# Network Shaping 101

by [redacted]

# This presentation is classified:

TOP SECRET//COMINT//REL TO USA, AUS, CAN, GBR, NZL

Derived From: NSA/CSSM 1-52
Dated: 20070108
Declassify On: 20320108

# What This Will Cover

- Caveats
- Example network we will work with
- What shaping would look like for that network
- Basic shaping problems
- A bit more advanced shaping problems

# Initial caveats

- To understand how to do shaping, and why it does/doesn't work sometimes, you have to go back to networking basics
- To get the most of this presentation, you should already understand how IP's, CIDR's and Autonomous Systems (ASN) work
- Some ips/facts are just made up. This presentation uses Yemennet as our target network. This info is outdated and incomplete. Don't use any of this information for any real analysis

# You’re gonna talk about Layer 2 shaping right?

- No
- It is extremely situational and only worth talking about if you are in a position where you have the right kind of access.
- Until then, Layer 3 shaping is where it’s at (in my opinion)

# Example network - Yemen

- Yemen has 1 ASN (AS12486)
- We'll pretend it has 6 upstream providers
  - Mobily (AS35819)
  - TATA (AS6453)
  - FLAG (AS15412)
  - PCCW (AS3491)
  - STC (AS39386)
  - SPRINT (AS1239)

# AS12486

This network owns the following IP ranges:

46.35.64.0/19, 89.189.64.0/19, 46.35.72.0/21, 109.74.32.0/20, 46.32.80.0/21, 109.74.40.0/21, 63.168.168.0/23, 109.200.160.0/19, 63.171.18.0/23, 109.200.168.0/21

So, when we reference AS12486, you can assume it includes any IP address that falls within any of the above ranges.

# Connectivity for this network could be viewed like:

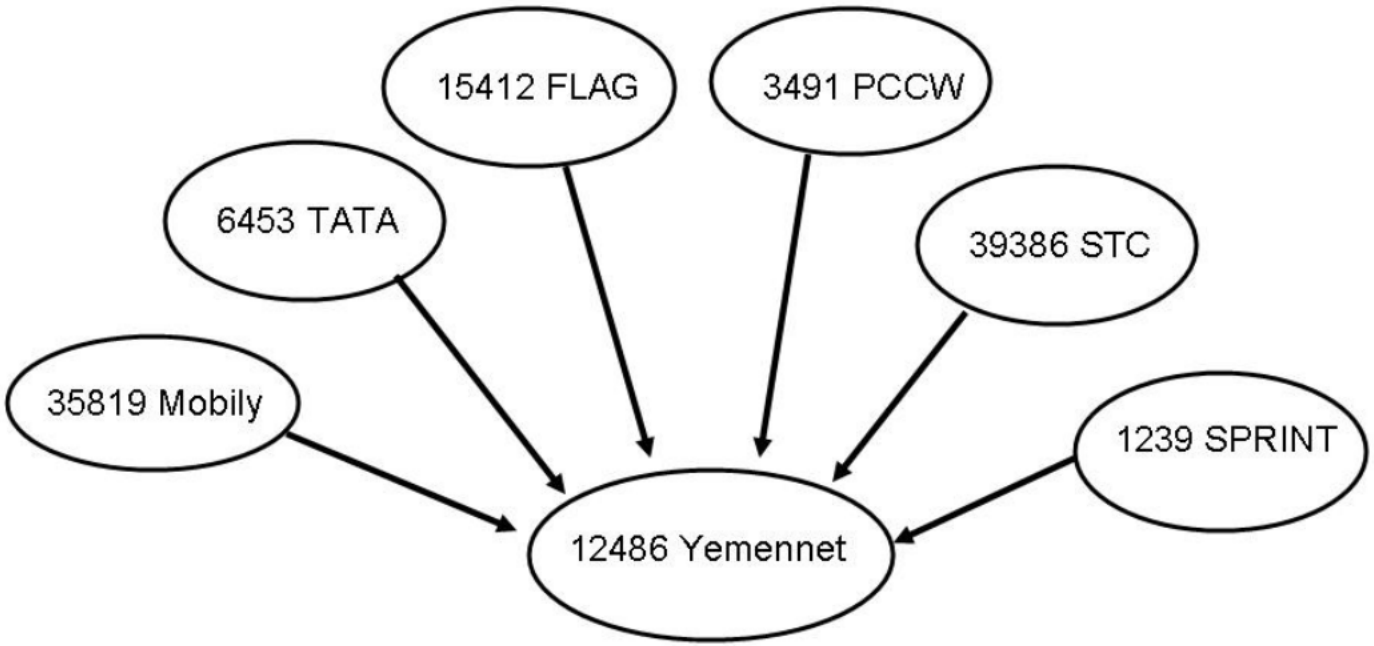

So, for traffic to get from Yemen to the rest of the Internet (or from anywhere on the Internet to get to Yemen), it *HAS* to go through one of those 5 upstream providers. Which could be viewed like this…

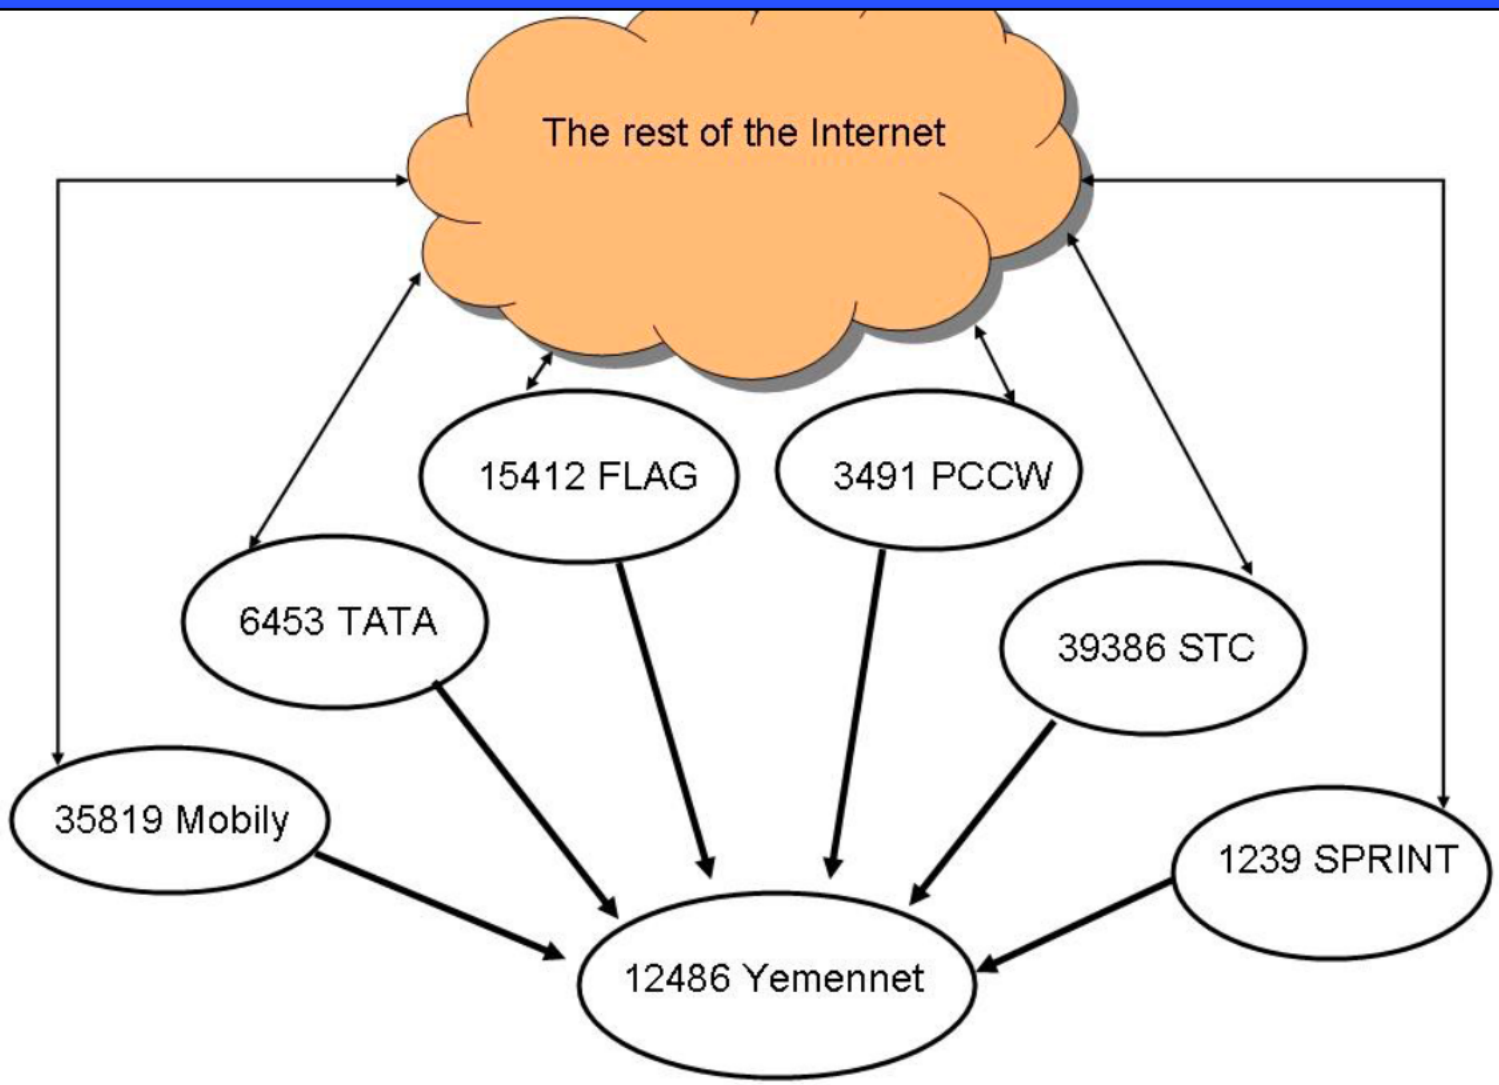

# Okay, so traffic for Yemen has to go through 1 of 6 providers, so?

Armed with this high-level knowledge of Yemen's connectivity, think about what that means:

- Yemennet has to have a router that connects it's own network with it's upstream providers. That router is going to have a unique interface and IP address for each connection.
- That router has to use physical cables to connect between Yemennet and each upstream provider (think big Transnational undersea fiber cables).
- Yemennet *CAN* control which upstream provider it sends data *OUT* of the country through (because it controls the router that's sending the data out).
- Yemennet *CAN NOT* control which provider the data comes back *IN* to the country through, because that is left to BGP routing tables out on the Internet.

# So, to visualize the last 2 points…

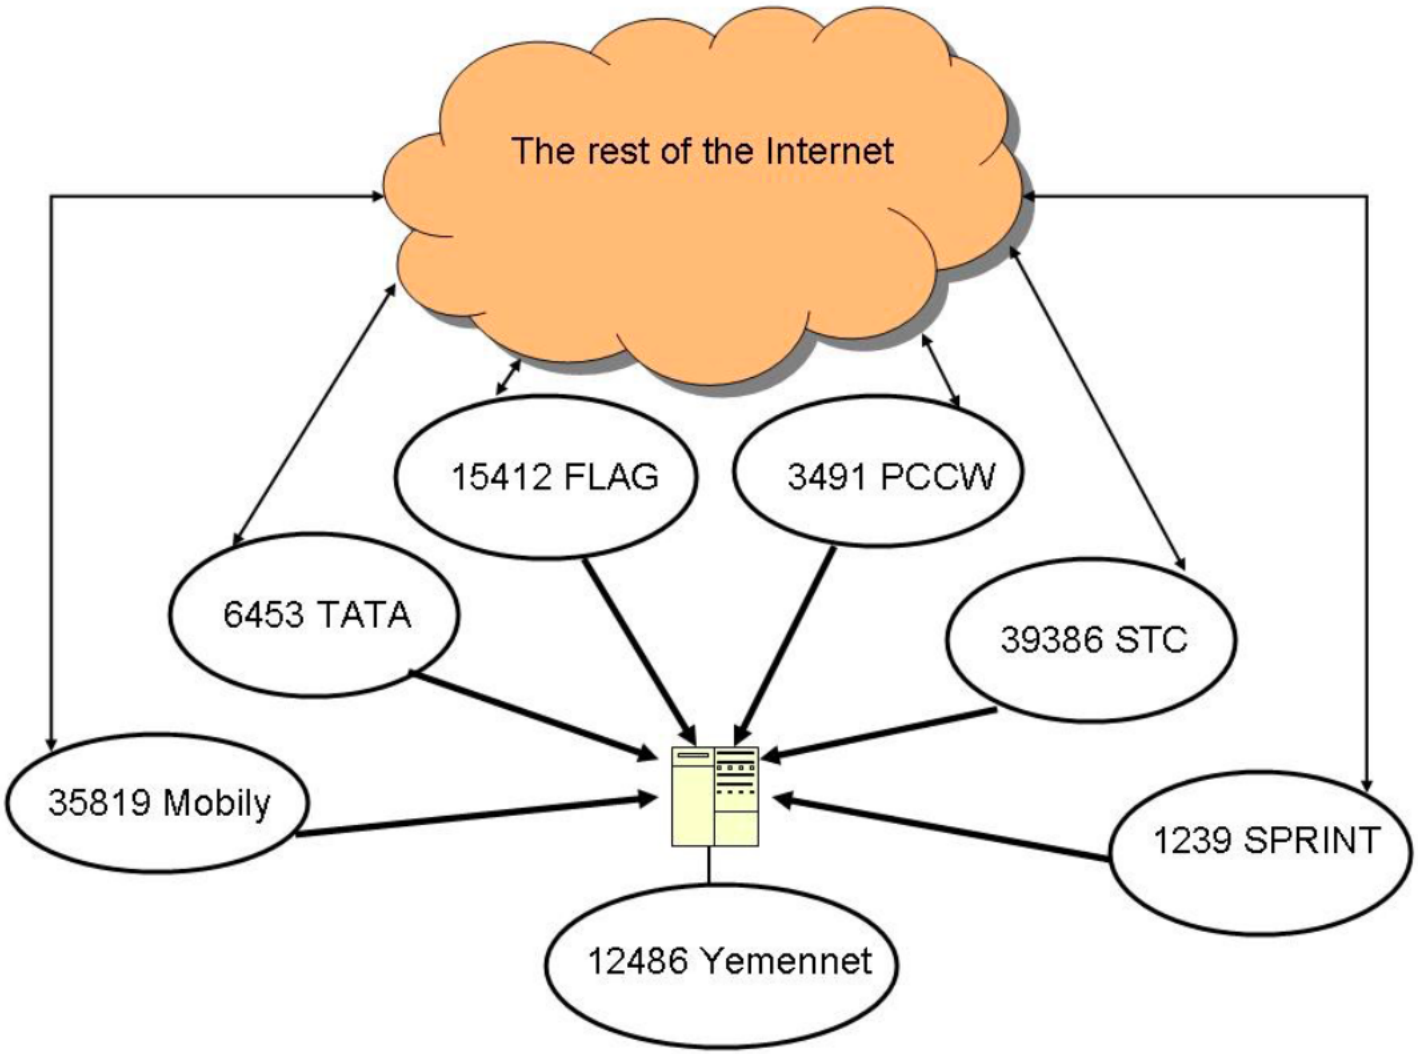

# So, to visualize the last 2 points…

Yemennet can control how data gets out…

Yemennet does NOT control how data gets back in…

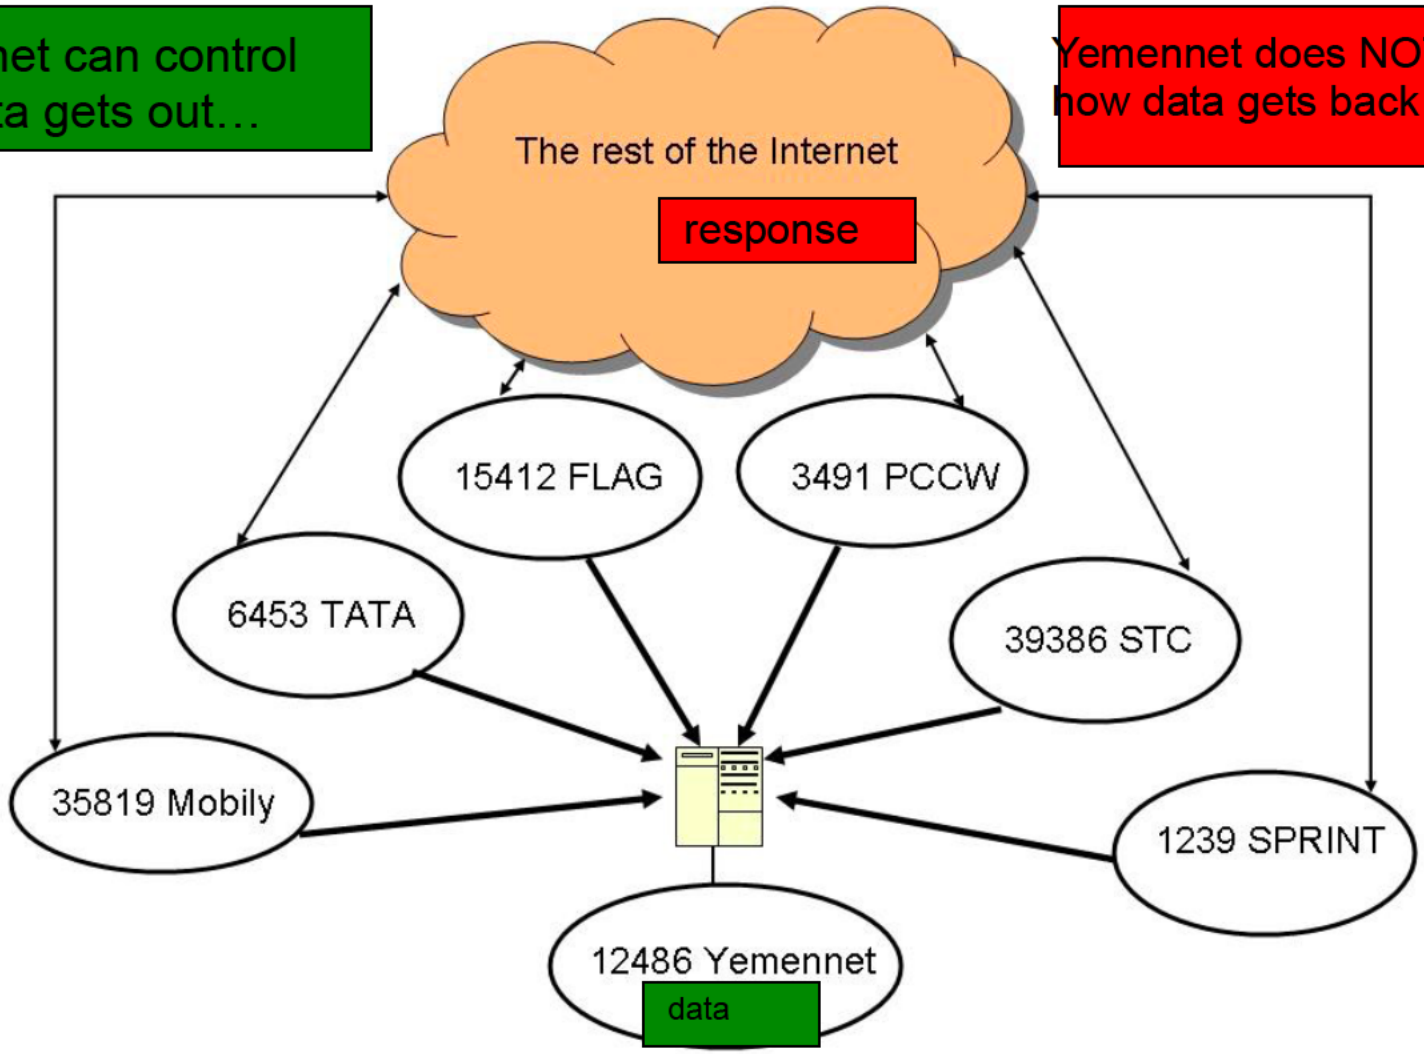

Next, let's visualize the physical connections between Yemennet and its upstream providers. You can see here which cables are used.

TOP SECRET//COMINT//REL TO USA, FVEYS

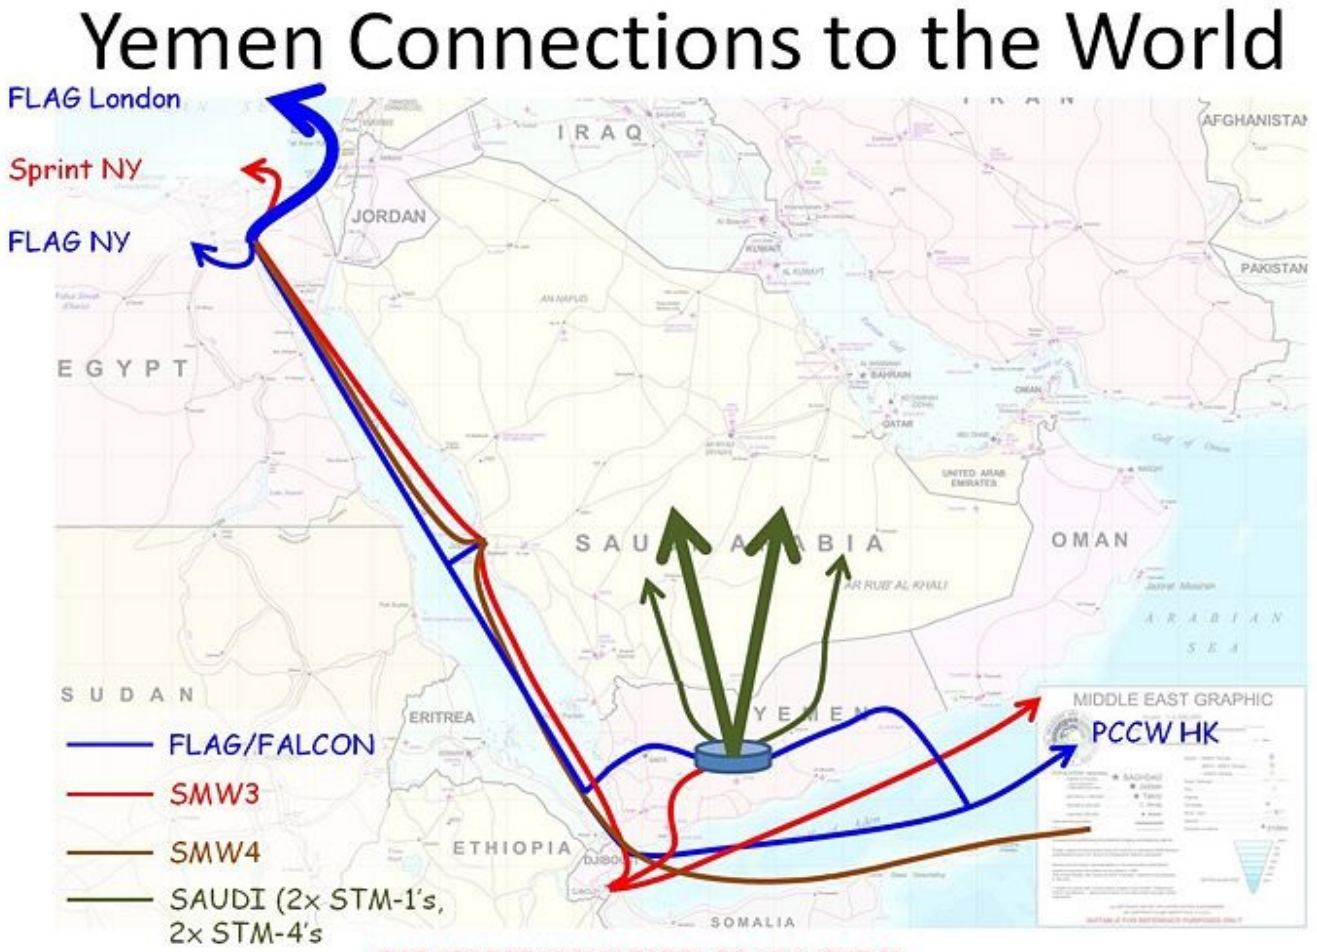


TOP SECRET//COMINT//REL TO USA, FVEYS

# So to recap:

- You understand the logical connectivity of Yemennet (who it has to go through to get to the Internet)
- You grasp the physical connectivity of Yemennet (you know which fiber cables physically connect it to the rest of the world)
- You know that Yemennet can choose which provider it sends data *OUT* through
- Big Internet BGP routing tables can dynamically choose which link data comes back *IN* to Yemennet through

- There are a couple more things to know before we talk about shaping…

# About that router that connects Yemennet to it's peers…

- Remember how I mentioned that router has a different interface and unique IP address for each upstream provider? That router will have at least 7 interfaces (one for each upstream, and one connected to the rest of it's network).
- The connection between the router and an upstream provider has to use IPs that are in the same subnet (normally it's a /30 subnet, which consists of 2 usable IPs).
- This means that one of the two networks will have to sacrifice an IP address to put on the other end of the connection (most of the time it's the bigger network that gives up an IP address to assign to the customer side's router).
- So, if we were to use the connection with SPRINT for example, here's what it might look like…

Note that the IP address on Yemen's side of it's connection with Sprint, is an IP address that belongs to SPRINT. This is because one of the two networks has to use one of their IP's on the other end of the connection.

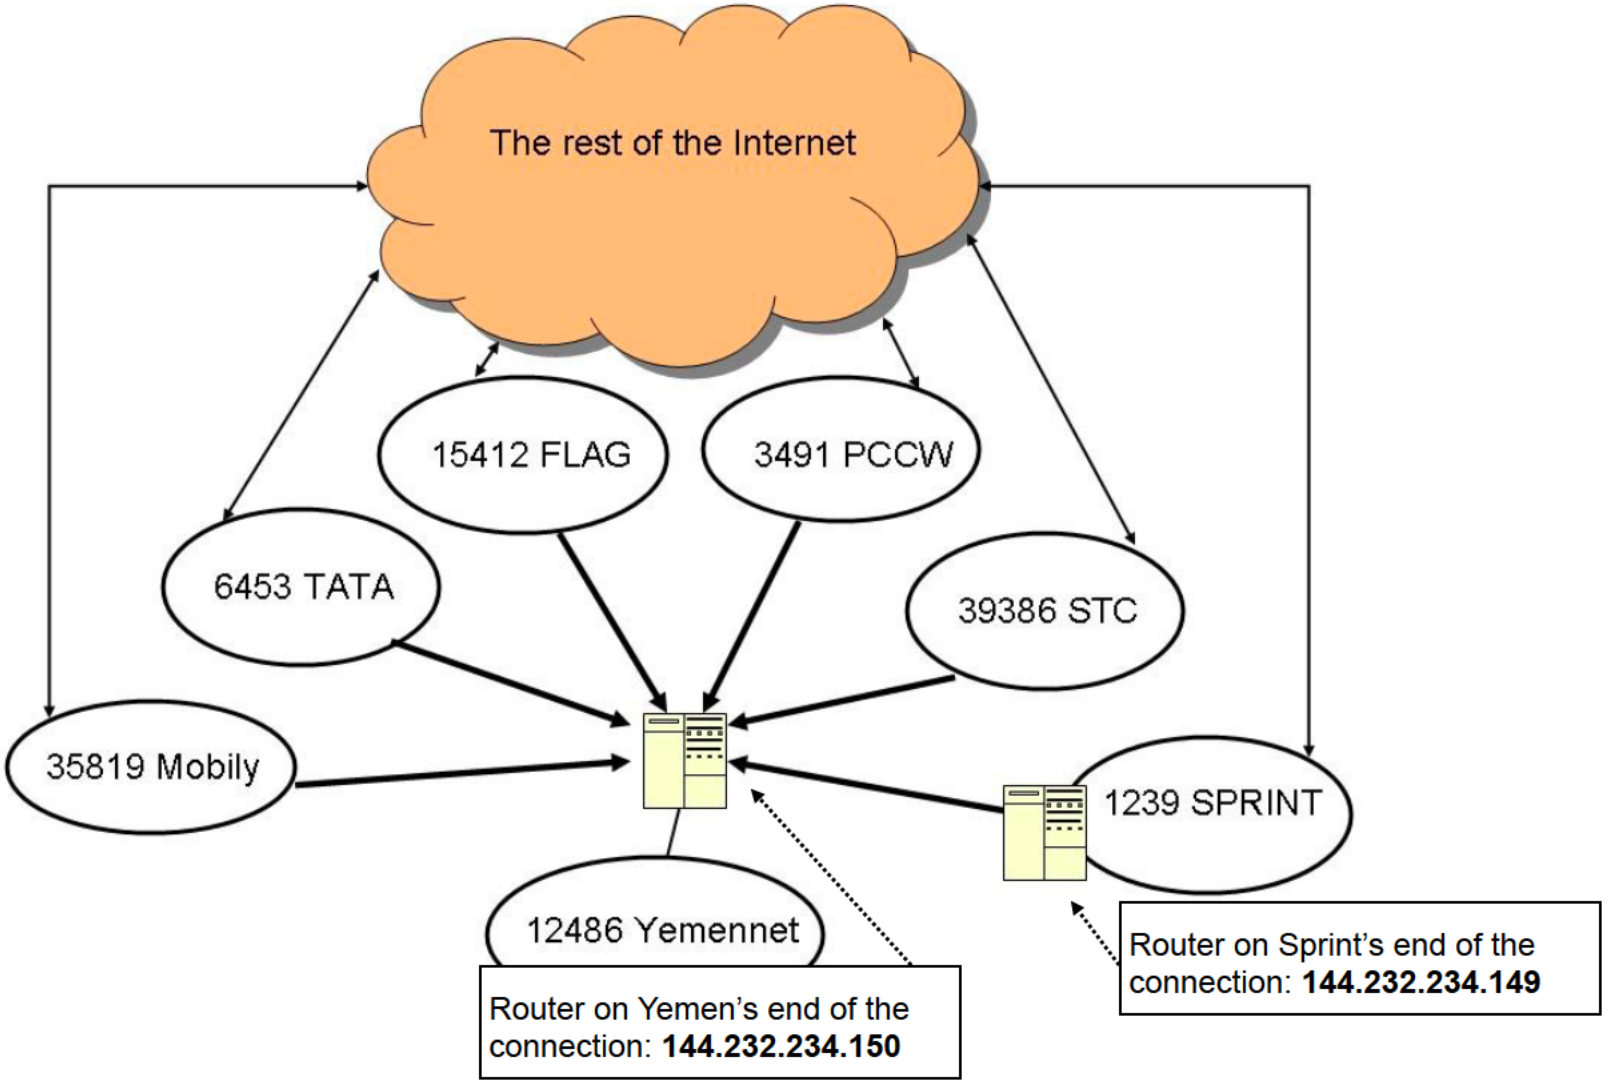

PRO TIP: If you do a traceroute from somewhere in the Internet to anywhere in Yemen, if one of the hops before it gets into the country is **144.232.234.150**, then you can assume your trace went through SPRINT's network (and over the SMW-3 cable if you remember the map 3 slides ago) to get into the country.

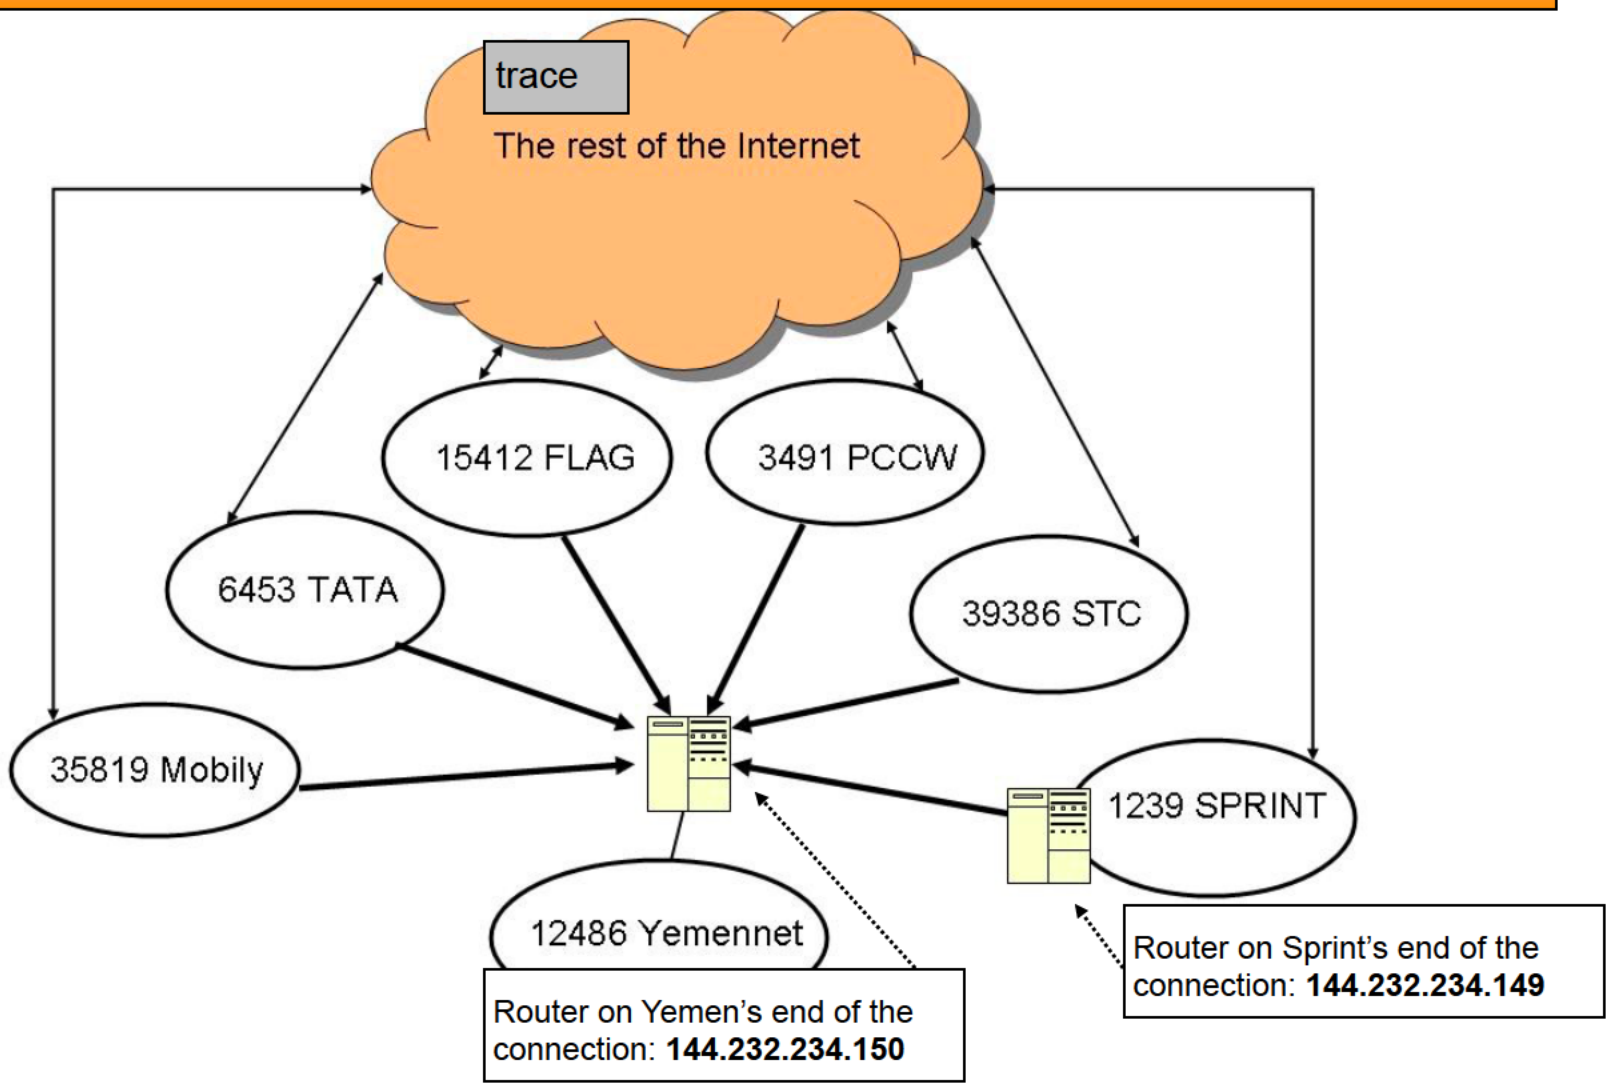

# For the sake of completeness, adding in the rest of the routers…

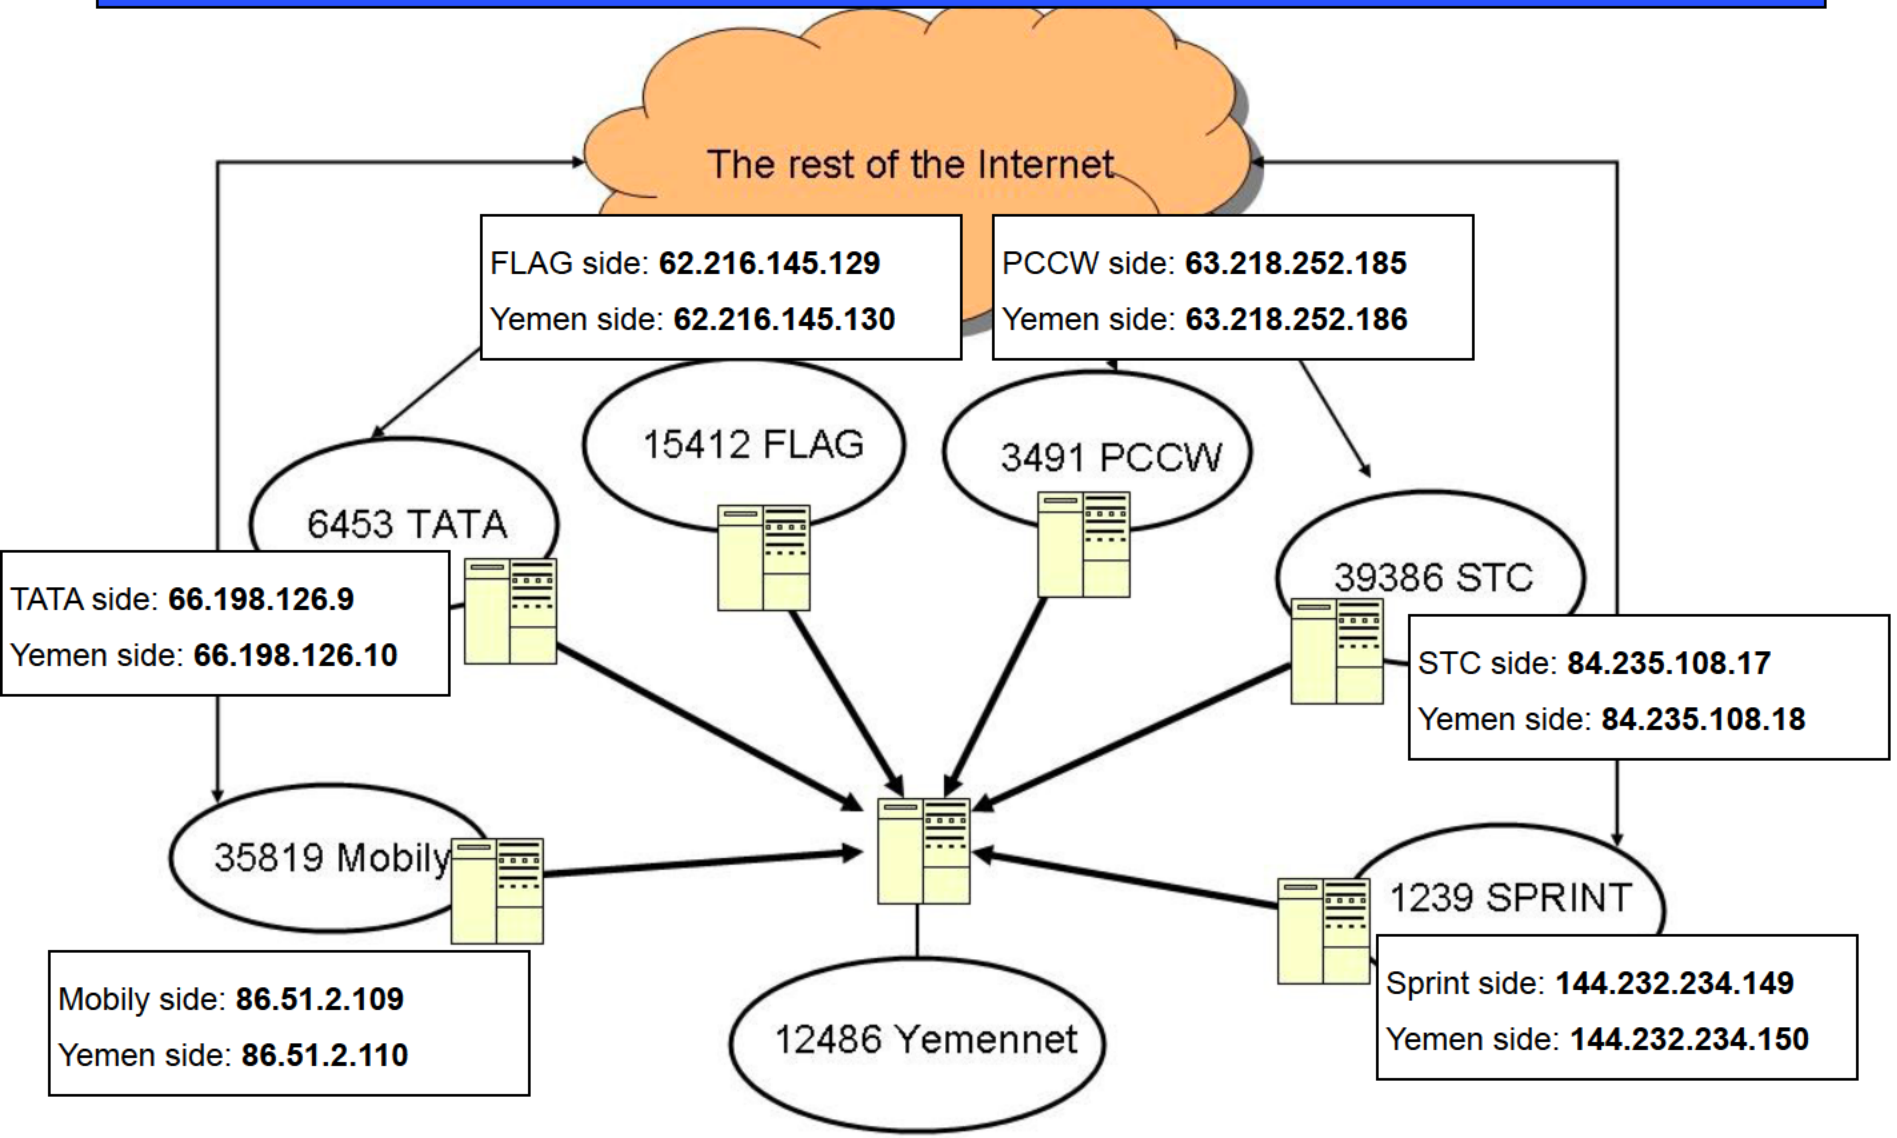

# WARNING! WARNING!

- In the following slides when I talk about SSO collection capabilities, I am completely MAKING UP:
  - SIGADs
  - Case notations
  - Which cables are collected
  - Where SSO's collection capabilities are
- I am MAKING UP this info for the sake of this lesson.
- For info on what SSO's capabilities are for your own target, you will have to go talk to them yourself.

# WARNING 2!  WARNING 2!

- For the sake of this example, I am assuming that all of Yemennet's International links are equal.  By that I am making the assumptions that:
  - An equal amount of traffic is going in/out each link
  - Yemennet is not doing anything to manipulate traffic going over specific links
  - All links are actually active, and are not just backups or down due to maintenance or cable breaks…
- With that out of the way…

# So now you have a good idea about Yemen's connectivity…

- Now time to overlay it with SIGINT collect…
- Without going into how to do this yourself, work with SSO to determine which of those links we can passively collect.
- Let's pretend that they have capabilities to collect the Yemen-Sprint link and the Yemen-FLAG link, but have no capabilities on the rest.
- Once again, this is only PRETEND for the sake of this lesson.

# Here is what our passive collection capabilities would look like for Yemen…

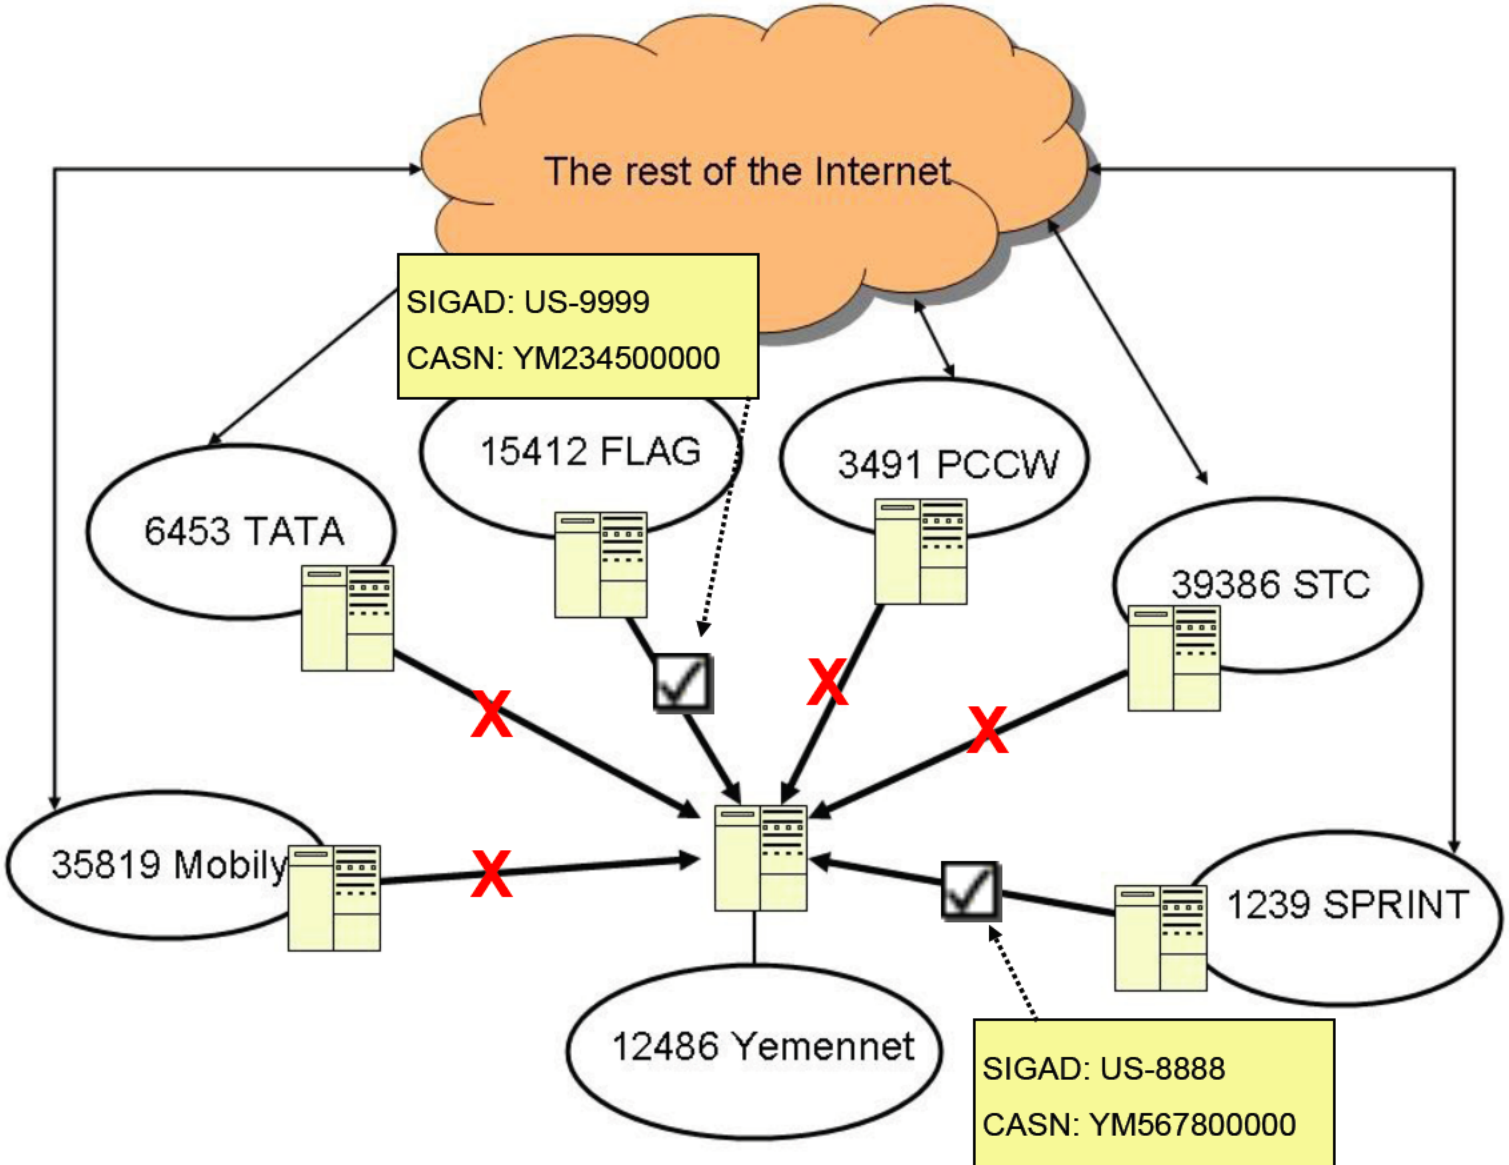

# So, remembering what we talked about in the beginning, here's what we can tell…

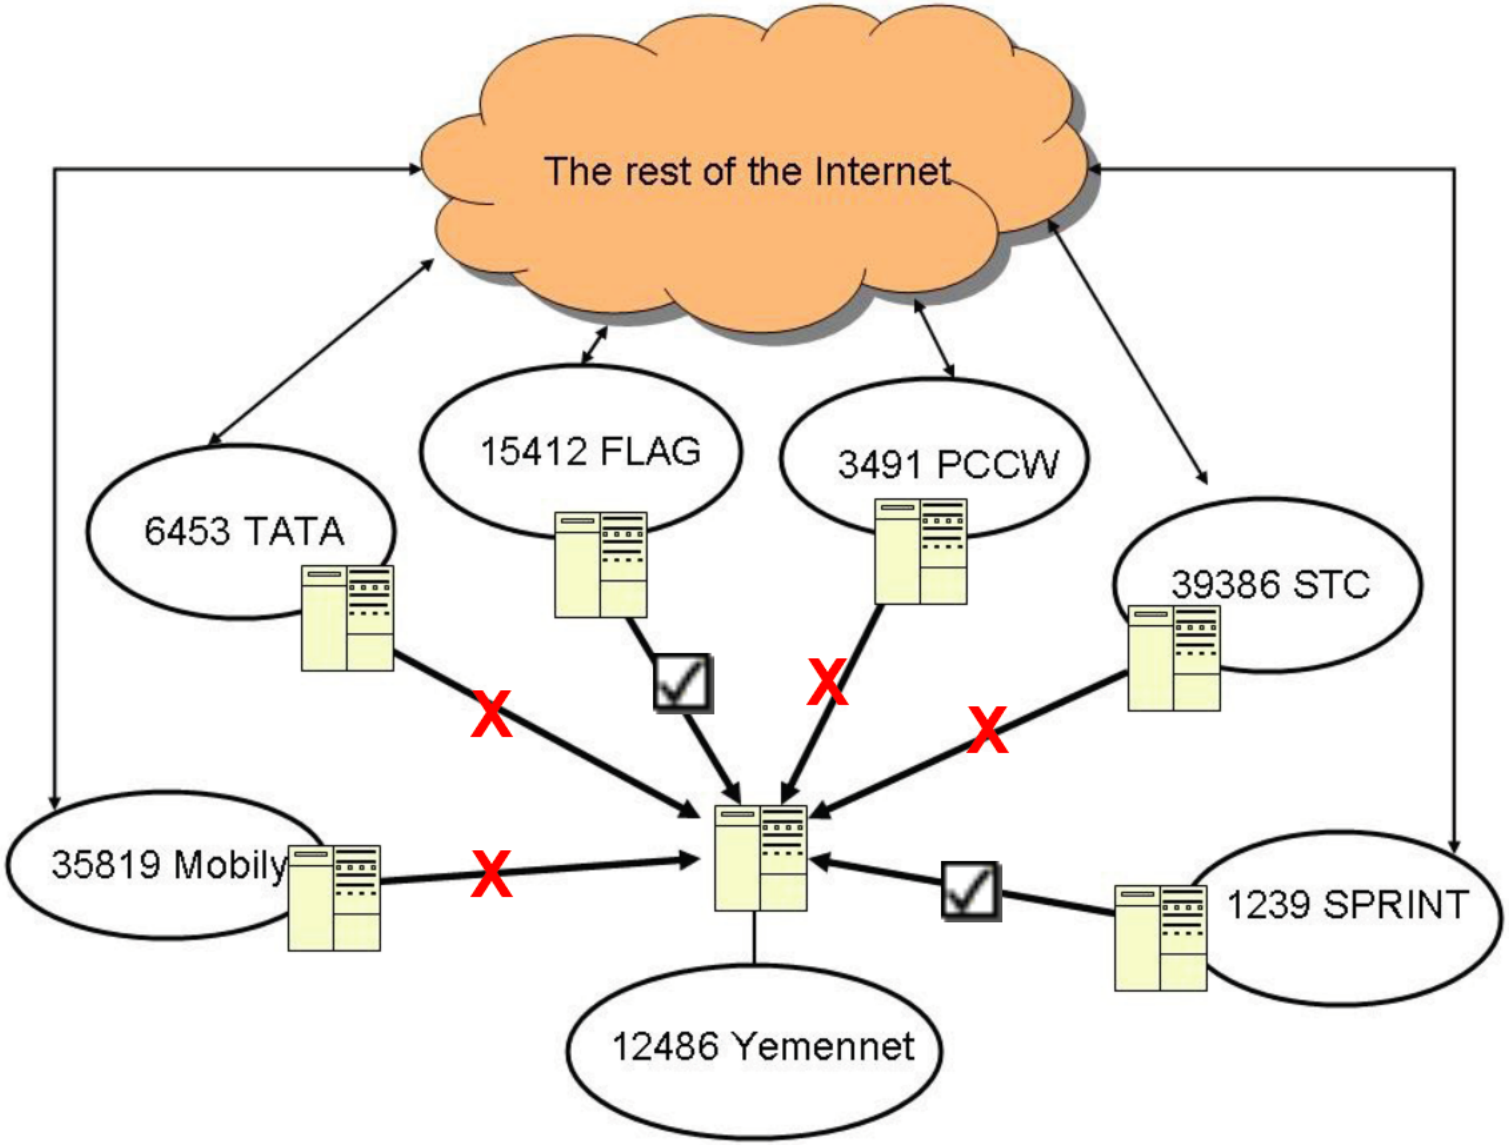

If data *happens* to go in or out the FLAG or Sprint links for Yemen, we will collect those comms.

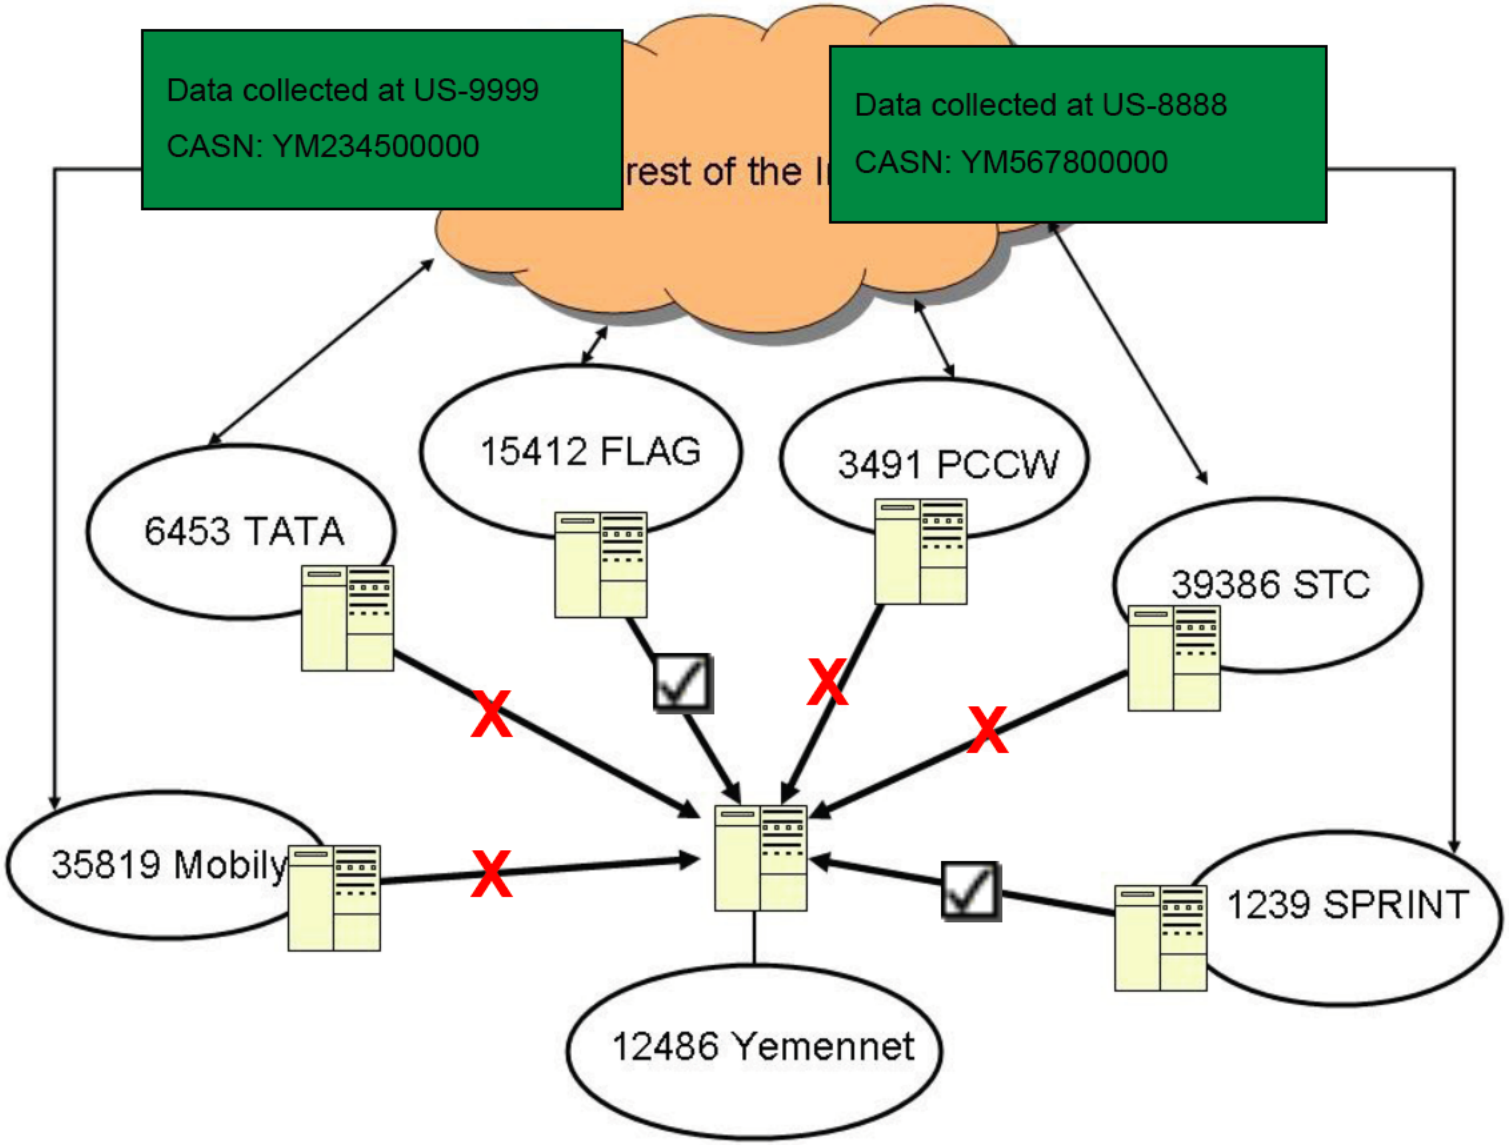

If data does *NOT* go in or out the FLAG or Sprint links for Yemen, we will *NOT* collect those comms.

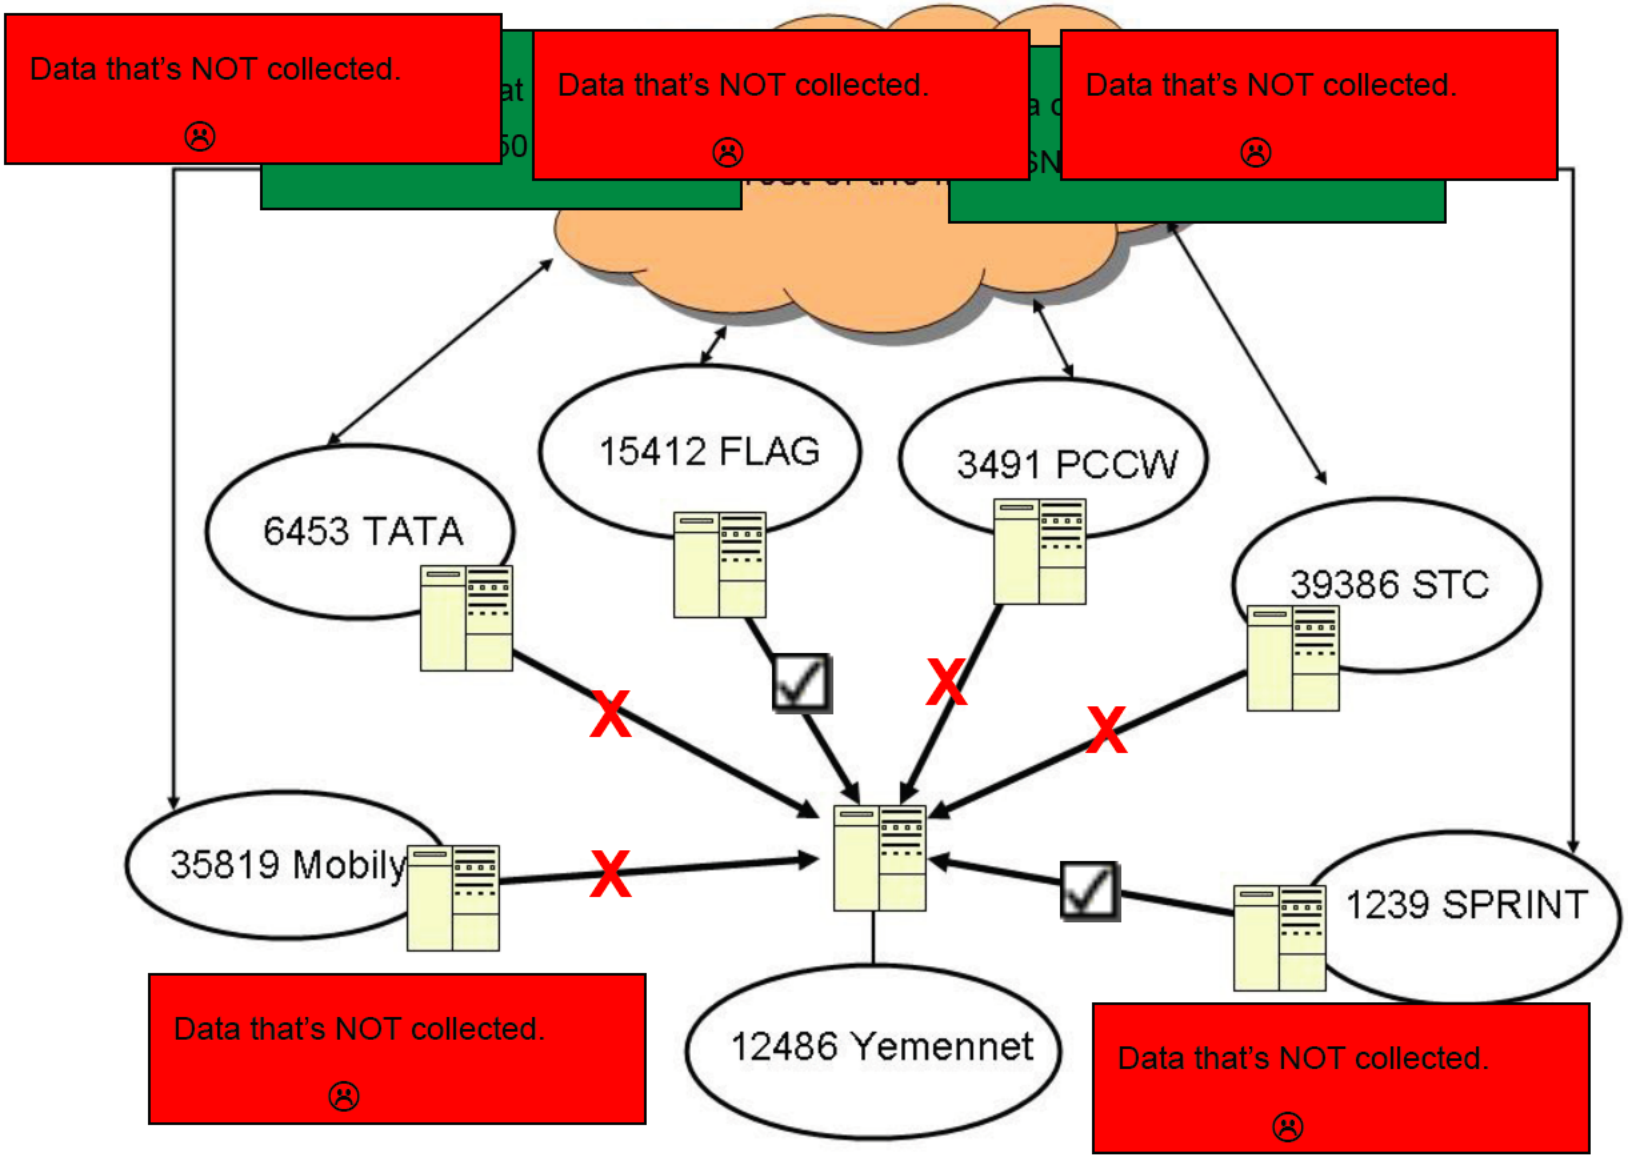

So, now you know what our passive collection posture against Yemennet is like… we may only collect about 33% (2 out of 6 links) for all of the country's traffic (this is assuming they send and receive equal amounts of traffic over each of the links).

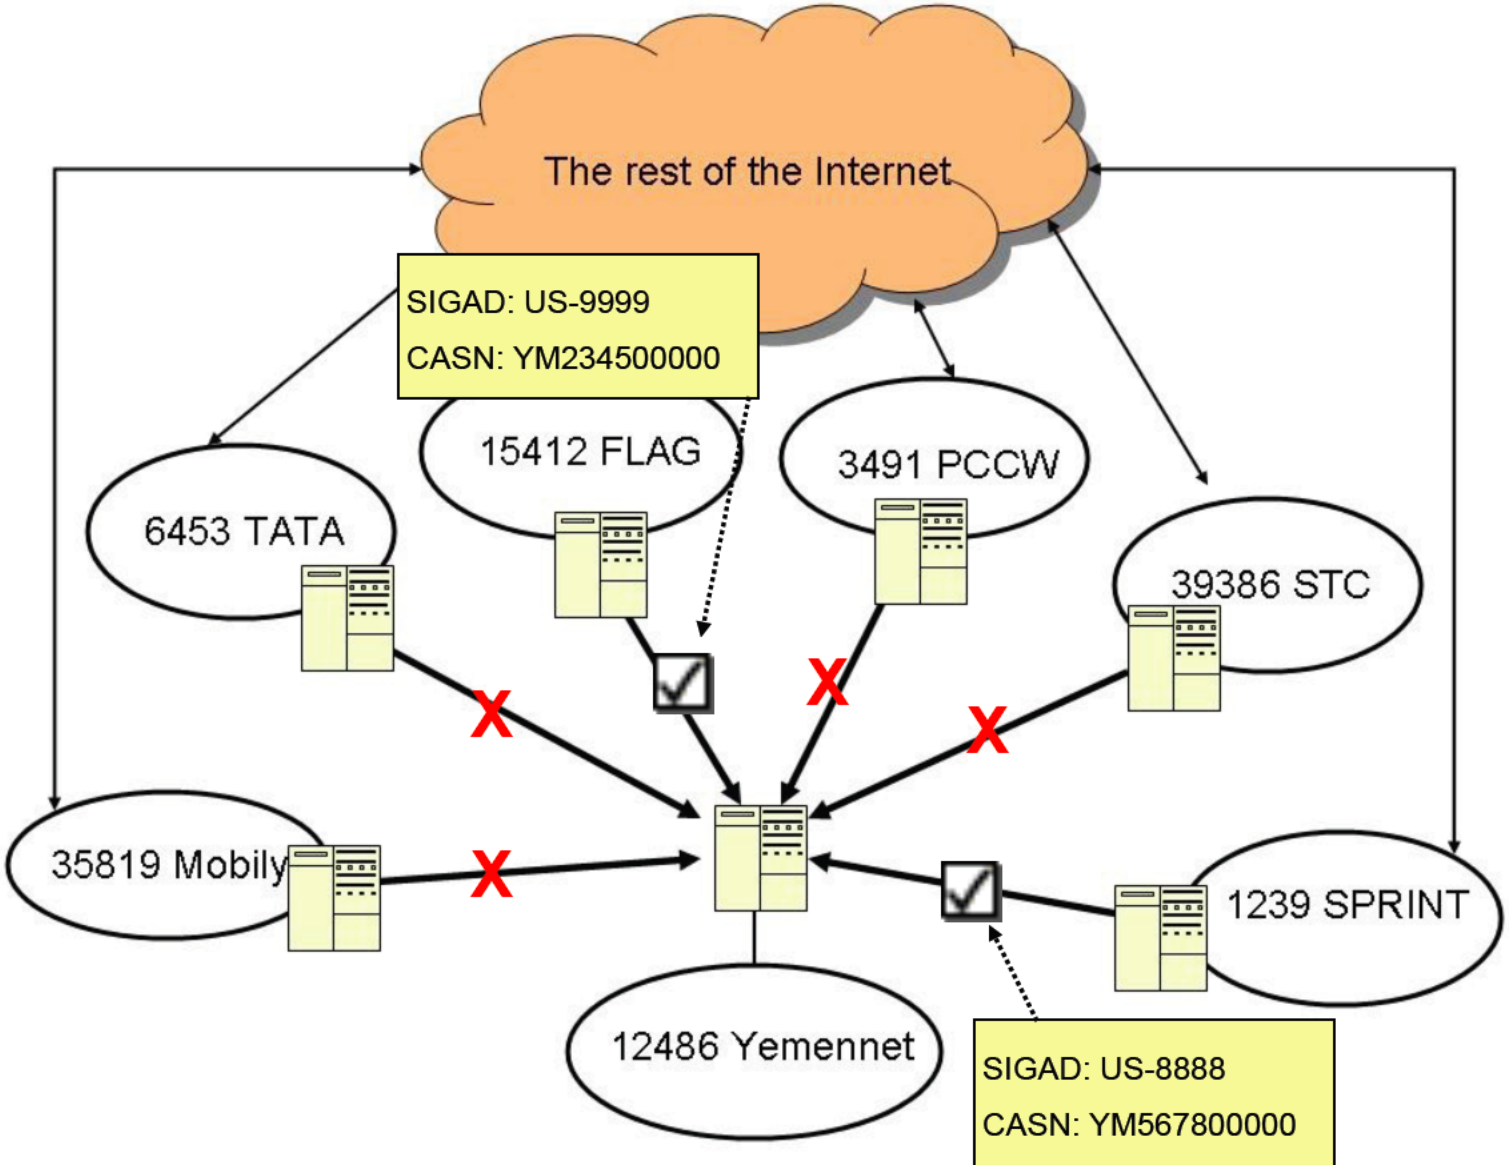

# Now that you're an expert on Yemennet, let's talk about shaping

- The purpose of "shaping" is taking traffic that wouldn't normally go through one of our passive links, and *making* it go through one of our passive links, so we can collect it and get it into the SIGINT system.
- Before we talk about how to shape traffic on Yemennet, let's explore a couple different scenarios in which we would consider shaping as a solution…they will be shaping traffic OUT of Yemennet, and shaping traffic INTO Yemennet.

# Shaping traffic OUT of Yemennet

- For this scenario, you have an access (probably CNE) inside of Yemennet, and you want to make that access send traffic, but make sure it goes out over a link that is passively collected by SSO.
- You need a DESTINATION on the Internet where you can send data to, where you know it will go over 1 of the 2 links we can collect.
- Earlier I mentioned that Yemennet can control which links they send data OUT.  This is true, Yemennet has that control, however, you, as an end-user on their network do *NOT* have that control.
- So, how can you control which link your traffic will go out through?

If you send traffic from somewhere inside Yemennet to some random place on the Internet, you are at Yemennet's mercy as to which link it will send the data out through.

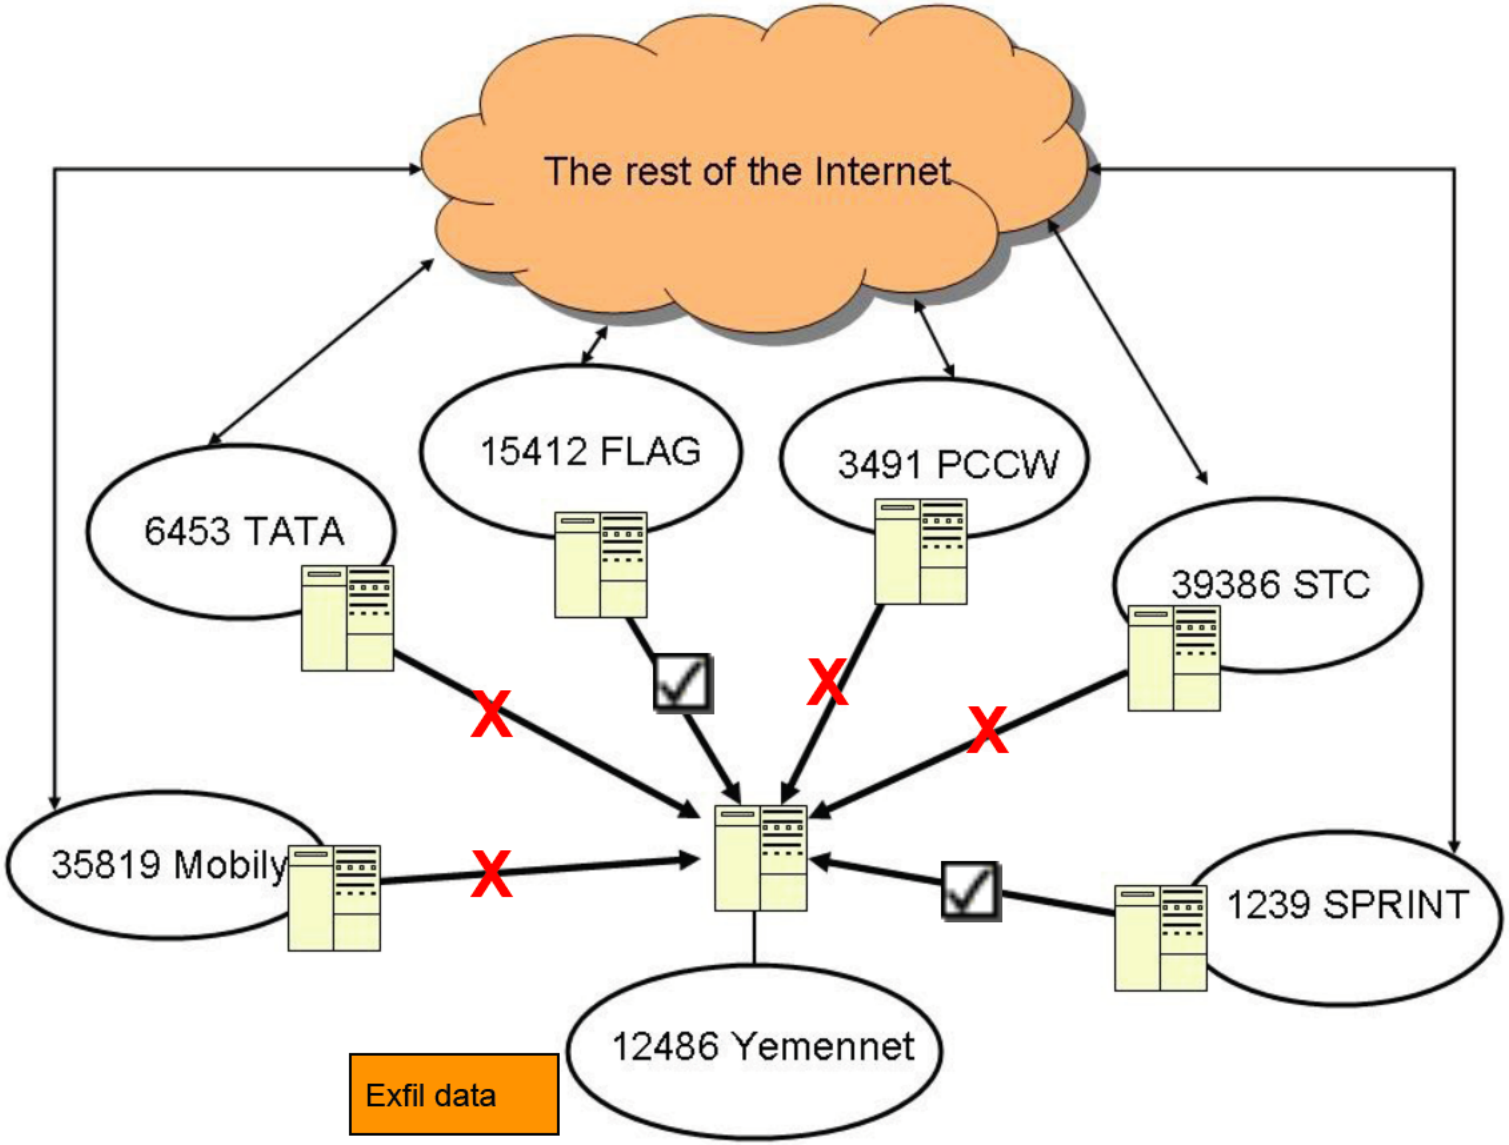

Uh oh, you tried to exfil traffic, but Lady Luck did not shine her favor upon you, and the traffic went out a link we could not passively collect…

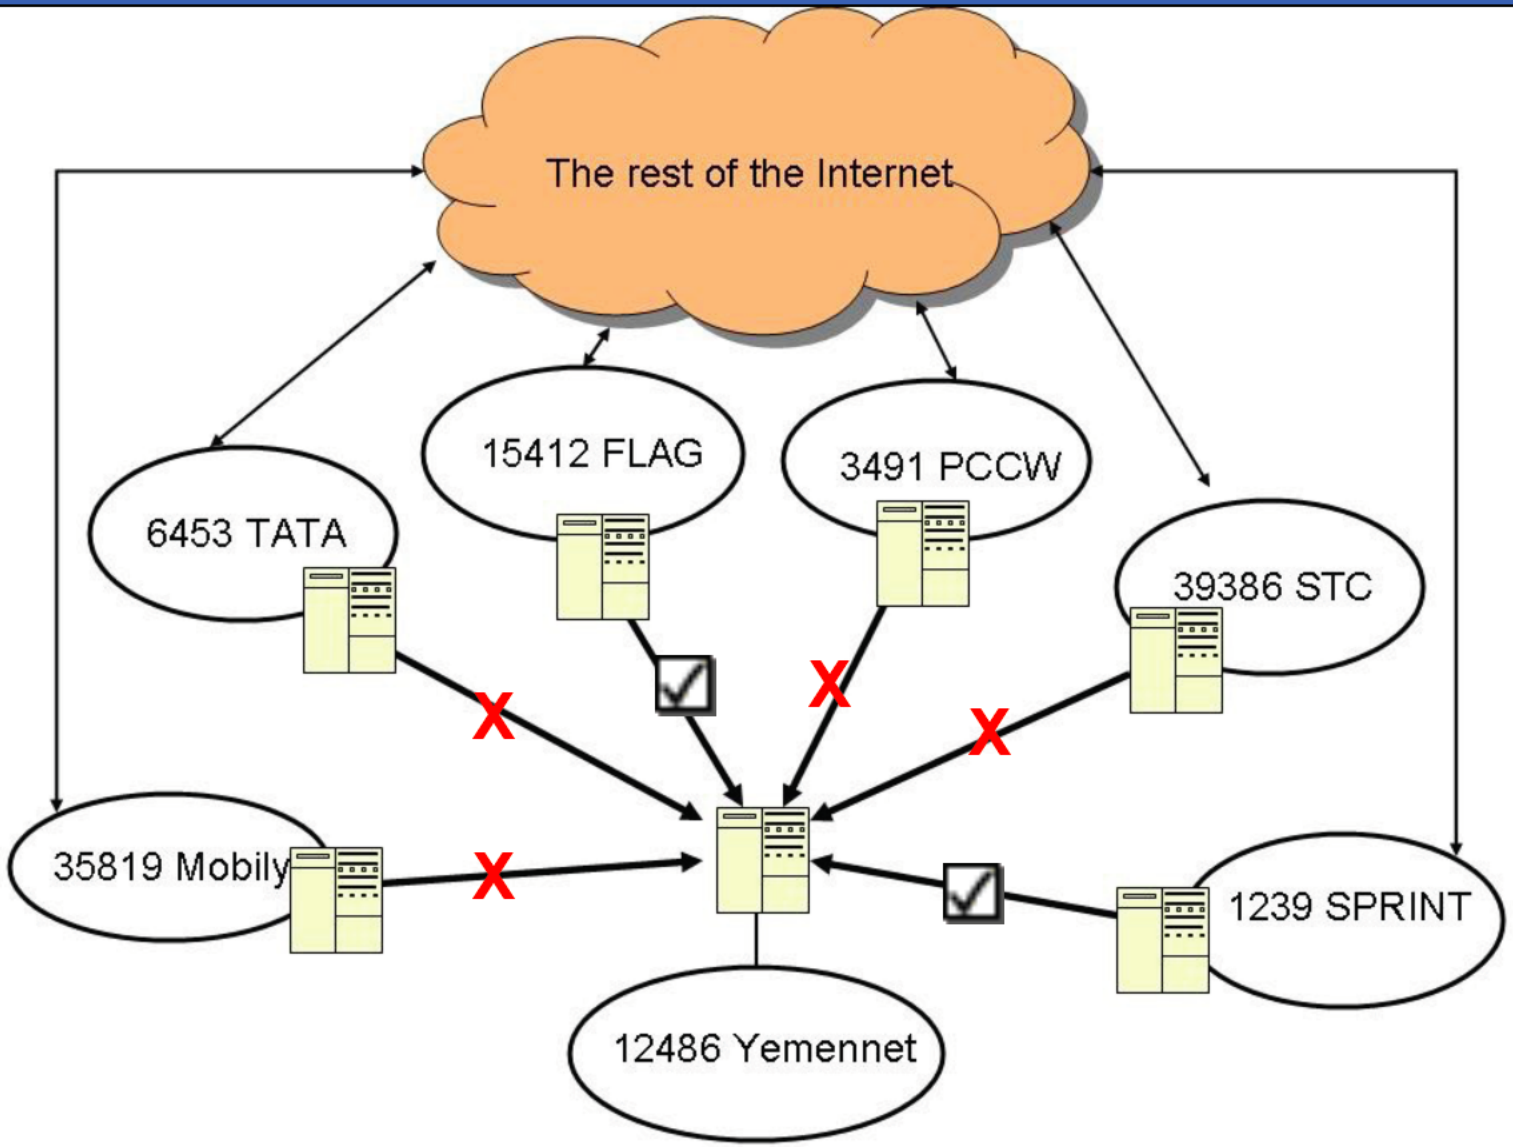

SUCCESS! The exfil will probably die somewhere in Sprint's network, but we don't care, because we accomplished our goal of collecting it along the way. Now we can go look at collect from US-8888 CASN: YM567800000 for our exfil!

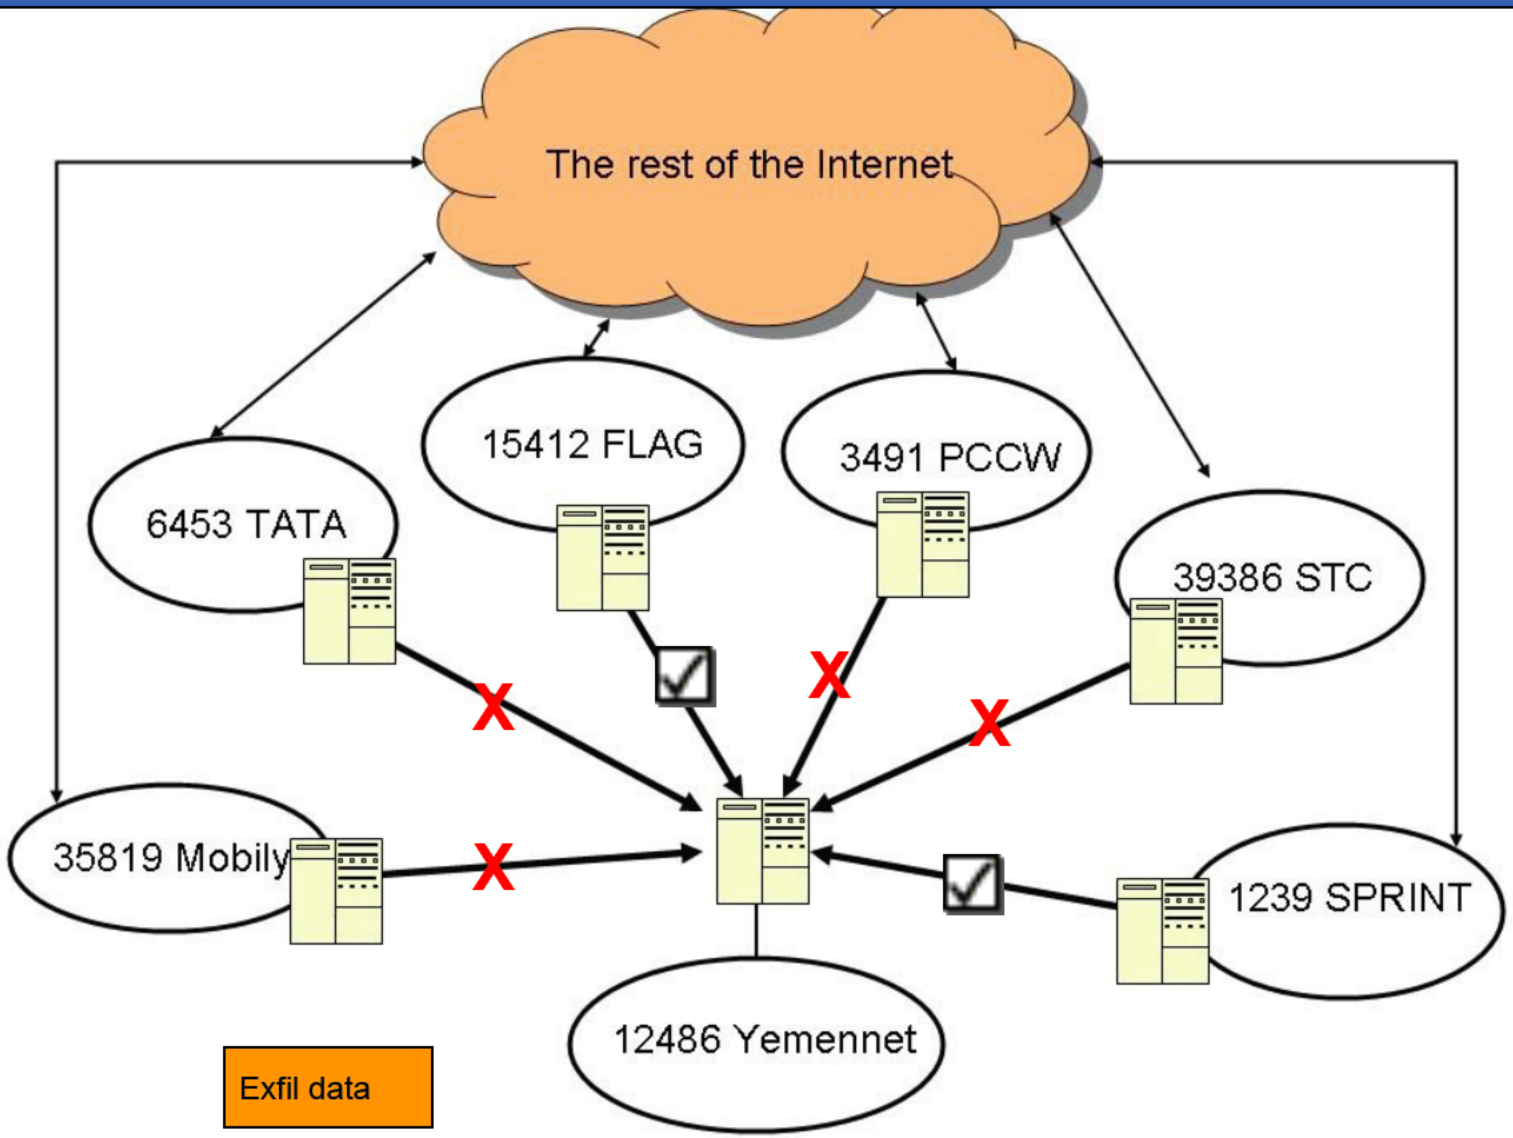

# So in order to shape traffic OUT of a network

- You need to have an understanding of the network that you are starting in, who it's upstream providers are, and what the collection capabilities are against that network.
- Then you can find a destination IP address directly on the other end of that link (by looking at any of the IP ranges in that provider's ASN).
- From there, you have a higher probability that traffic will traverse a link you can passively collect.

# Shaping traffic INTO a network

- This is a whole different animal, and probably more relevant to what people traditionally think of as “shaping” in the SIGINT sense.
- There is only 1 feasible way (that I can think of) to make this work reliably.
- But first, let’s go back and look at our Internet connectivity…

In order to shape traffic into a network, we'll be starting from some random place on the Internet. Somewhere in here…

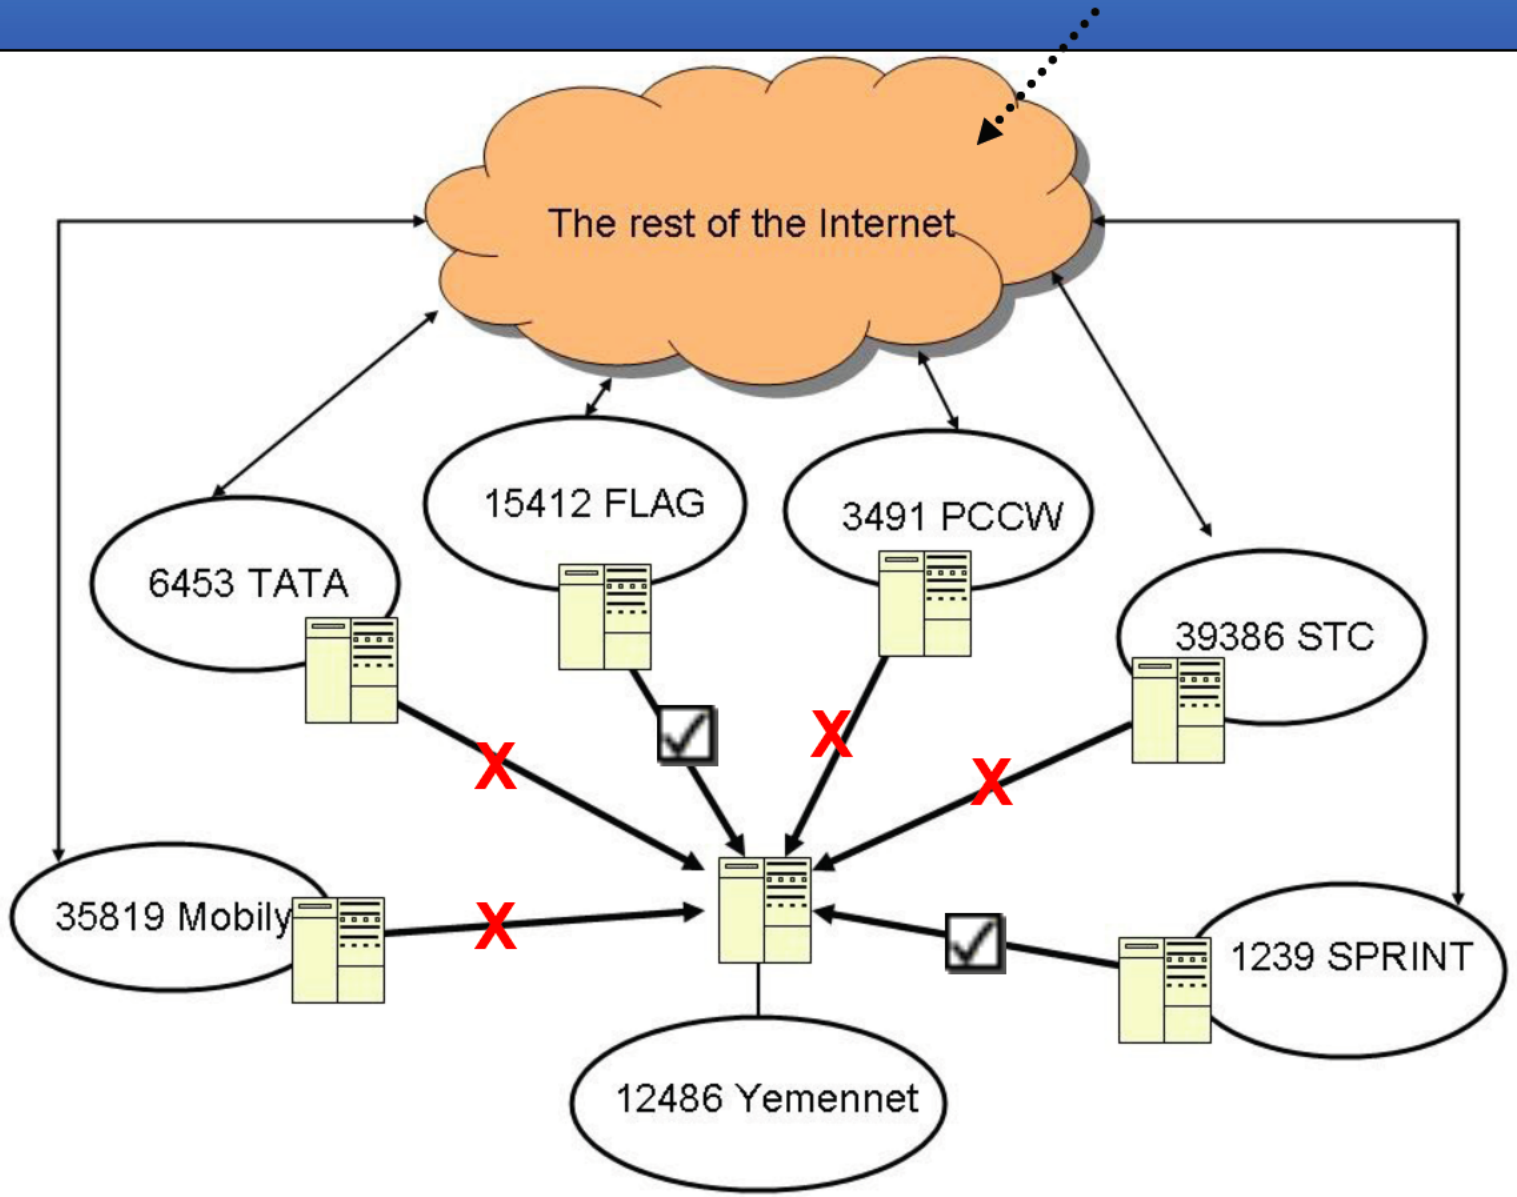

And our goal, will be to send traffic into Yemennet with the express purpose of making it go through 1 of the 2 links we can passively collect…because the entire point is getting our exfil into the SIGINT system through those passive access points.

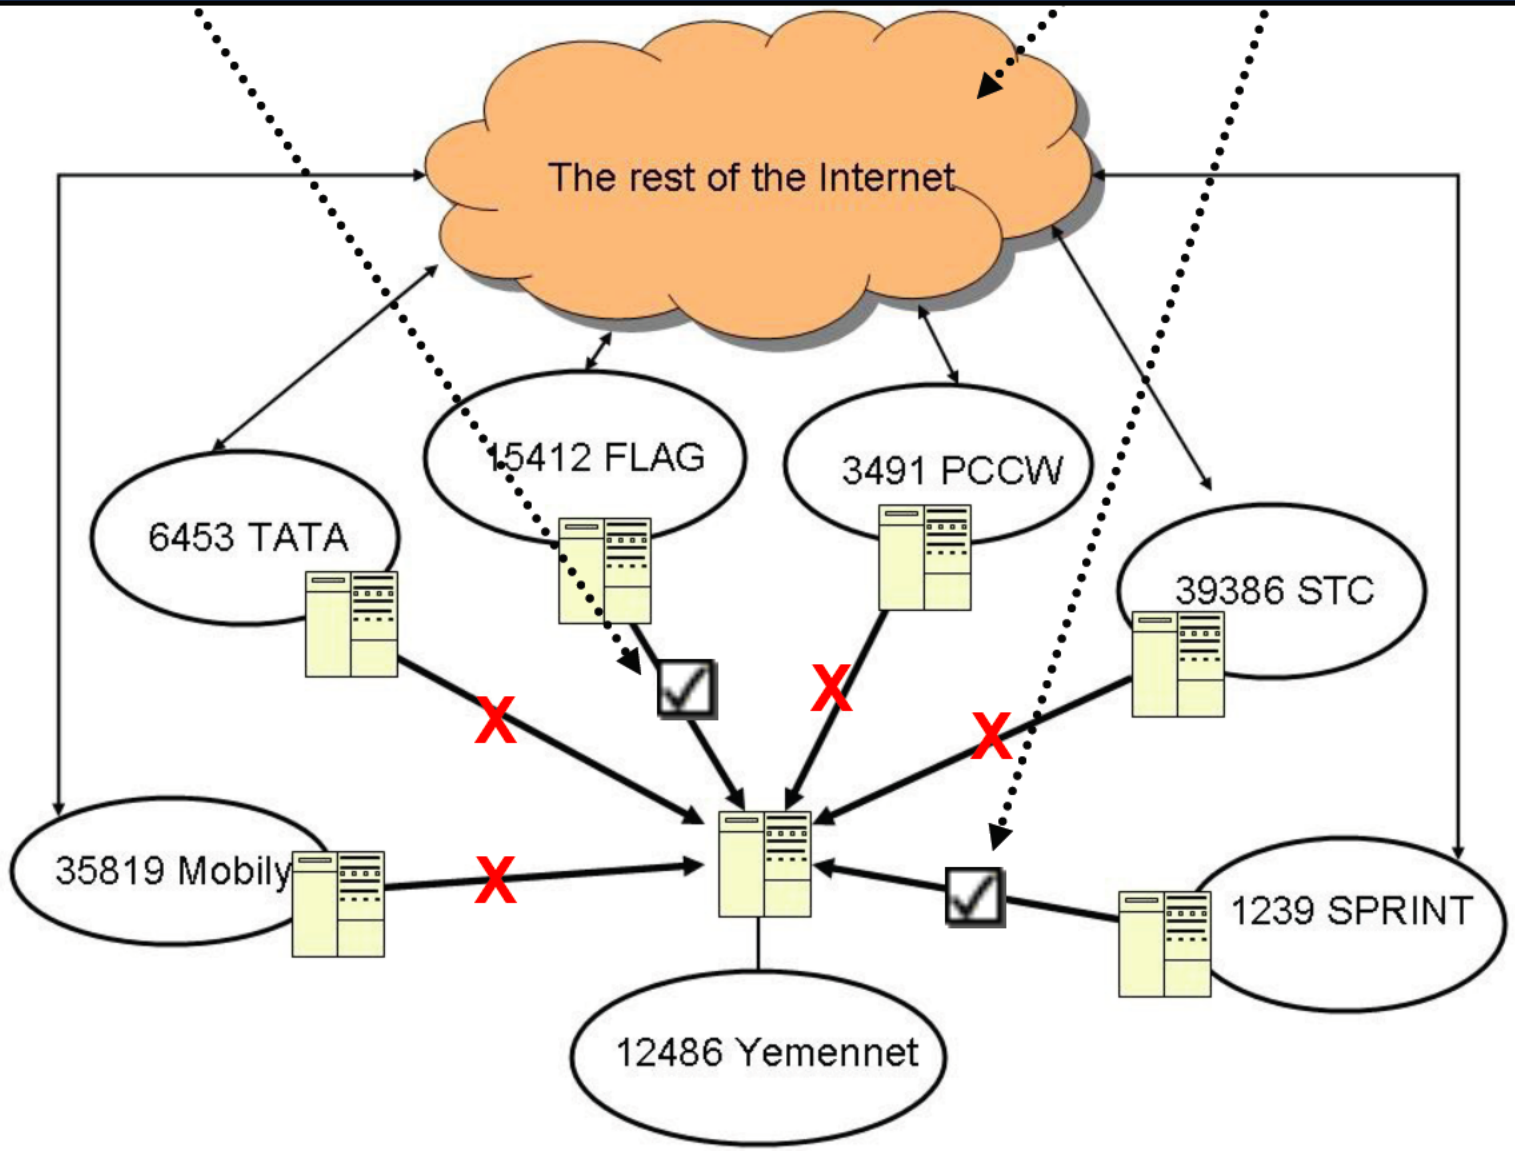

But remember, the route for traffic going into Yemennet is completely at the mercy of BGP routing tables out on the Internet. So there is no guarantee that it will go through one of the links we can collect (once again, only a 33% shot).

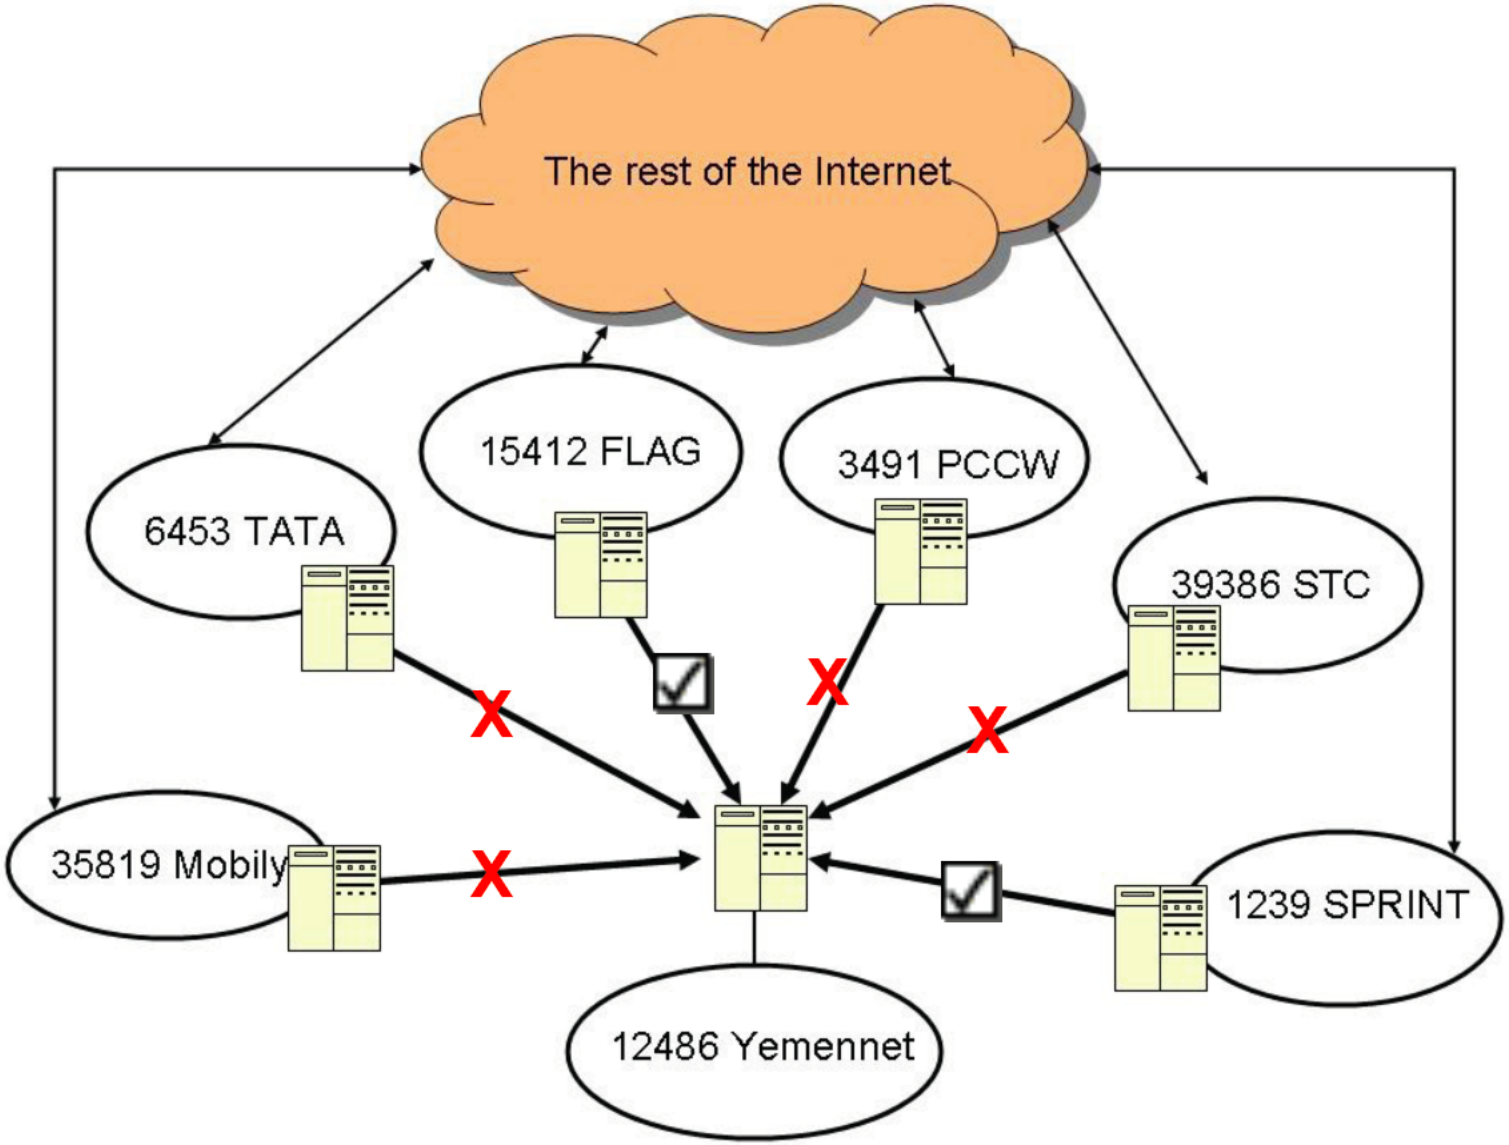

But remember, the route for traffic going into Yemennet is completely at the mercy of BGP routing tables out on the Internet. So there is no guarantee that it will go through one of the links we can collect (once again, only a 33% shot).

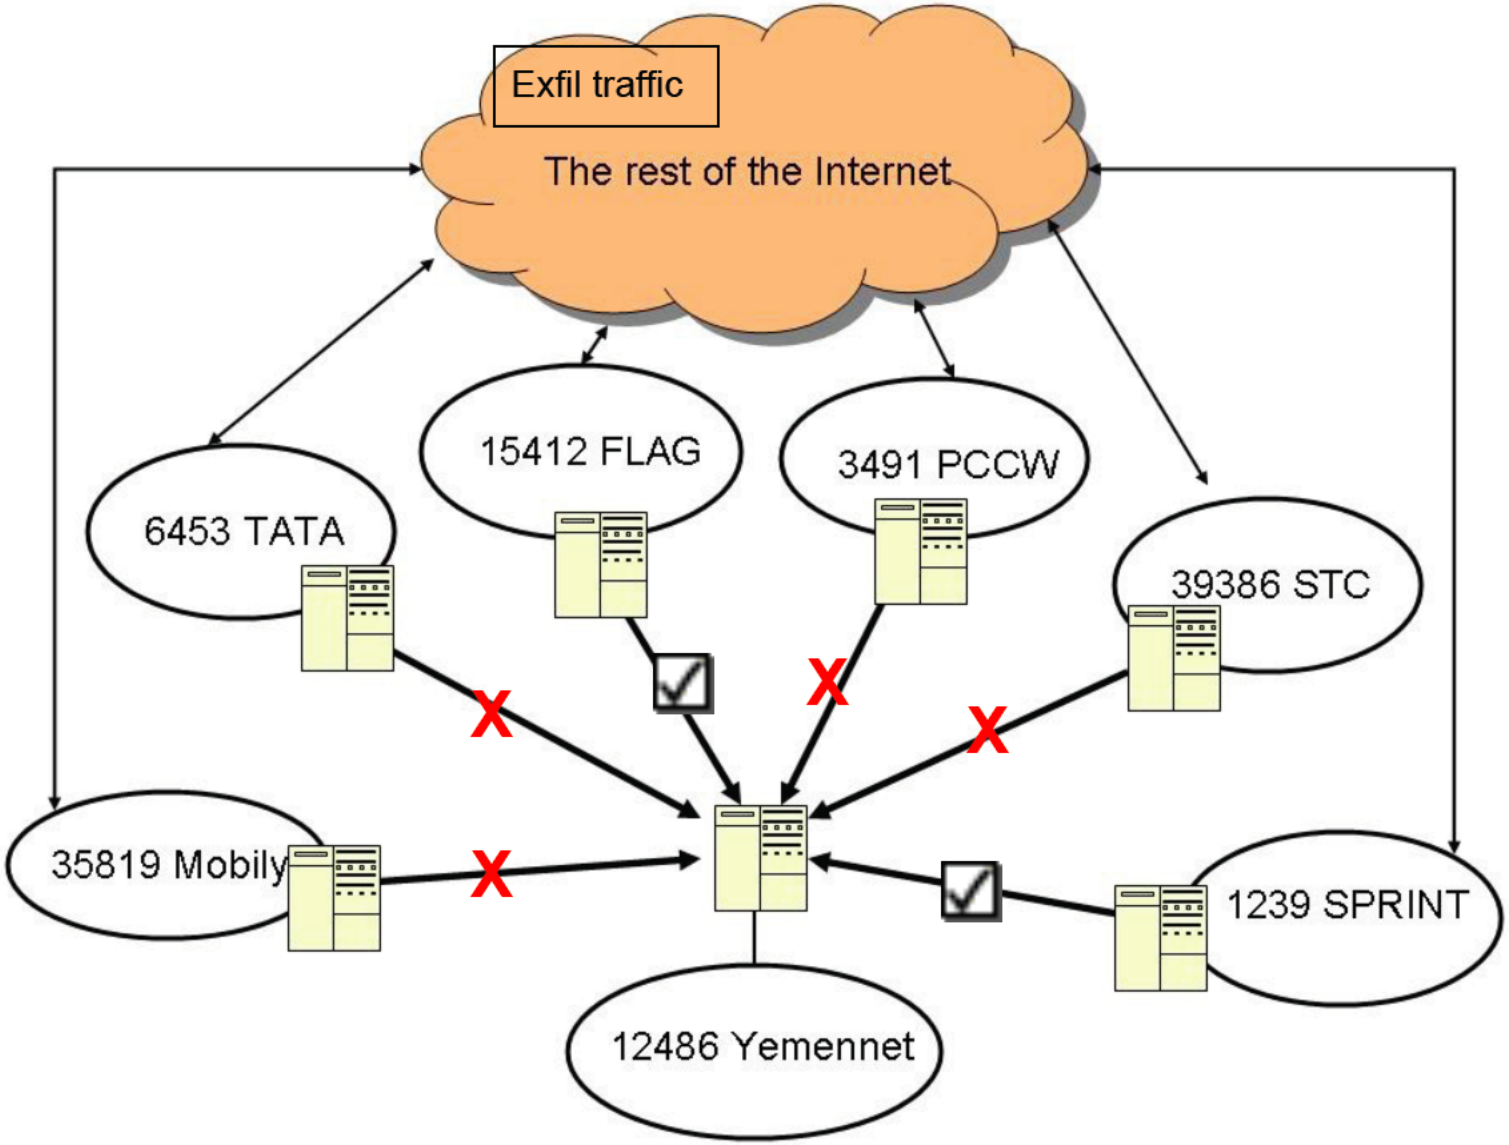

Wow, that didn't work. ☹ So, if we are completely at the mercy of big Internet routing tables, how could we reliably send traffic from ANYWHERE in the Internet into Yemennet while guaranteeing it goes through passive collect?

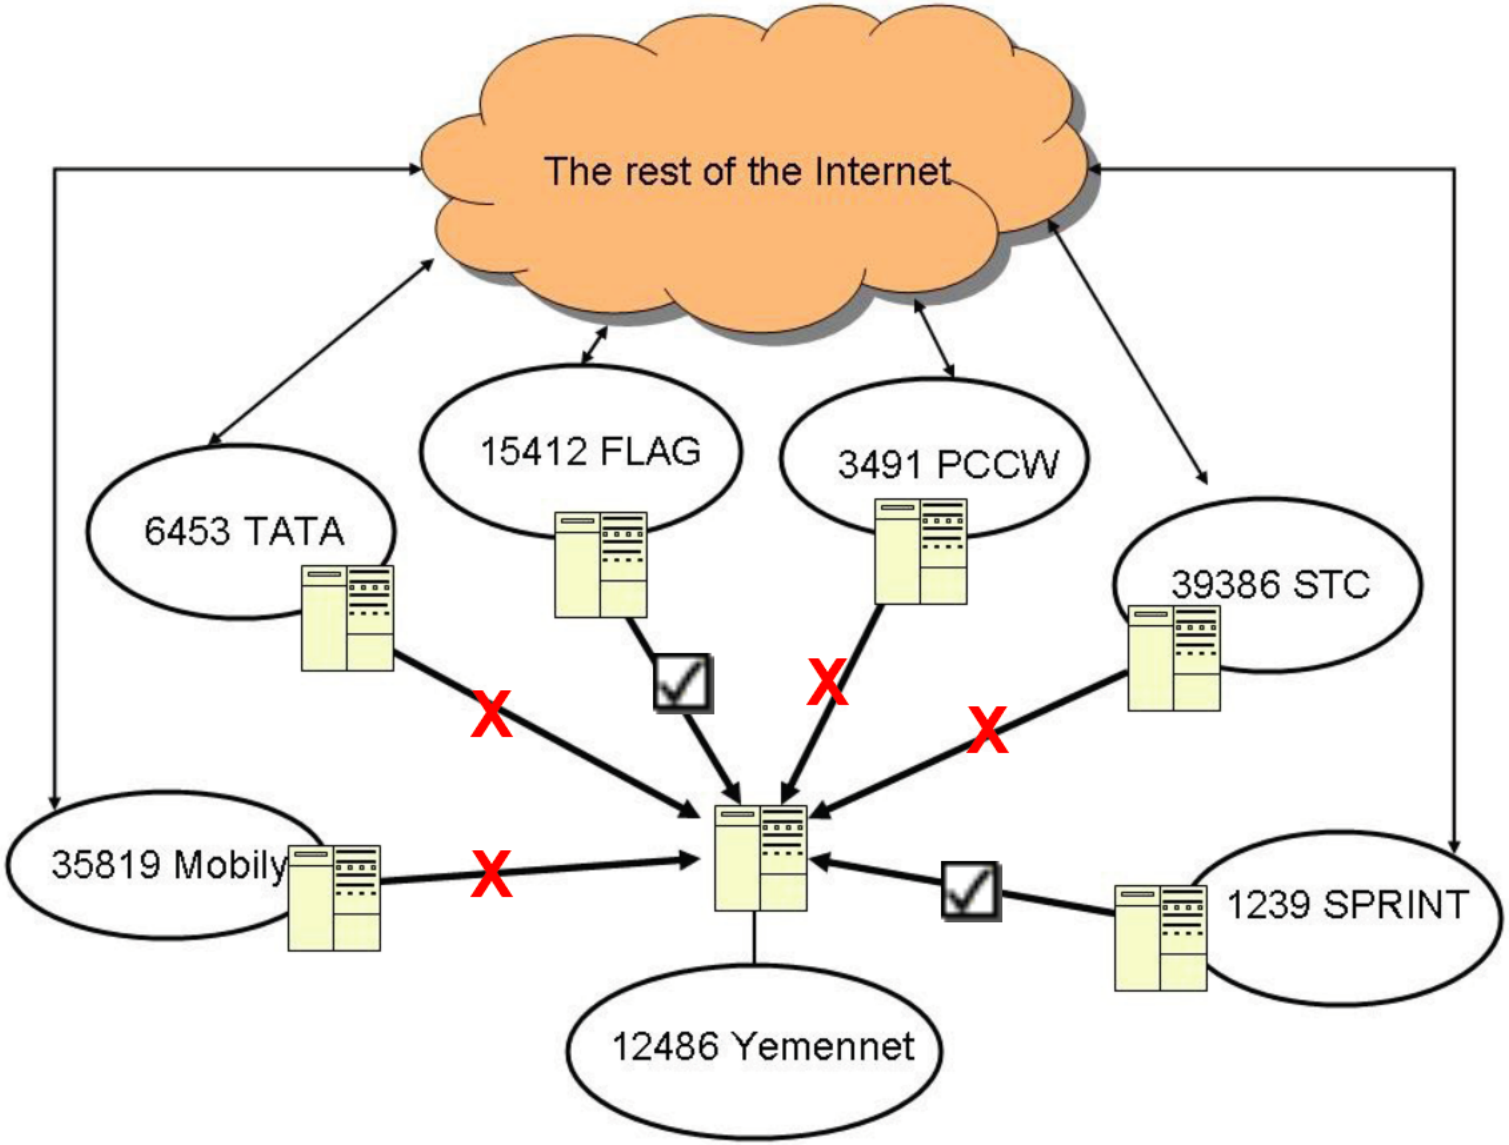

There are a few options you can use to try to make that happen…

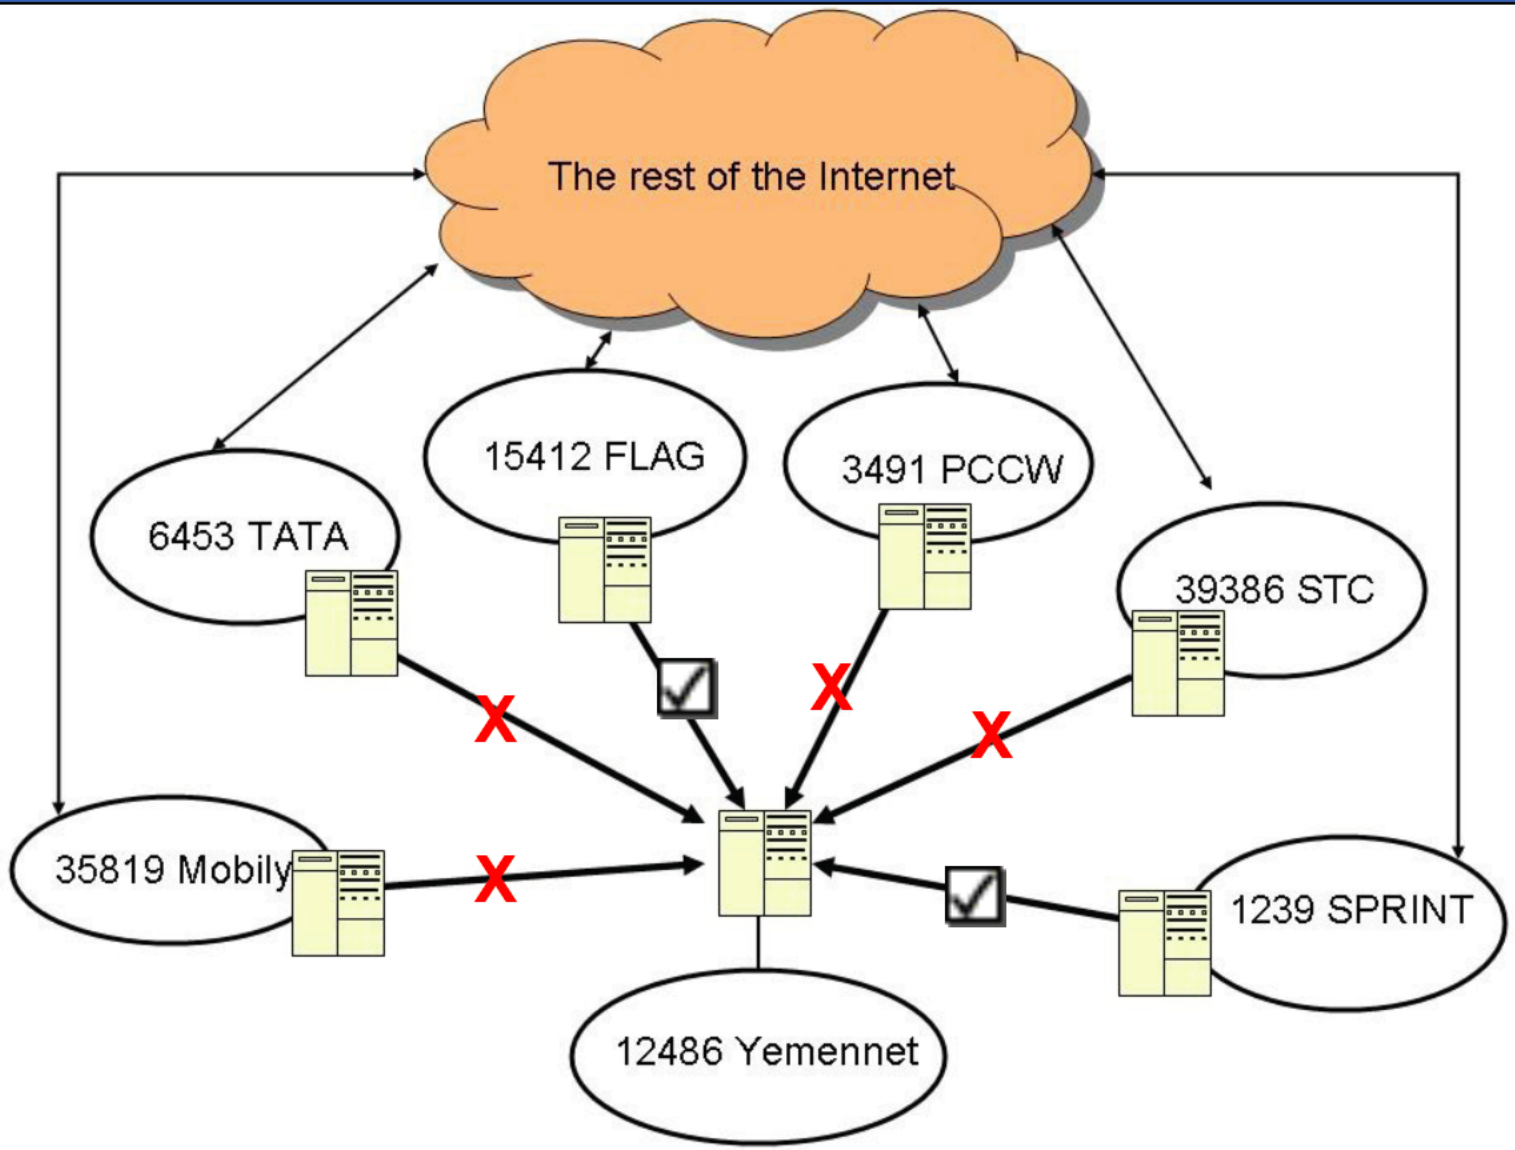

1 – You can try to tweak the BGP routing tables to make your 2 links the most attractive for inbound traffic…

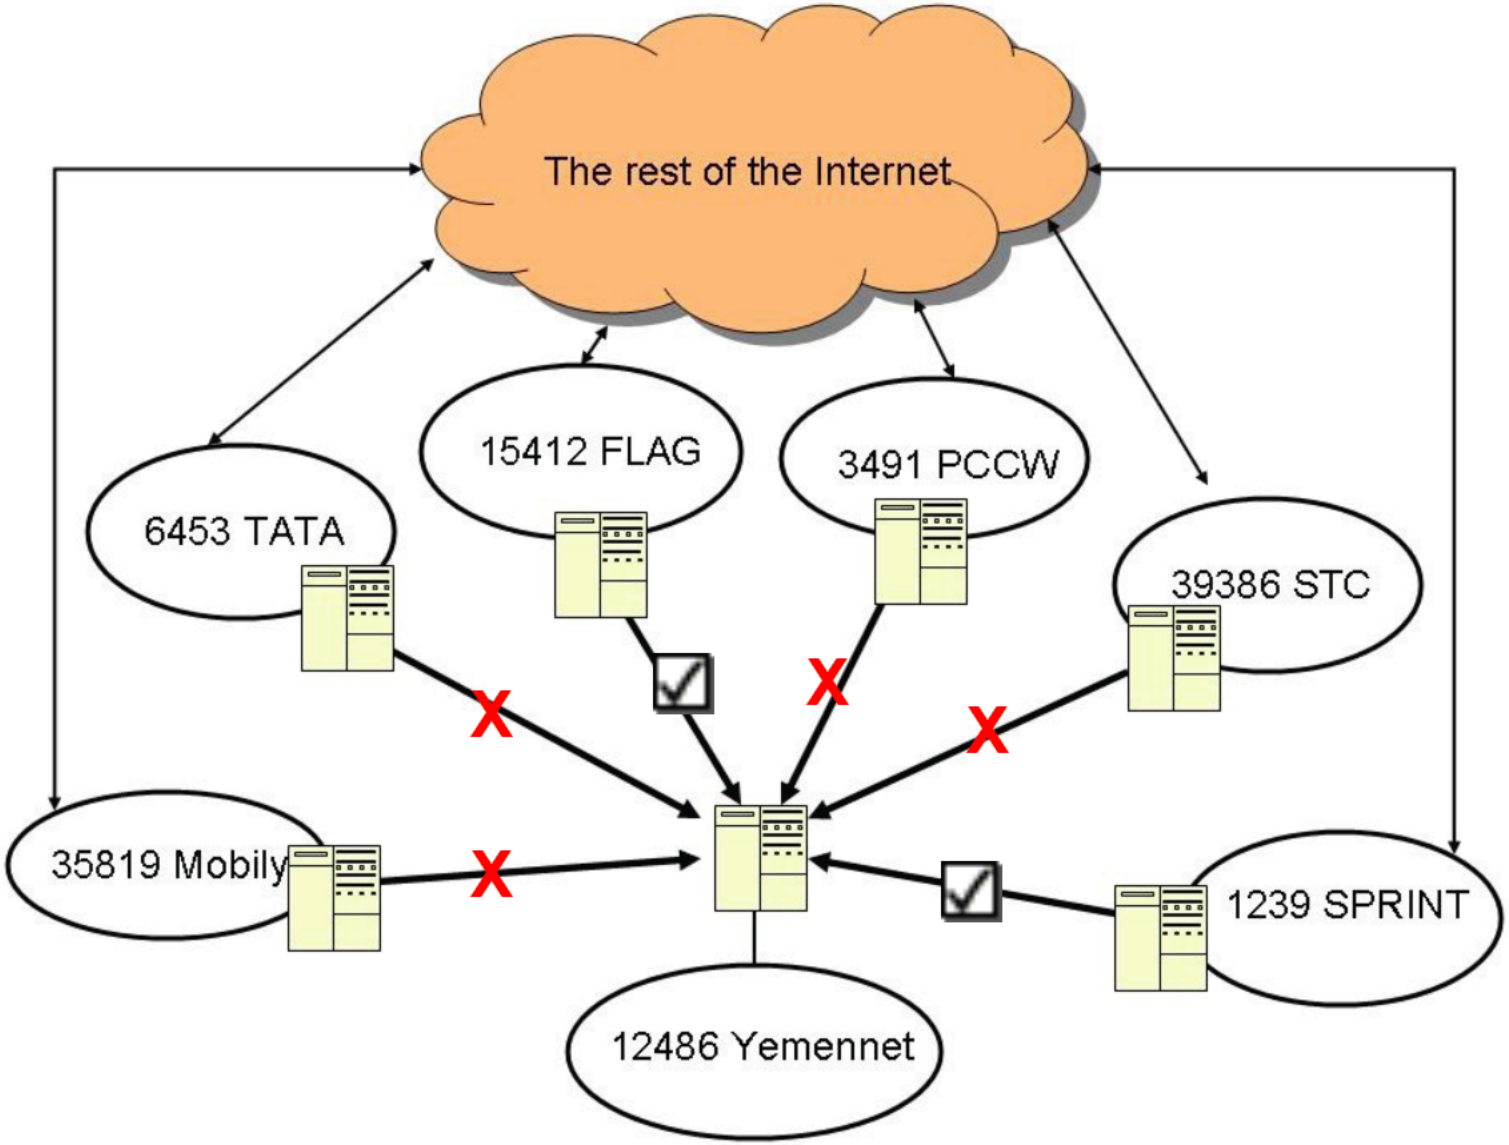

That *could* work, but 2 bad things would result from this:

1 – It's a pretty noisy thing to do on the Internet. People would notice bad BGP updates.

2 – You would throttle ALL Internet traffic through those 2 links, which Yemennet would probably notice.

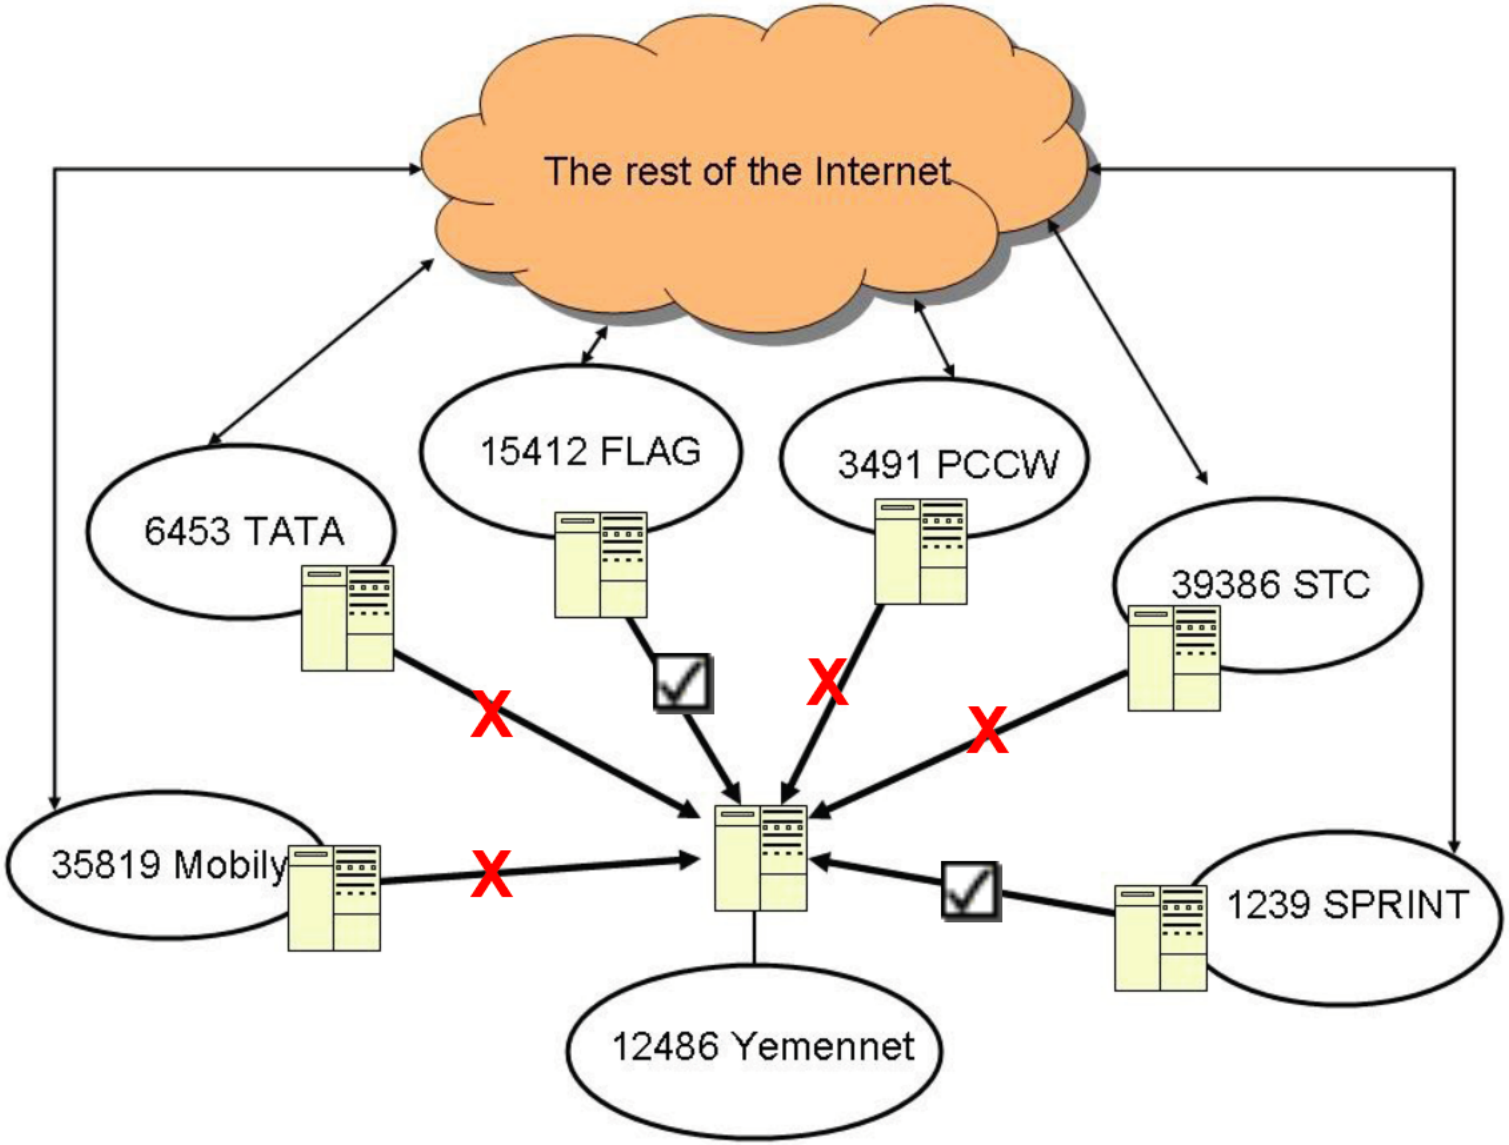

Or…you could blow up or cut all of the International links that we can't collect! ☺

That's fun to think about, but not very reasonable. ☹

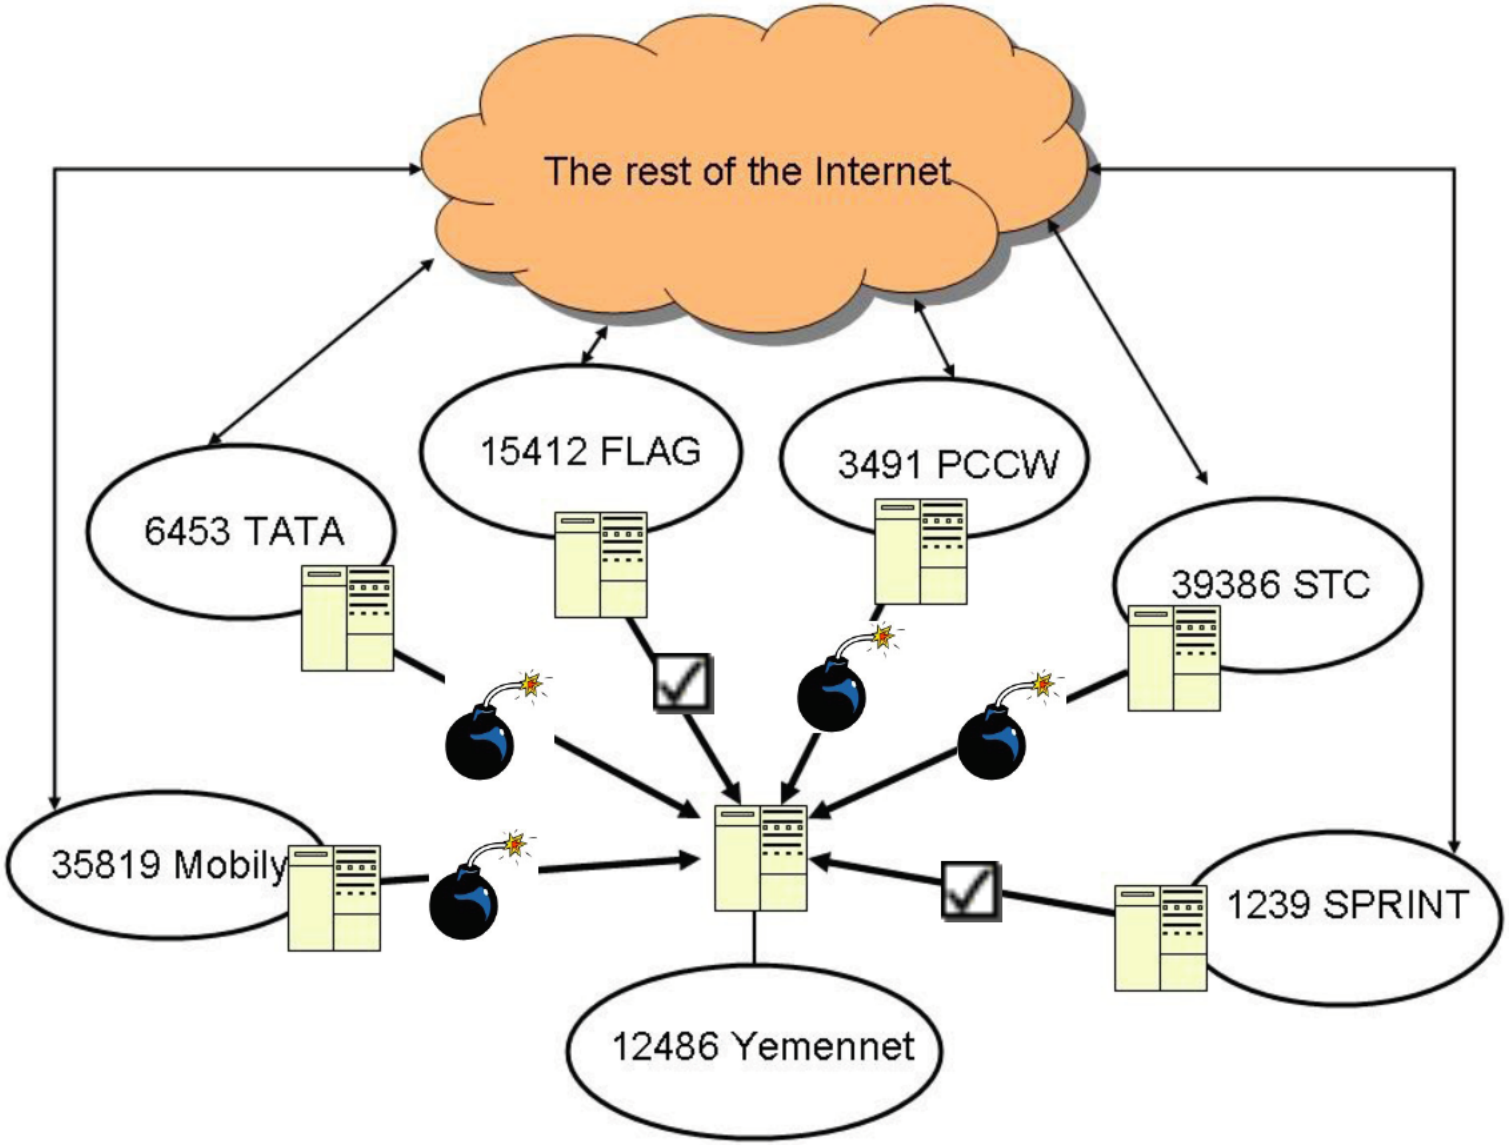

There is another way that might work…we'll have to go back to one of our previous maps though…

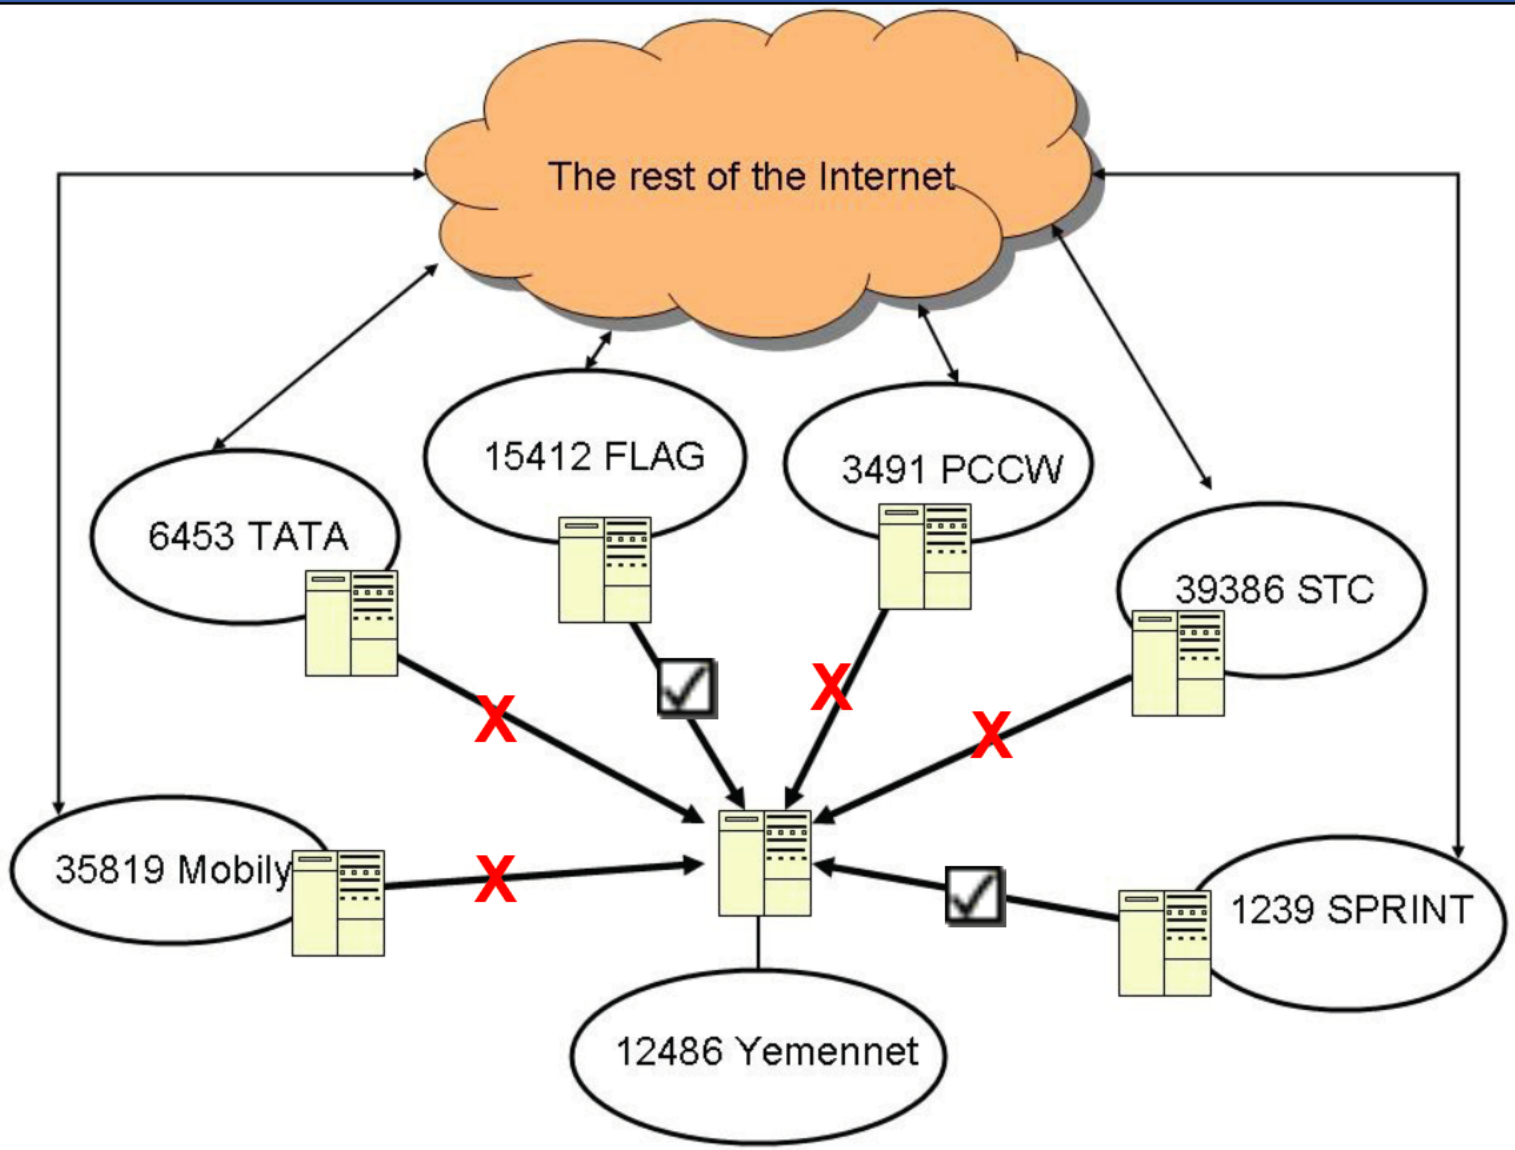

# Remember the /30 connections between Yemen and each of the upstream providers…

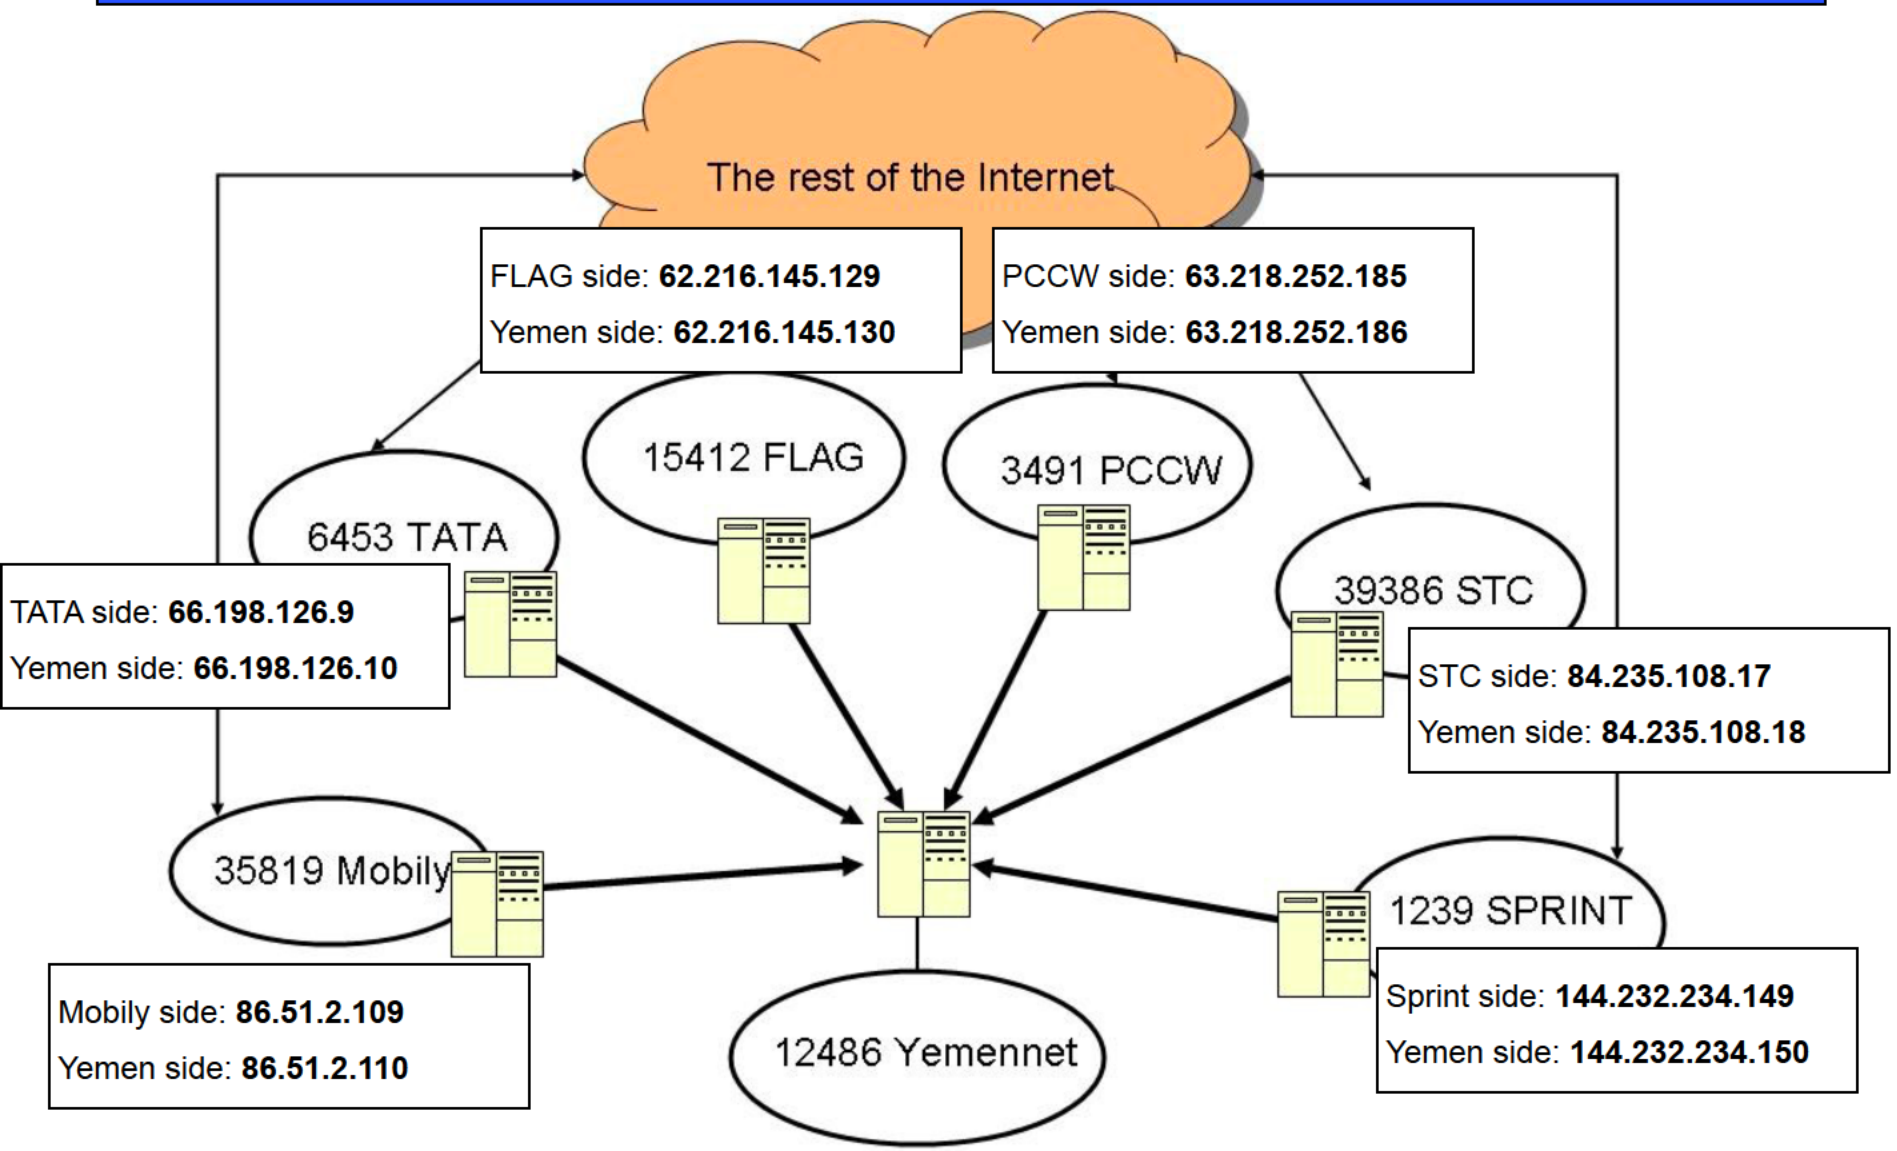

# It actually looks a bit more like this…

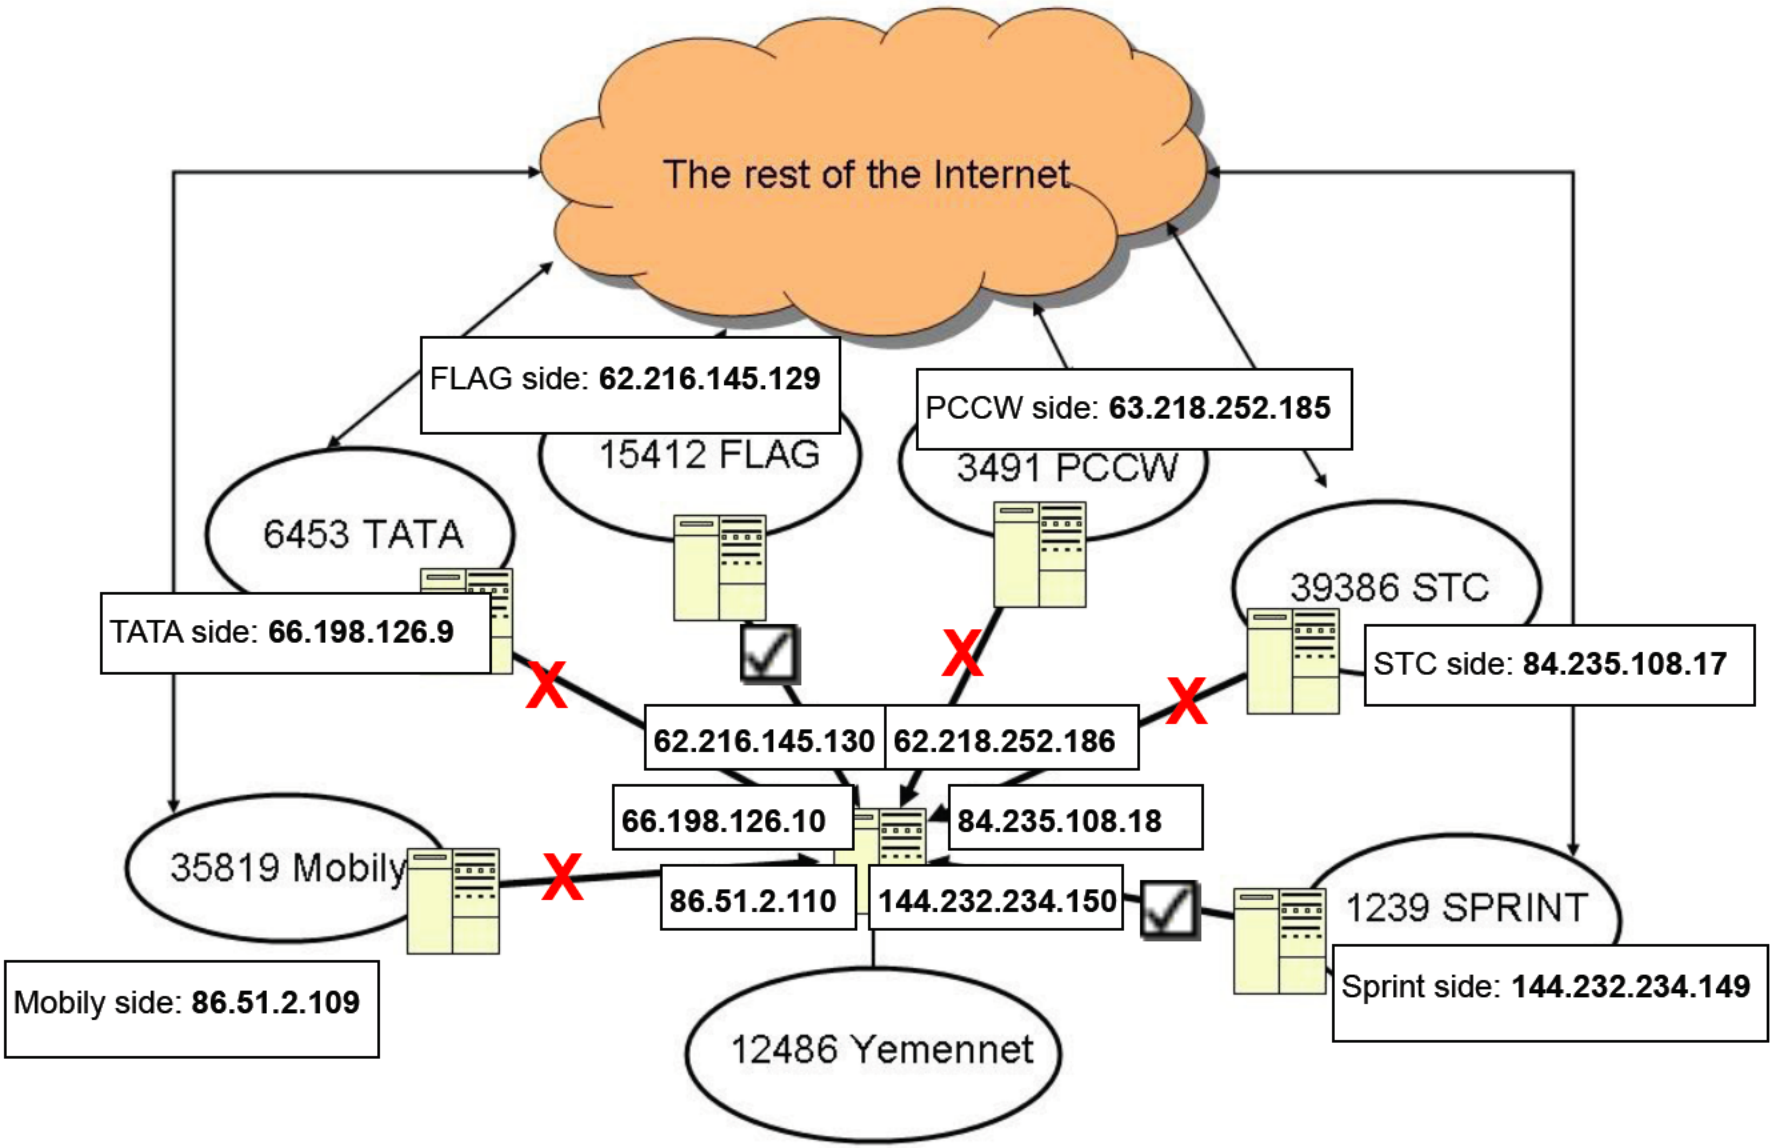

- Whoa, that looks kinda messy, what do I really need to take away from this slide?

- Good question, just keep in mind the fact that each provider has put one of *their own* IP addresses in Yemen for those connections. Here's why it matters…

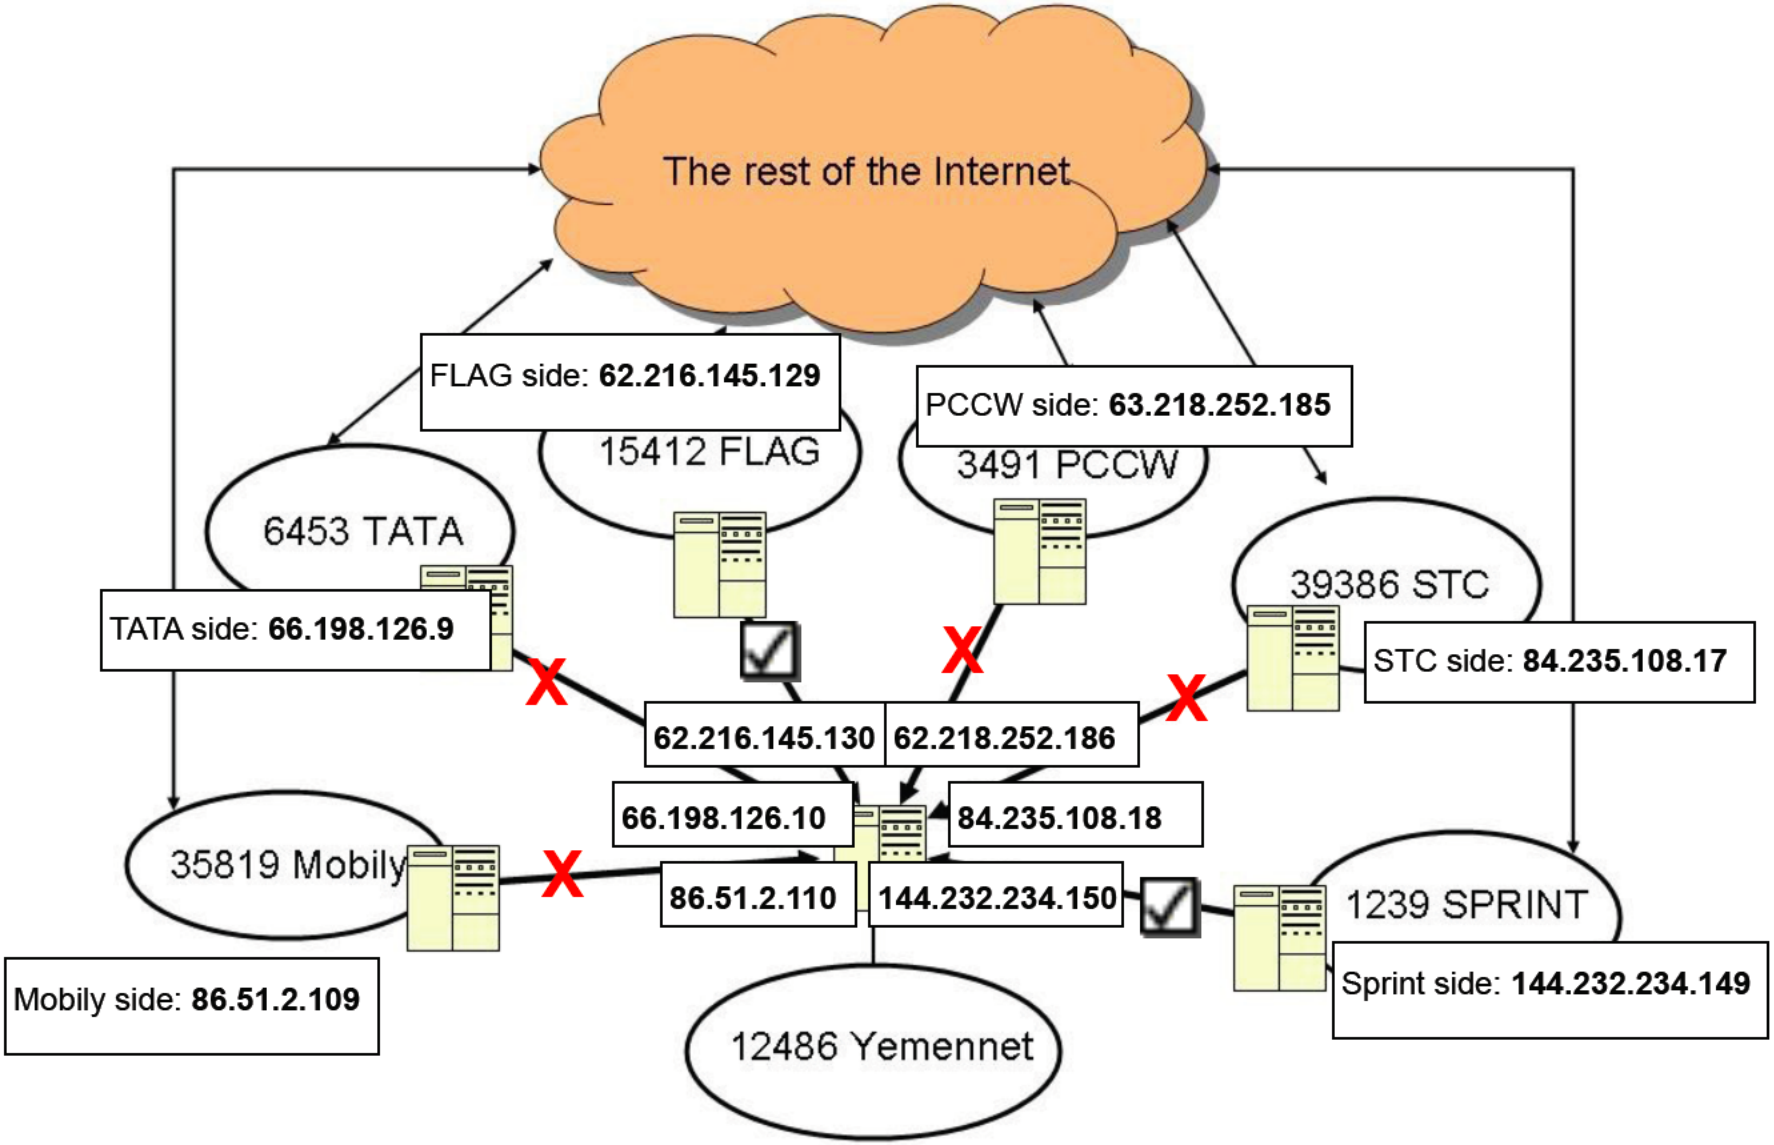

Remember, if you do a traceroute to a random IP address in AS12486, you won't know which link it will go through to get there…

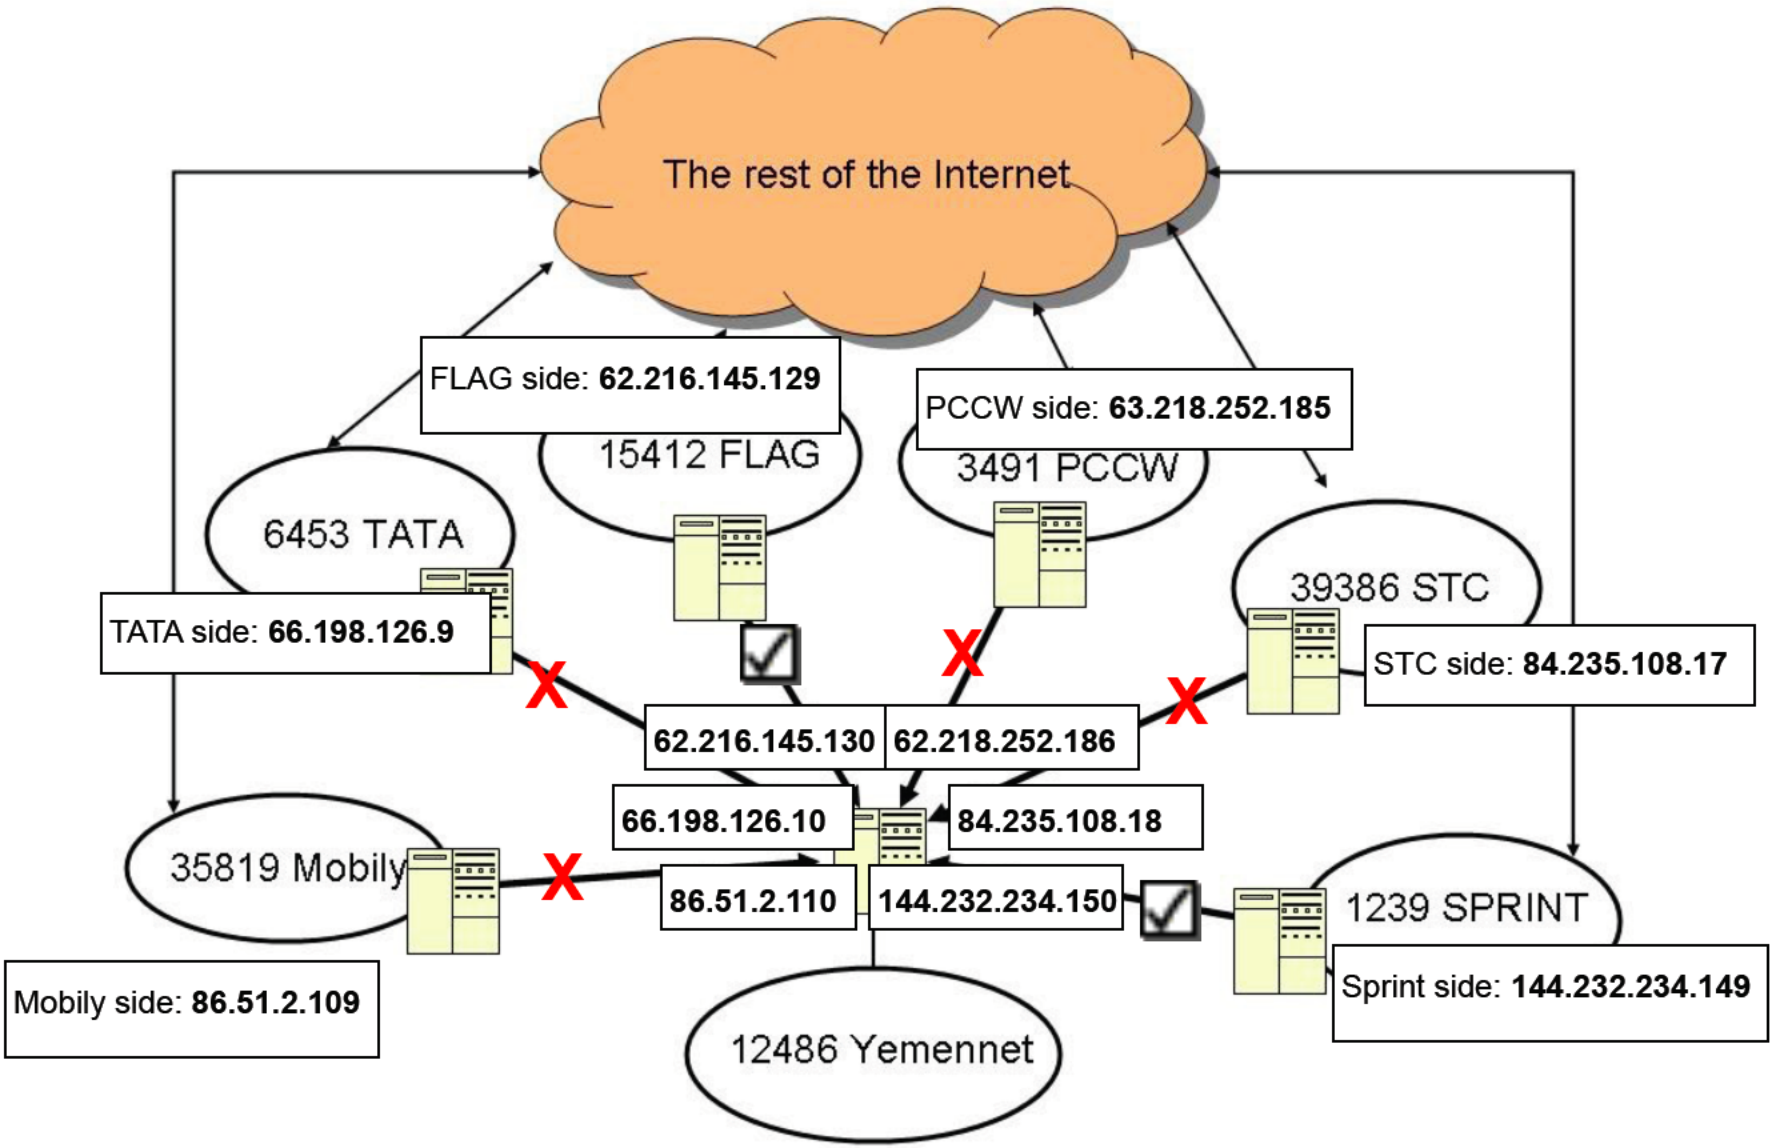

But, hypothetically, what happens if you do a traceroute to 86.51.2.110 (the Yemen side of its Mobily connection) from some random place on the Internet?

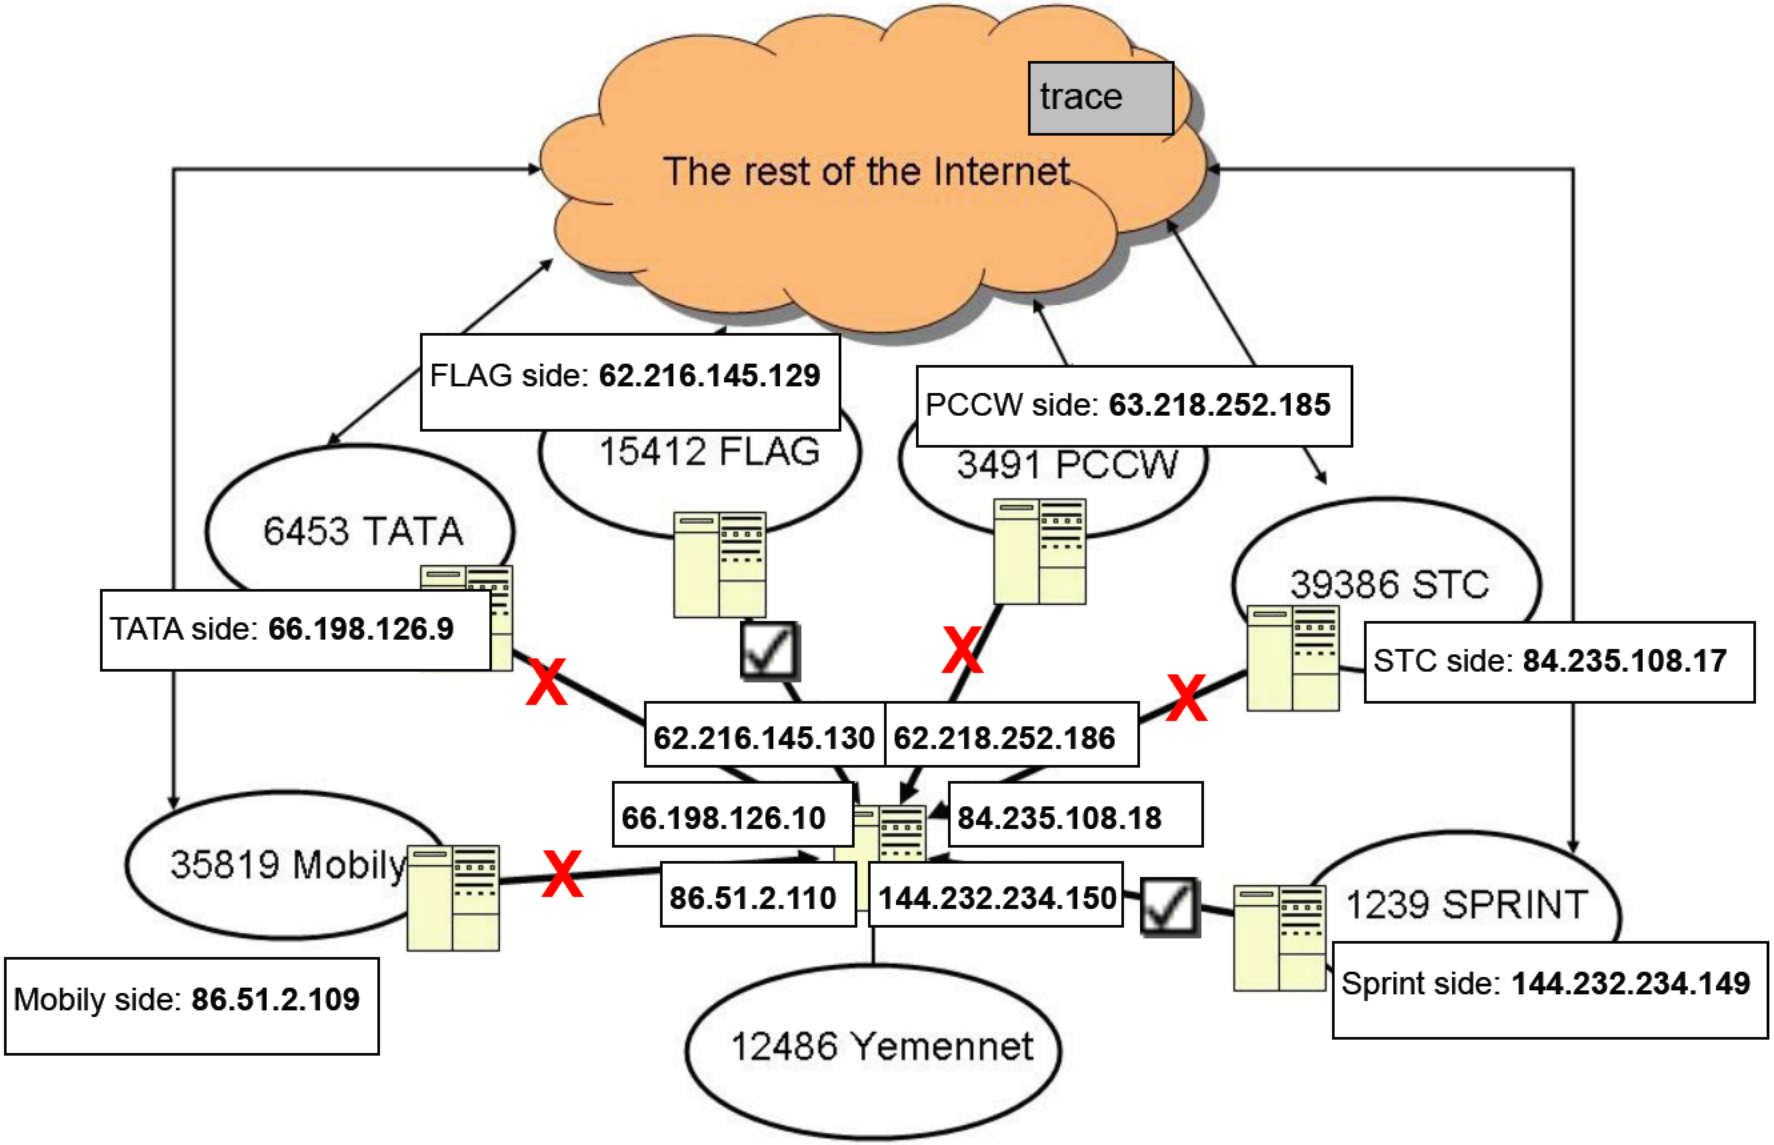

-Ok, it went where I expected it to go…what's so special about that?

- First, keep in mind the IP we traced to belongs to Mobily, not Yemennet. So the traffic will first get routed to Mobily's network. From there, it happens to reside on a router belonging to Yemennet.

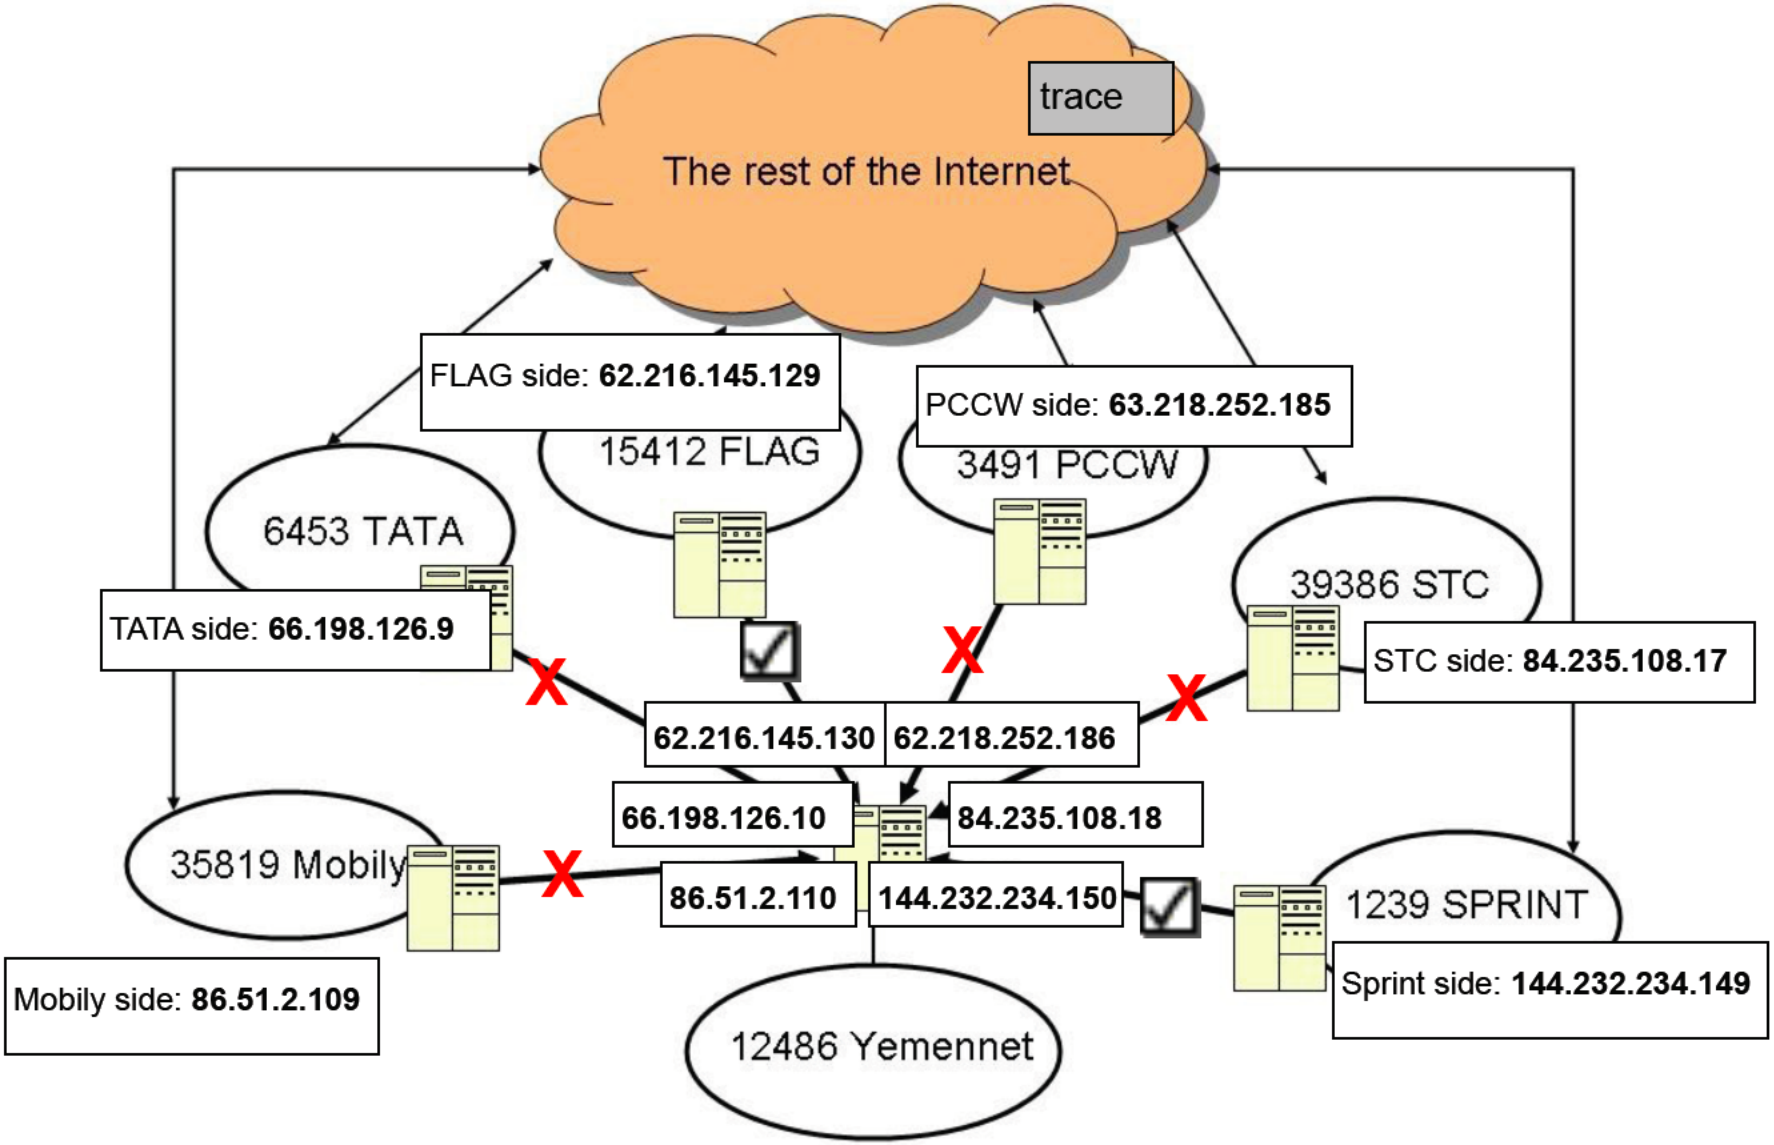

-AAAAAND?

- By tracing to the Mobily IP on Yemennet's router, we forced the trace to go through Mobily to get there. What happens if we did that for FLAG's IP address?

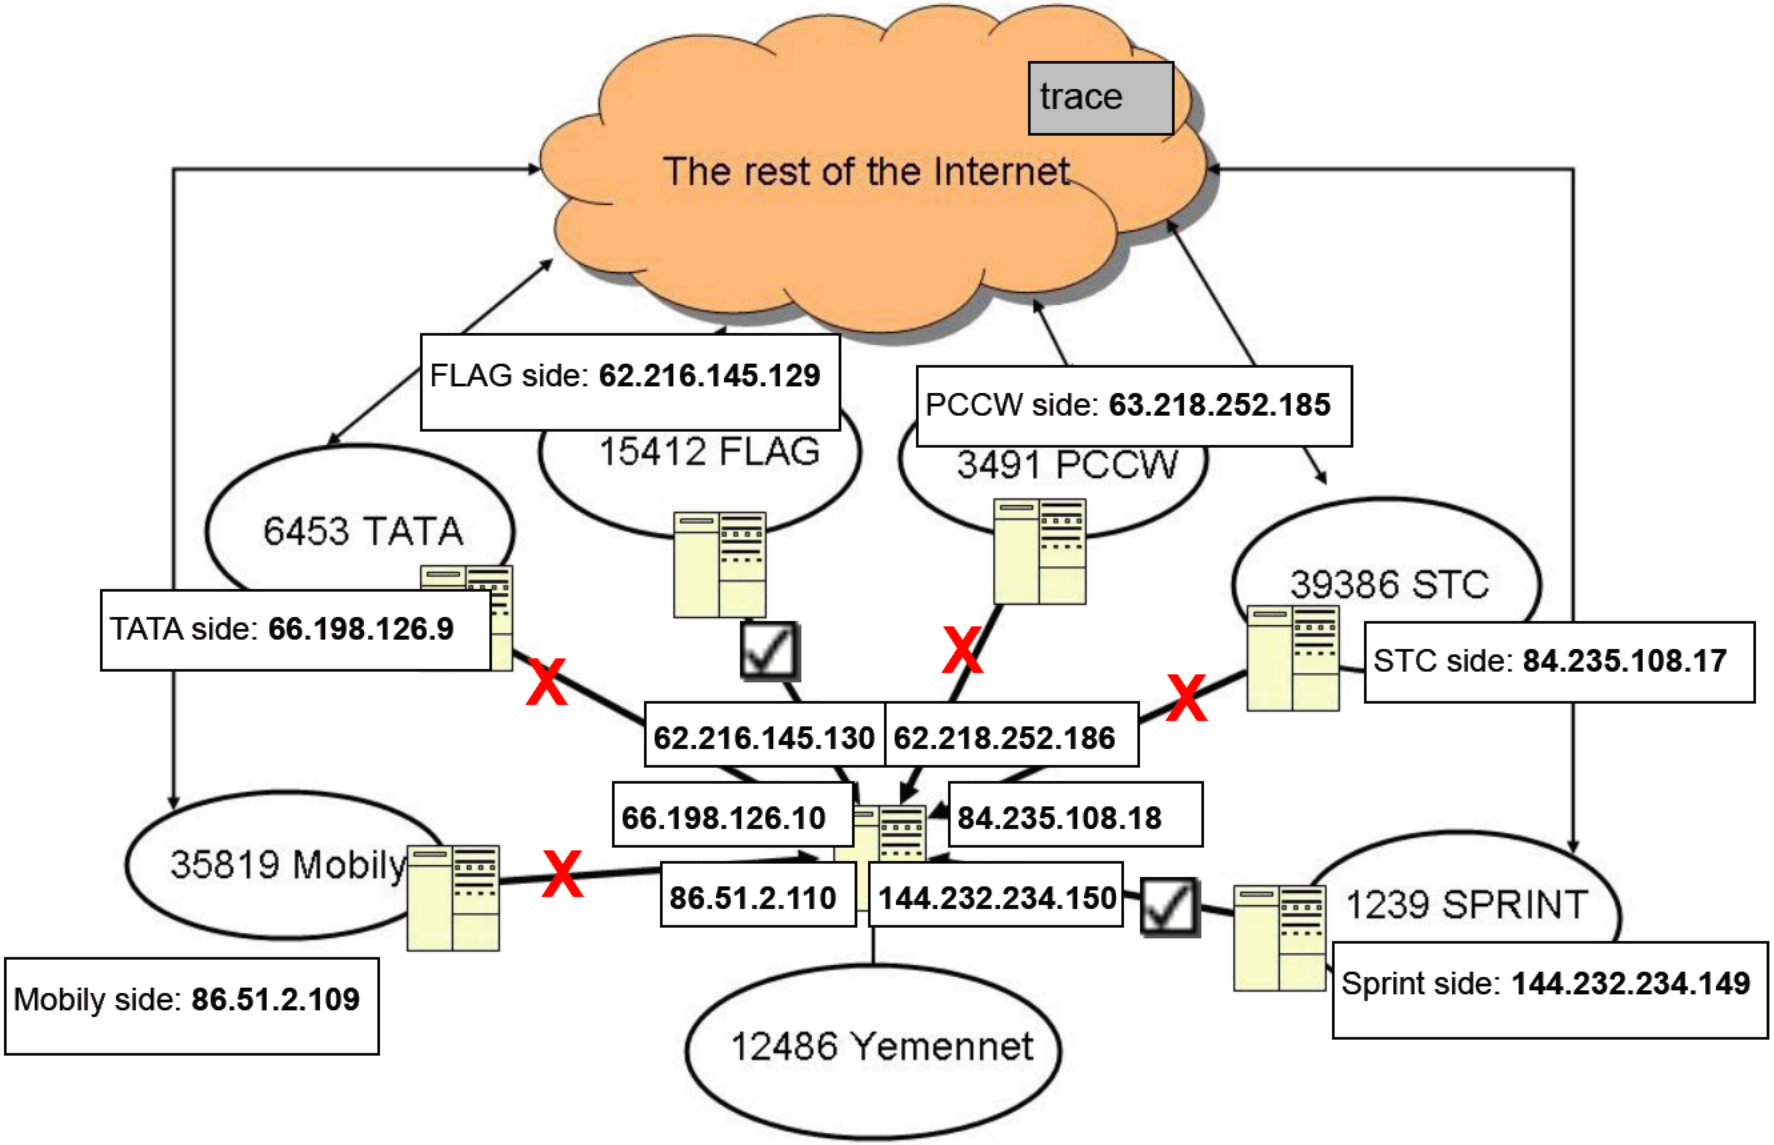

# Starting traceroute from anywhere on the Internet to 62.216.145.130…

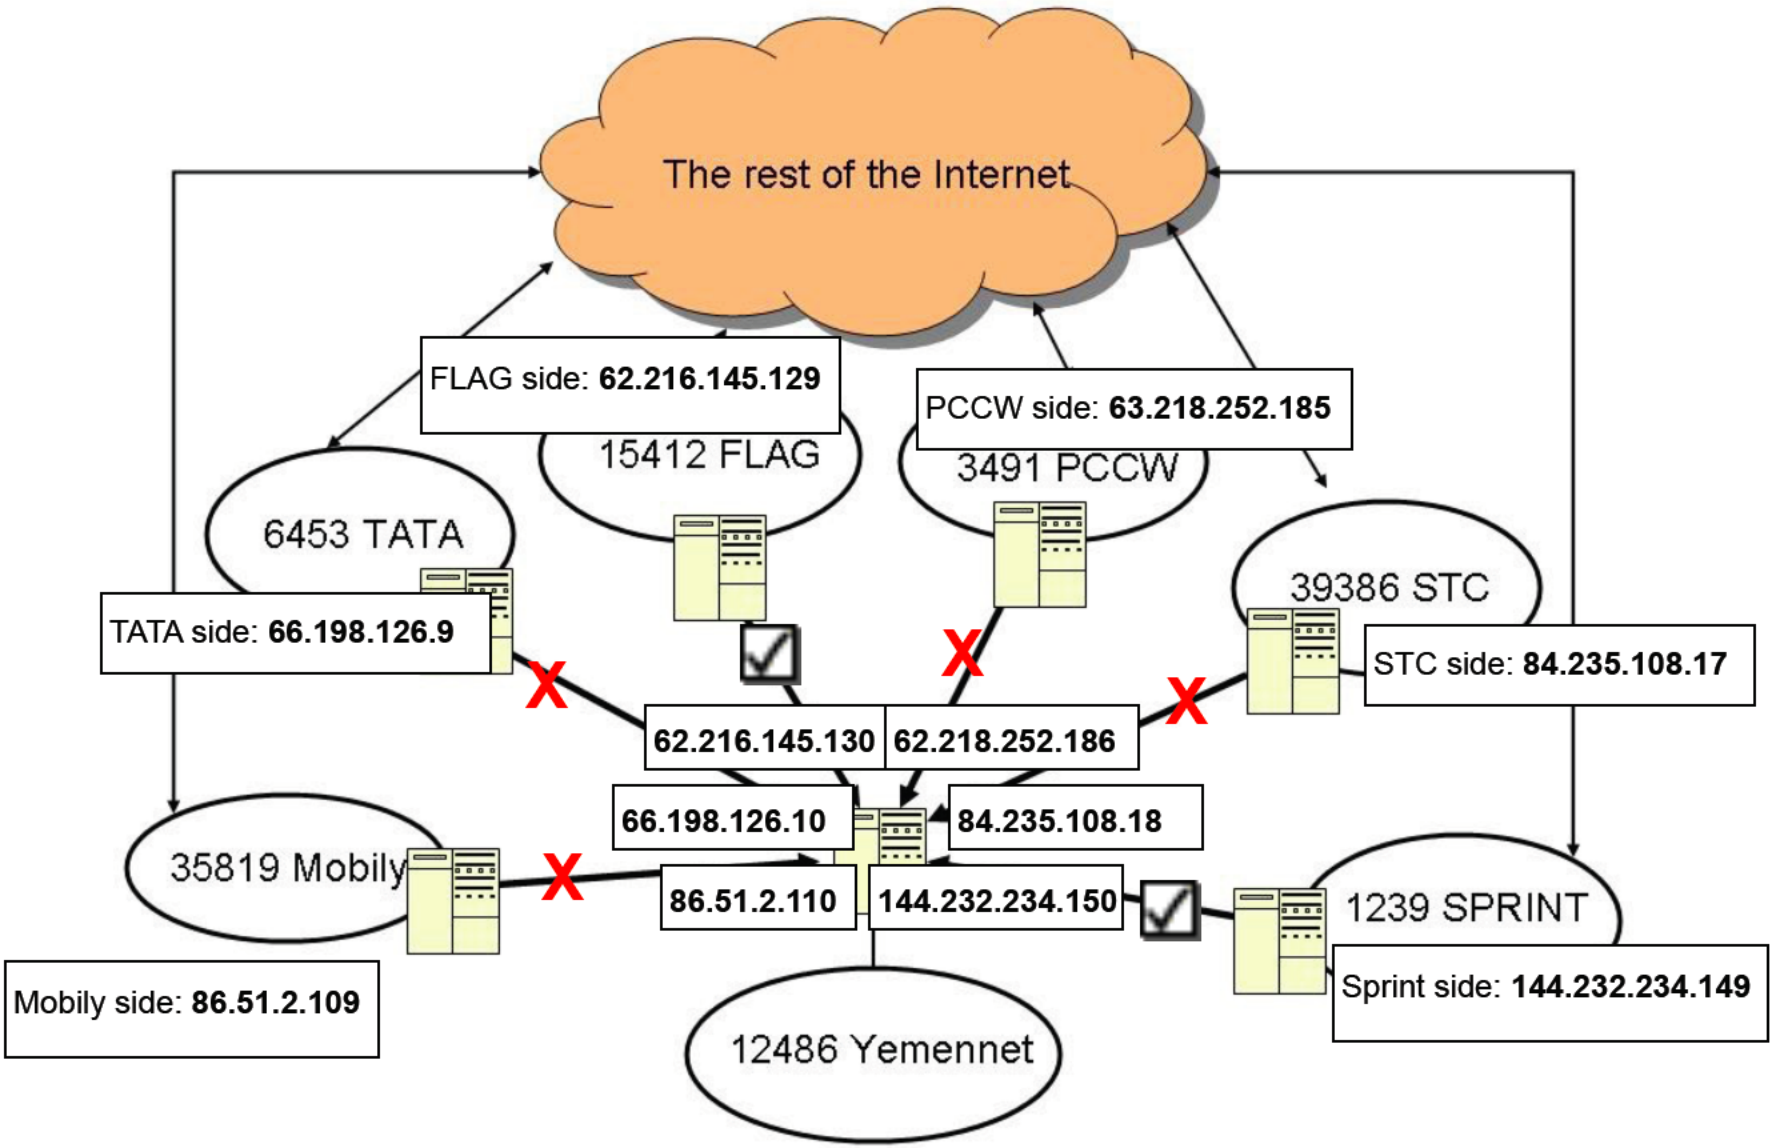

# Starting traceroute from anywhere on the Internet to 62.216.145.130…

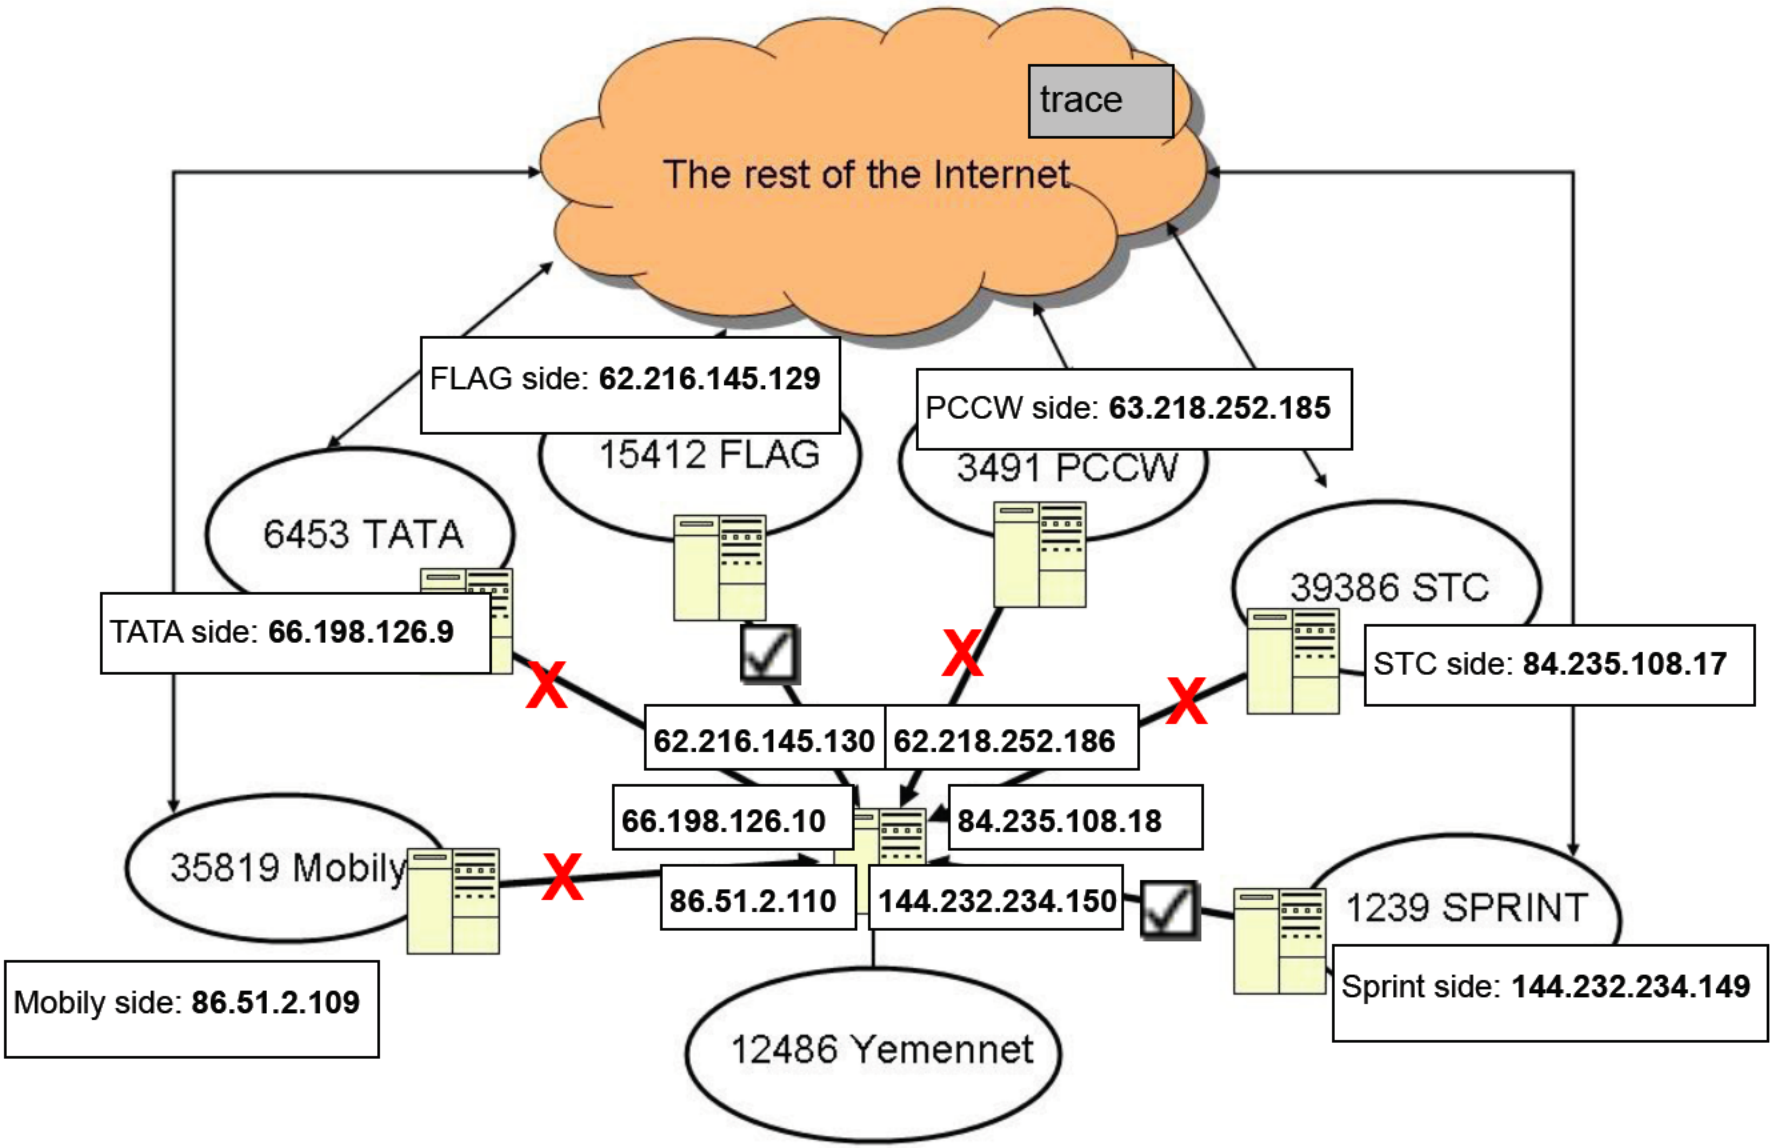

We have just successfully forced traffic to go through FLAG's network and over a link that we can collect! <throws confetti in the air>

So, it seems we have found a good candidate to shape traffic to from anywhere!

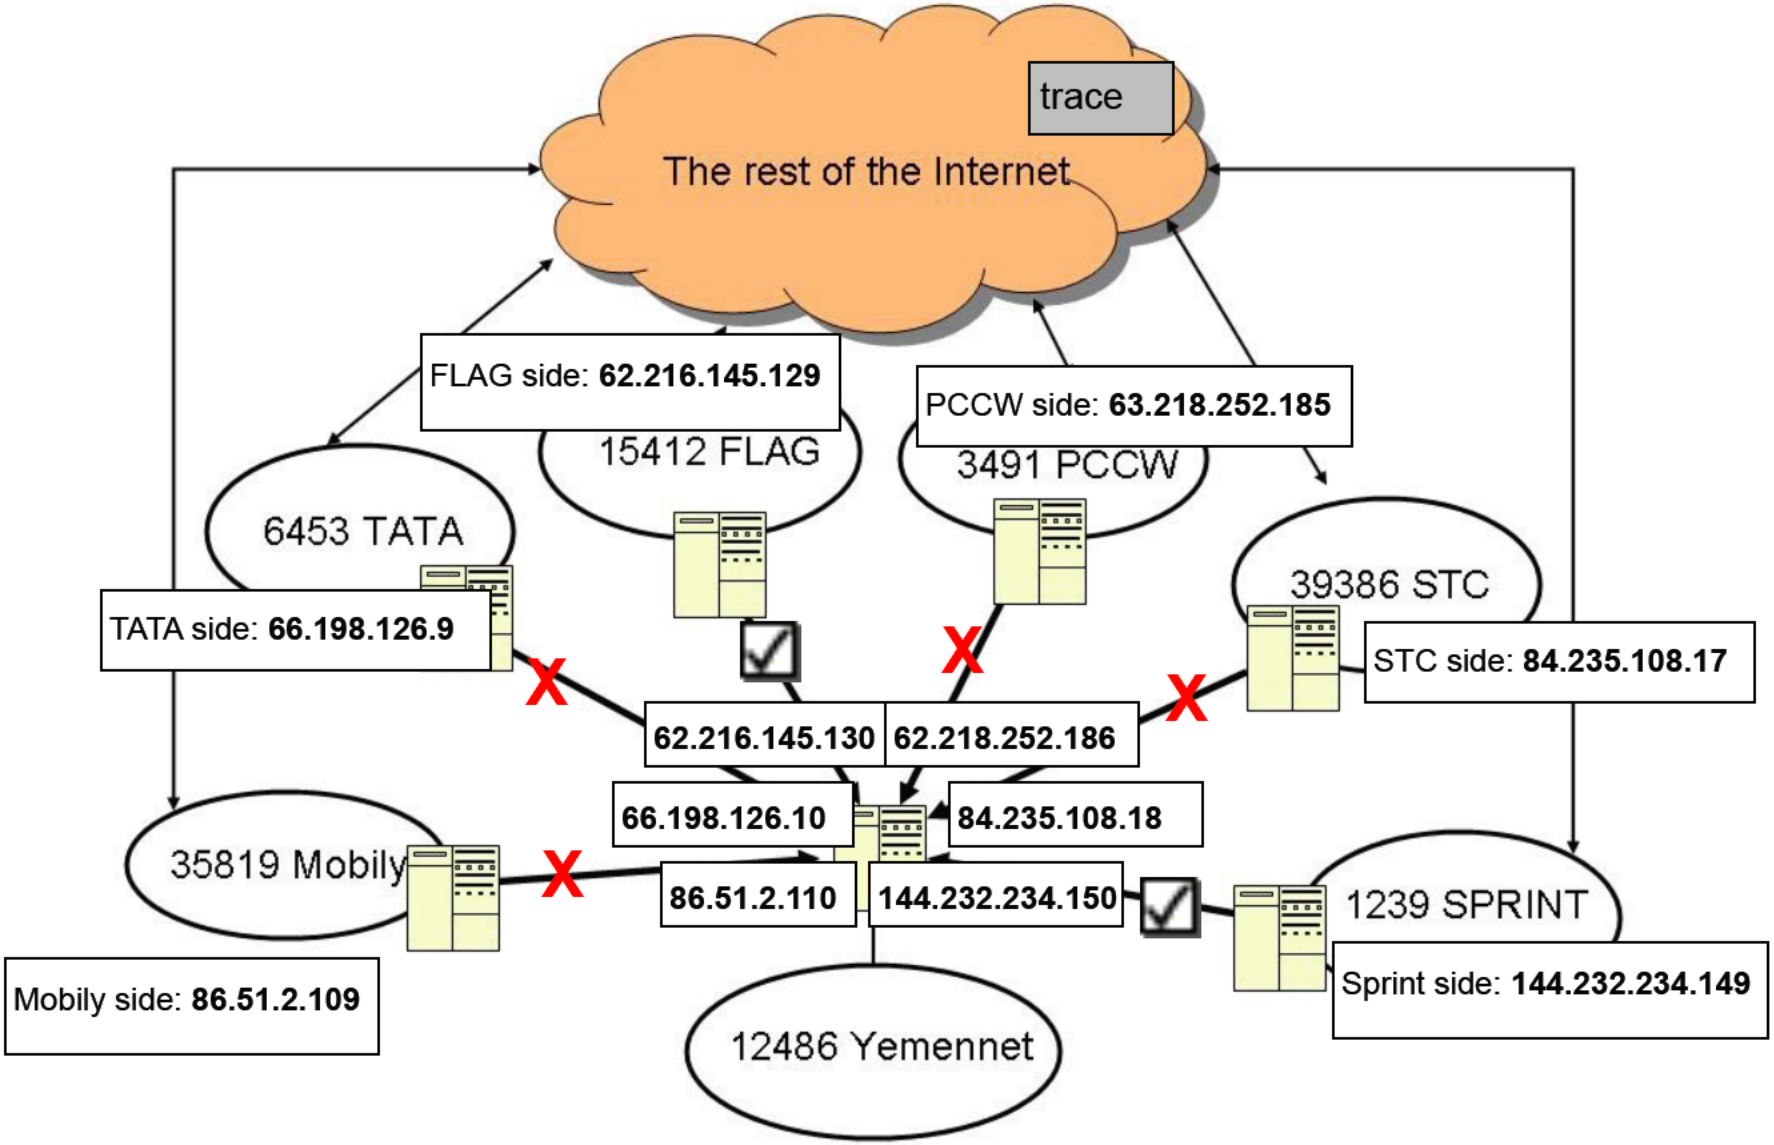

# Scenario – “I tried shaping one time and it didn’t work…”

- As we’ve seen earlier, there are many facets that make shaping efforts unreliable
- It matters whether you are trying to shape traffic OUT of a network or whether you are starting at a random place on the Internet and trying to shape traffic INTO the network
- So what steps could you take?

# So, you might consider the following train of thought…

- First you say, “I want to do shaping through SIGAD US-9999”
- Then look at all of the links collected at that site (probably in BLACKPEARL)
- Find a World-to-Geekistan link over CASN GE010100000 (we’ll assume you also know that this is actually a Level 3-to-Geekistan link)
- Look at the IP space on the dest side of the link and say, “I will send my exfil to that IP space, and it should go through US-9999, CASN GE010100000.”
- Then you are left sorely disappointed when your exfil isn’t reliably collected.
- What went wrong?  Let’s consider what we know so far…

This isn't a bad start for network knowledge, but there's still some missing pieces…

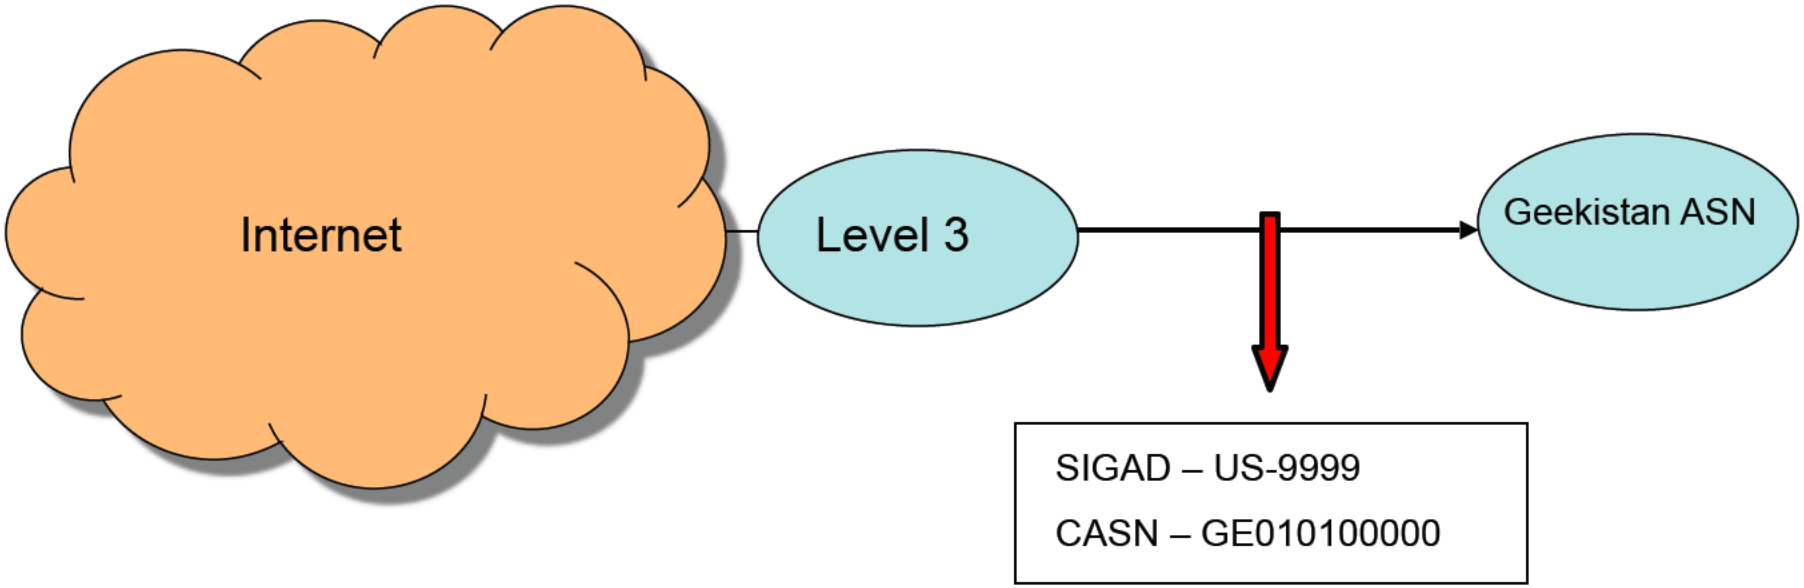

For example, do you know if Level 3 is the *ONLY* upstream provider for Geekistan? Or are there other ways for traffic to get in and out of that network? How would you find out?

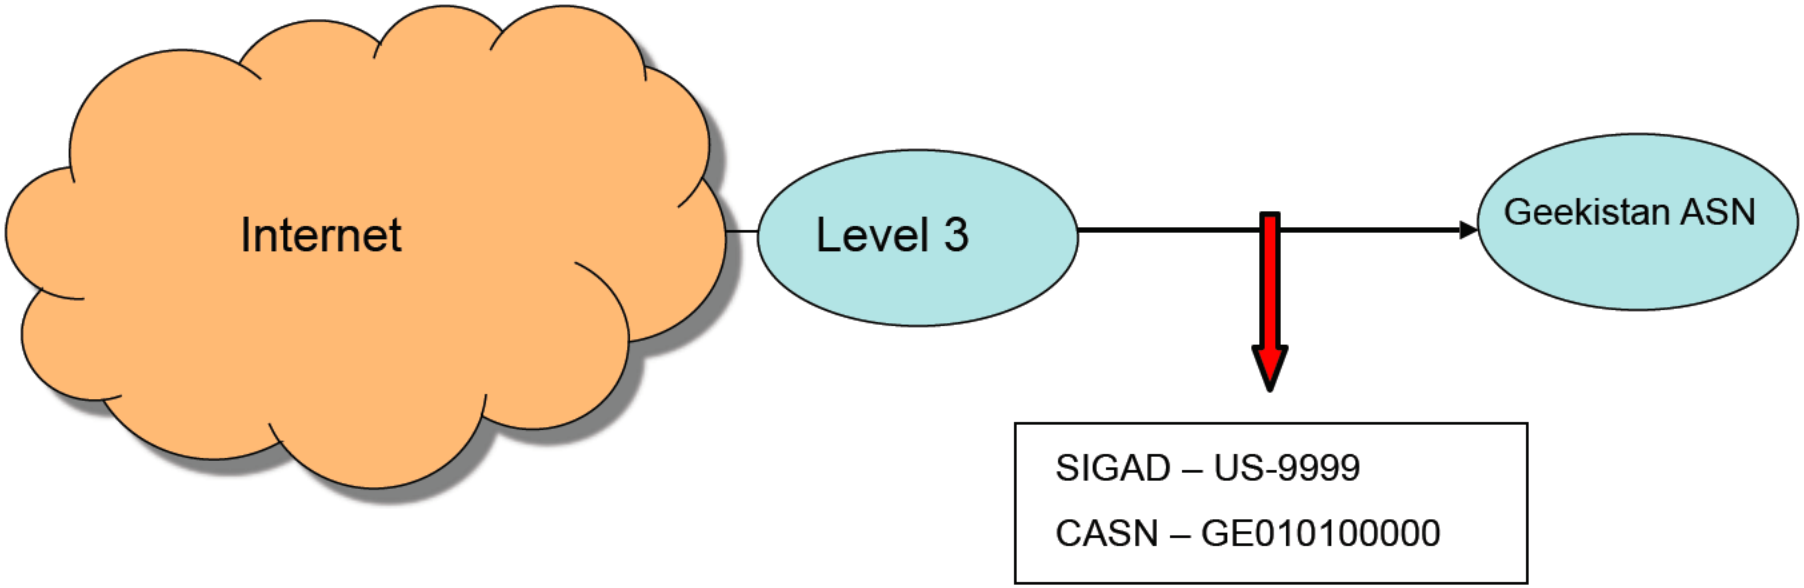

You could look at BGP to find upstream providers for Geekistan ASN.

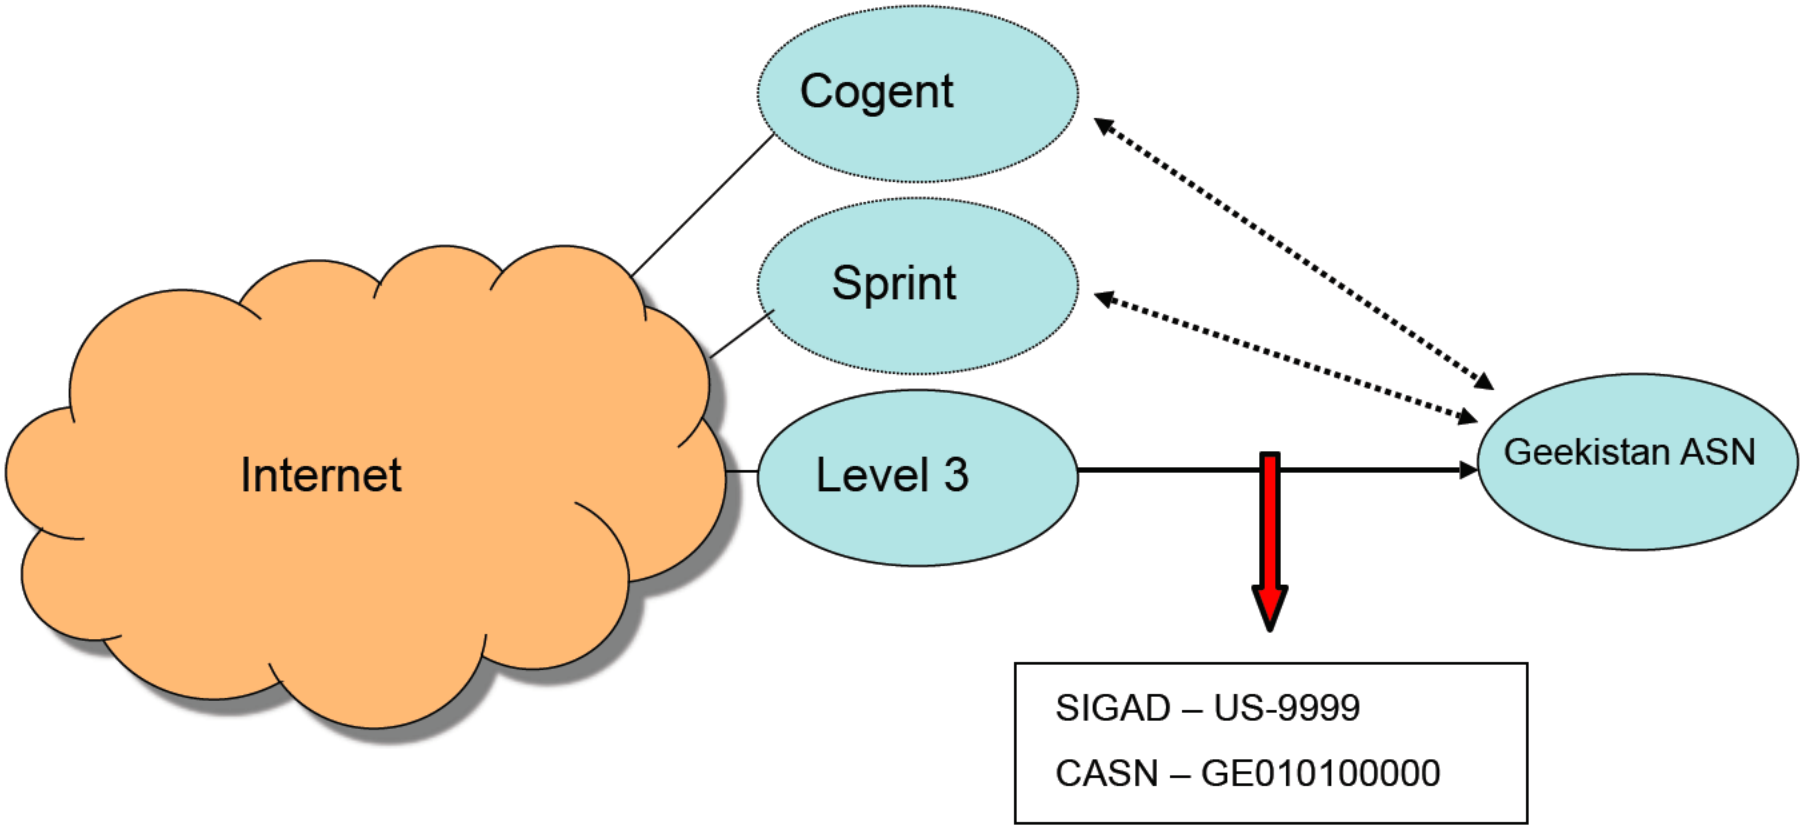

You could do traceroutes from random places on the Internet to IP's in Geekistan's network and see who it goes through to get there.

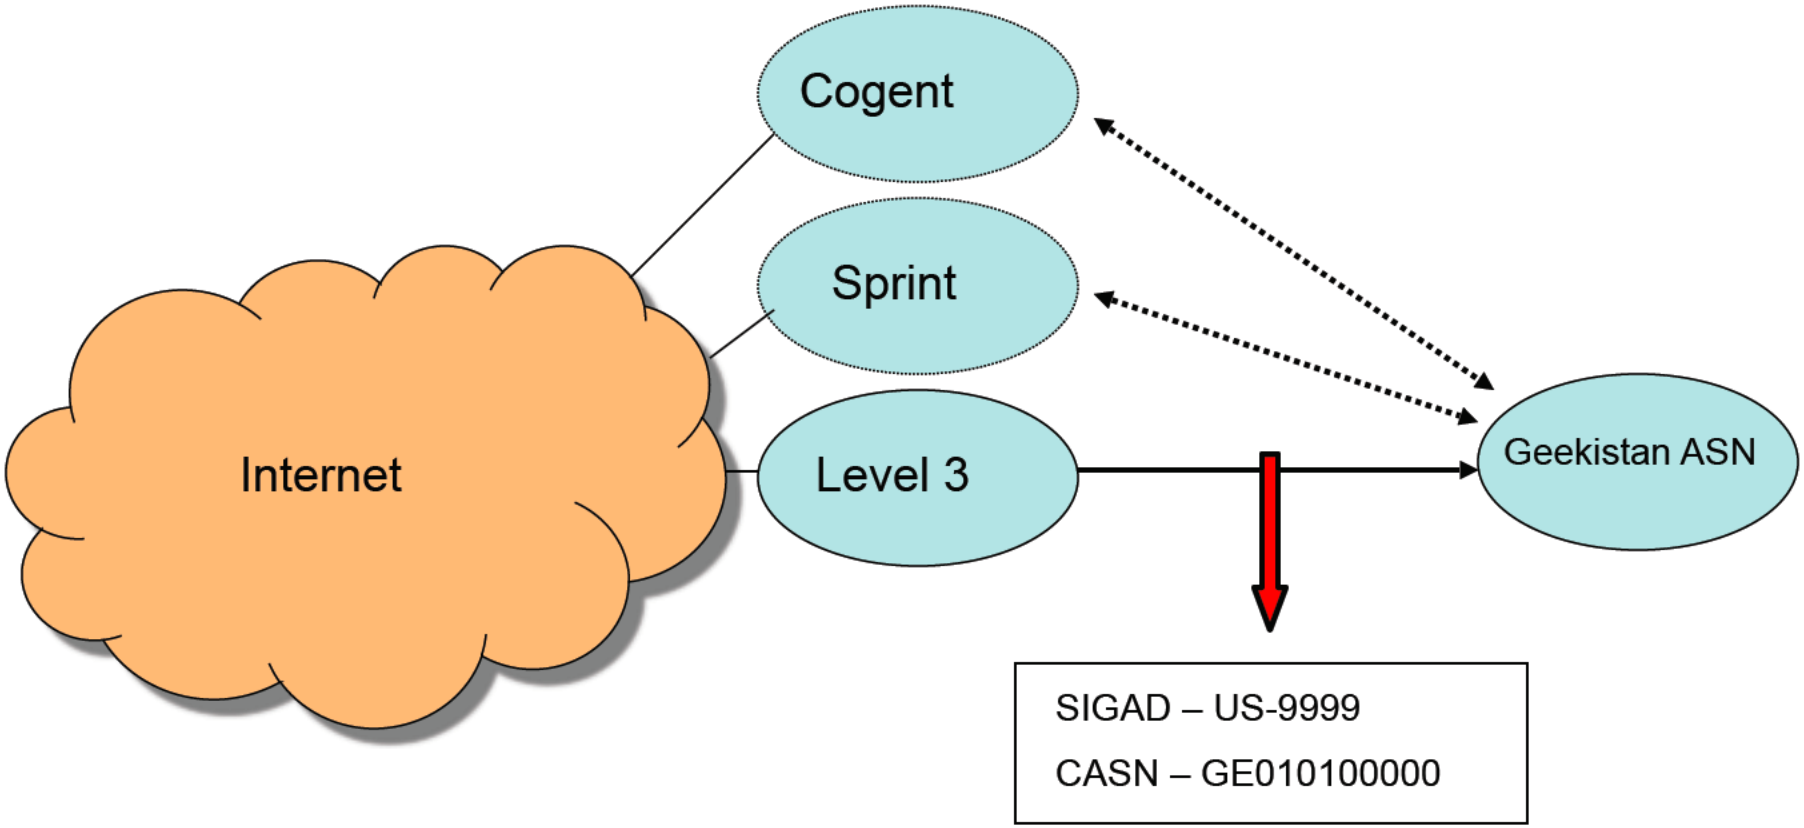

You could do traceroutes from random places on the Internet to IP's in Geekistan's network and see who it goes through to get there.

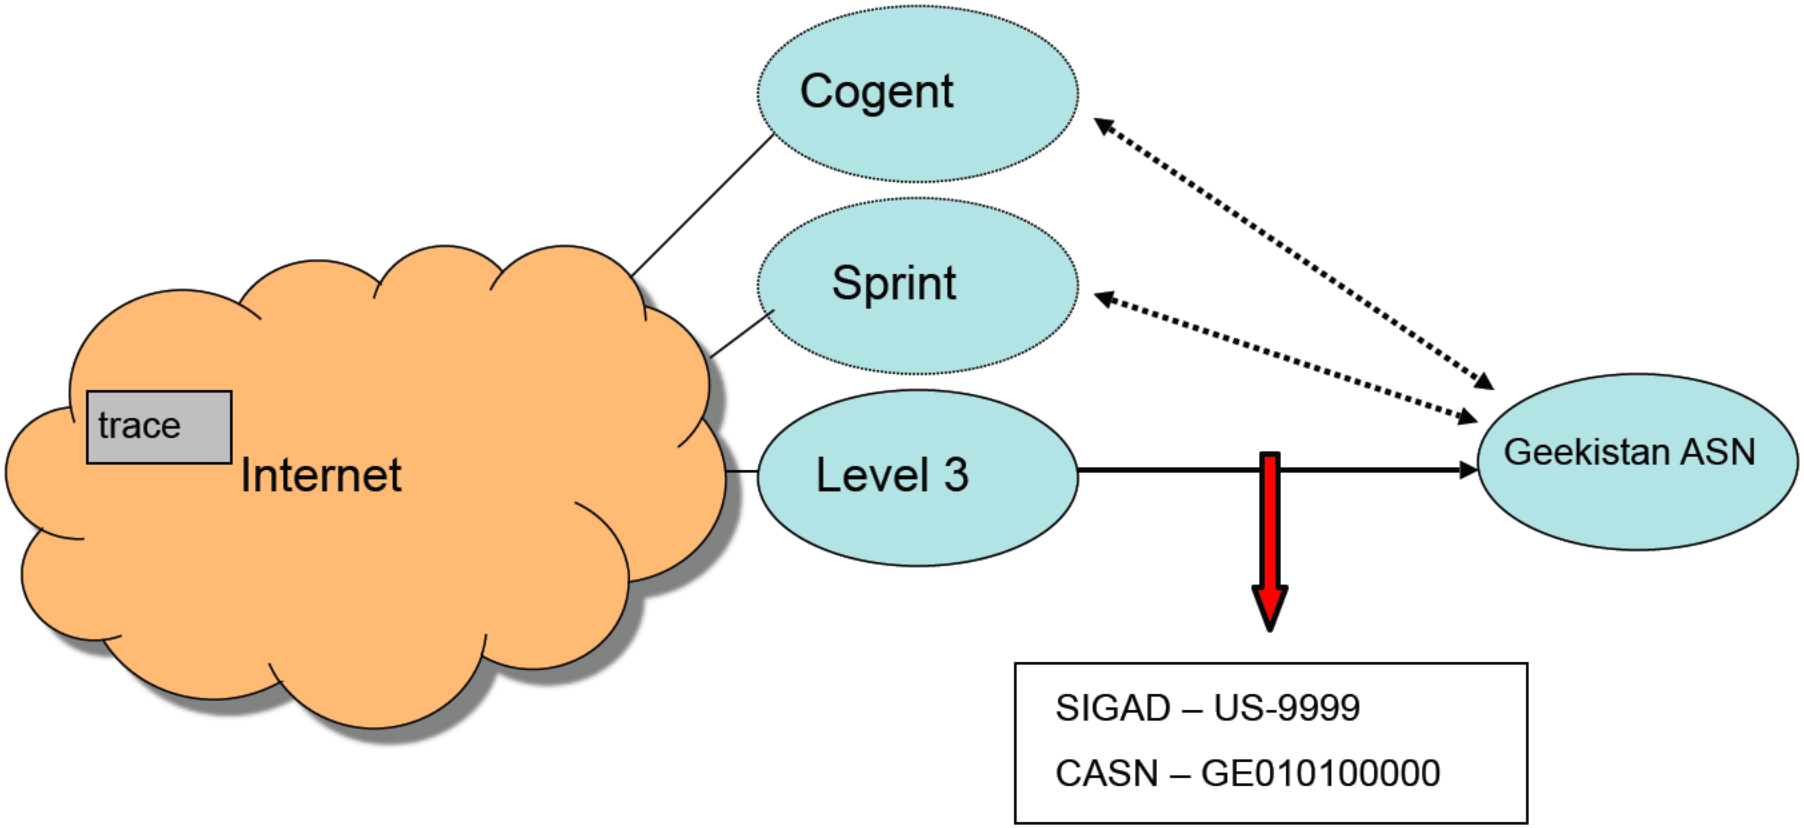

Now with this information, what do you do next?

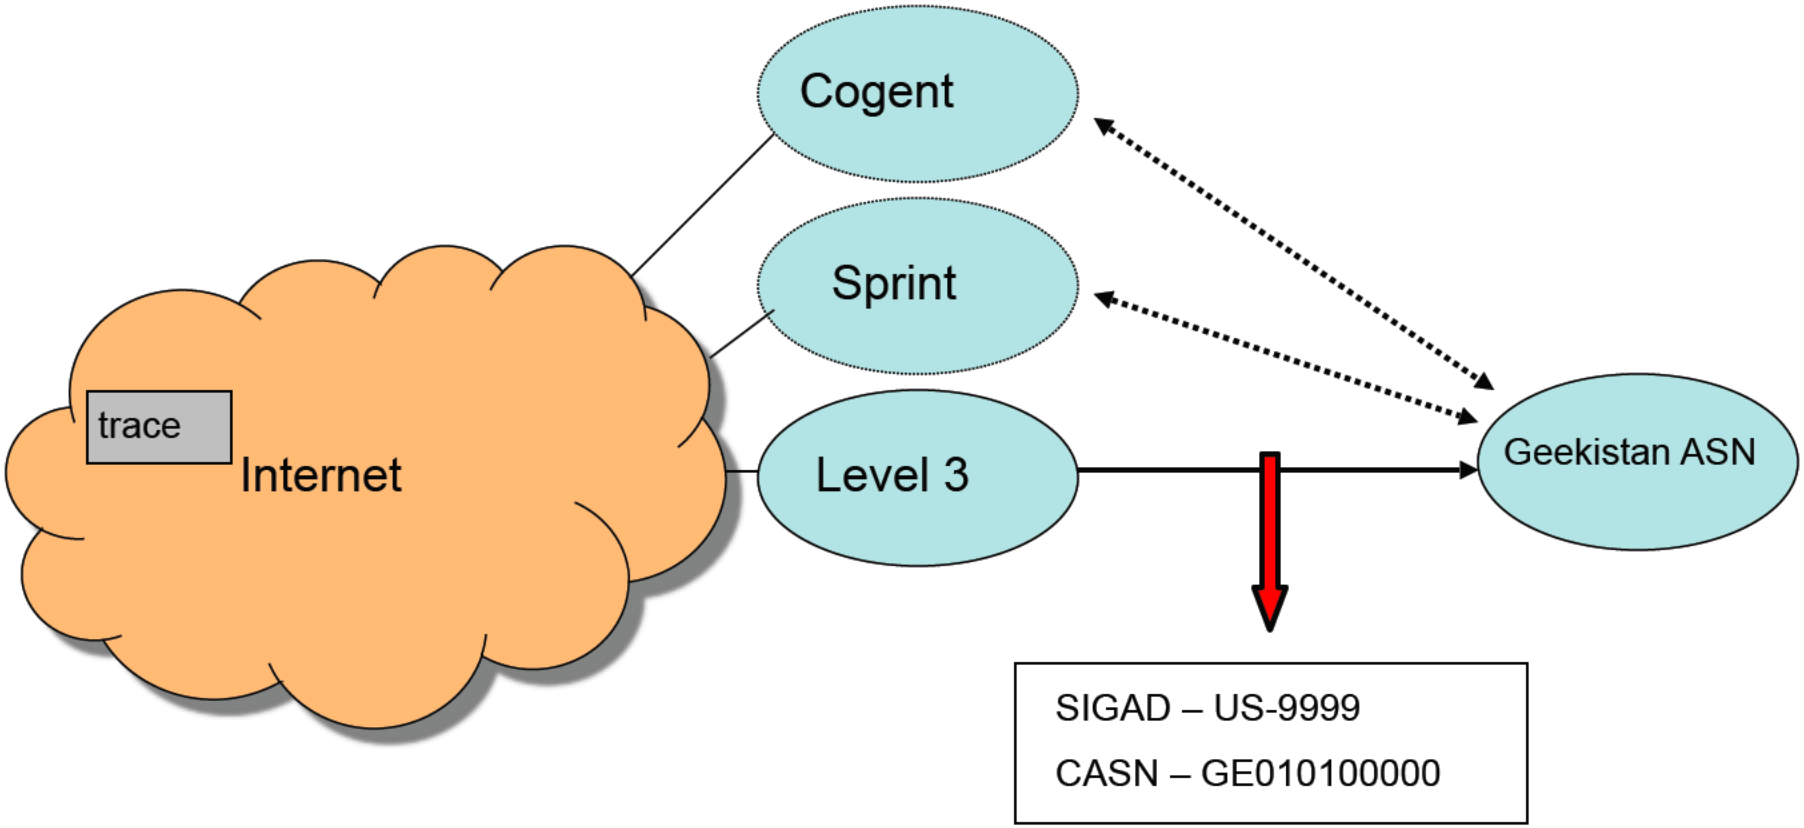

If you are bound and determined to shape traffic into Geekistan, you need to do 1 of 2 things…

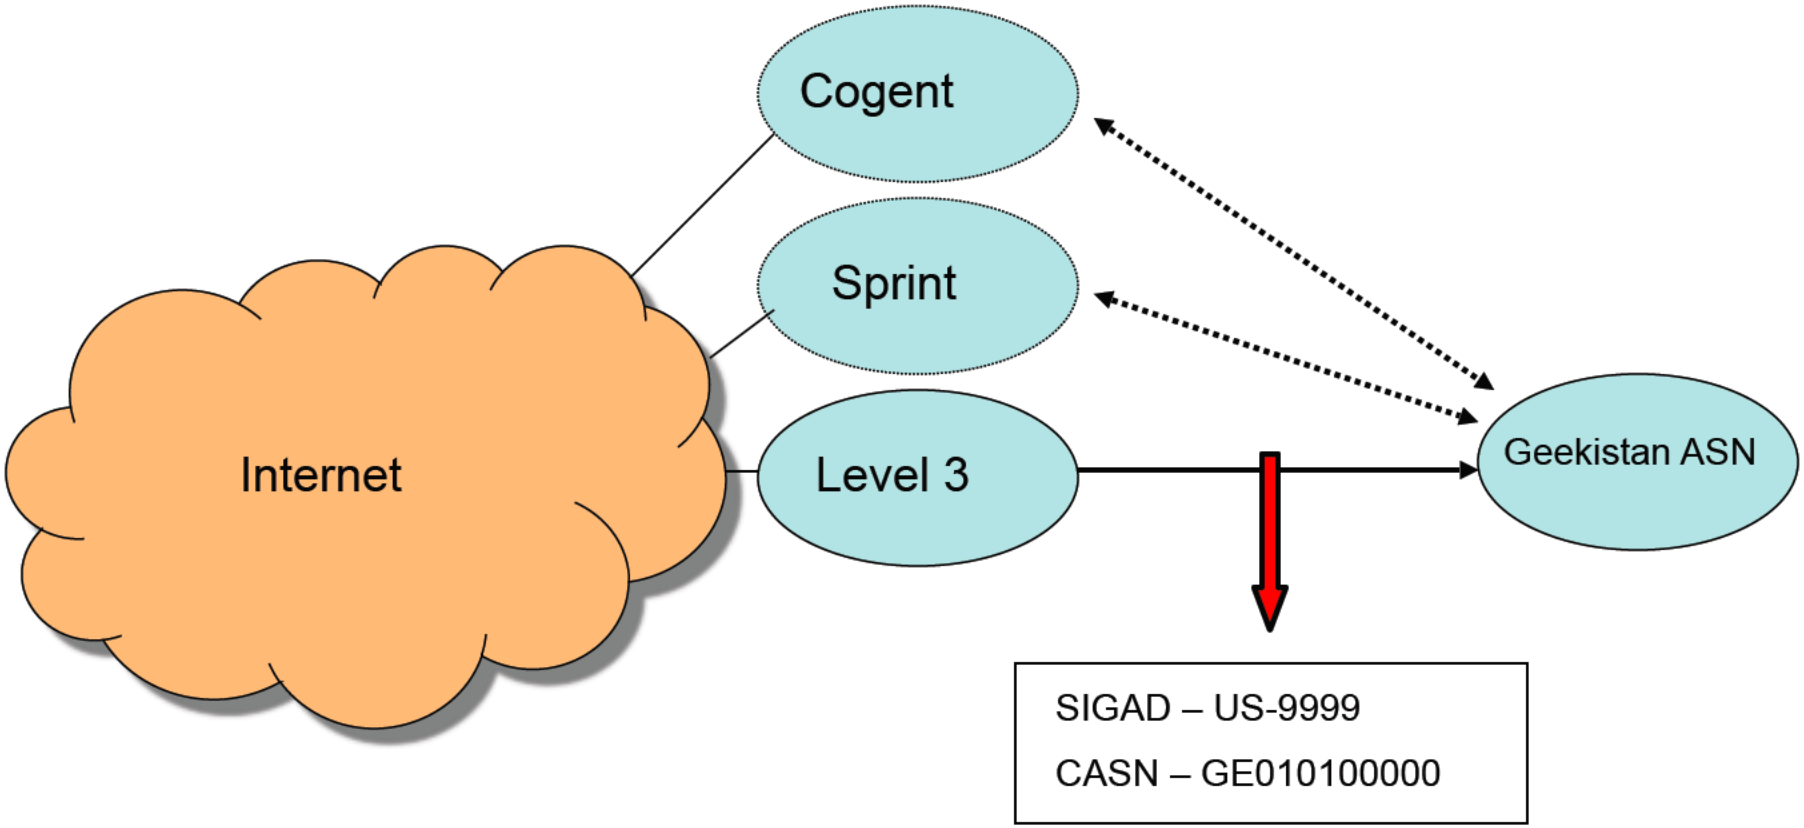

1 – You'll need to find passive collection on both the Cogent and Sprint links with Geekistan. If you do, then you can have confidence in collecting exfil as you shape it into that network.

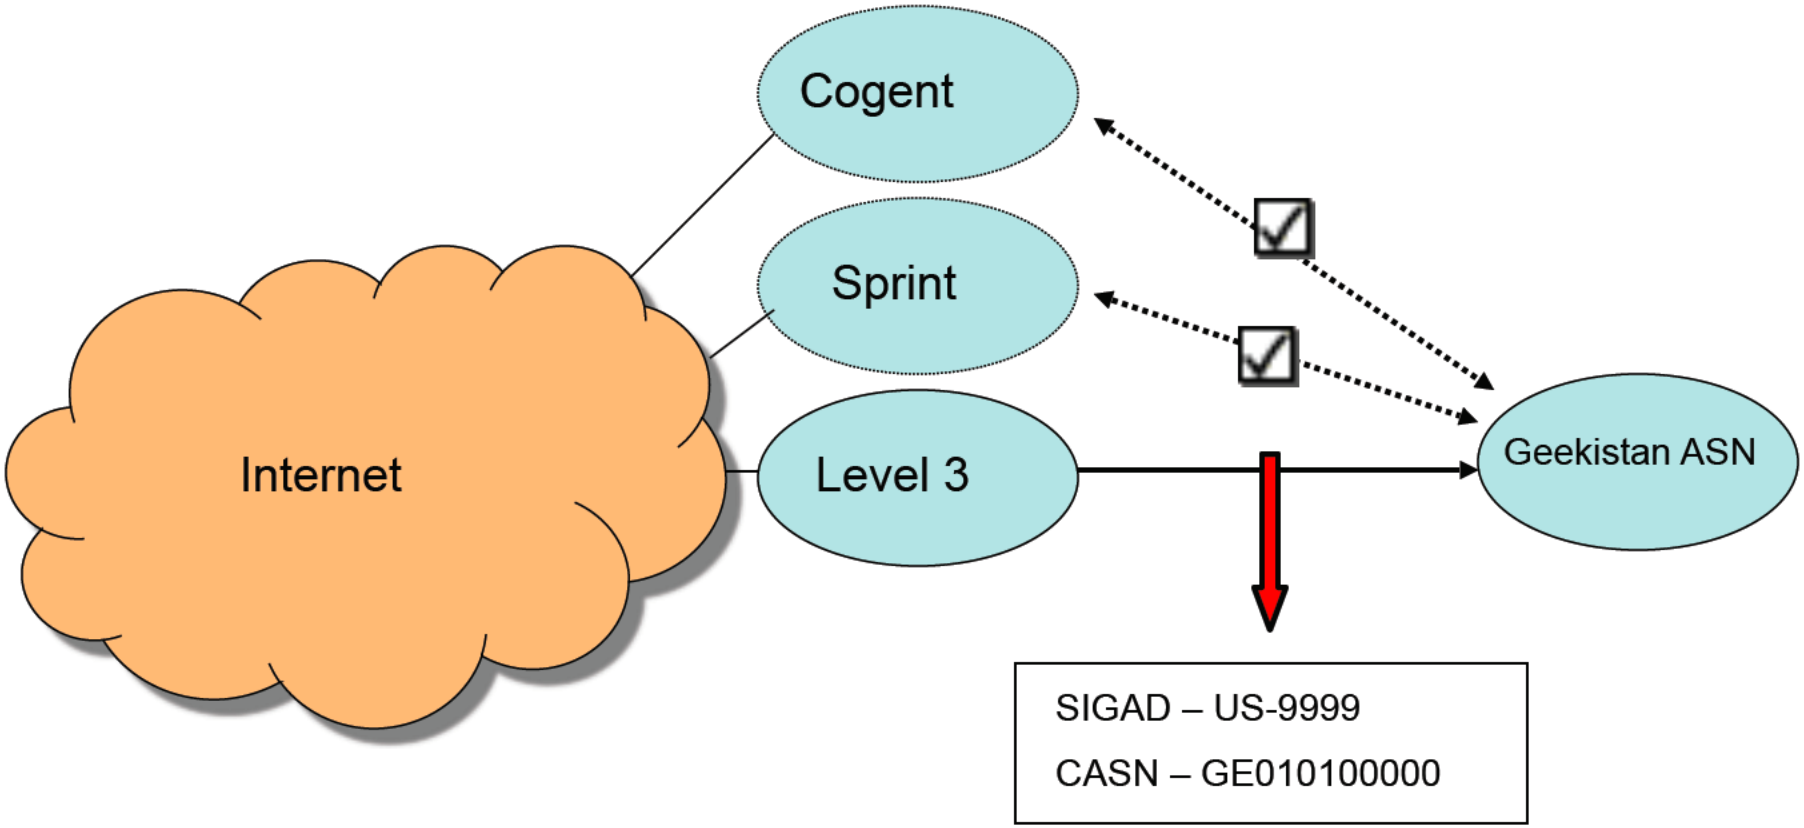

2 – You'll need to identify the Level 3 IP address that is on the Geekistan side of the connection, and send exfil directly to that IP address…

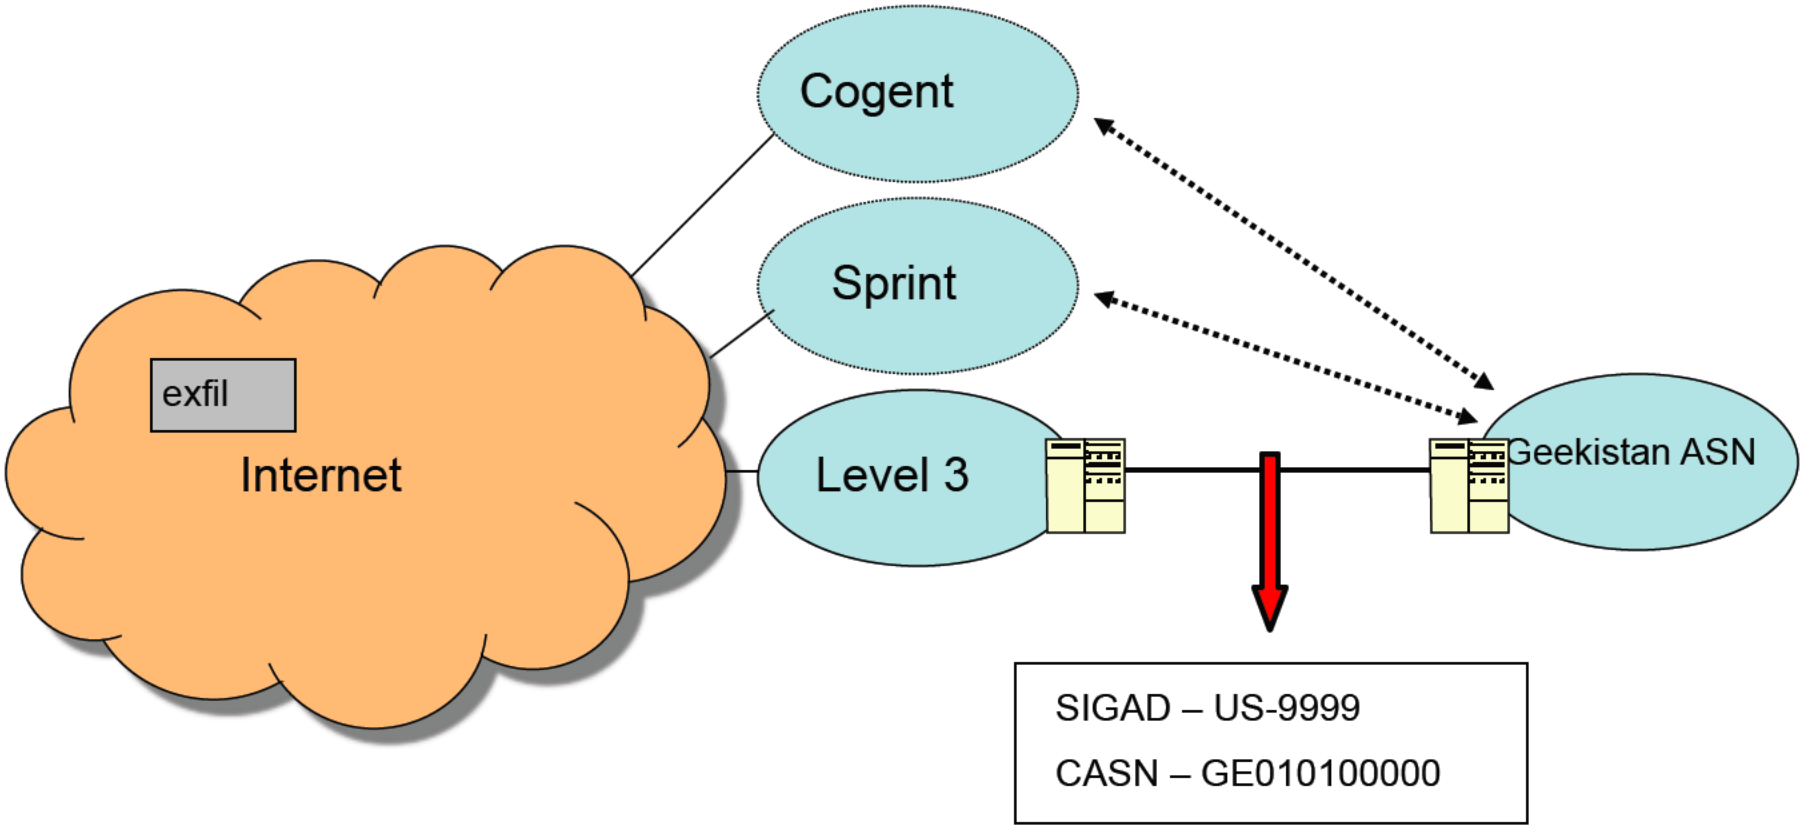

If you do your research, and realize that Level 3 is the *only* connection Geekistan has to the rest of the Internet, and you are collecting that link, then you're safer to assume your exfil will get collected if you send it to anywhere in that network…

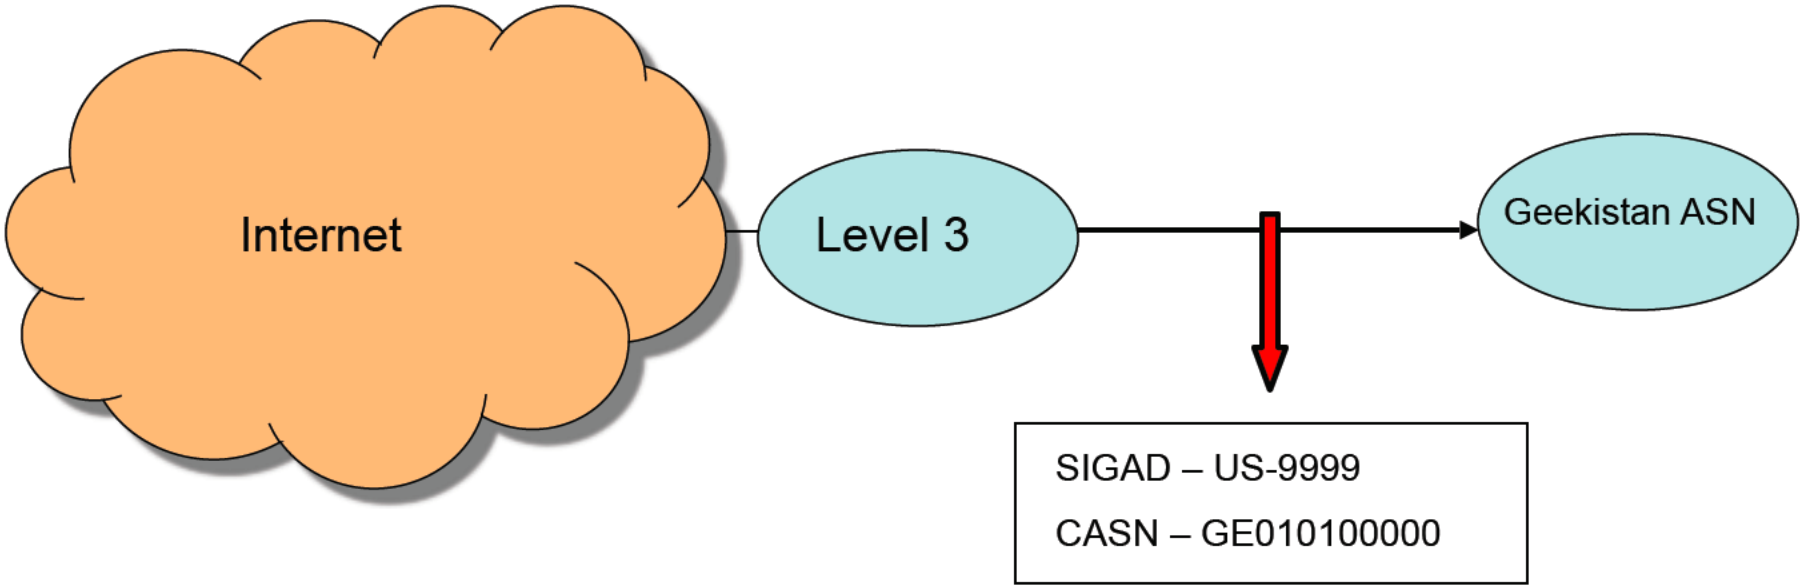

If you do your research, and realize that Level 3 is the *only* connection Geekistan has to the rest of the Internet, and you are collecting that link, then you're safer to assume your exfil will get collected if you send it to anywhere in that network…

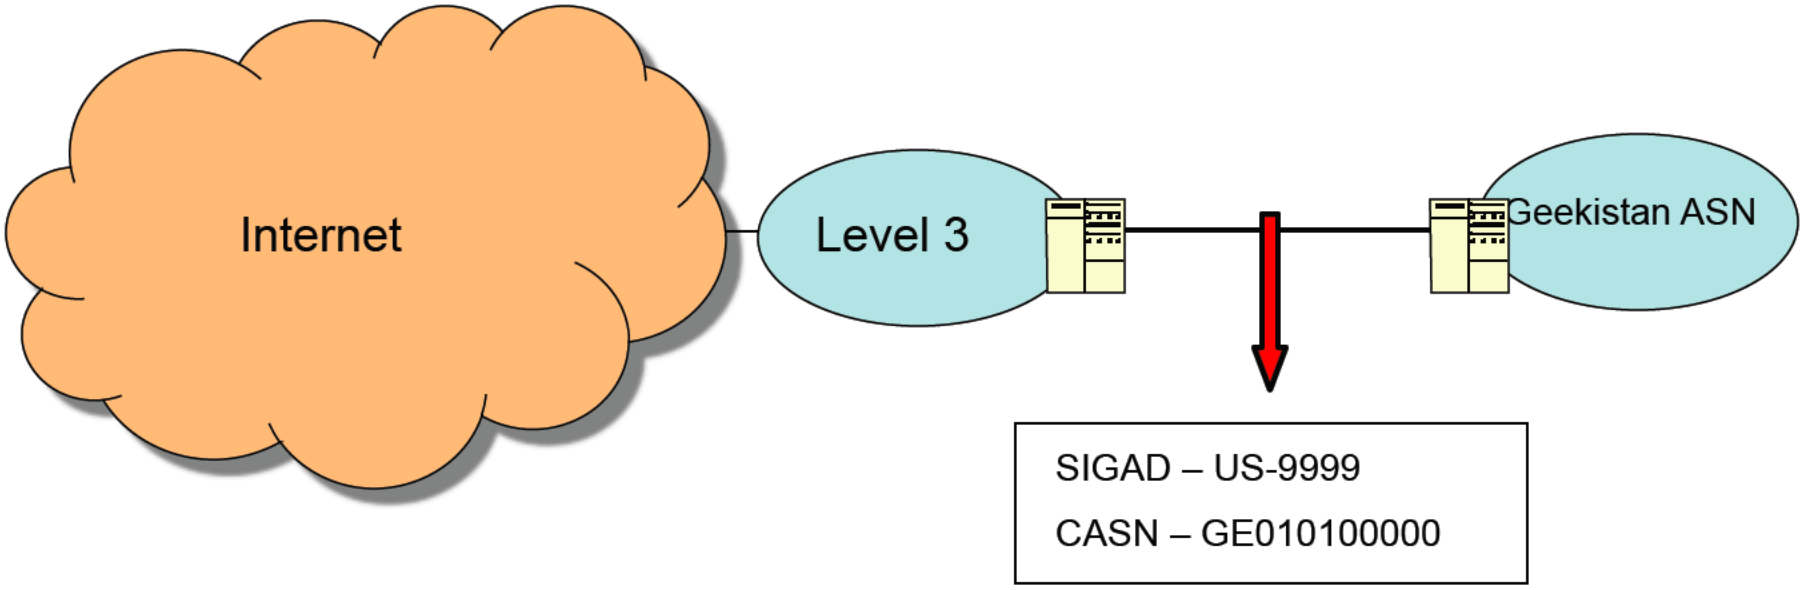

# Now to the nittier-grittier…

- This following section could also be renamed the “I’m pulling my hair out in the fetal position while screaming ‘Why didn’t it work?!’” section.
- The previous slides described how shaping should work at a theoretical level, following are a few reasons why it doesn’t always work in the real world.
- The following issues are not all-encompassing of why shaping might not work, just a few examples.

First, the multiple links problem.

Below, the graph looks really pretty with 1 link per upstream provider, but that isn't always the case. Sometimes it's like this:

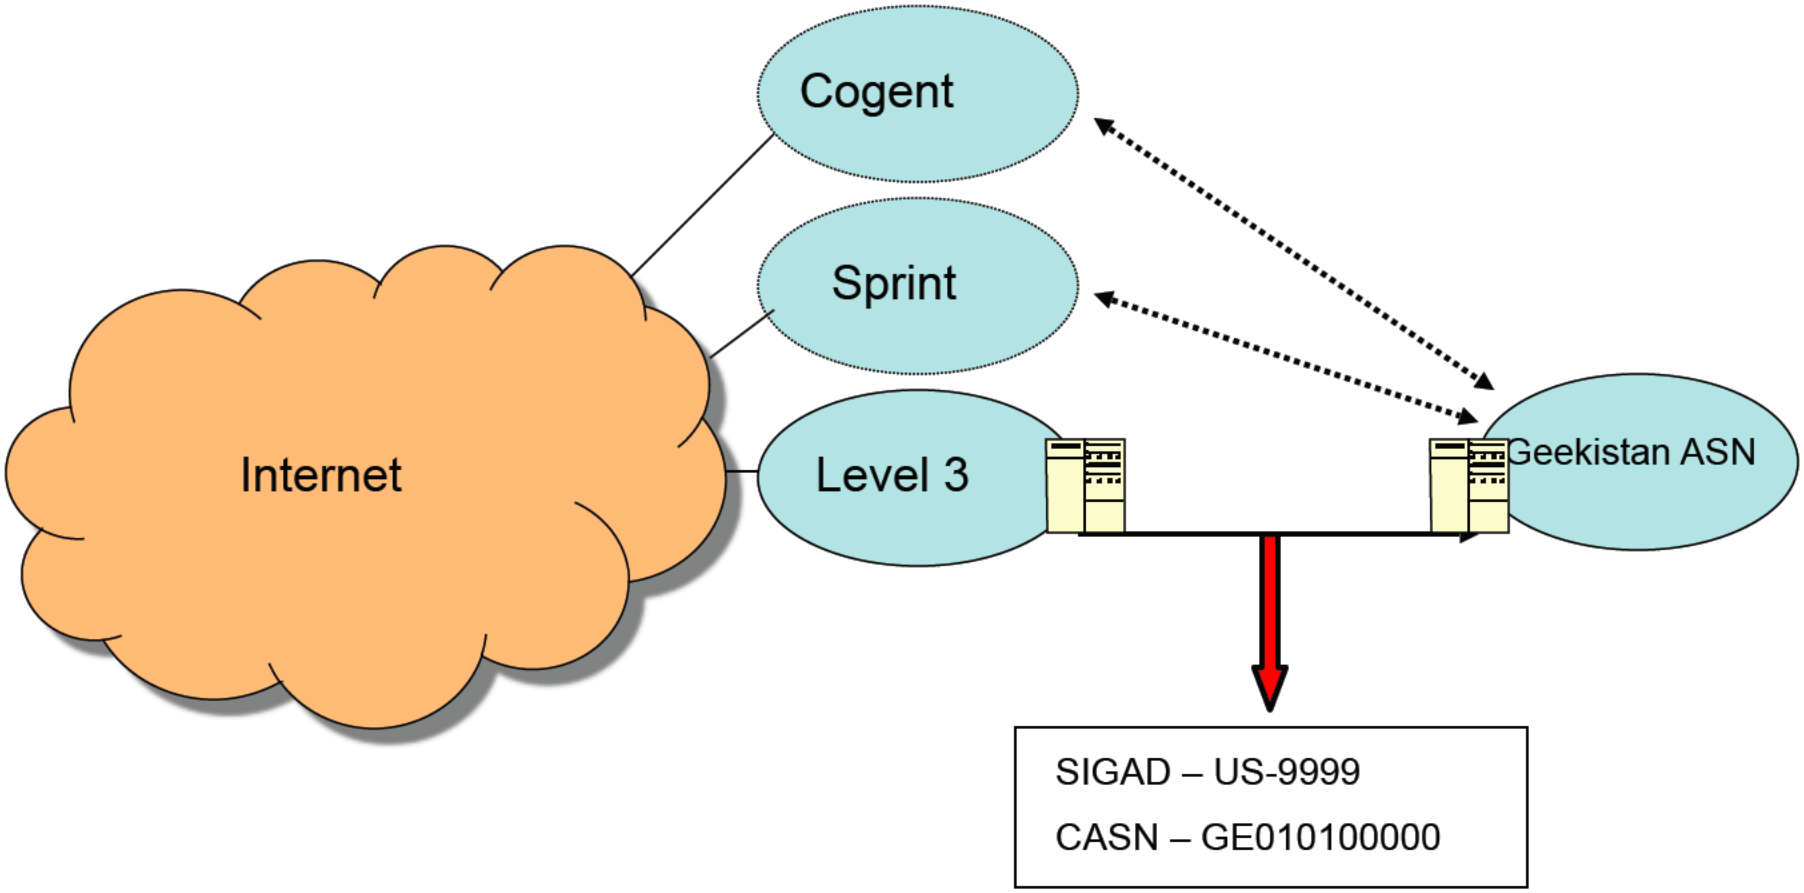

Notice how there are 2 links between Geekistan ASN and Level 3. It is not abnormal for ASNs to have multiple links to eachother for redundancy reasons.

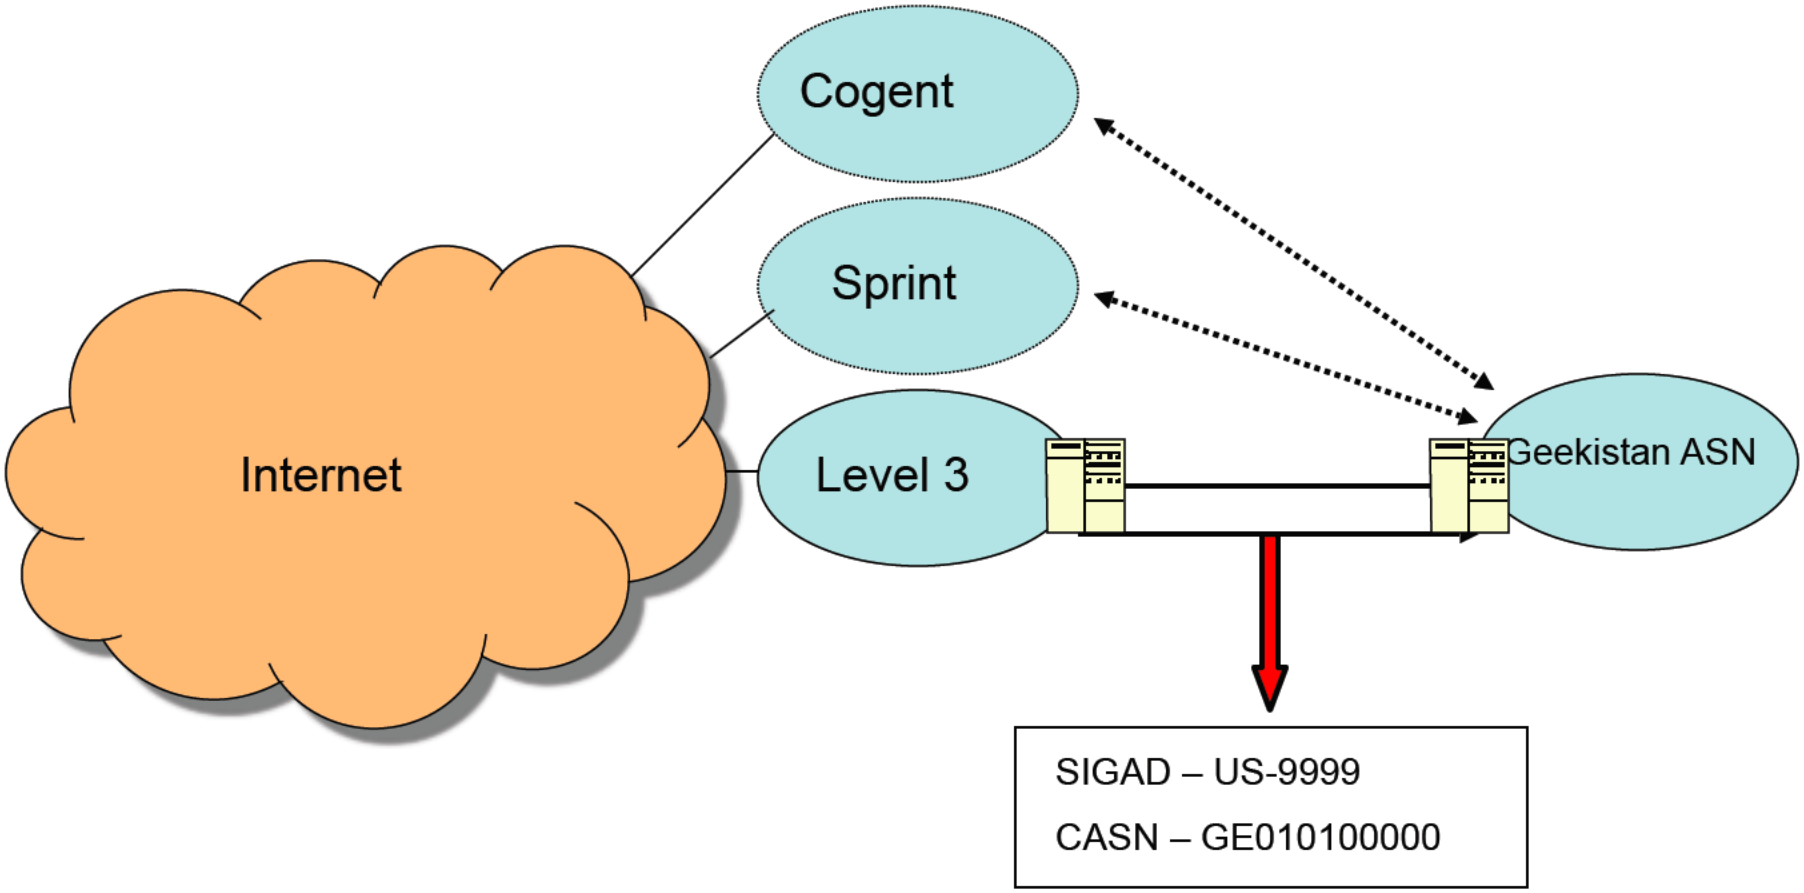

Why this matters, is if we are only able to collect 1 of the 2 links, even if we try to force our exfil into/out of the Level 3-Geekistan ASN link, we can't guarantee it will traverse the link we collect.

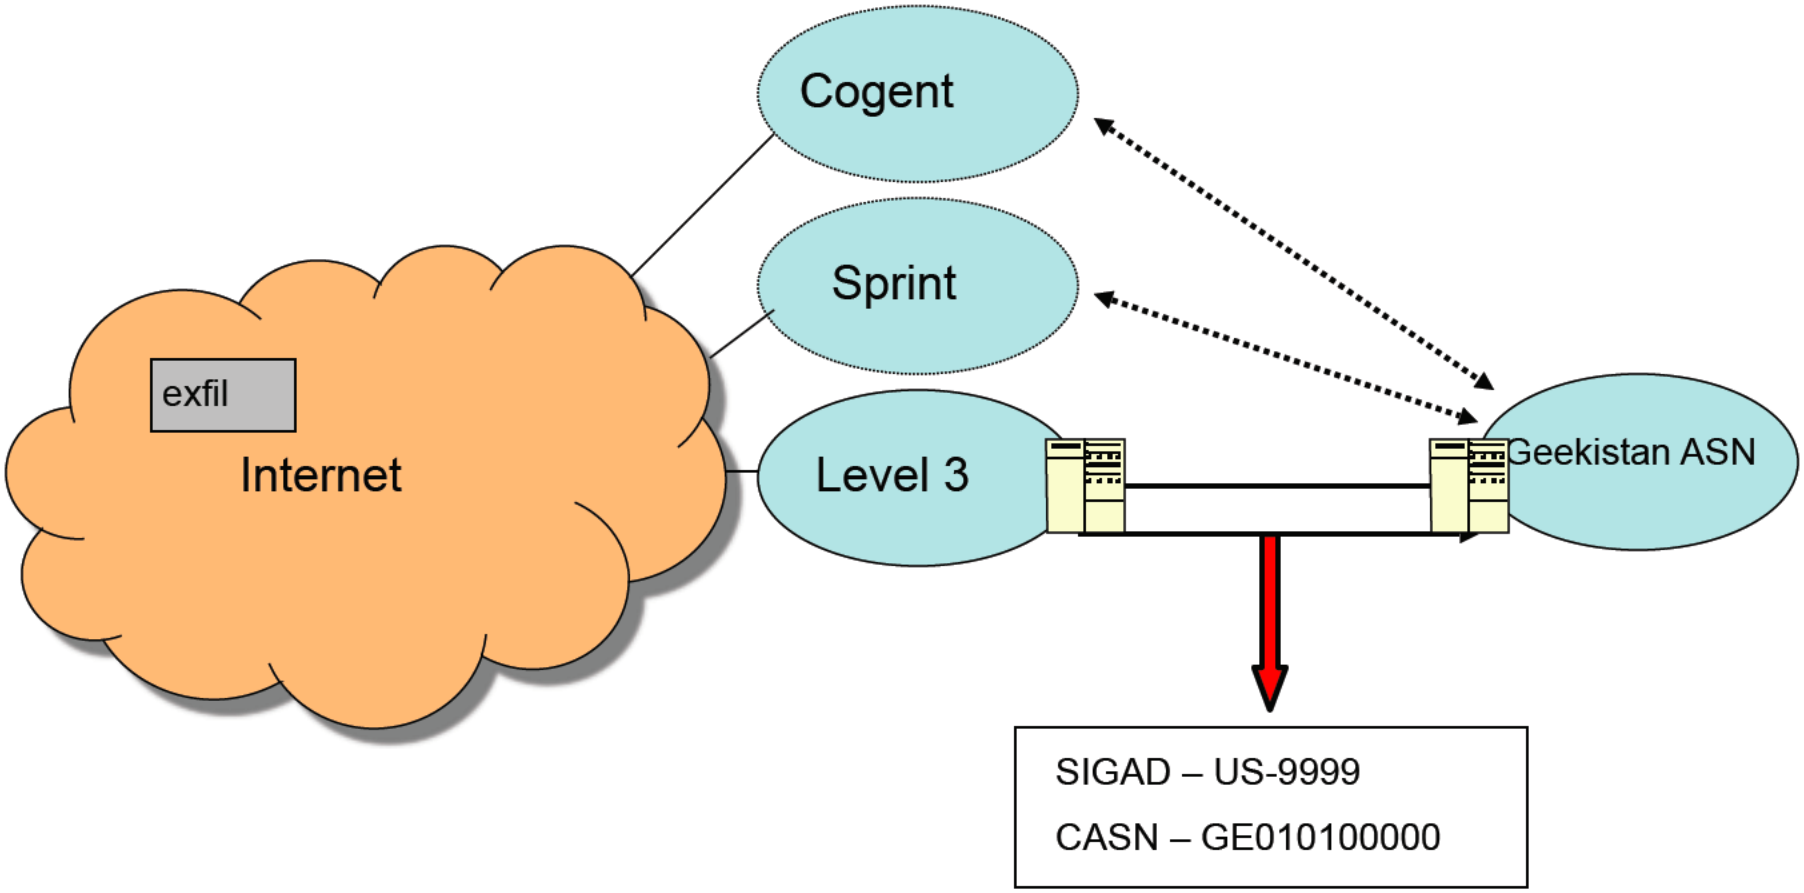

Even if we try to shoot directly at the Level 3 interface that we *can* collect on, there's a chance it could go over the *other* Level 3 link to get there…

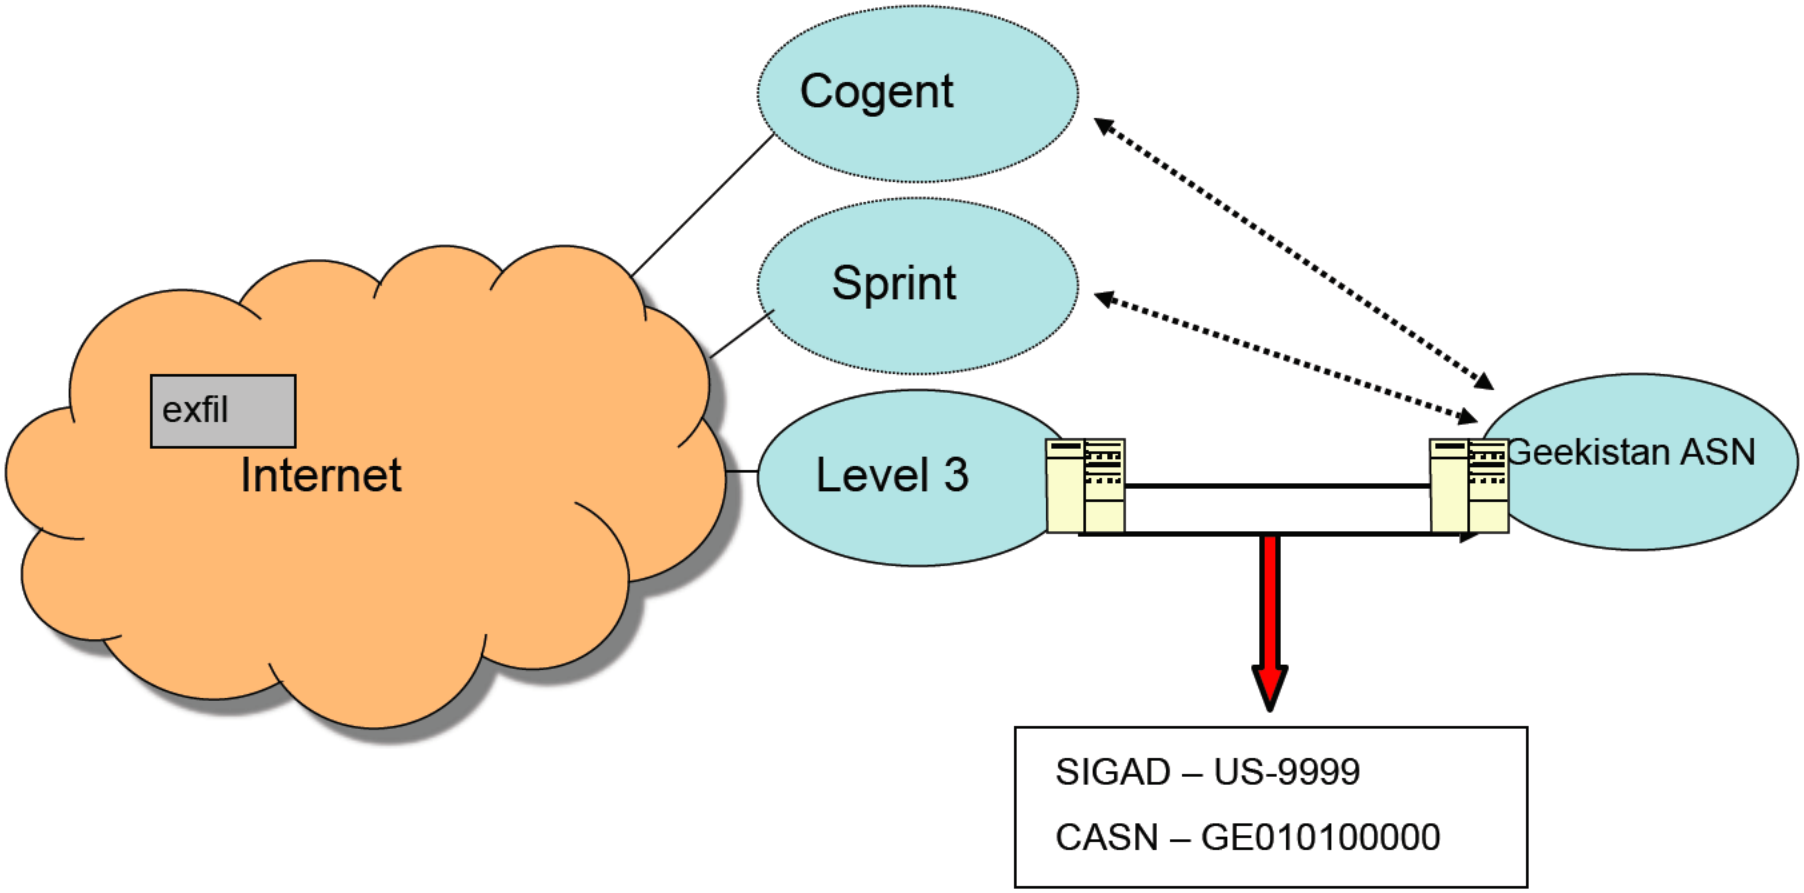

Next, there is the “back-up” link problem. In the ISP world, bandwidth is money, and some people charge more for letting your data traverse their networks…

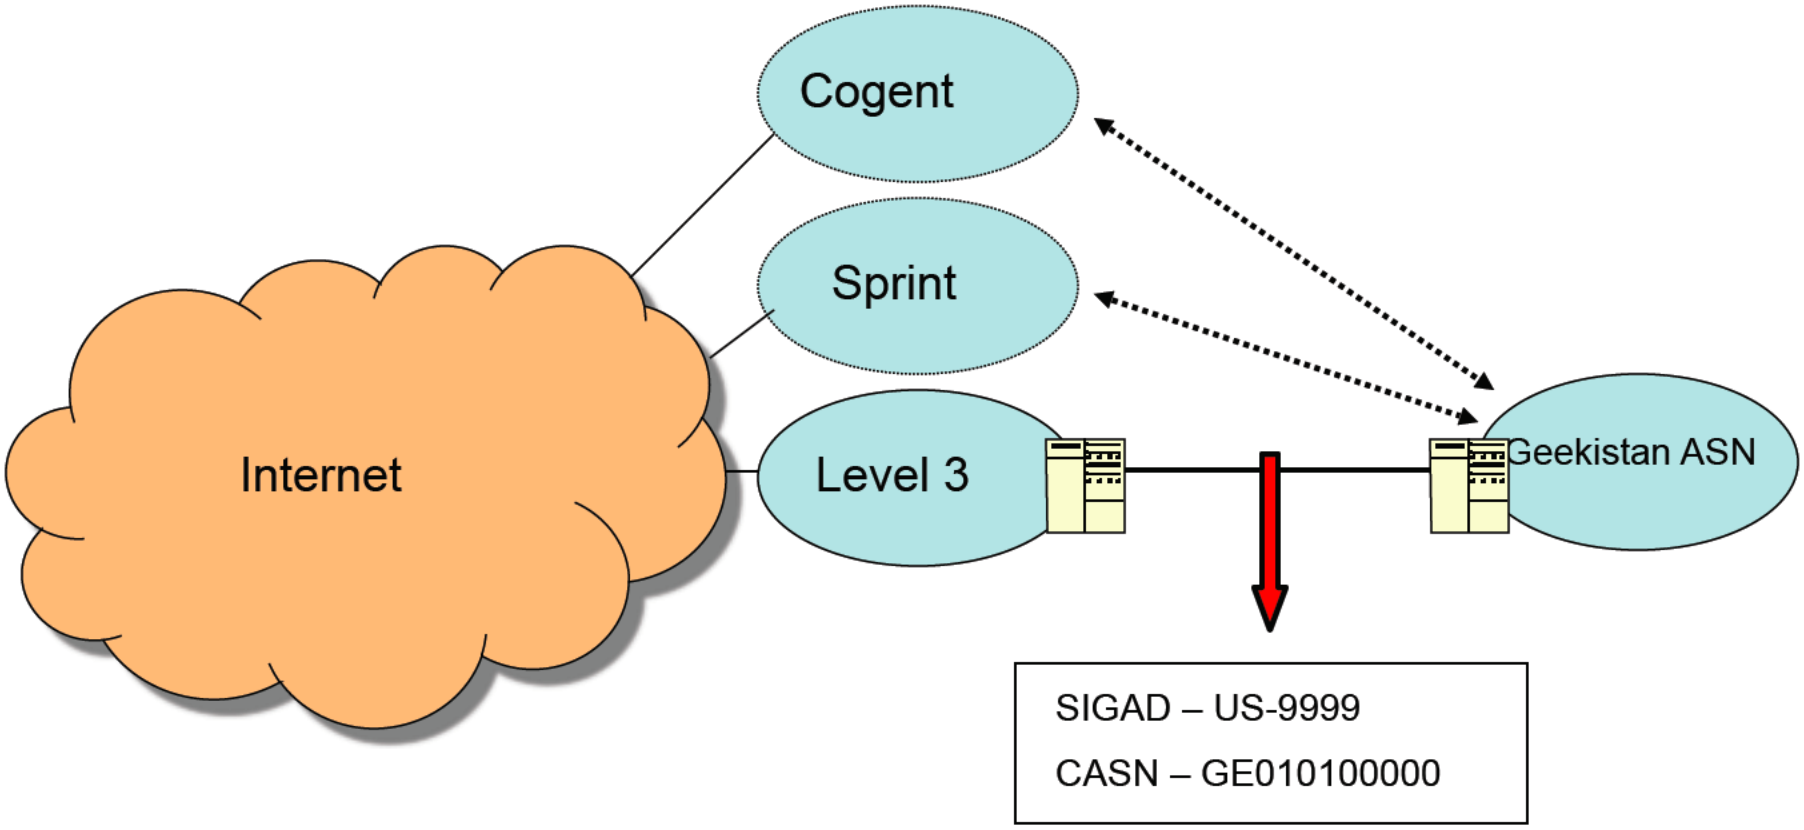

Say, for example, it was *twice* as expensive for Geekistan to send data through Level 3 as it is for Cogent or Sprint. They would use Level 3 as little as possible (as in, only when necessary), while using the other 2 as much as they can.

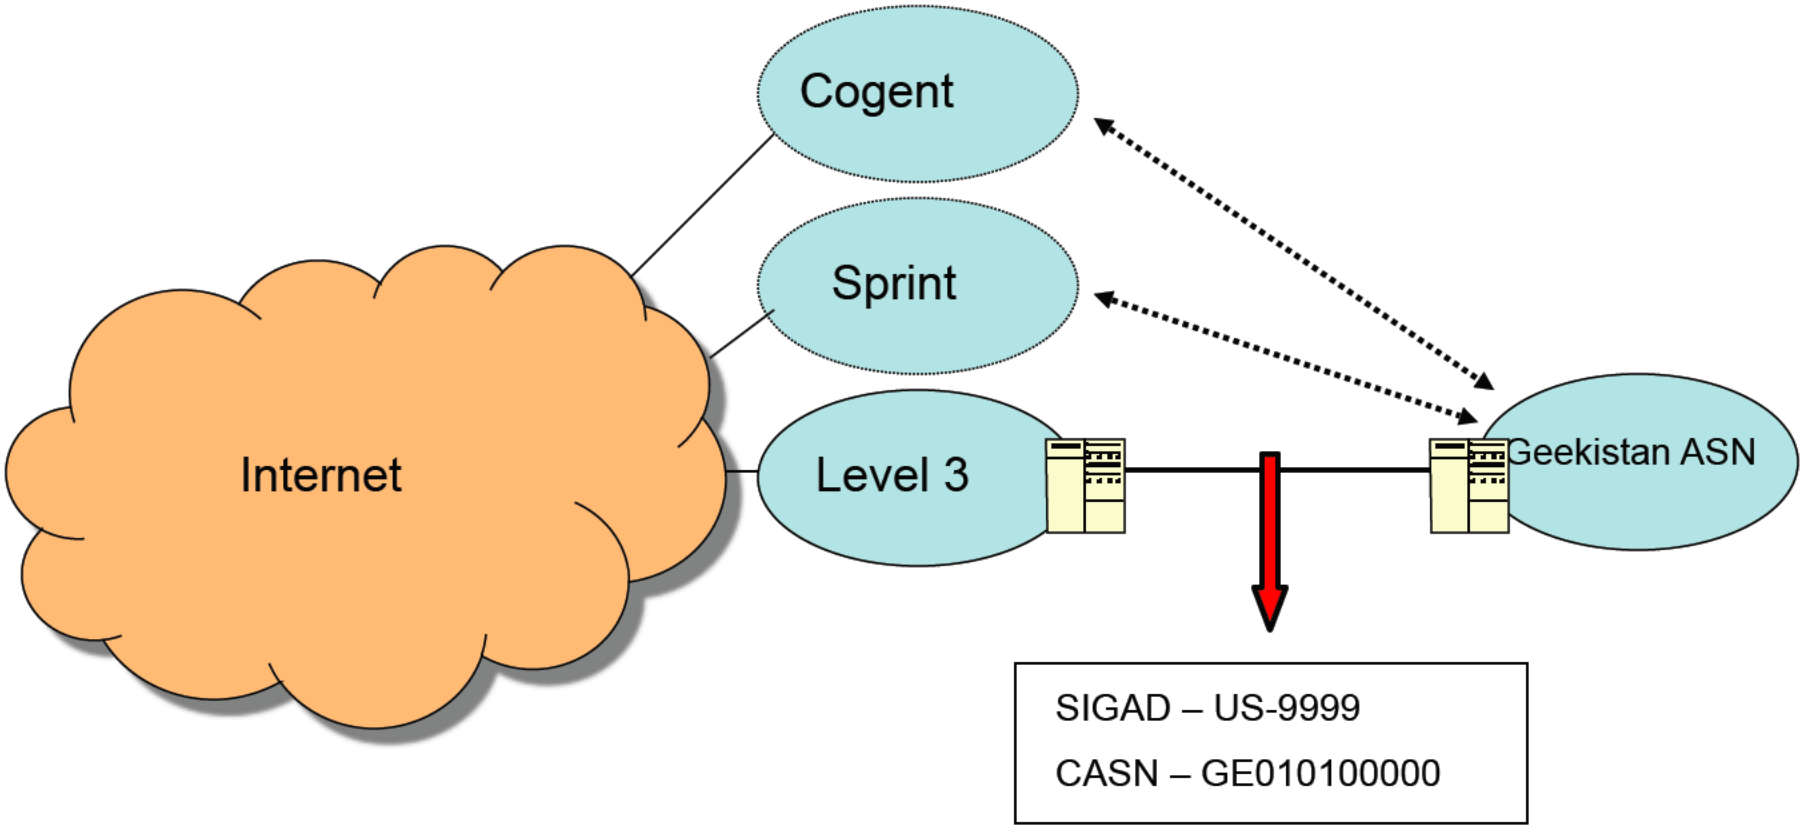

If they try to force traffic in/out Cogent and Sprint, that could spell bad news for our shaping efforts, we might see something like this:

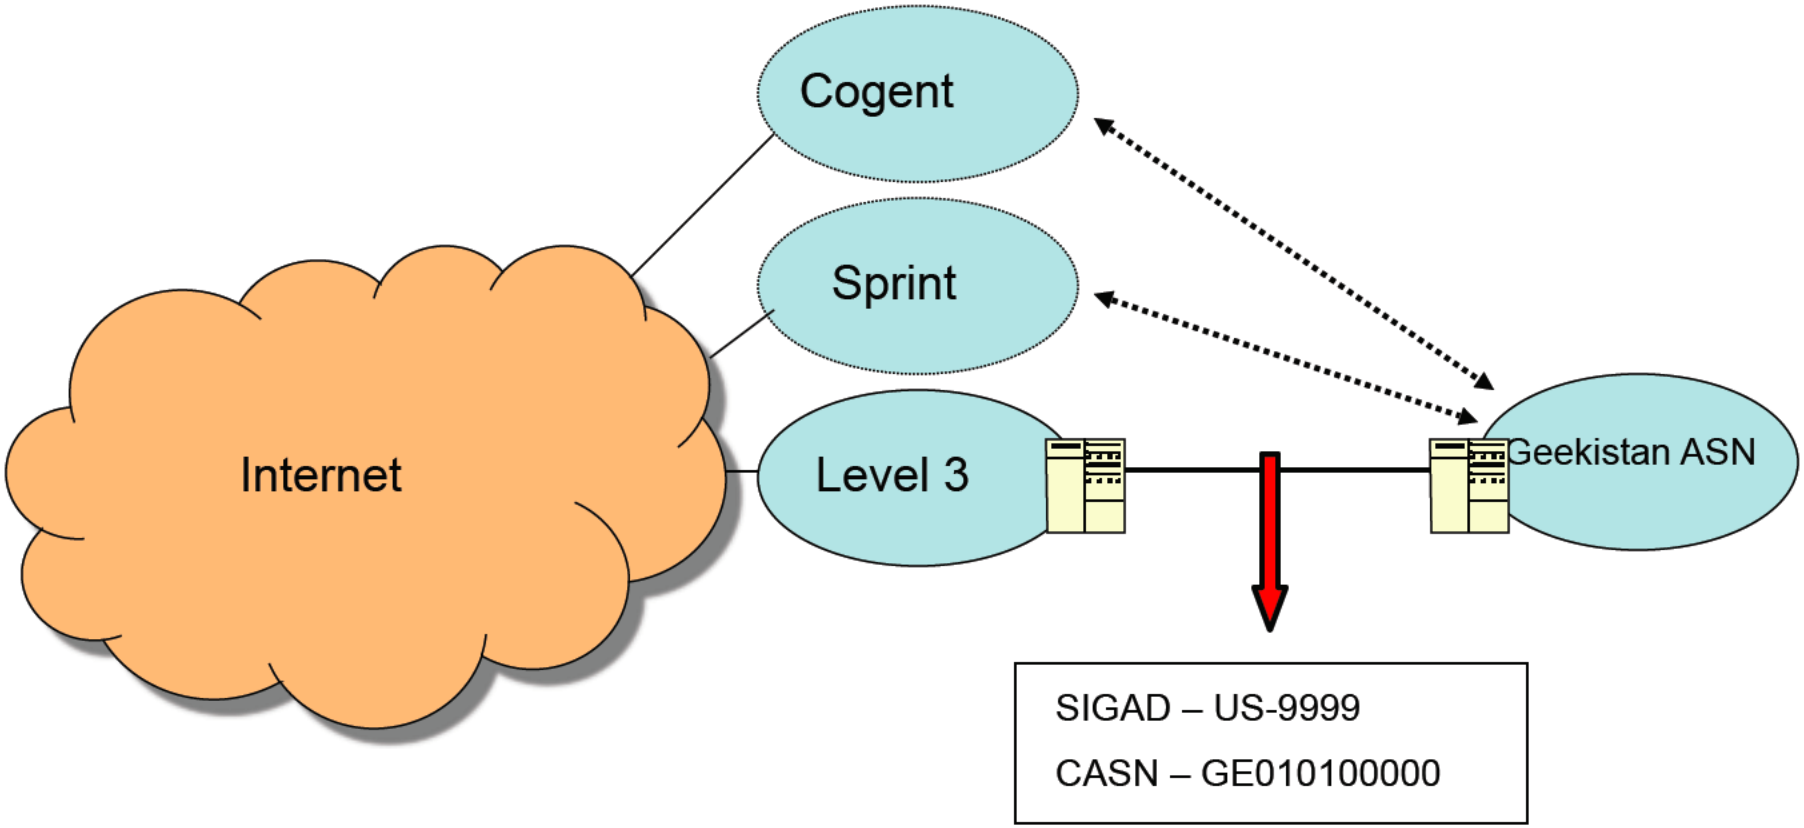

If we try to send data directly over the Level 3 link, we might be sorely disappointed with Geekistan forcing us over a different link…

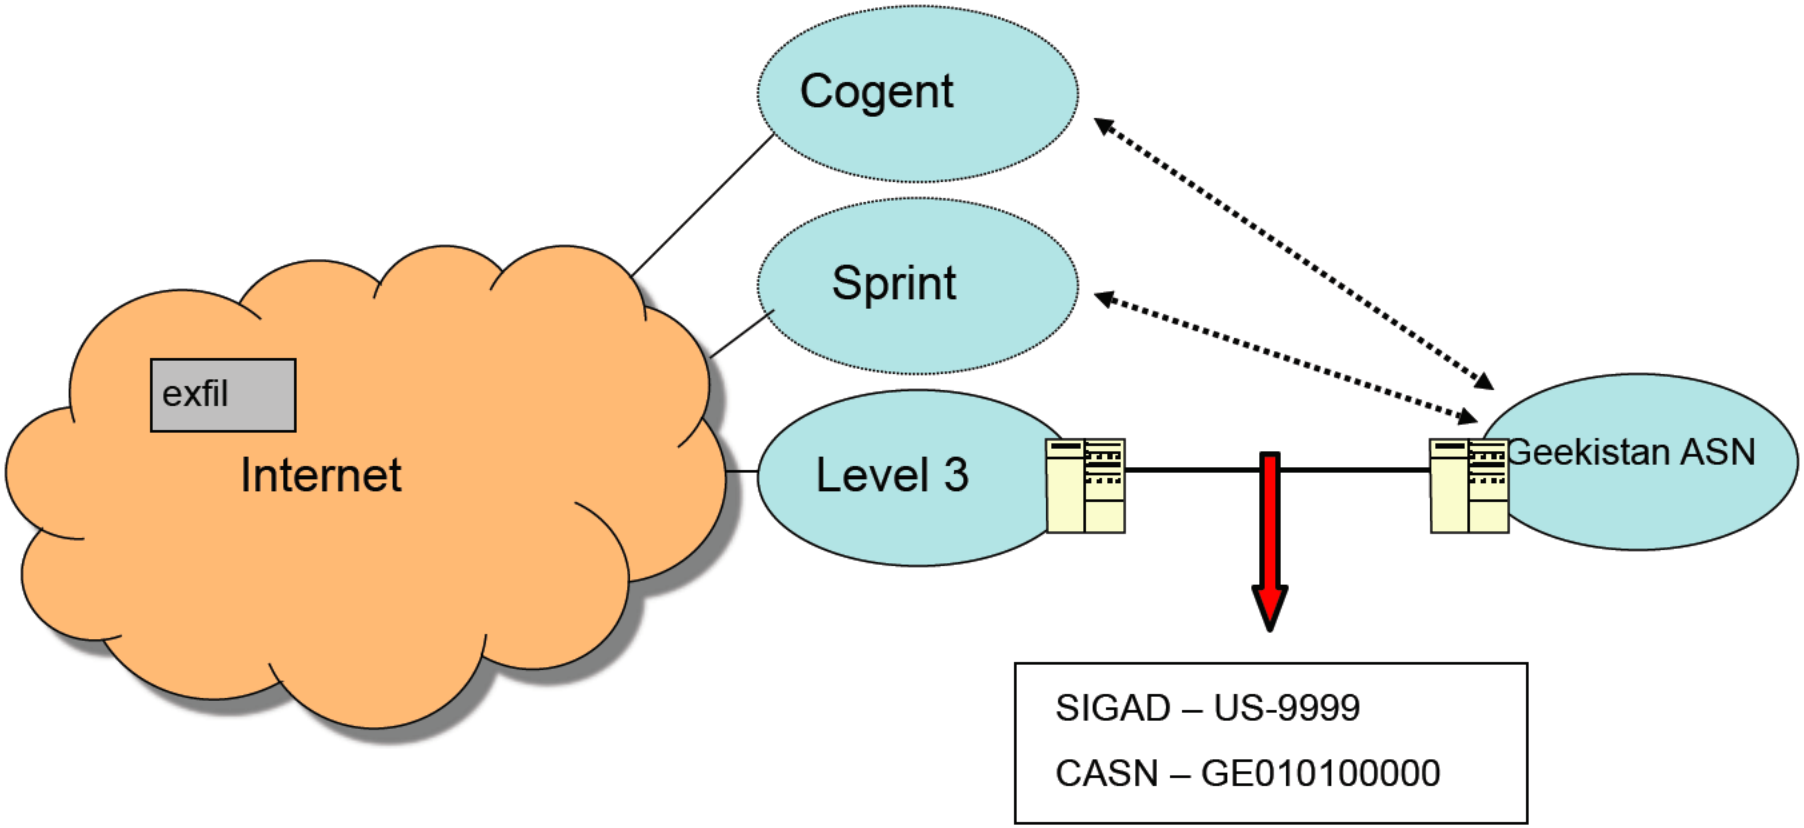

Or it might look something like this if we tried to exfil ***out*** directly into Level 3's network…

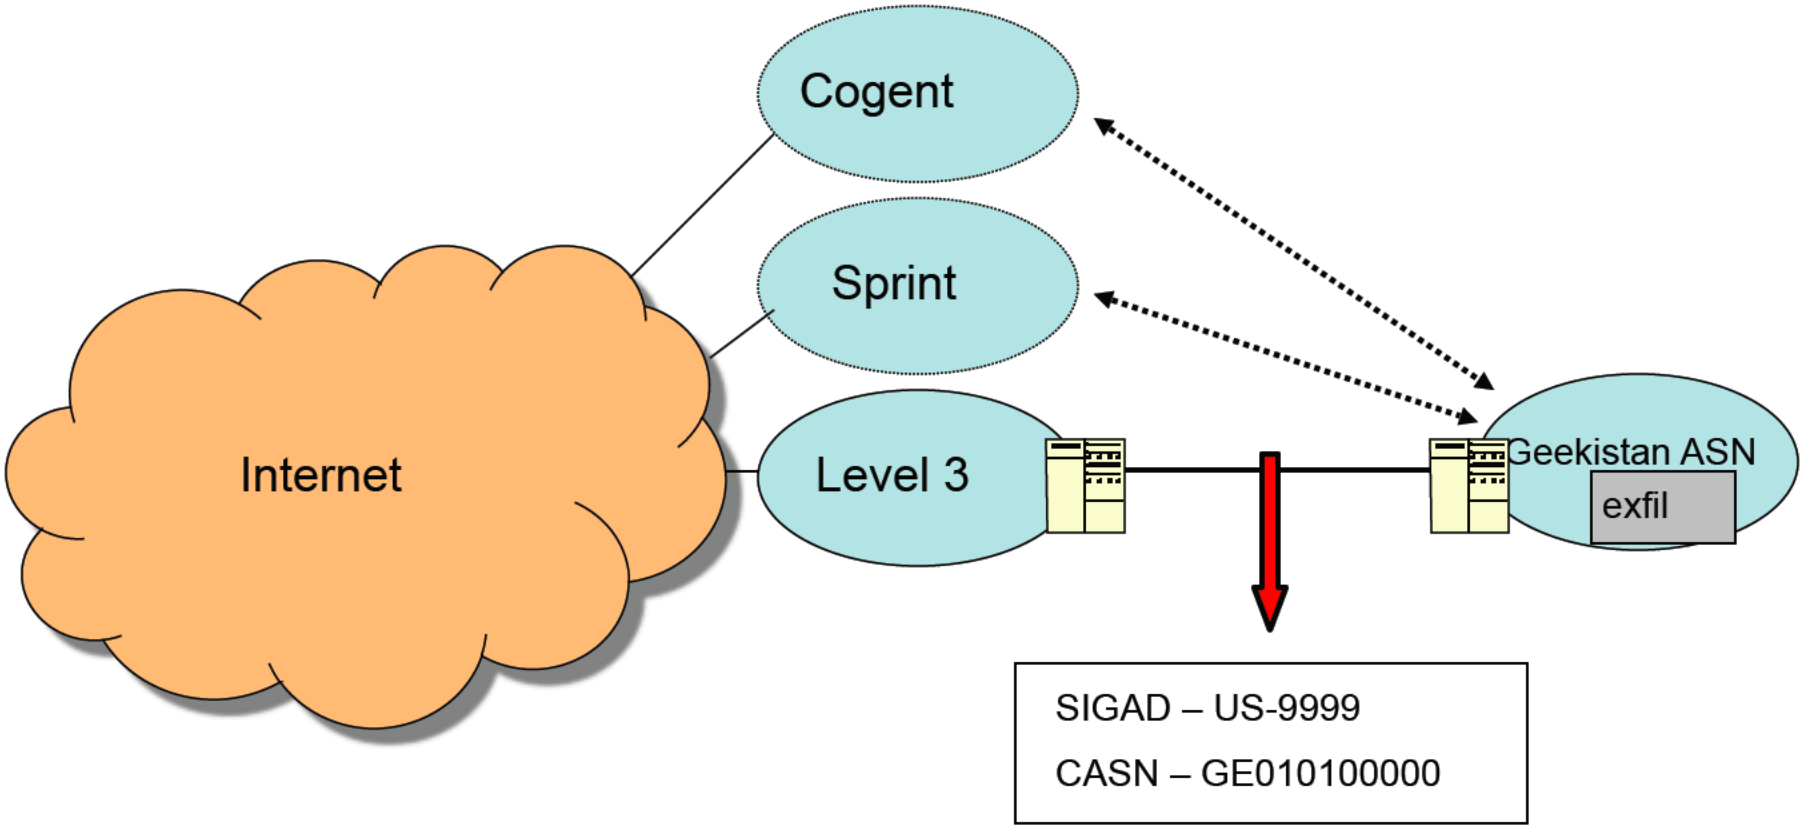

Lastly (for now), is the infamous collection problem…

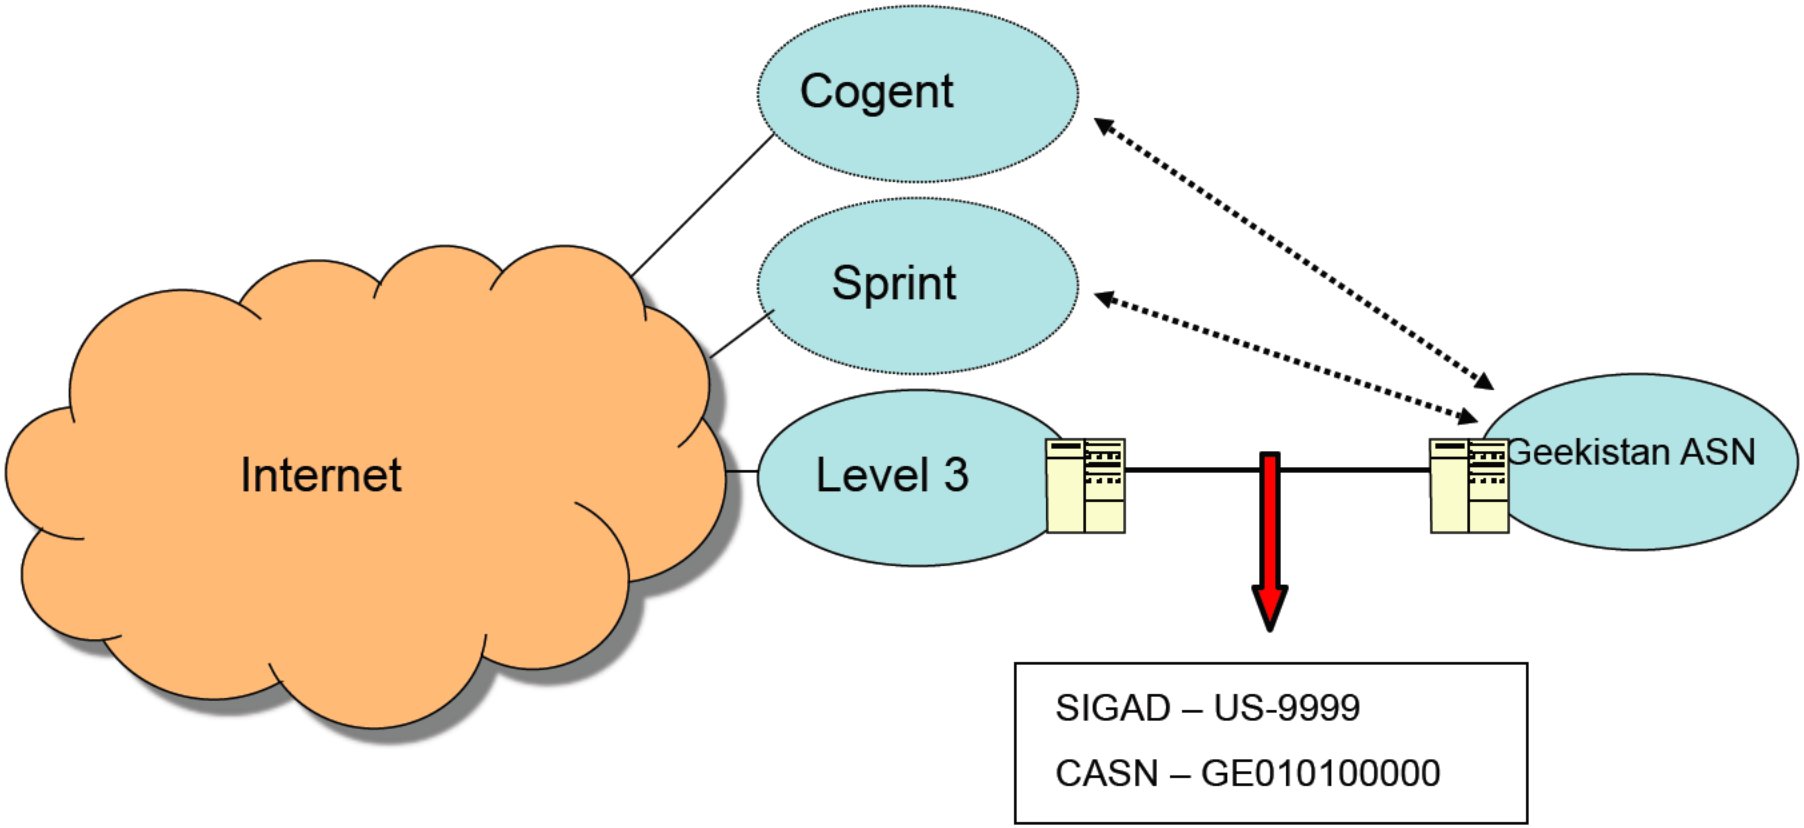

Pretend all the stars aligned…there is only 1 Level 3 link, we *can* send data over it, and we *do* have access to it at SIGAD US-9999…

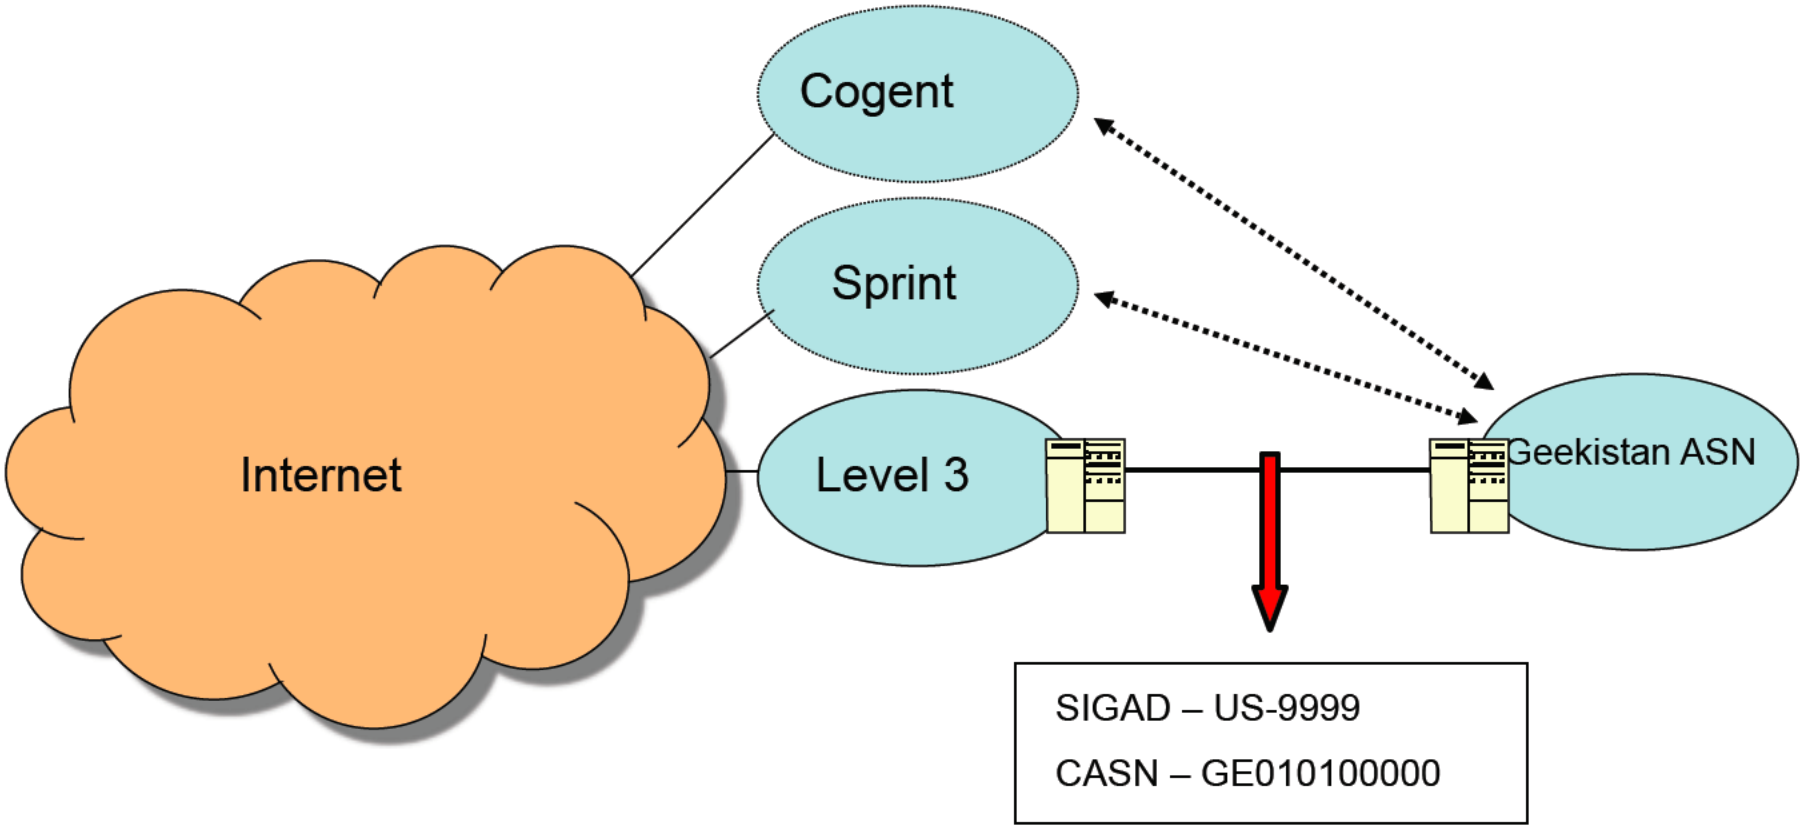

In a fit of joy we start sending exfil over that link. Then we go look at US-9999, GE010100000 for our exfil, and find that ***it isn't there***!

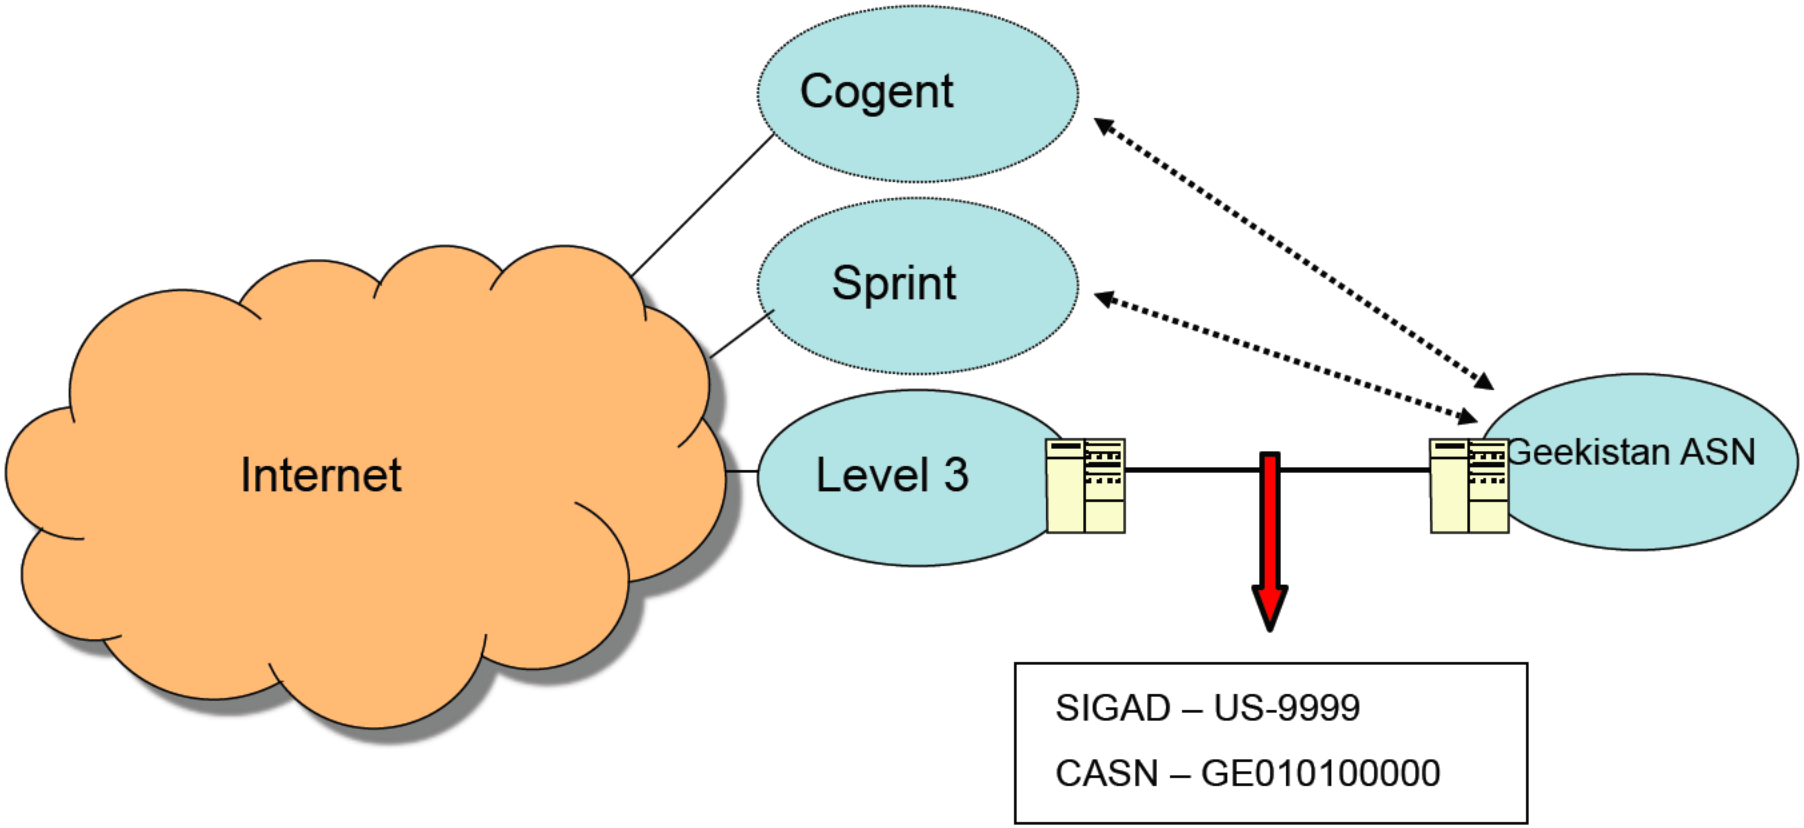

Without going into the gory details, US-9999 may be dropping your exfil (as in, ‘able to see it, but not collecting and processing it’). If you think this is the case, please coordinate with SSO to make sure appropriate IP’s/protocols are promoted at the site.

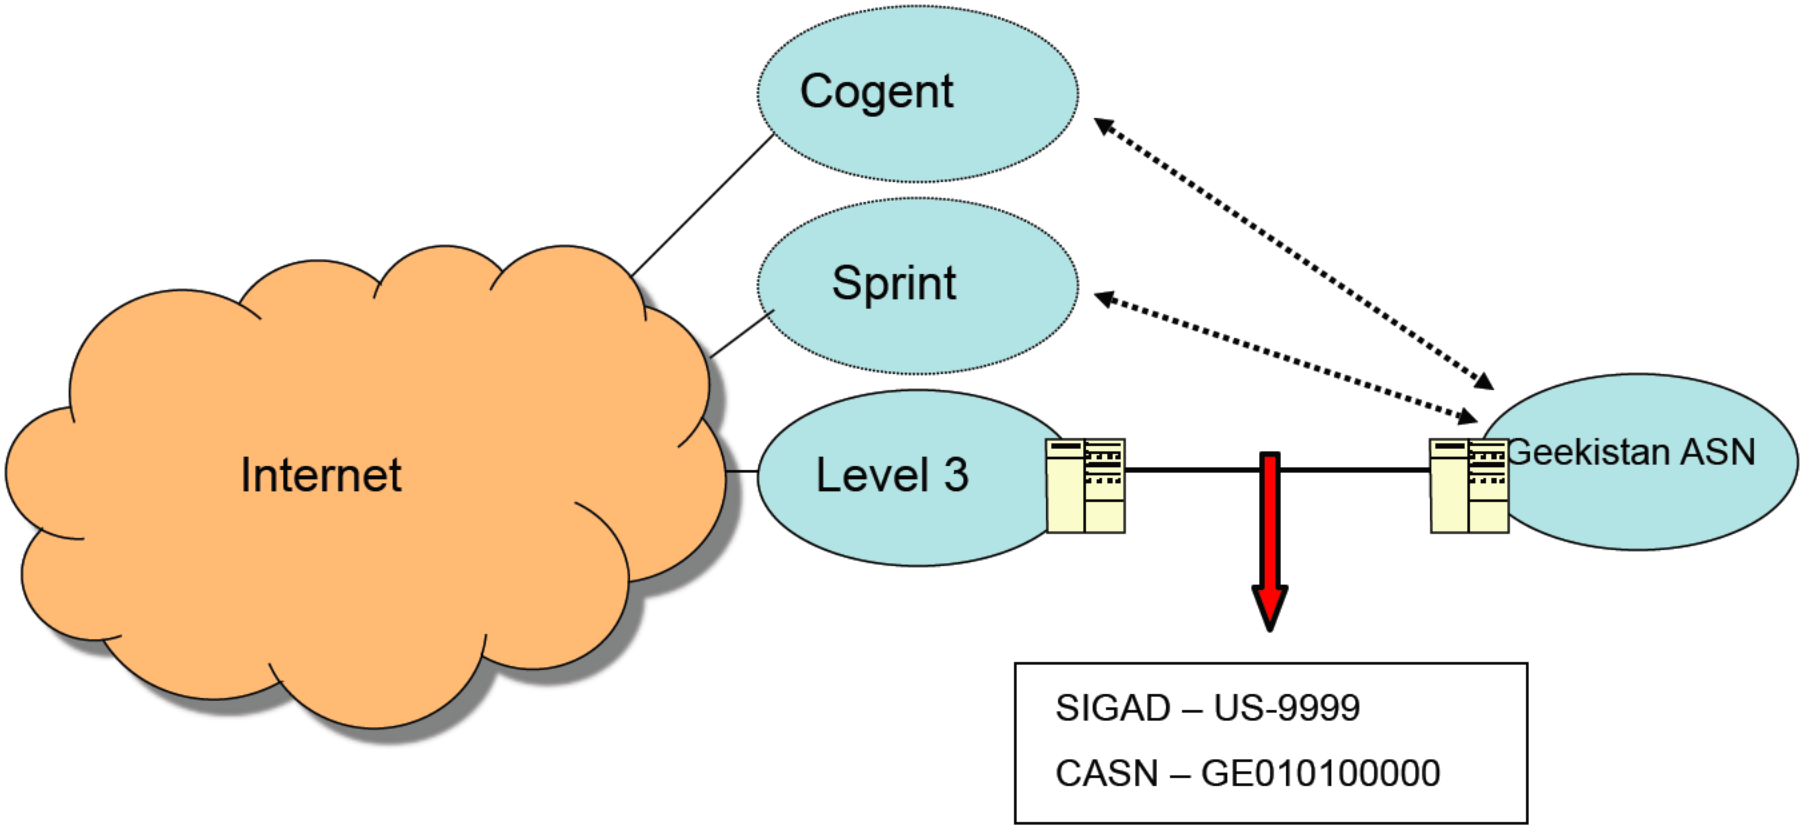

# Contact deets…

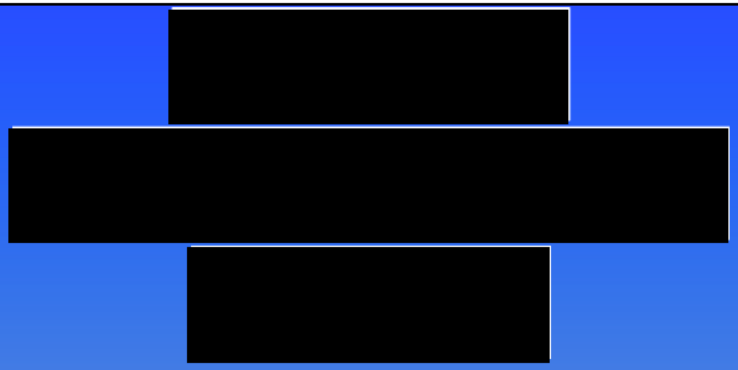


National Security Archive,

Suite 701, Gelman Library, The George Washington University,

2130 H Street, NW, Washington, D.C., 20037,

Phone: 202/994-7000, Fax: 202/994-7005, nsarchiv@gwu.edu