Embracing SLA's

How Customer Requirements Drive Innovation

Most common TLD Anycast SLA's

DNS – Zone Propagation Times

- Time between first NOTIFY and cloud updated

DNS – Availability

- Percentage of queries responded to
- Percentage of servers and/or sites up in the cloud

DNS – RTT

- Average regional/global response time

WWW.CIRA.CA

Zone Propagation

Zone Propagation – Old Mindset

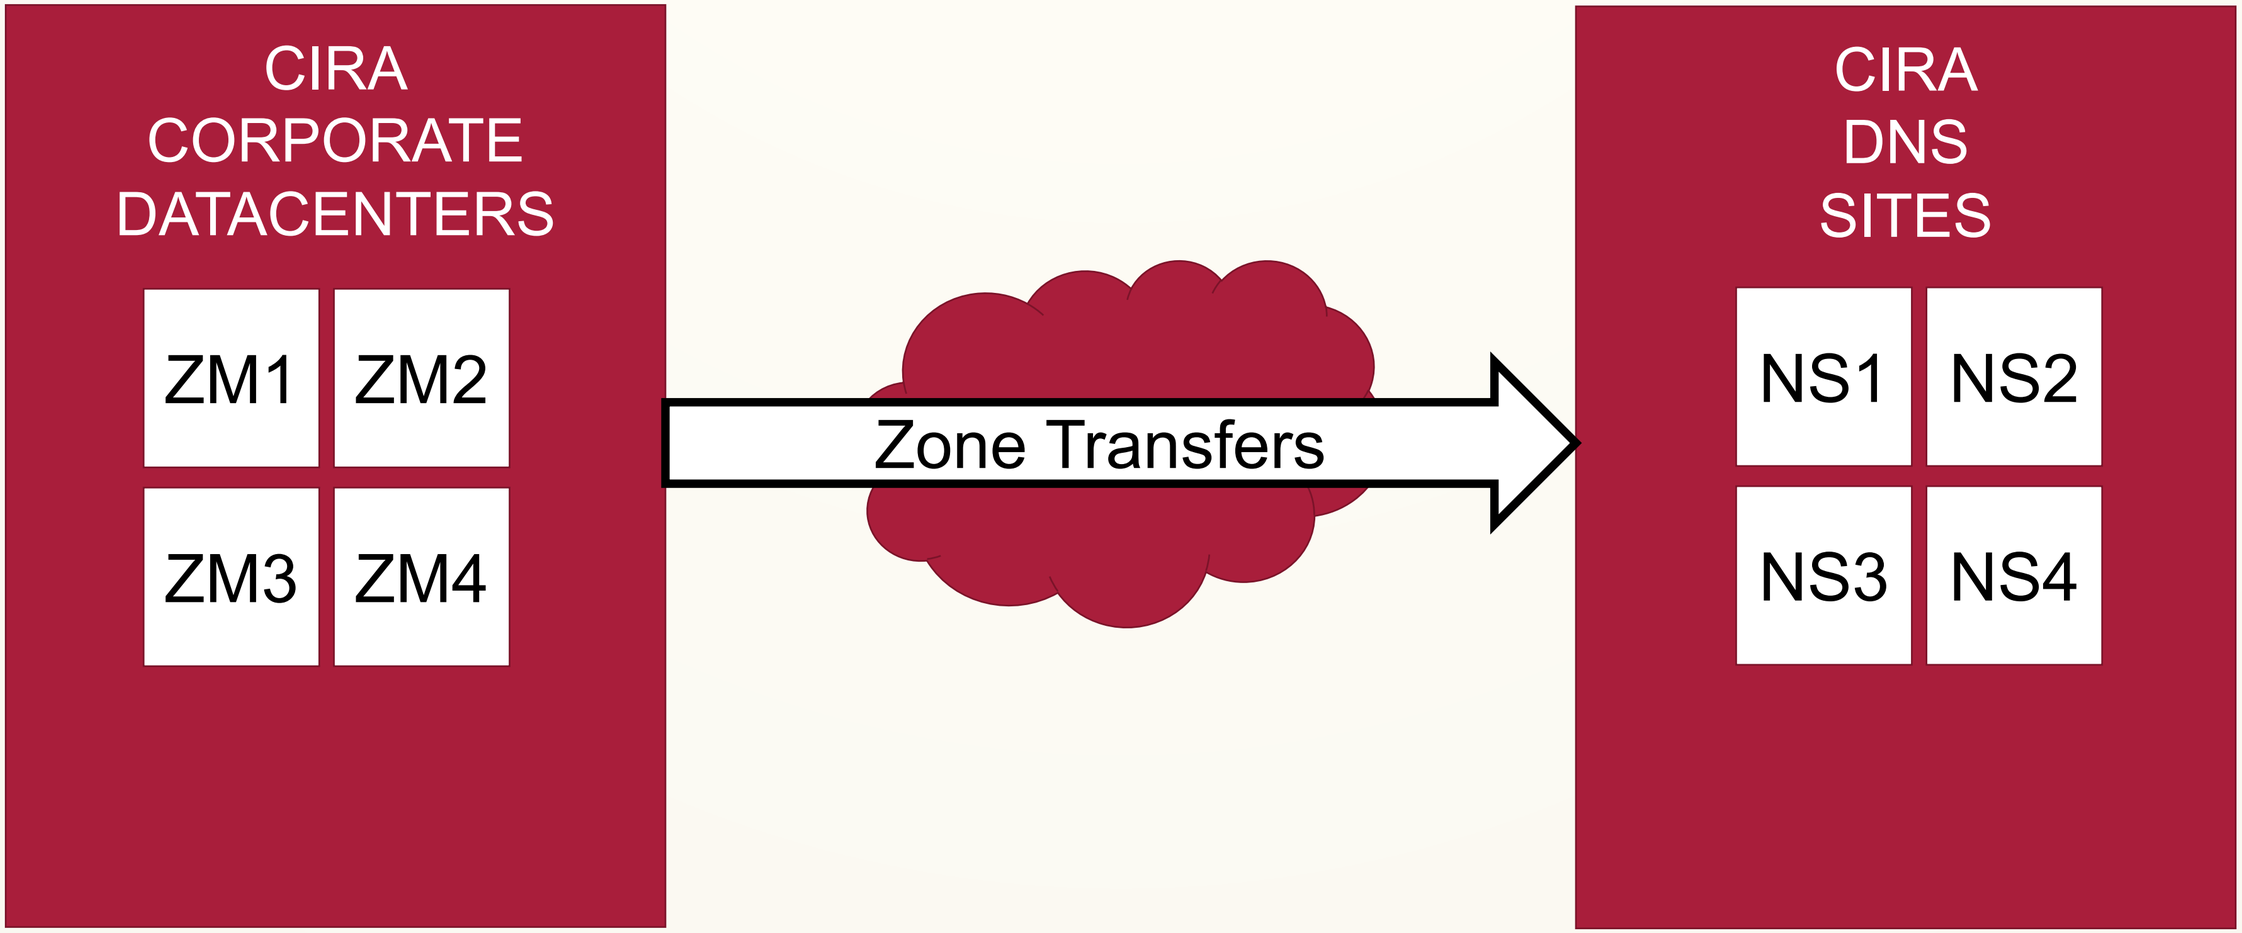

Zone Propagation – Old Mindset

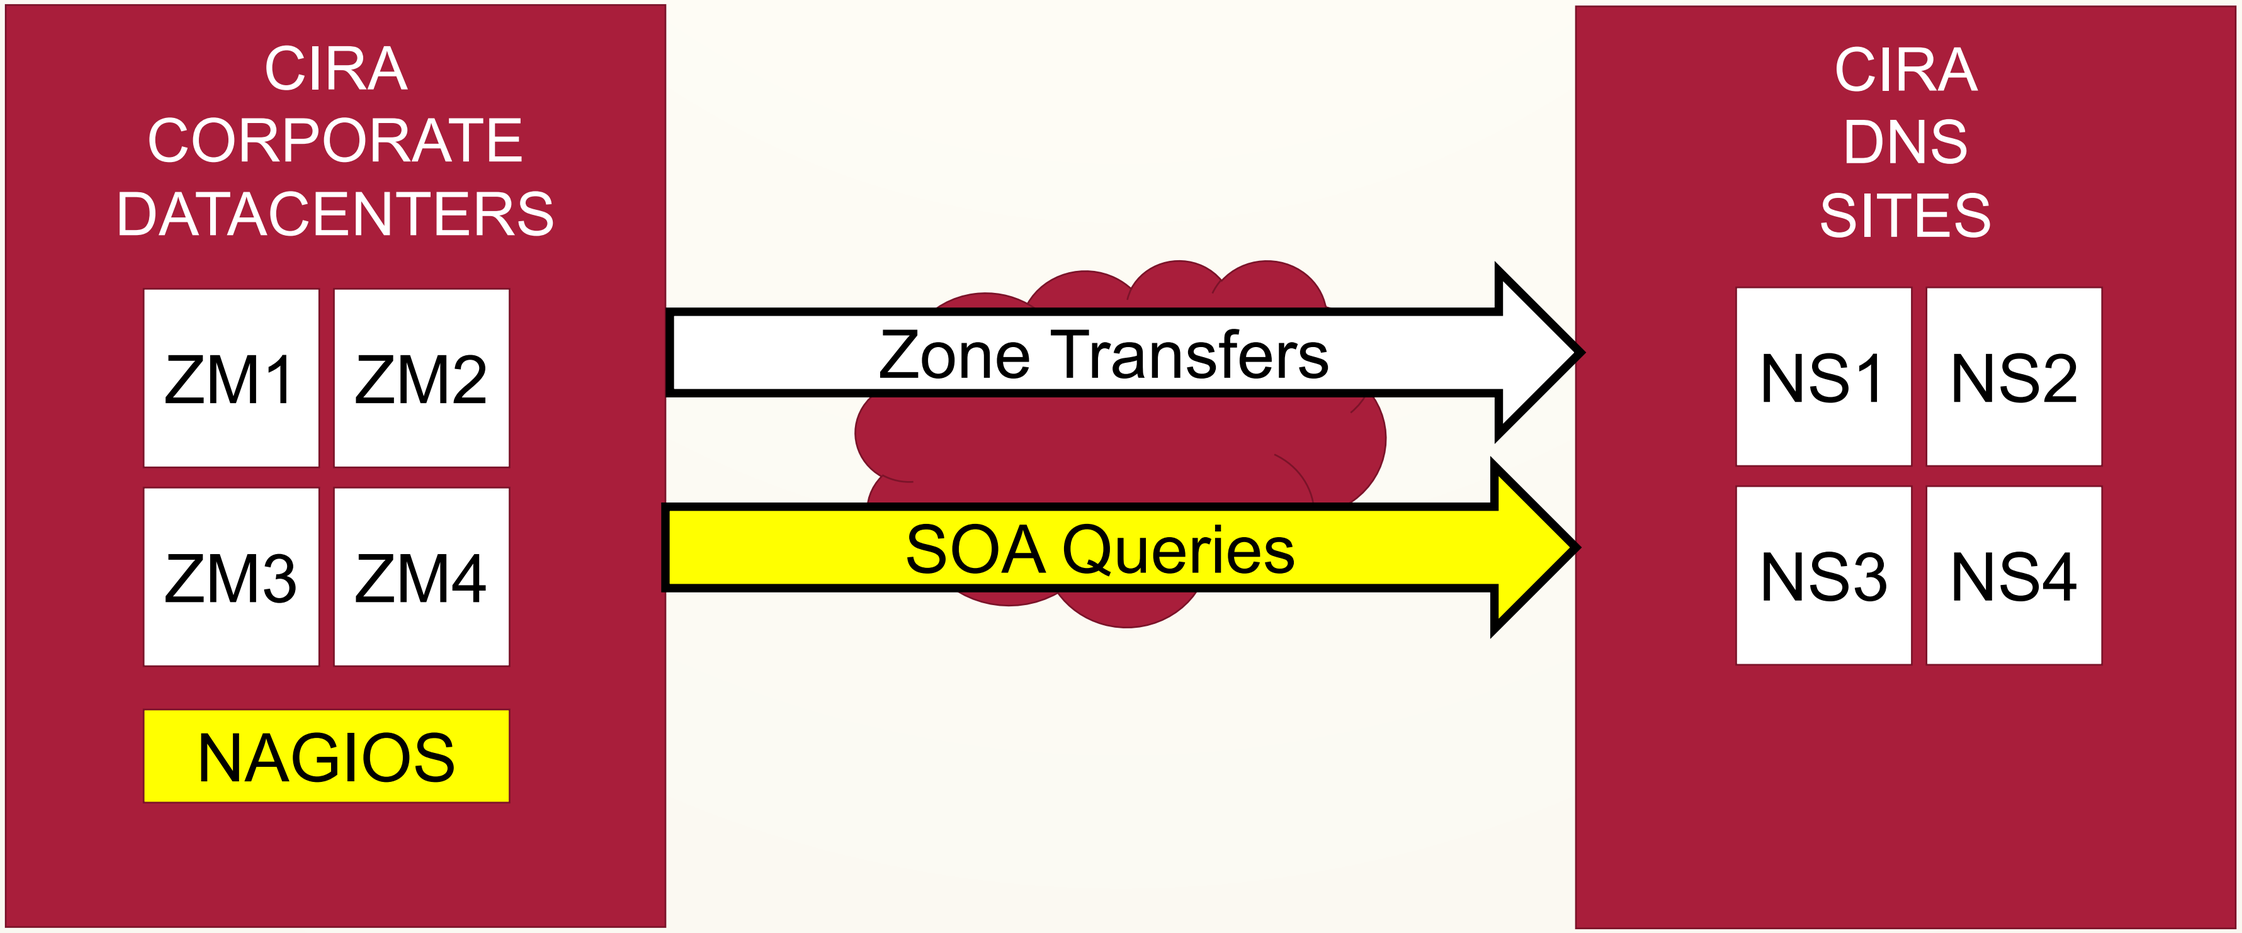

Zone Propagation – Old Mindset

- Predictable Serial Number - YYMMDDHHMM
- Compare to $localtime
- Send alert if $localtime - $serialnumber is > $threshold minutes
- On-Call SysAdmin triage
- Problems:
  - Missed Notify – Trigger a refresh
  - Packet loss – Stop BGP advertisements
  - Transit down – Can’t manage, have datacenter pull power?

cira

Zone Propagation – New Mindset

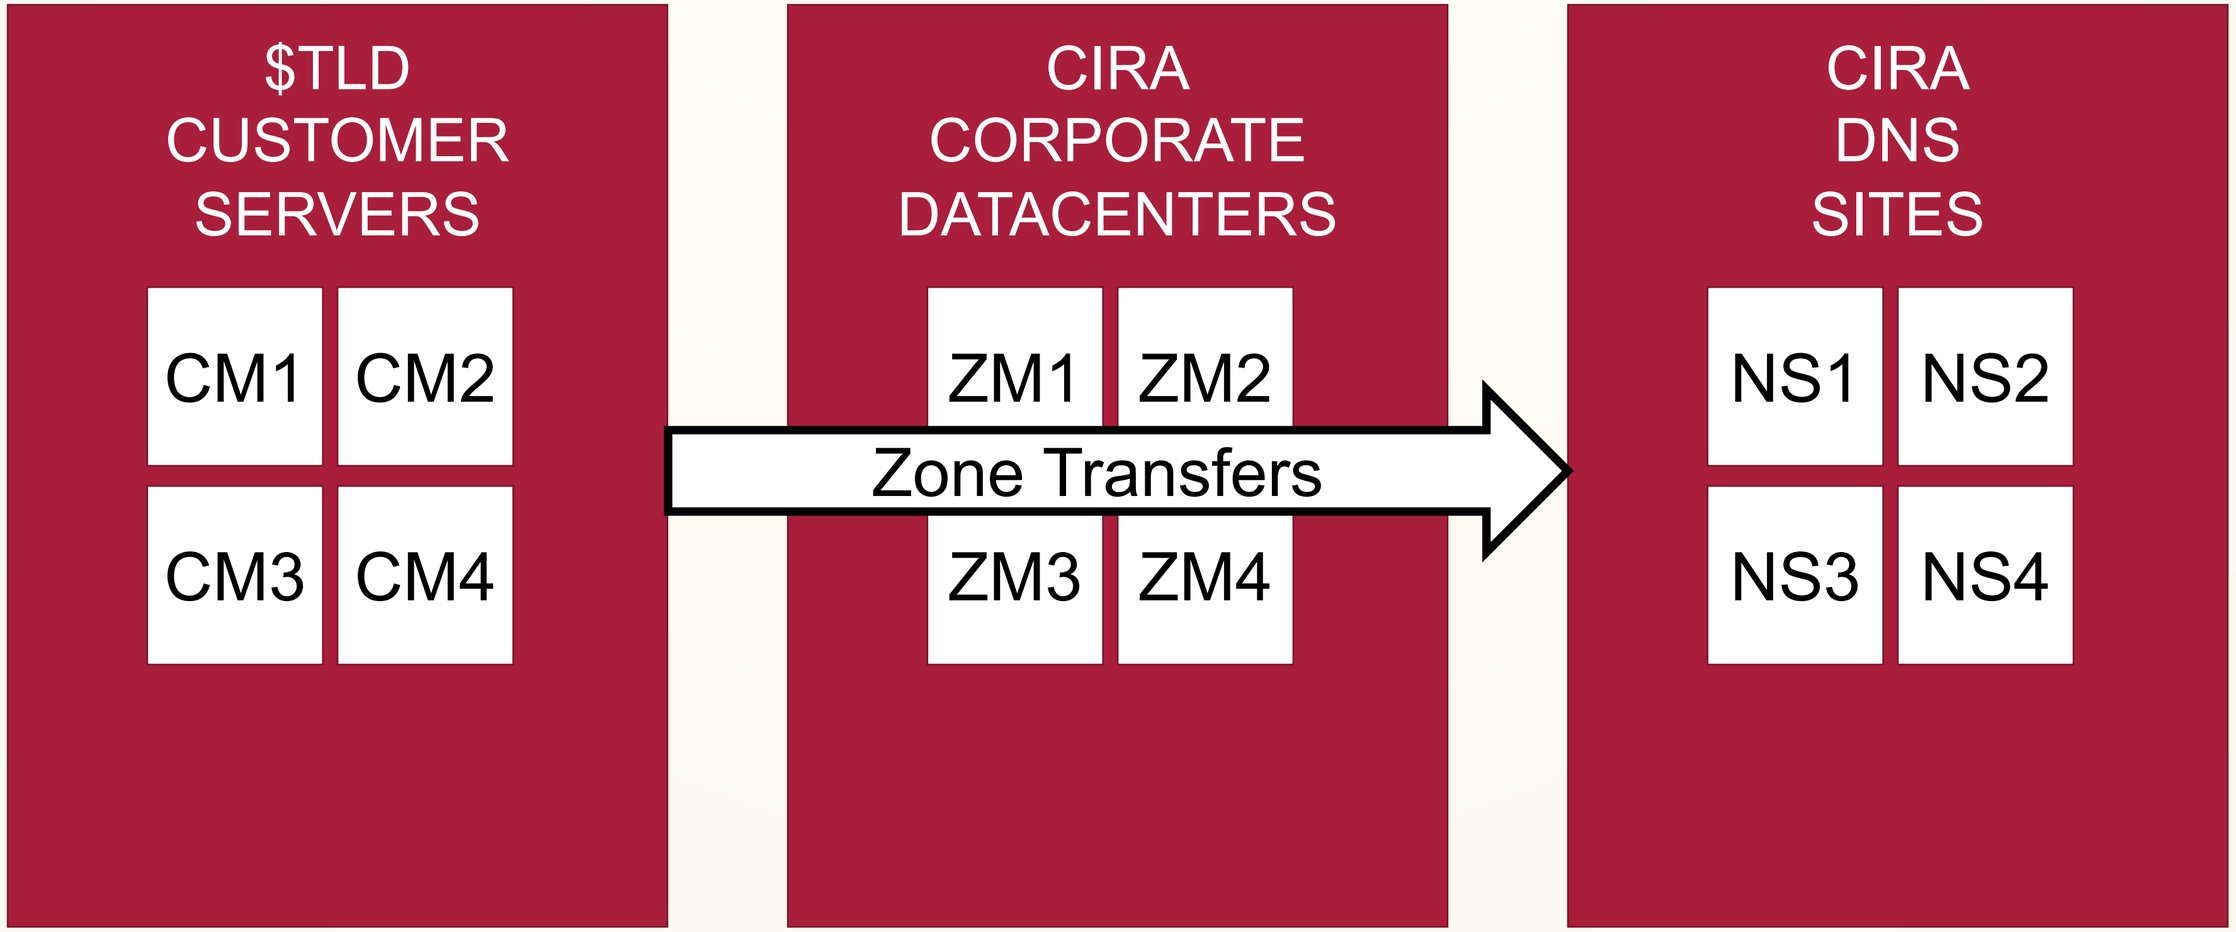

Zone Propagation – New Mindset

- New Requirement:
  - Customers require < 5 minute update times

- Challenges:
  - Serial numbers are no longer predictable
  - Many hundreds of zones
  - On-Call intervention and resolution in a 5 minute time frame

- User Stories:
  - We must architect to avoid late-transfer situations
  - We must know a zone is out-of-date before the customer
  - We must self-heal this situation without Operator intervention

Zone Propagation – Architecture

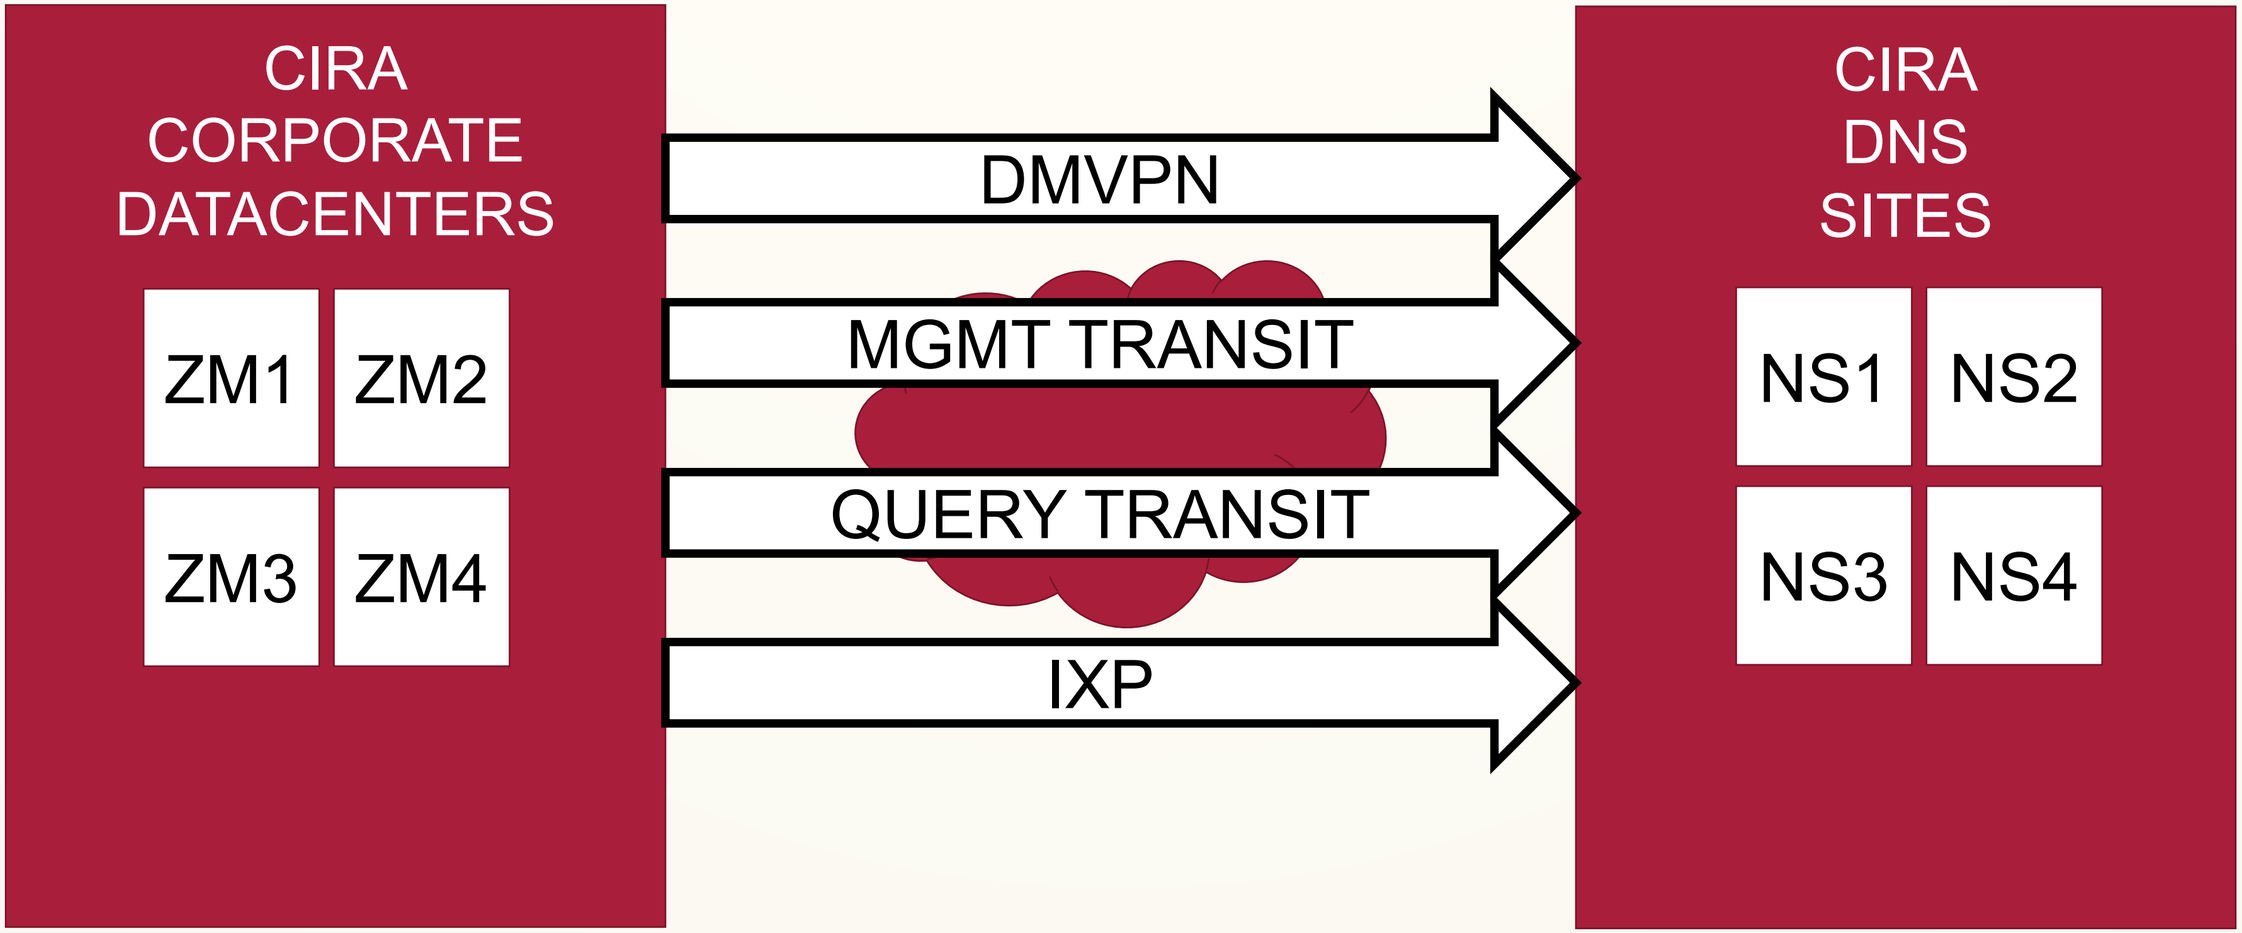

Zone Propagation – Detection & Alerting

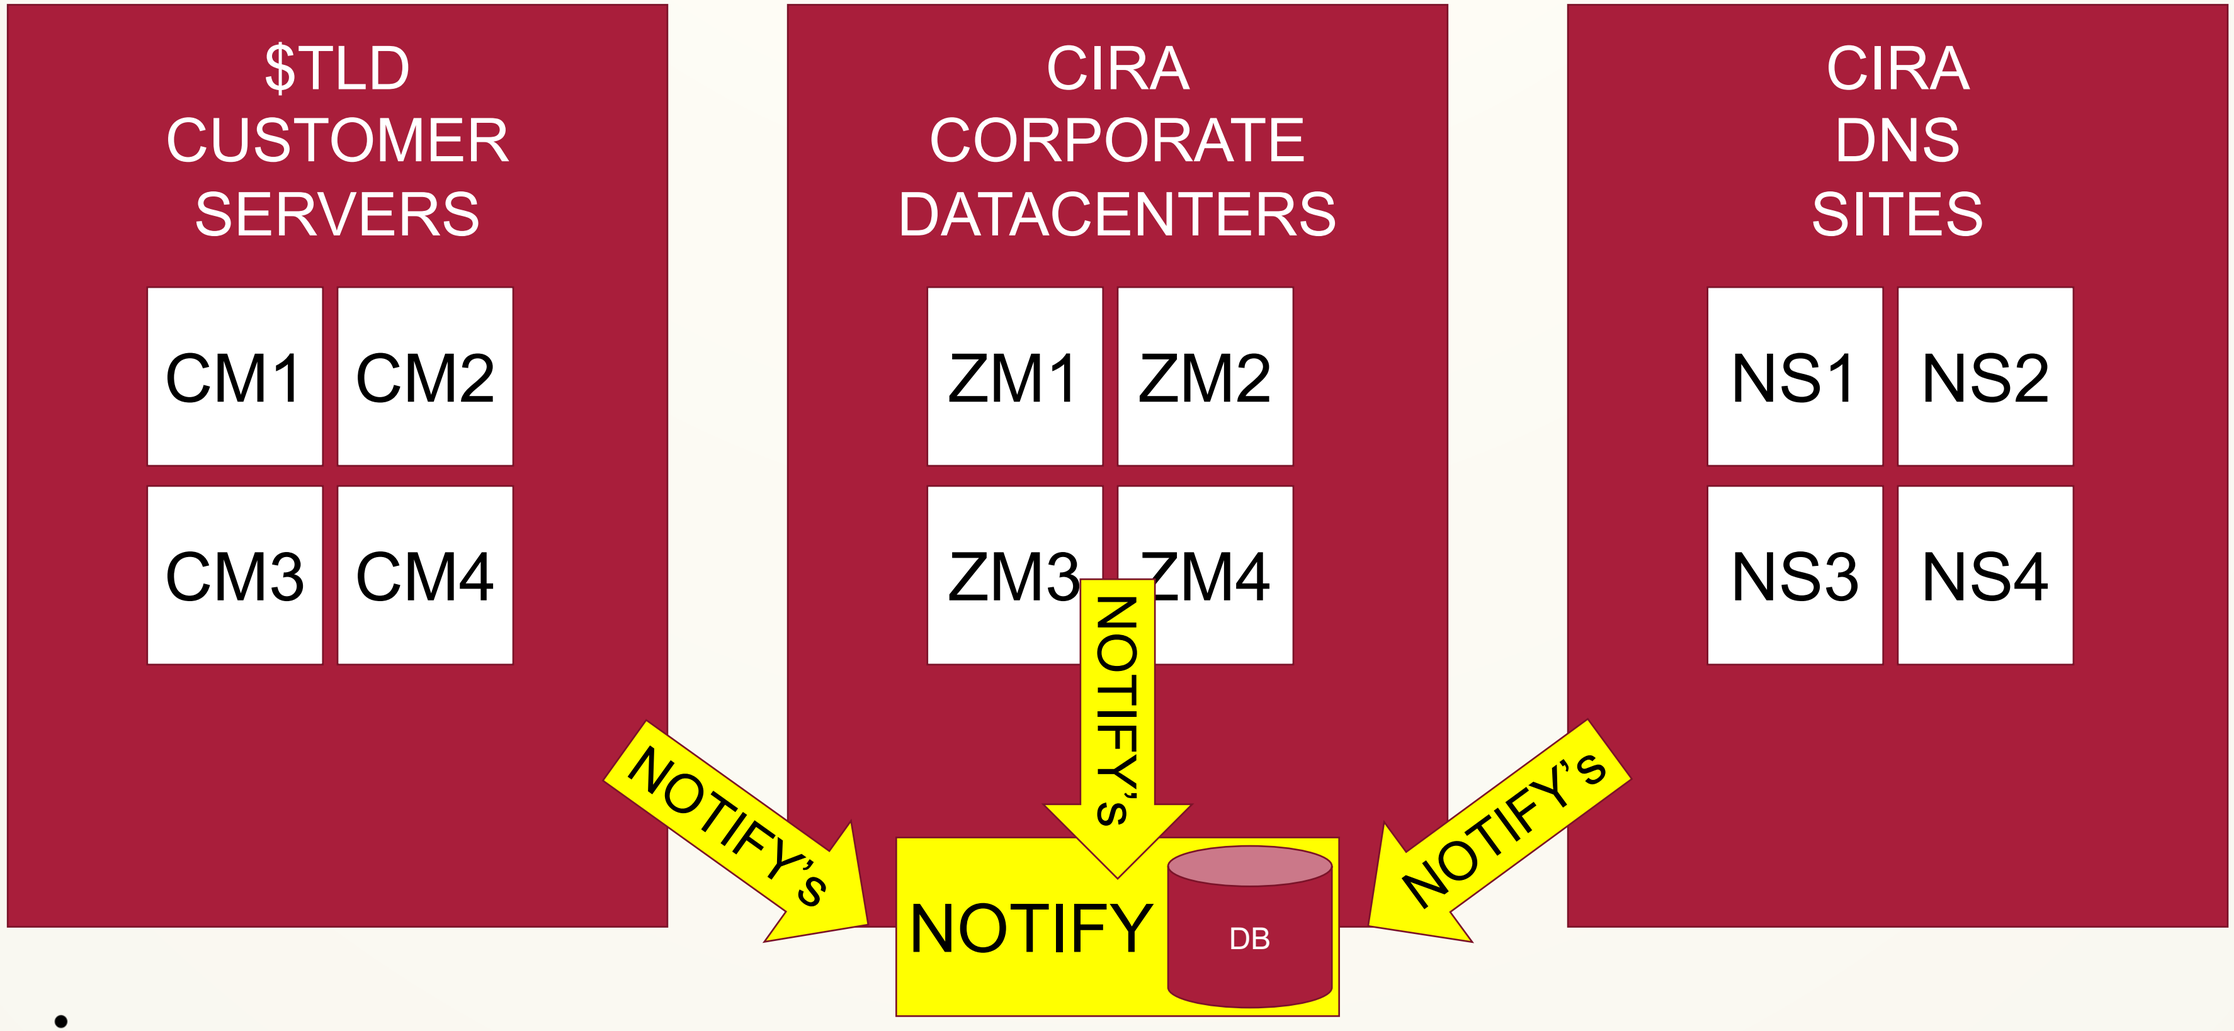

Zone Propagation – Detection & Alerting

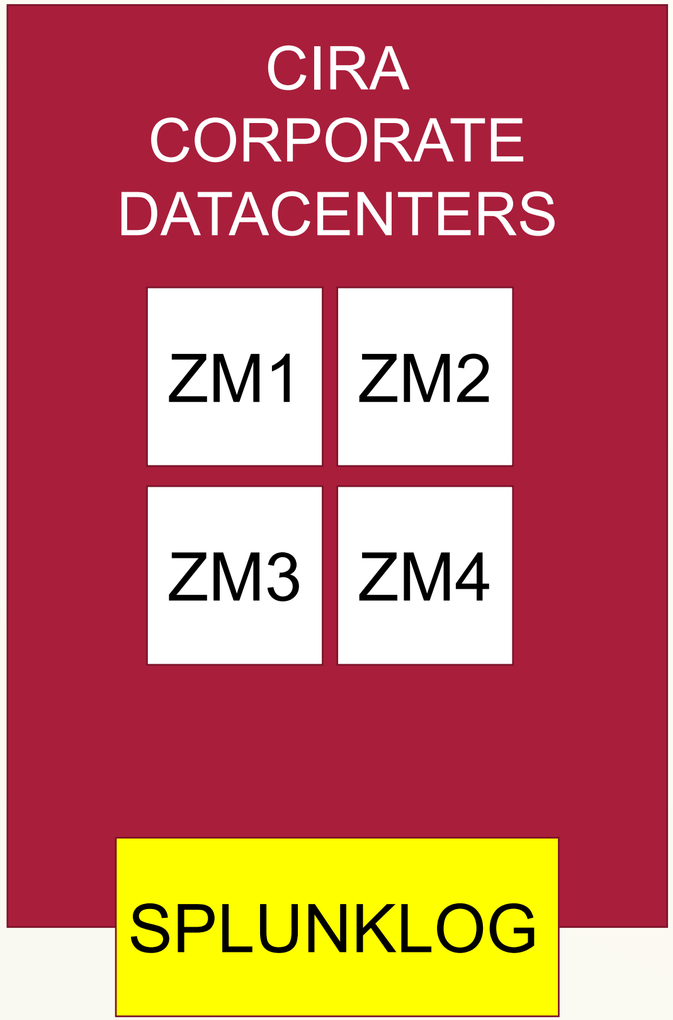

From BIND logs via Splunk

- Customer – “Connection refused”
- Customer – “Connection timed out”
- Customer – Lack of “received notify”
- Customer – Lack of “transfer complete”

- Customer – Received notify for unconfigured zone

…above threshold? Alert!

Zone Propagation – Detect/Alert/Heal

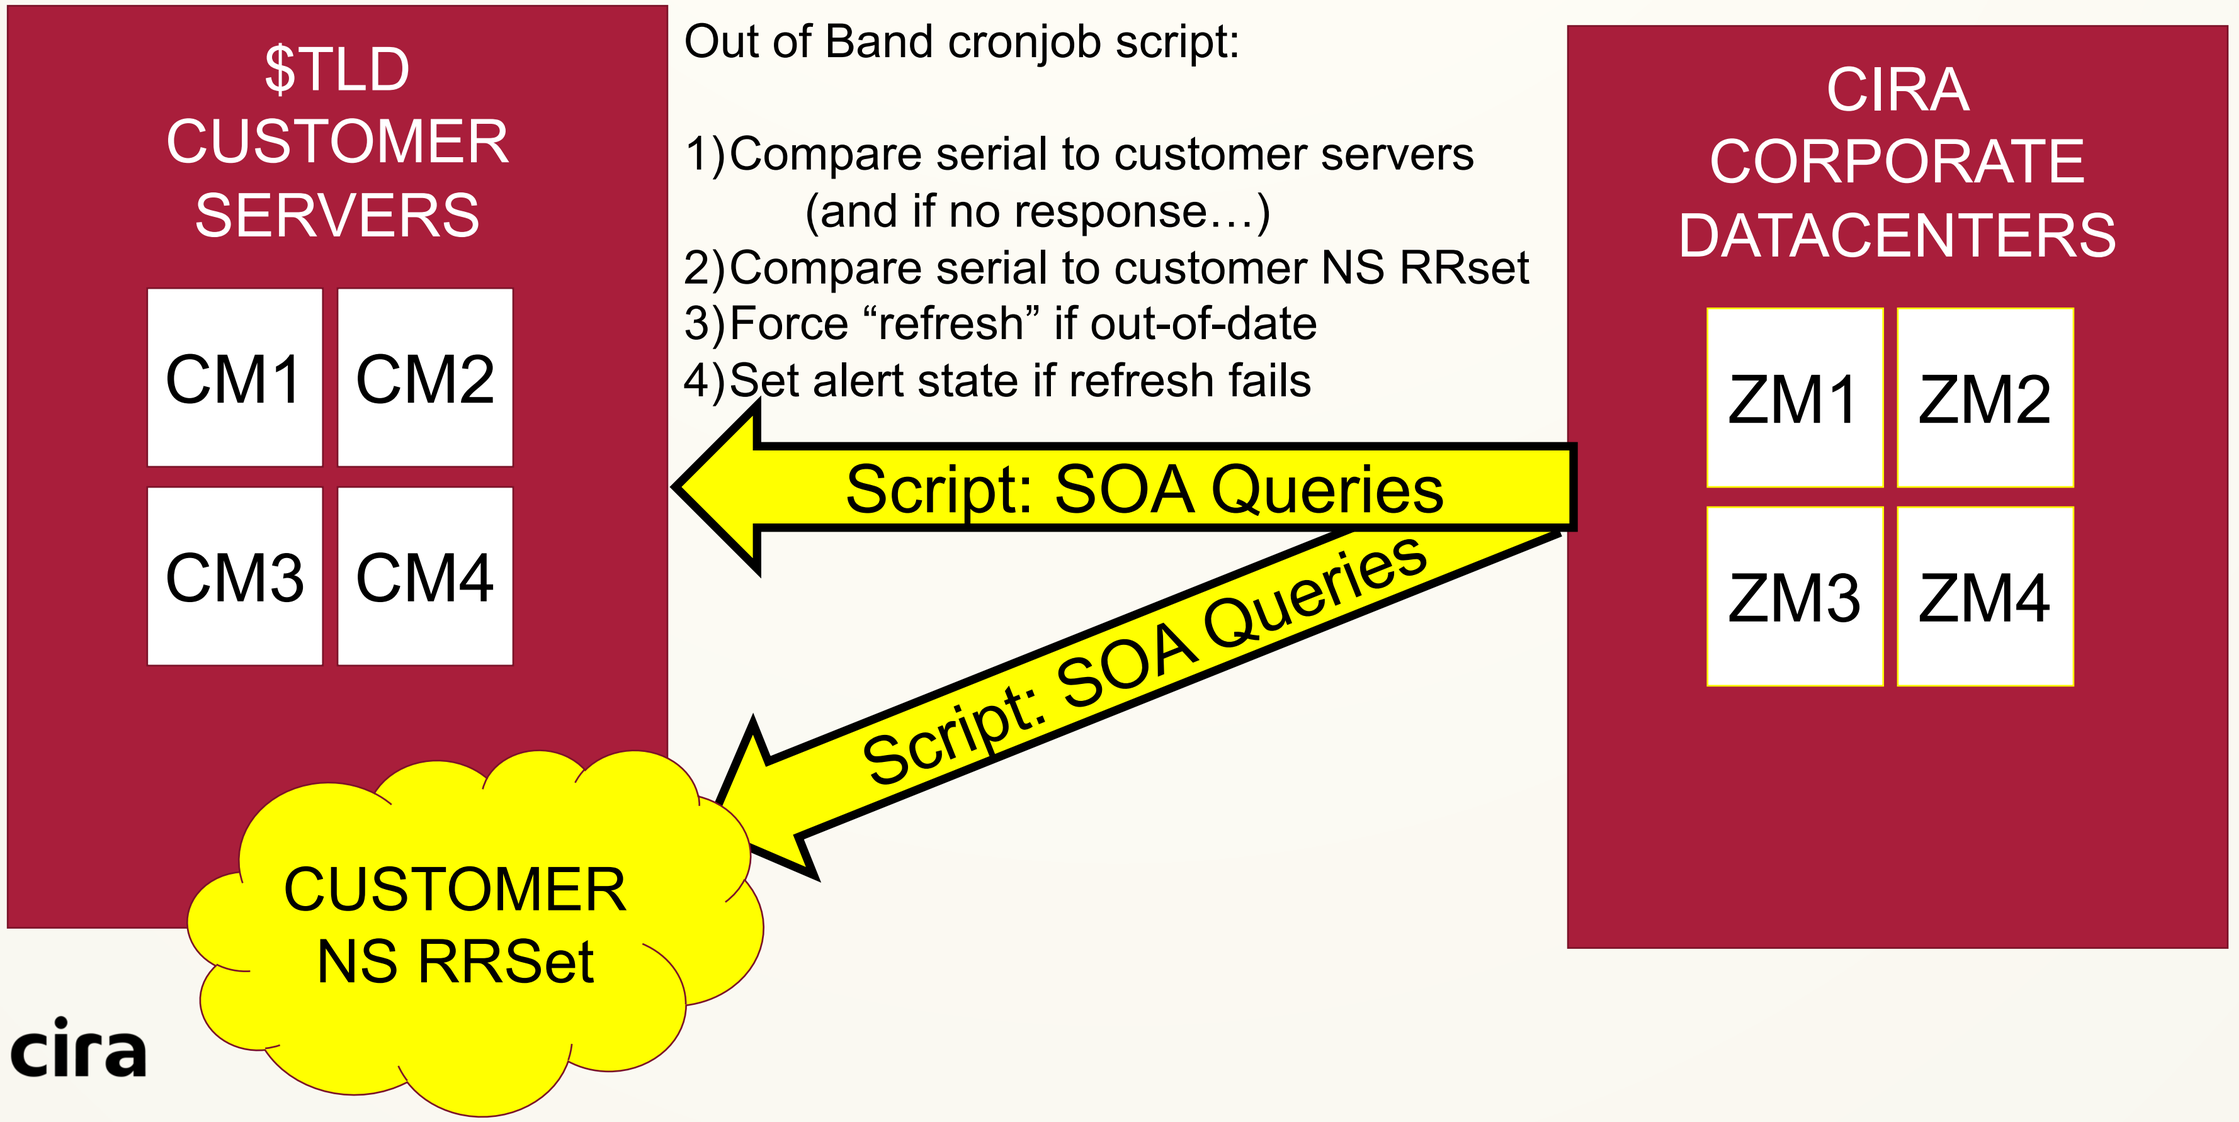

Zone Propagation – Detect/Alert/Heal

CIRA CORPORATE DATACENTERS

ZM1 ZM2

ZM3 ZM4

Out of Band cronjob script:

1) Compare serial to distribution servers
2) Did zone update in the last minute?
3) Force “refresh” if out-of-date
4) Set alert state if refresh fails

Script: SOA Queries

Optional:

If number of outdated zones is above $threshold, pull the route(s) and/or ask permission to pull the route(s).

CIRA DNS SITES

NS1 NS2

NS3 NS4

Zone Propagation - Outcomes

Beneficial Outcomes:

- Multiple levels of monitoring throughout the chain
- Auto-detect late zones within 2 minutes
- Auto-repair late zones immediately if possible
- Auto-withdraw lagging/lossy servers quickly

- Far less customer service requests
- Far less incidents
- Confidence

WWW.CIRA.CA

DNS Availability

Availability – Old Mindset

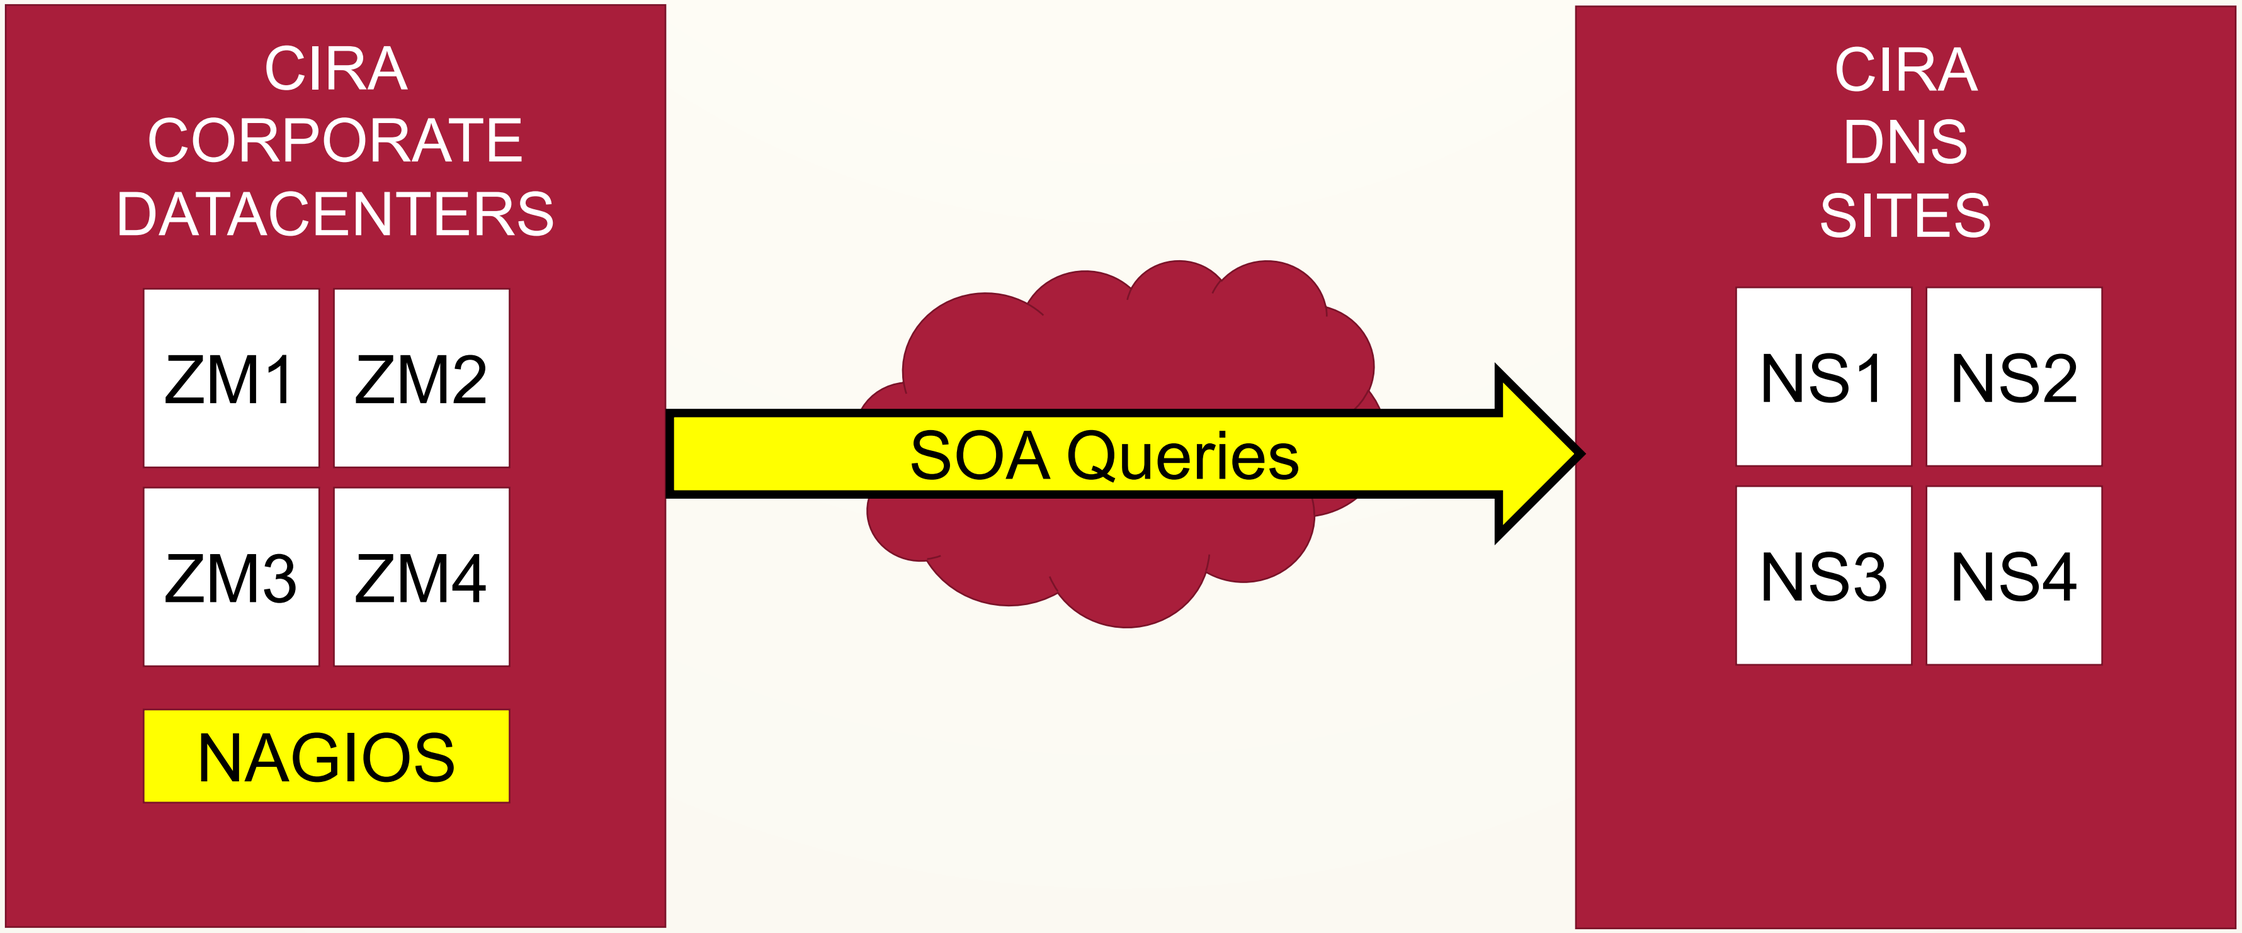

Availability – New Mindset

- New Requirement:
  - Anycast requires us to monitor from around the planet

- Challenges:
  - Love RIPE ATLAS – but…
    - Testing v4/v6 UDP/TCP every minute feels like abuse
  - Industry monitoring leader quoted ~USD $10k/month…
    - Only a few of our clouds, every 15 minutes…

- User Stories:
  - We need a low cost global remote monitoring network
    - Custom dashboards
    - Customer alerts/reports

Availability – Monitoring Architecture

cira

Availability – Monitoring Architecture

Probe	TESTS	UDP4	TCP4	UDP6	TCP6
johannes1	1440	100.00%	100.00%	100.00%	100.00%
johannes2	1440	100.00%	100.00%	100.00%	100.00%
bangalore1	1440	100.00%	100.00%	100.00%	100.00%
hongkong1	1440	100.00%	100.00%	100.00%	100.00%
hongkong2	1440	100.00%	100.00%	100.00%	100.00%
mumbai1	1440	100.00%	100.00%	100.00%	100.00%
mumbai2	1440	100.00%	100.00%	100.00%	100.00%
osaka1	1440	100.00%	100.00%	100.00%	100.00%
seoul1	1440	100.00%	100.00%	99.93%	100.00%
seoul2	1440	100.00%	100.00%	100.00%	100.00%
singapore1	1440	100.00%	100.00%	100.00%	100.00%
singapore2	1440	100.00%	100.00%	100.00%	100.00%
singapore3	1439	100.00%	100.00%	100.00%	100.00%
singapore4	1440	100.00%	100.00%	100.00%	100.00%
telaviv1	1440	99.93%	100.00%	99.93%	100.00%
tokyo1	1440	100.00%	100.00%	100.00%	100.00%
tokyo2	1440	100.00%	99.93%	100.00%	100.00%
tokyo3	1440	100.00%	100.00%	100.00%	100.00%
amsterdam1	1440	100.00%	100.00%	100.00%	100.00%
amsterdam2	1440	100.00%	100.00%	100.00%	100.00%
amsterdam3	1440	100.00%	100.00%	100.00%	100.00%
athens1	1440	100.00%	100.00%	100.00%	100.00%
barcelona1	1440	100.00%	100.00%	100.00%	100.00%
belgrade1	1440	100.00%	100.00%	100.00%	100.00%
brussels1	1440	100.00%	100.00%	100.00%	100.00%
bucharest1	1440	100.00%	100.00%	100.00%	100.00%
copenhagen1	1440	100.00%	100.00%	100.00%	100.00%
frankfurt1	1440	100.00%	100.00%	100.00%	100.00%

Probe	UDP4	TCP4	UDP6	TCP6
amsterdam1	LHR6	LHR5	LHR7	LHR5
amsterdam2	FRA6	FRA3	FRA4	FRA3
amsterdam3	FRA3	FRA5	FRA6	FRA5
ashburn1	ASH3	ASH3	ASH3	ASH3
athens1	LHR5	LHR7	LHR8	LHR8
atlanta1	MIA4	MIA4	MIA4	MIA4
bangalore1	FRA4	FRA3	HKG3	HKG4
barcelona1	FRA3	FRA4	FRA6	FRA3
belgrade1	FRA3	FRA6	LHR6	LHR6
brisbane1	SYD1	SYD1	SYD1	SYD1
brussels1	LHR7	LHR6	LHR8	LHR5
bucharest1	LHR5	LHR5	LHR9	LHR7
chicago1	YYZ3	YYZ3	YYZ4	YYZ3
chicago2	YYZ4	YYZ3	YYZ4	YYZ3
copenhagen1	FRA3	FRA5	FRA5	FRA5
dallas1	YYZ3	YYZ3	LAX4	LAX3
frankfurt1	FRA5	FRA6	FRA3	FRA6
frankfurt2	FRA5	FRA3	FRA6	FRA3
frankfurt3	FRA6	FRA5	FRA5	FRA5
hongkong1	HKG4	HKG4	HKG4	HKG4
hongkong2	HKG3	HKG4	HKG3	HKG3
johannes1	JNB1	JNB1	JNB1	JNB1
johannes2	JNB1	JNB1	JNB1	JNB1
kiev1	LHR6	LHR6	LHR10	LHR8
ljubljana1	LHR10	LHR7	LHR6	LHR9
london1	LHR8	LHR9	LHR6	LHR6
london2	LHR6	LHR7	LHR10	LHR10
london3	LHR8	LHR6	LHR7	LHR6

Availability – Monitoring Architecture

2042869152	2022-09-01 21:36:03	vienna1	25	Y	LHR3	23	Y	LHR1	38	Y	LHR4	23	Y	LHR2
2042860359	2022-09-01 21:35:05	vienna1	27	Y	LHR4	25	Y	LHR2	24	Y	LHR3	25	Y	LHR4
2042851700	2022-09-01 21:34:09	vienna1	23	Y	LHR3	28	Y	LHR1	24	Y	LHR3	27	Y	LHR1
2042845651	2022-09-01 21:33:27	vienna1	24	Y	LHR3	23	Y	LHR1	0	N		0	N	
2042836461	2022-09-01 21:32:27	vienna1	25	Y	LHR4	23	Y	LHR3	0	N		0	N	
2042827277	2022-09-01 21:31:27	vienna1	23	Y	LHR3	23	Y	LHR4	0	N		0	N	
2042815233	2022-09-01 21:30:11	vienna1	121	Y	LHR1	125	Y	LHR3	133	Y	LHR4	124	Y	LHR2
2042805118	2022-09-01 21:29:03	vienna1	24	Y	LHR3	24	Y	LHR3	31	Y	LHR1	24	Y	LHR4
2042795904	2022-09-01 21:28:03	vienna1	23	Y	LHR1	24	Y	LHR3	32	Y	LHR3	23	Y	LHR3
2042786694	2022-09-01 21:27:03	vienna1	24	Y	LHR4	24	Y	LHR3	24	Y	LHR2	23	Y	LHR2
2042777552	2022-09-01 21:26:04	vienna1	23	Y	LHR1	24	Y	LHR2	33	Y	LHR3	24	Y	LHR1
2042768332	2022-09-01 21:25:04	vienna1	35	Y	LHR2	23	Y	LHR4	23	Y	LHR4	23	Y	LHR3
2042759098	2022-09-01 21:24:03	vienna1	23	Y	LHR1	23	Y	LHR3	26	Y	LHR3	24	Y	LHR3

Instant visibility with granular start/stop times

- DNS server issues
- Network issues
- External routing events

Availability – New Mindset

Beneficial Outcomes:

- Per-minute real-time granularity for all protocols
- Custom alerting/reports
- Immediately identifies problematic circuit/server/site

- Detects:
  - Misconfiguration (ex: IP address missing/blackholed)
  - Server down or packet loss
  - Transit down or packet loss
  - IXP’s down or packet loss

WWW.CIRA.CA

Round-Trip Time

RTT – New Mindset

- New Requirement:
  - Two independent clouds/ASN's (gTLD requirements)
  - Constant RTT feedback loop
  - Site failure in any location doesn't impact RTT SLA's
  - Option to provide both clouds in a "unified" format

- Challenges:
  - Routing – What's true today may be different tomorrow
  - Routing – Asia is a very very very large place

CIRA TLD Anycast Locations

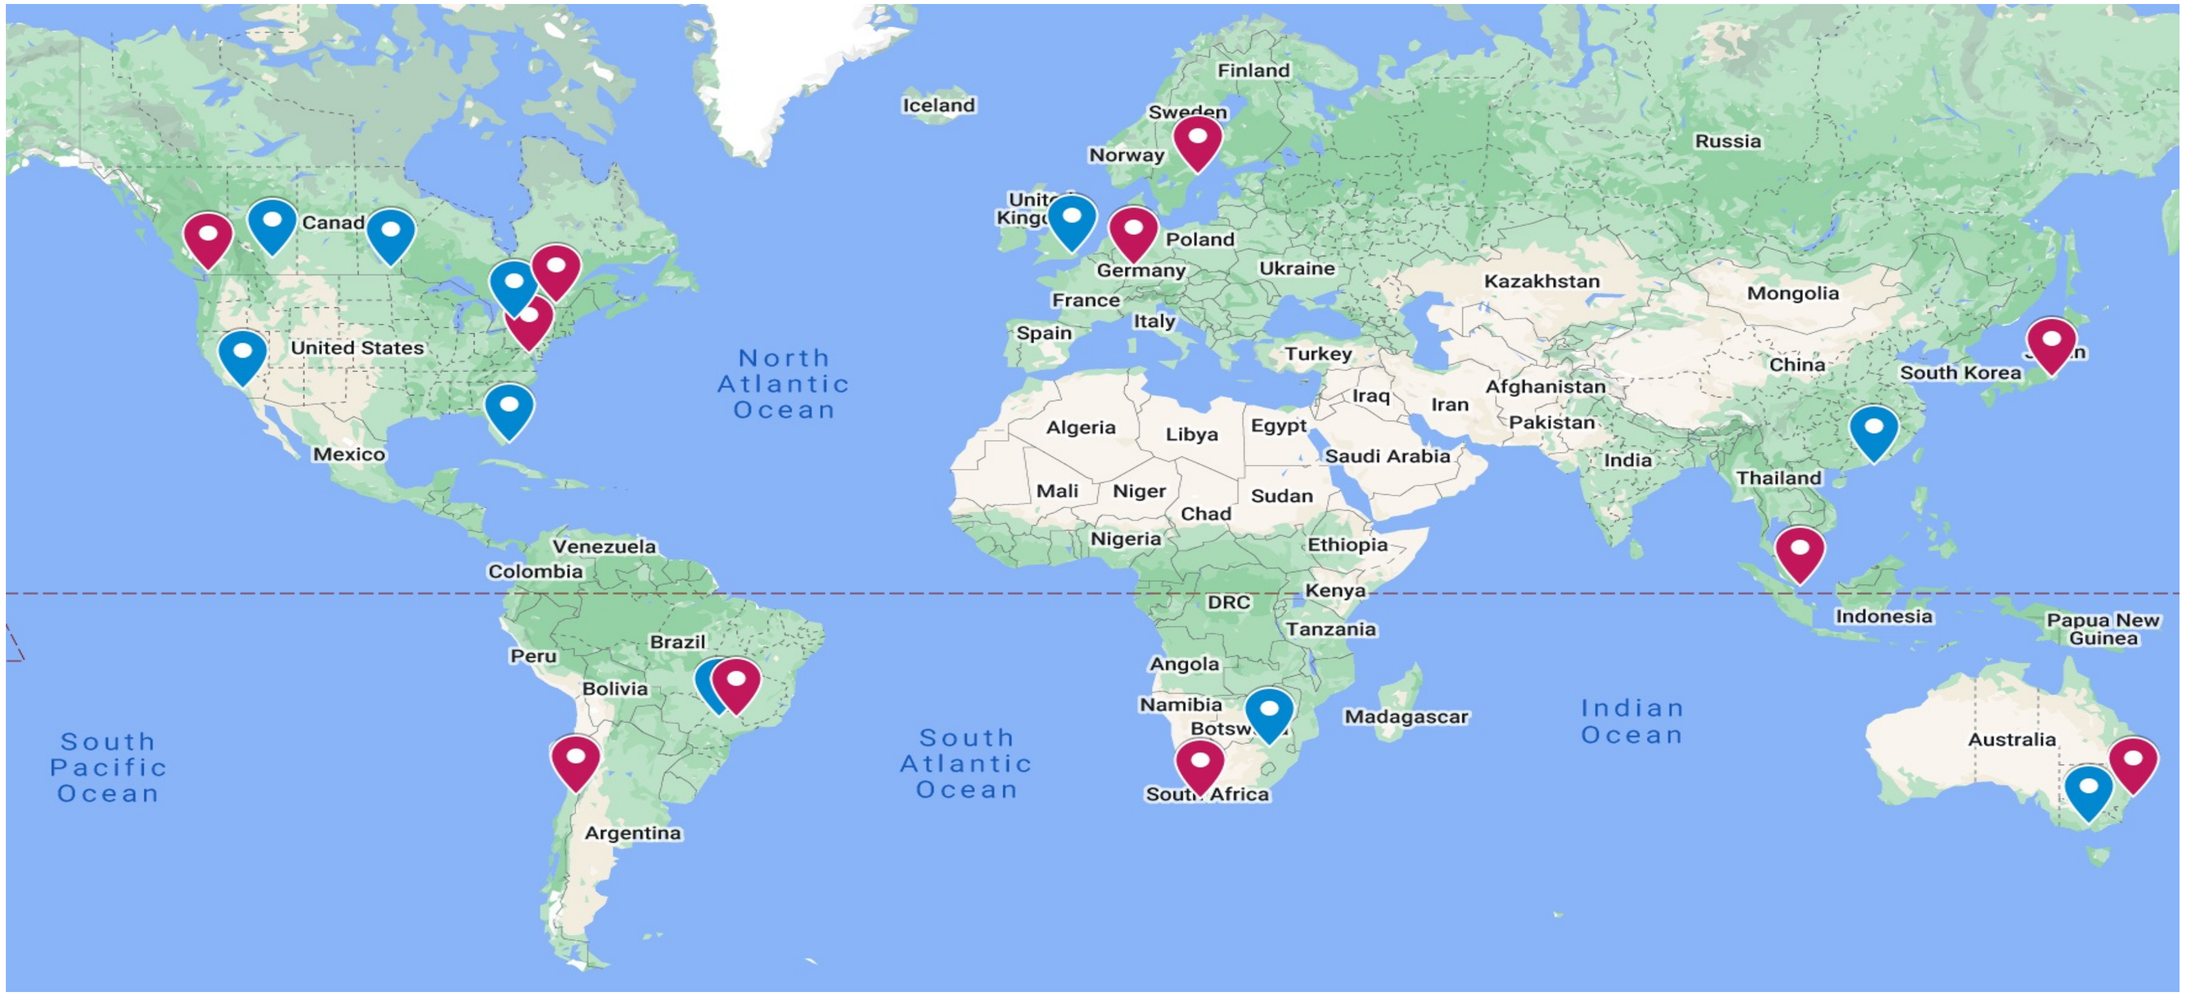

WWW.CIRA.CA

cira

Bandwidth Selection Considerations

- Tier 1 level global transit providers
  - Regional transit providers/aggregators are bad!

- Cloud 1 has consistent transit provider(s)
- Cloud 2 has consistent transit provider(s)

- Routing failover testing!
  - Know where your traffic goes when a site is down!

- Not having an overall strategy for this results in traffic crossing oceans!

CIRA's Ideal in-continent routing

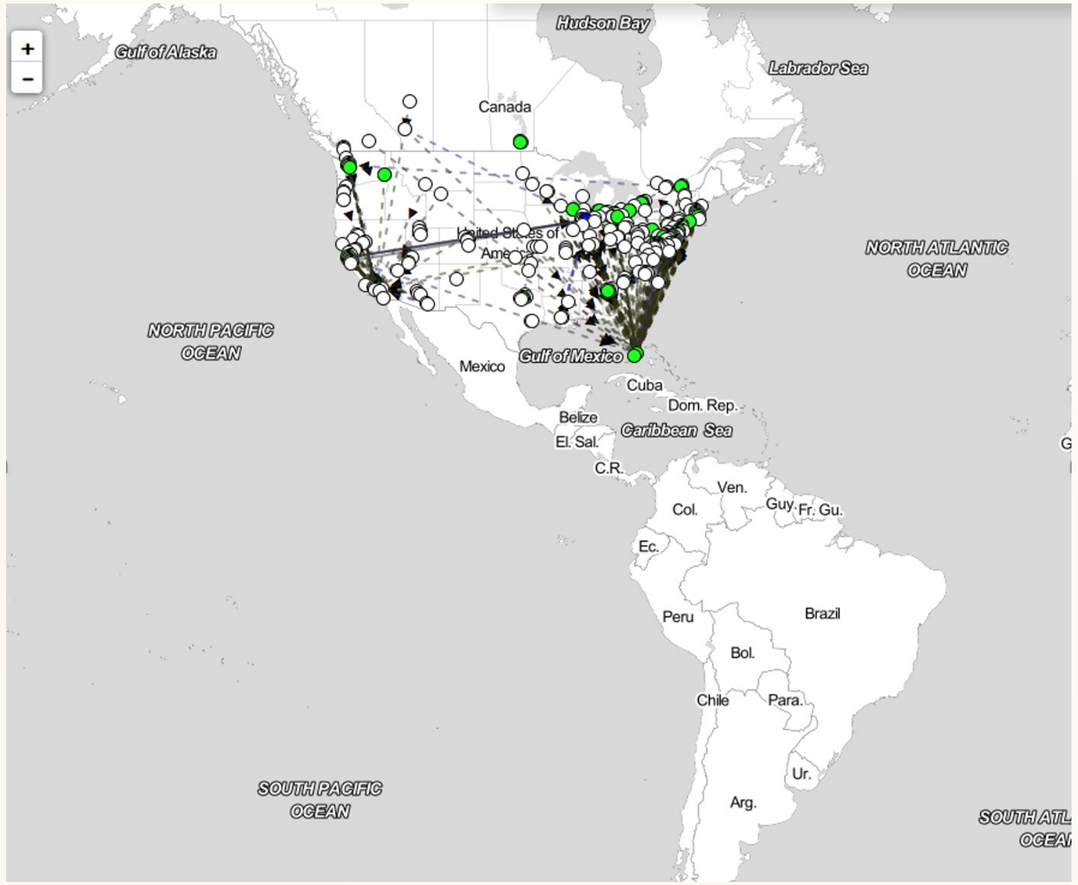

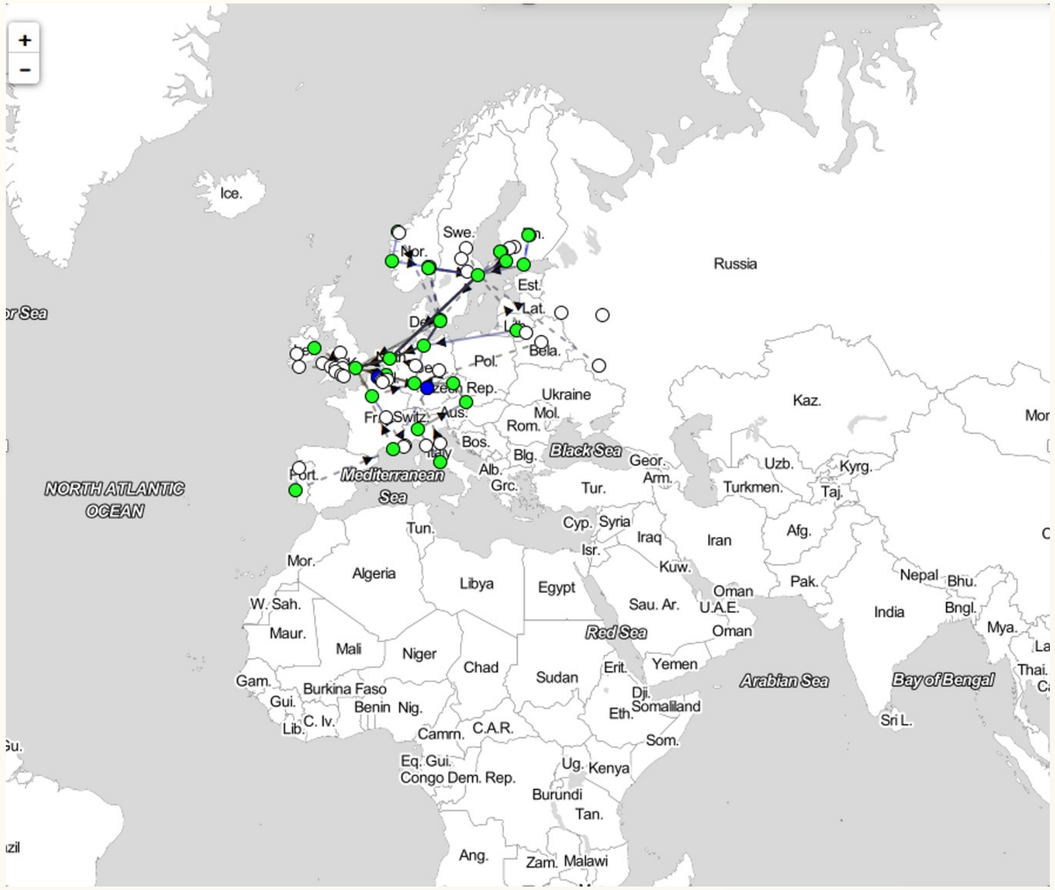

What you wish to avoid…

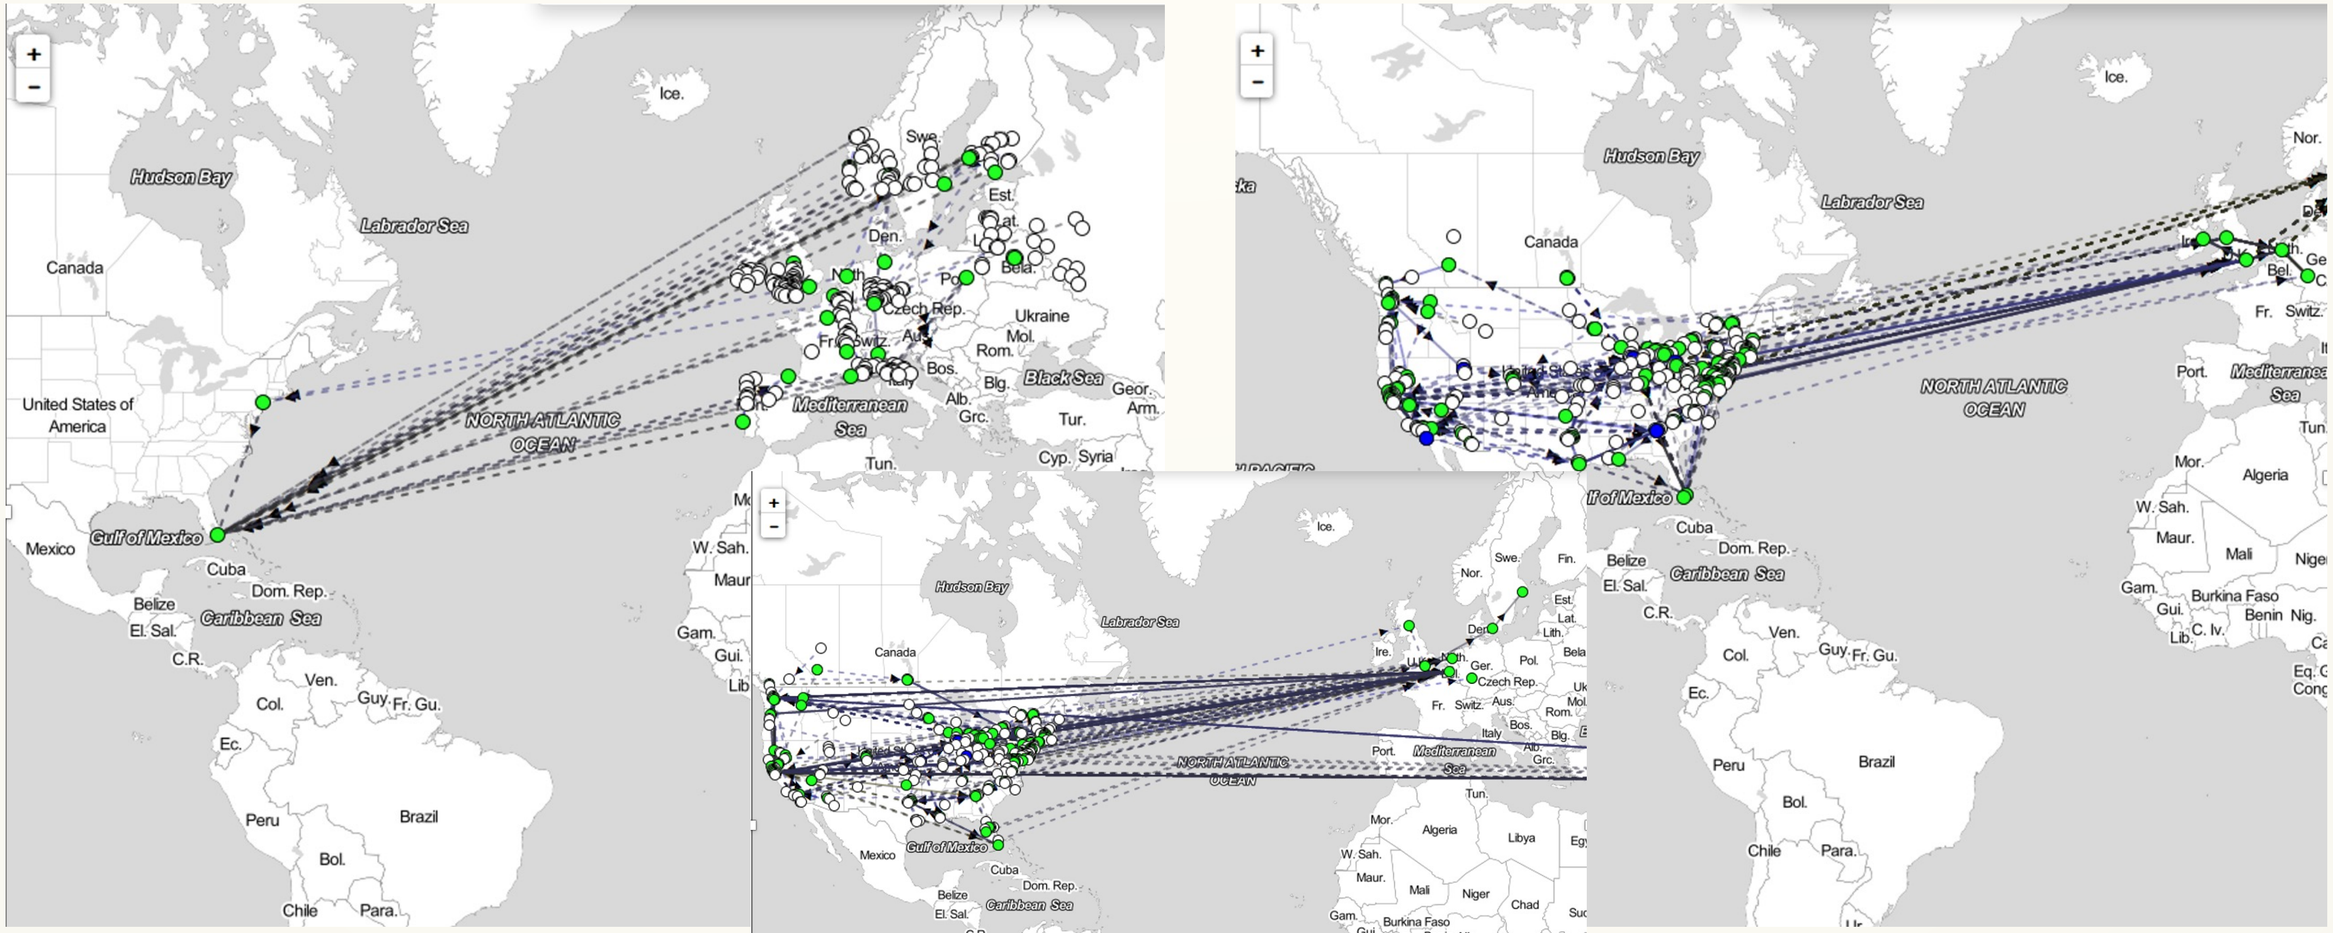

IXP Selection Considerations

- Largest number of peers and/or traffic!

- Locations with multiple IXP's in-building!
  - London: LINX and LONAP!
  - Frankfurt: DE-CIX and NL-IX!
  - Sydney: Equinix-Sydney, MegaIX-Sydney, IXAustralia-NSW

  ...give yourself expansion options!

Tools:
PeeringDB, PCH IXP DIR, BGP.HE.NET, etc…

…and stop only peering with route-servers!

Know your strengths and weaknesses

- RIPE ATLAS
  - Global, per-region, or per-country averages!
  - Focus on the worst RTT's!

- APNIC ASPOP
  - Know the population of every ASN in the world!

- PeeringDB
  - Know which ASN's are present at every IXP!

- Calculate which exchanges will remediate the most issues!
- Maintain an index of the lesser-served large ASN's!

Where we stand today…

- Approximately 1/3rd of all TLD's in the root zone file!
- 484 total TLD's
  - 448 gTLD's
  - 36 ccTLD's

- ~51 billion DNS queries/day (average)
  - 65% of traffic over local IXP's (8000+ IPv4 peers)
  - 35% of traffic over global transit