Wave is an International Peering Fabric
  - US, Canada, Europe, South America
  - Distributed IP peering points:
    - NYC, WDC, ATL, MIA, SPB
- SURA, FIU-AMPATH, MAX, SoX, MANLAN, FLR, and in partnership with the Academic Networks of Sao Paulo (ANSP) are combining efforts to establish AtlanticWave
- A-Wave is an integral component of the NSF IRNC WHREN-LILA project to create an open distributed exchange and transport service along the Atlantic rim.

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

There were the usual national peering pictures with lots of lines

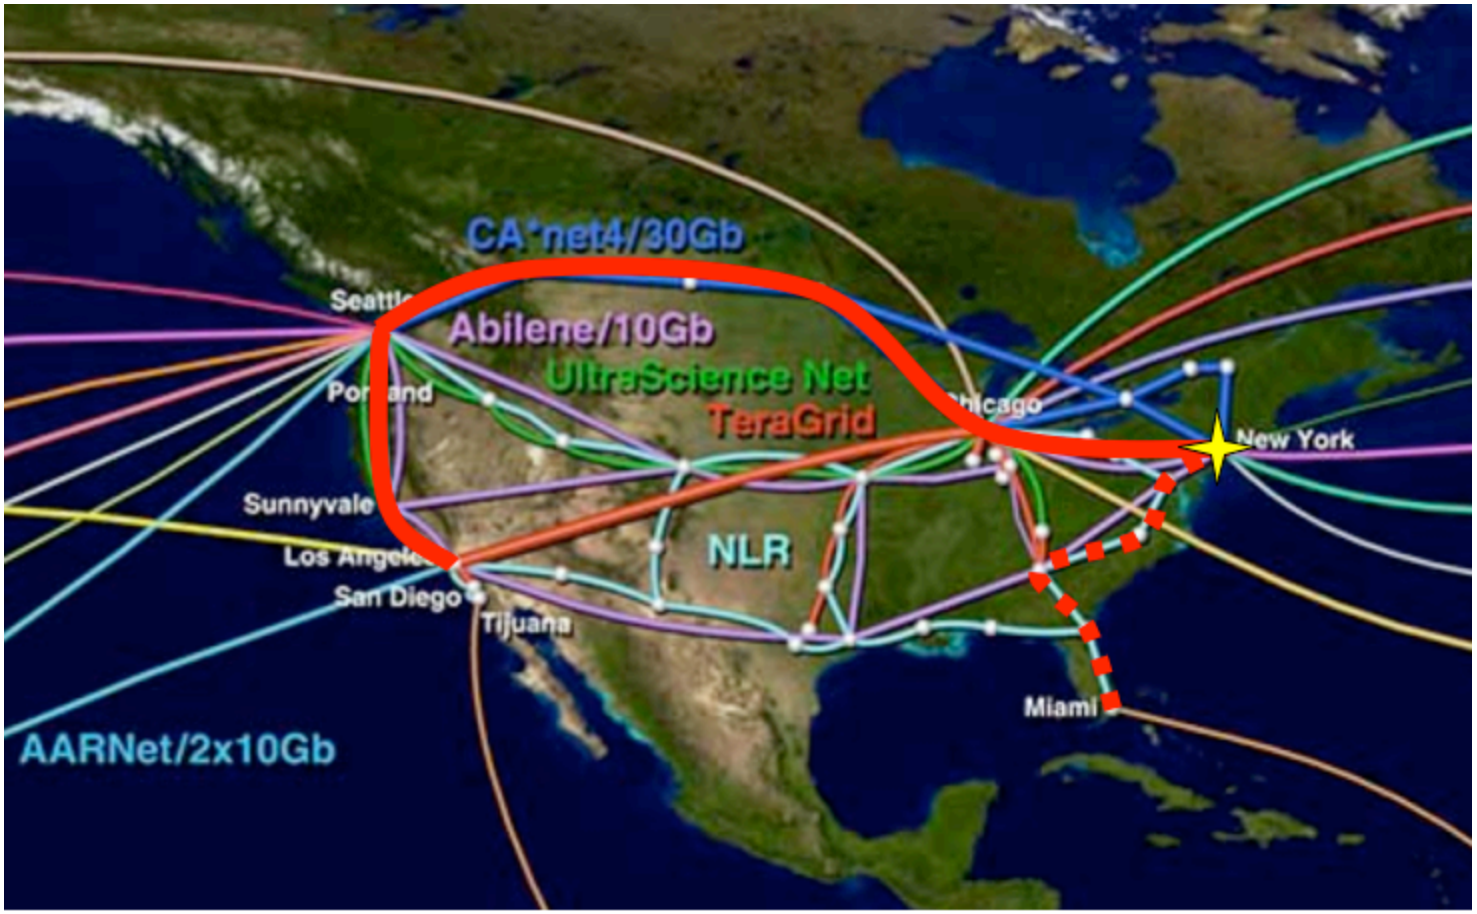

And international pictures talking about A-Wave integration with GLIF with even more lines

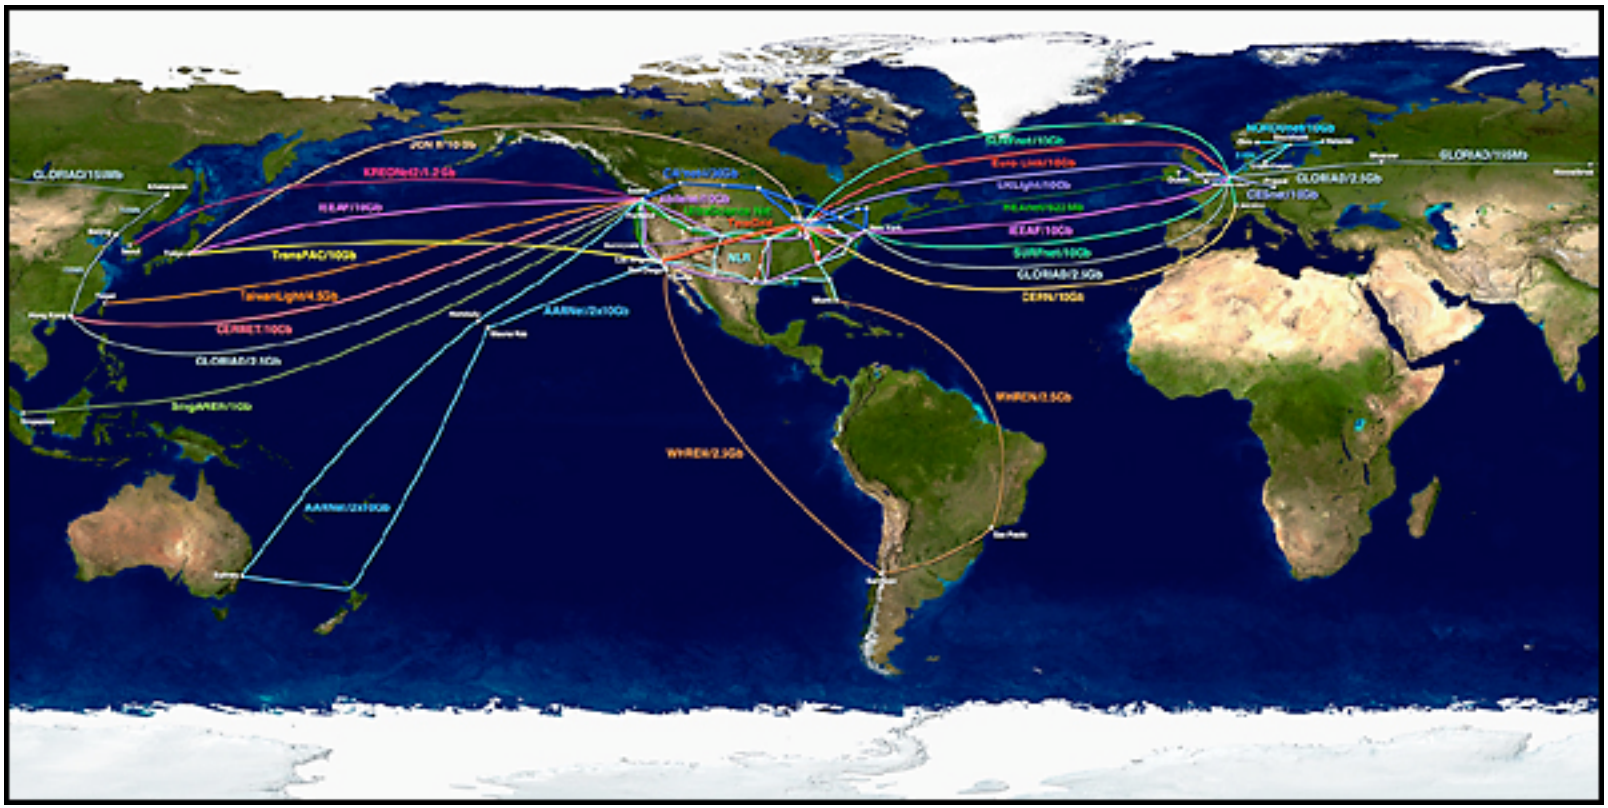

There were baby pictures about The Strategic Picture and how we intended to use Next-Gen Sonet

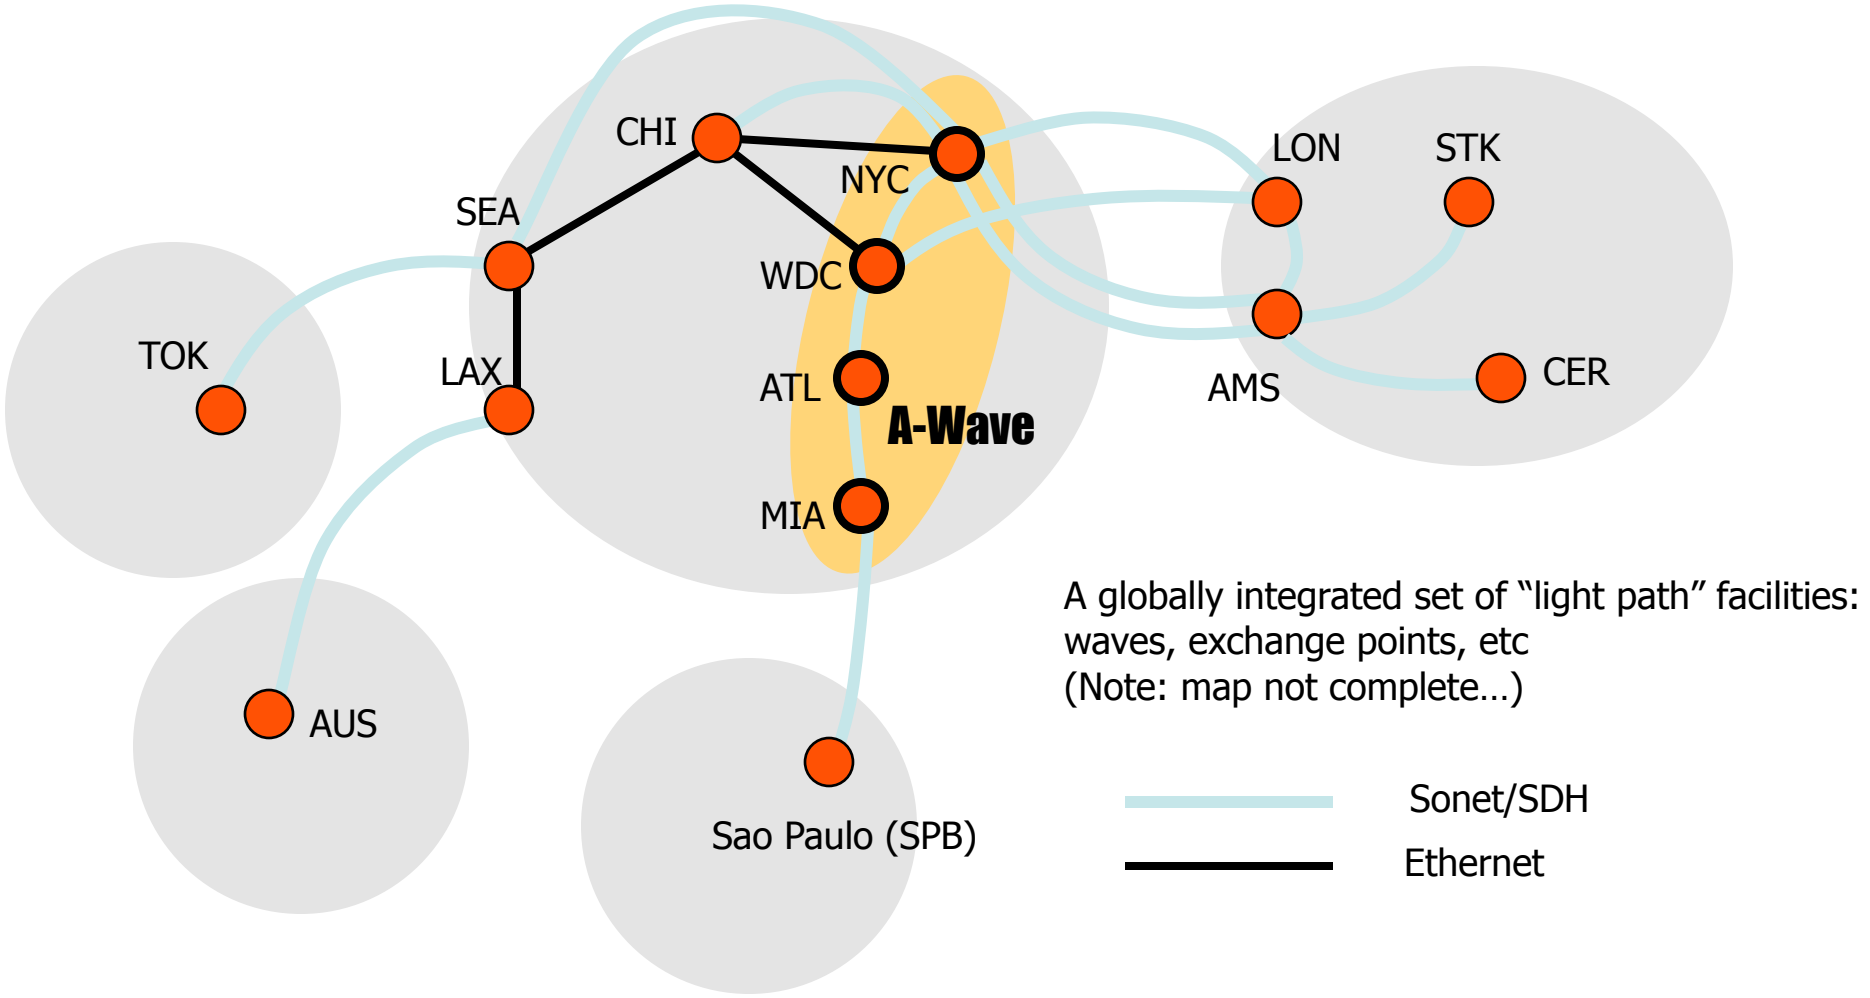

Rationales for why next-gen sonet was the right approach

- **Compatibility with existing network links**
  - **The vast majority of international links (all?) are presented at the exchange points as Sonet.**
    - Trans-oceanic links are all Sonet/SDH (IEEAF OC192, EuroLink OC192, NetherLight, South America, Japan and Australia…)
    - Canadian links are Sonet
    - Commercial services are [still] mostly Sonet or SDH
  - **With a superSonet framed backbone, A-Wave can transit VCG light paths directly from the inbound Sonet circuit to the outbound sonet VCGs without adding [unnecessary] decap/ encap steps & cost, without inserting [unnecessary and poorly understood] switch buffering, and preserving the synchronous and deterministic flow characteristics across the core.**
  - **Some such links may require re-configuration**
    - E.g. OC192c reconfigured to 4x OC48c, OC48c to 2x GFP-F GbE
    - Or: OC192c front ended with VCAT/LCAS capable switching gear
- **$$$ New generation of Sonet/SDH switches and DWDM optical gear are no more expensive than Ethernet**
  - **Most 10Gbs transponders/tranceivers for DWDM applications are "UNI PHY" – I.e. software configurable for LAN, WAN, or Sonet service – so the cost is the same.**
  - **Most of the major manufactures are already offering either rate selectable 2R transponders or GFP encap/decap of 1GbE for 2.5 Gbs interfaces**
  - **Note: integrated Sonet/Ethernet switches are just now reaching the market.**

Details on an A-Wave Node Architecture to accomplish that

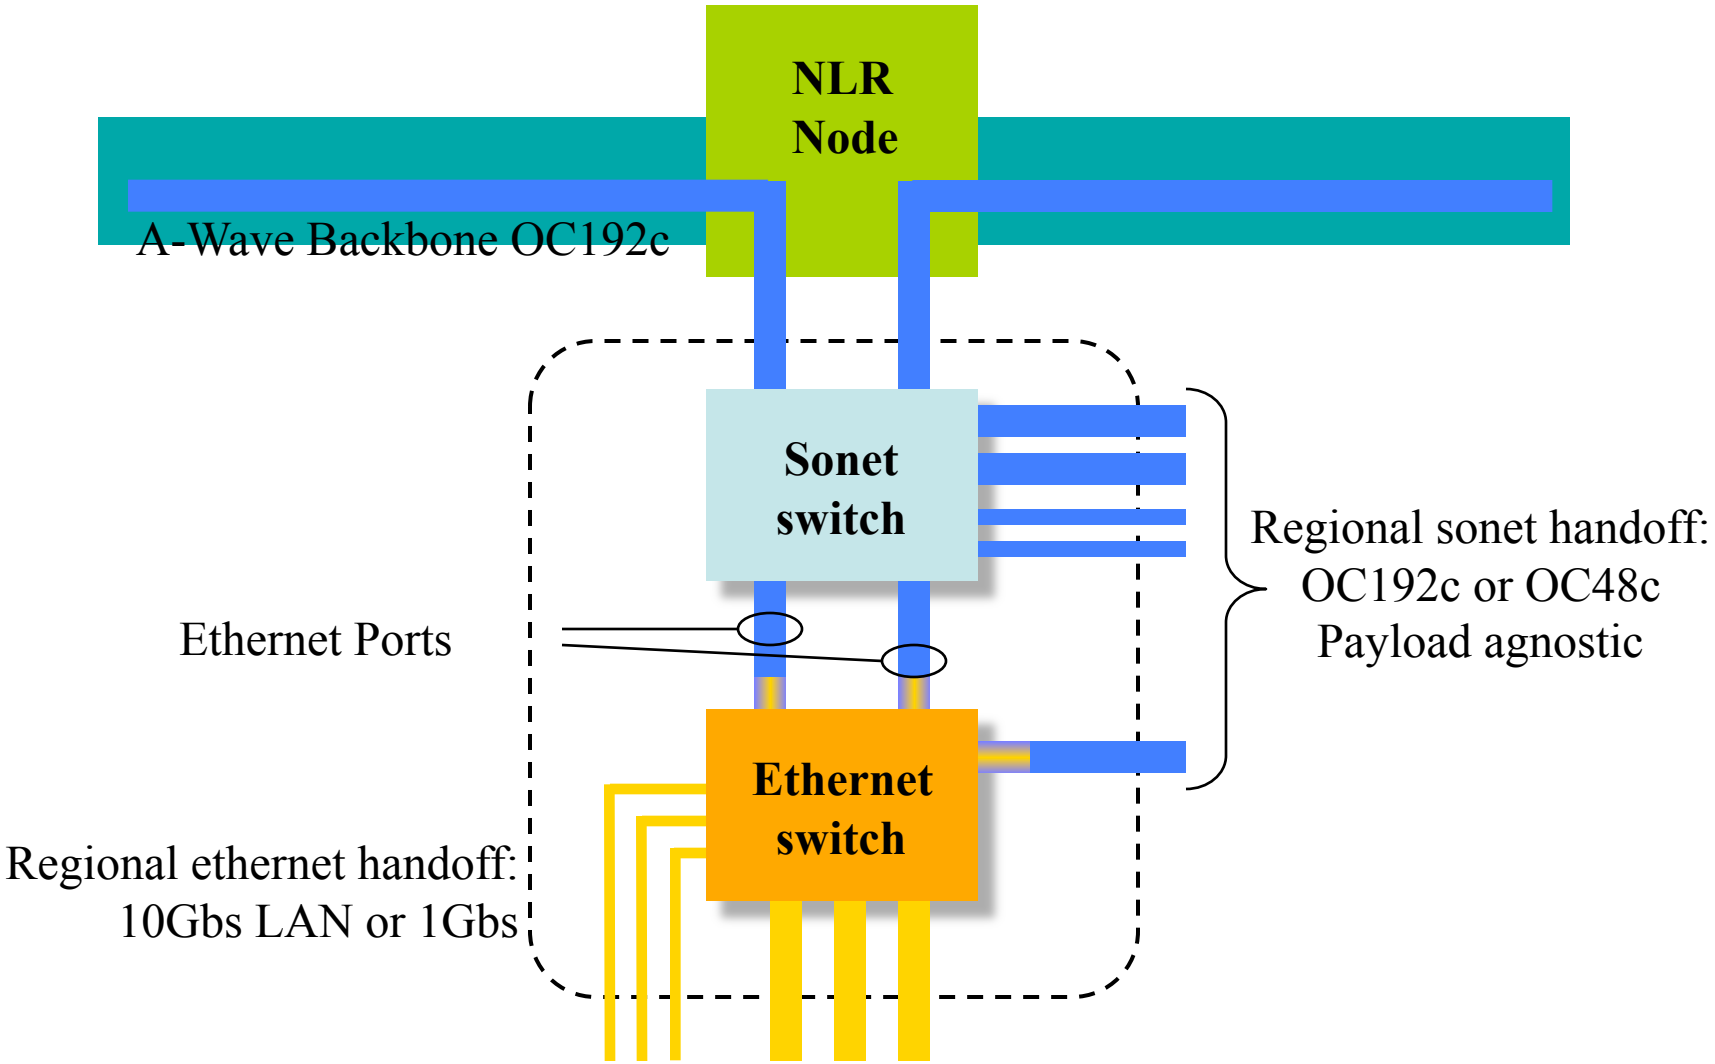

So what happened to this picture, and why did it take so long?

It turns out that herding 5 cats is a lot harder than herding one or two!

- Especially when they're academic cats
- SURA nicely paid the $$ for the NLR lambdas, but nothing else was funded.
- So the existing exchange points didn't really see the need to pay lots of $$ for sonet switches when they had perfectly good ethernet ones already.
- Even doing the MOUs with each of the institutions to accept the SURA lambdas turned out to be complex and very time consuming

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

So where are we today?

- The MOUs finally got (almost, almost) done!
- The 10G NLR lambdas finally got provisioned.
- The FLR lambdas to get from Jacksonville to Miami finally will be turned up this week
- We're Iperfing to test end-to-end performance NYC-Miami and on to Tampa for SC06.
- We have lots of folks queued up to use A-Wave for SC06 demos (everyone loves free bandwidth)
- And of course we're using ethernet with cool next-gen sonet lightpaths as phase II.

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

Revised Strategic Picture,
"initially ethernet"

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

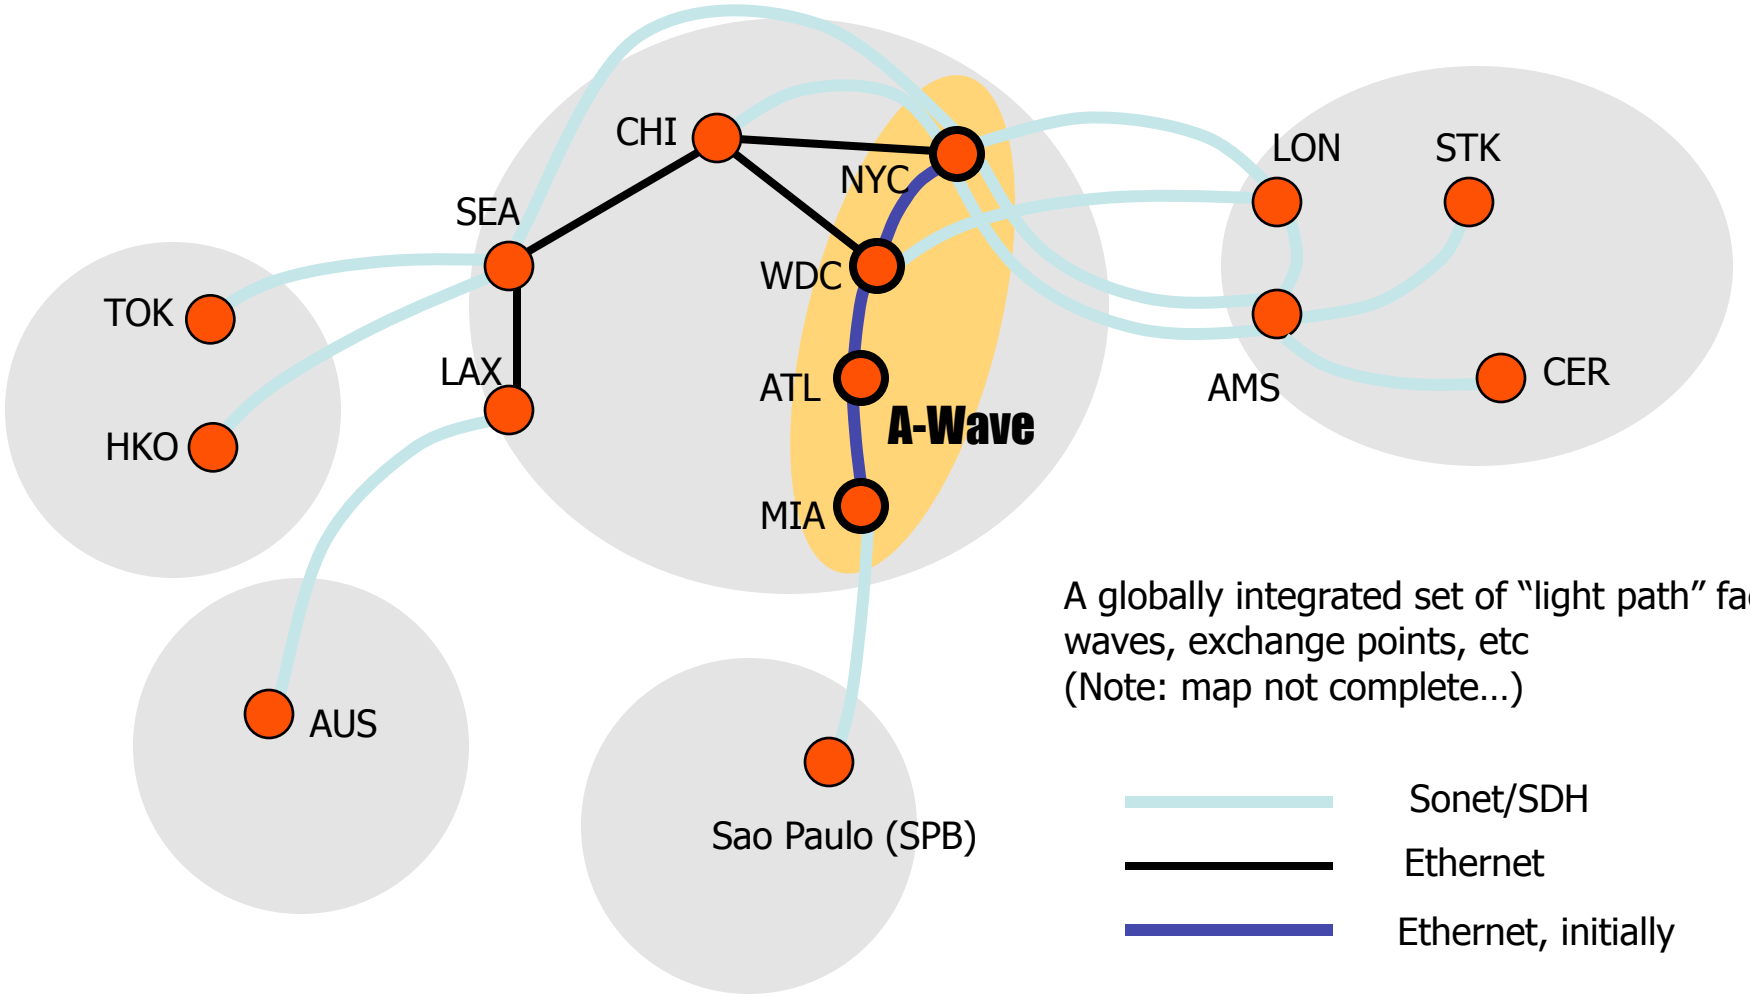

A globally integrated set of "light path" facilities: waves, exchange points, etc
(Note: map not complete…)

Sonet/SDH
Ethernet
Ethernet, initially

Of course we have a vlan plan, but we used p-t-ps instead of shared like P-Wave

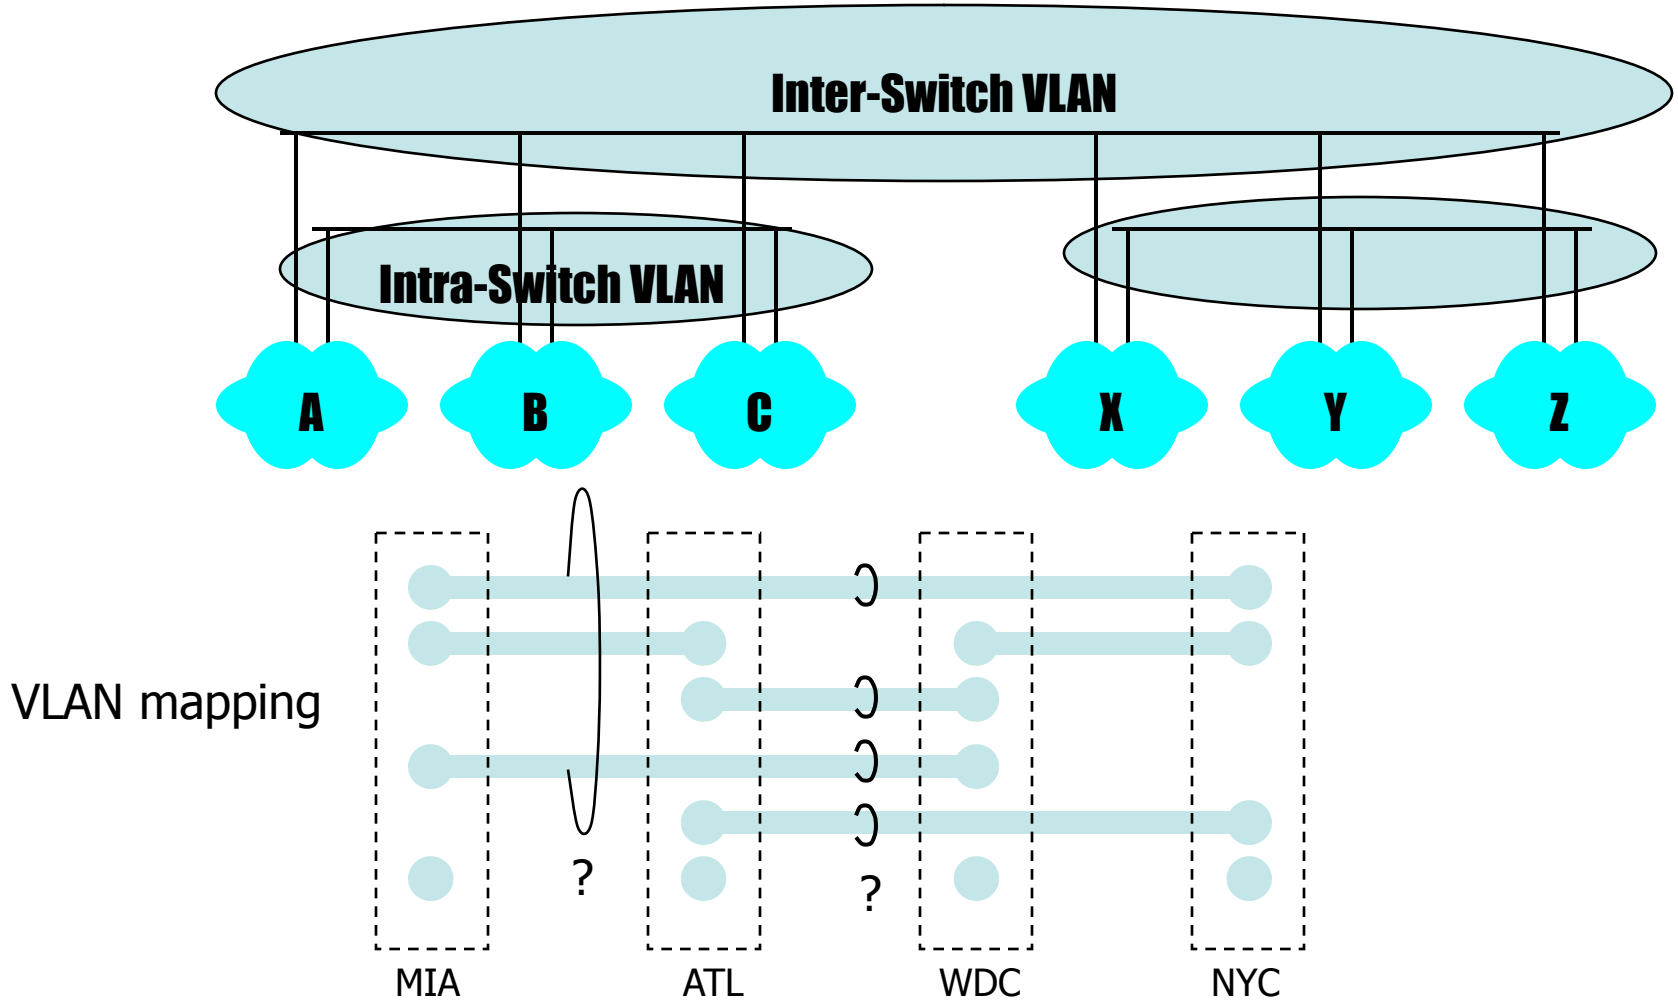

Timetable

- NYC MANLAN to NGIX/E DC segment will be announced to the peers at November JETnet mtg
- Internet2-CLARA peering over A-Wave scheduled
- Since many of the SC06 demos are Chicago – DC – ATL – MIA – Tampa, that part of A-Wave won't be open for peering business until after SC06.
- After SC06, the 10G connections to Chicago and CAnarie will stay in place, plus GEANT2.
- So this means the coast-to-coast and international "Strategic Picture" is happening pretty quickly.

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

What are the implications of this?

- Right now, everyone is talking about NLR and Newnet, sustainability/viability of two R & E backbones, & whether competition is good or bad.
- Eg, if a connector is paying $500k/yr for 10G I2 plus $800k/yr for NLR. How long can that go on?
- But consider: the lashup of STL with CHI with WDC that ties P-Wave with Starlight with A-Wave creates another "free" backbone, which I call "Ad-hoc net". No trustworthy central IU operation; all access, rules, business plans TBD.
- People might remember that this was the original "Quilt vision": people throwing lines to their neighbors, now happening 10 years later. RON-to-RON peerings: they are happening.

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

So what will “Ad-hoc net” backbones be used for? GLIF activities meet bean-counters

- Personally, I would like to see production peering traffic remain on well-managed well-funded backbones, and
- “Ad-hoc net” (MorphNet, but that name is taken) used just for 1) bilateral backup arrangements, and 2) optical testbed interoperability and research activities that don’t need “9s”.
- That sounds fine, but there is definite overlap with at least some of NLR and Newnet’s L1 and L2 offerings. Sure, one 10G pipe isn’t that much, but pipes are cheap these days. It could grow quickly topsy-turvy style with strong demand.
- Especially if it saves people on expensive “club costs” to join. We’ll see how much cost is a driver.

AtlanticWave

NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

This is a Peering Workshop, so what is A-Wave's Peering Policy?

- Maybe we should all do Peering Personals like at NANOG. "Do you have a pulse?"
- Seriously, with 4 exchange points, A-Wave initially is a no-cost value-add extension/new service for each XP's customers. Eg, we announce to NGIX/E peers that they can now peer with folks at MANLAN, SOX, AMPATH, etc.
- Beyond this, A-Wave is moving towards a "corporate front" so that people can approach A-Wave as an entity with a consistent interface, Consistent people – very soon (Nov). Common fees, business models, governance - later.
- Who will be A-Wave participants & peers 2 years from now? We don't know yet, but we're on the ground now and going!

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

Thanks to the cast of significant players making it happen

- Don Riley, for brokering the lambda deal with SURA and overall stewardship of the project.
- Julio, Ernie, Michael and the FIU folks for hosting and coordination of the project, esp for SC06, in addition to their role with their MIA XP.
- The XP operators: Brian and Cas at ATL, Bill and Christian at MANLAN/NYC; Dave, Chris, and Quang at MAX.
- Doubtless others I have forgotten.

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

Last techie bit: Let's talk about "vlan registries" (not A-Wave per se, but we tripped over it)

- "They don't exist. The idea is absurd, there are only 4k of them, it'll never scale". Say IP heads who favor layer 3 for doing things.
- But wait: aren't people creating extended layer 2 topologies anyway cross country and internationally? That seem to work?
- And there are whole new virtual ethernet protocols being deployed on carrier routers to make layer 2 heads happy and buy services.
- So vlans can be coordinated within one XP easily, within an enterprise (eg a campus) with some difficulty, between 4 independent XPs with more difficulty, but when we connect up Chicago on to Seattle and Canada, the chance of free transit vlans becomes almost nil without renumbering

  and re-doing peerings..

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

“Vlan Registries”, con’t

- Yes, we could avoid the problem with shared vlans. But p-t-p vlans are nice, we’re used to them and like them.
- What we need is vlan translation on ethernet switches without performance penalties. Ciena CoreDirectors can do it.
- Since we don’t have that, hence the need for a mechanism for coordinating vlan assignments.
- I have heard that with European RONs, the Germans are acting as the vlan number authority, and that other NRENs find them a bit autocratic. After doing this for A-Wave, I’m finding myself sympathetic with the Germans.
- We’re not talking about IRRs, it isn’t global. Each project has a particular scope. But we need better tool than what we have now. Not clear NDL will do this.

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO

Thanks!

www.atlanticwave.net

AtlanticWave
NEW YORK WASHINGTON ATLANTA MIAMI SAO PAULO