Anycast Authoritative DNS Service of MMIX

Saw Yan Paing

CCIE# 57007

Chief Engineer @ARK Premium Solutions Limited | @MMIX

Why is DNS Important?

- DNS ensures the internet is not only user-friendly but also works smoothly, loading whatever content we ask for quickly and efficiently.

- If a DNS cannot translate the domain name with the right IP address, we won't be able to access the website what we're looking for.

Without DNS.....?

How could you get any websites?

Types of DNS servers

1. **Recursive resolver**
2. **DNS root name server**
3. **TLD name server**
   - generic TLDs (gTLDs)
   - Sponsored TLDs
   - Non-sponsored TLDs
   - Country Code TLDs (ccTLDs)
   - Internationalized TLDs
4. **Authoritative name server**

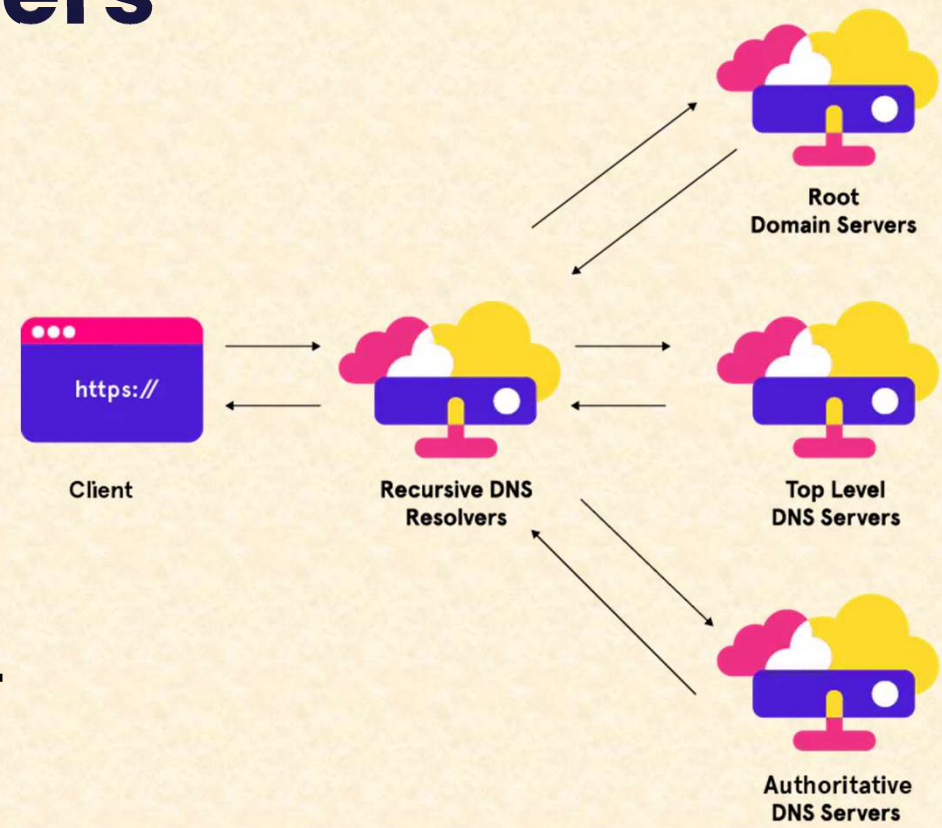

Naming Hierarchy

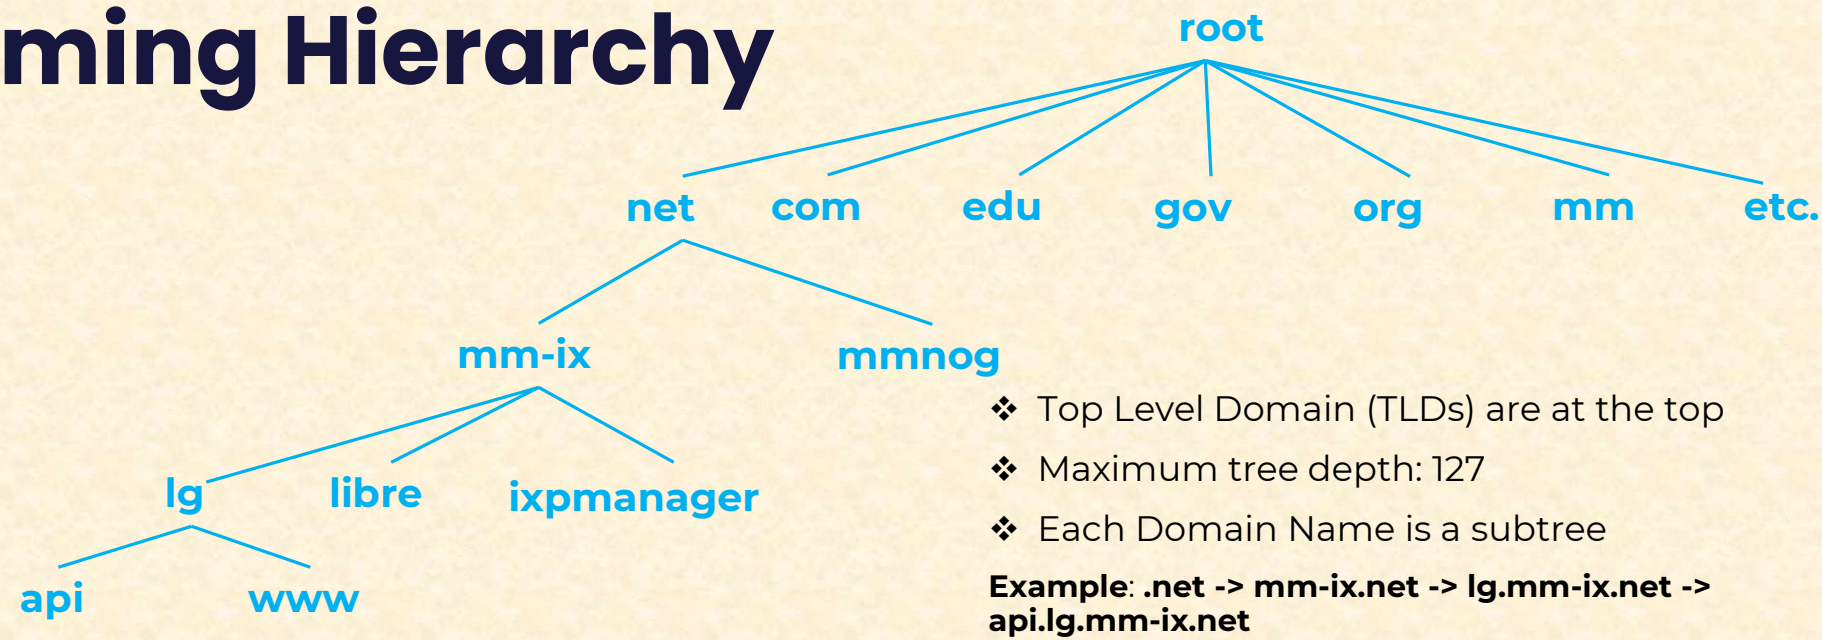

- Top Level Domain (TLDs) are at the top
- Maximum tree depth: 127
- Each Domain Name is a subtree

Example: **.net -> mm-ix.net -> lg.mm-ix.net -> api.lg.mm-ix.net**

- Name collisions are avoided

- Tree is divided into Zones
  - Each zone has an administrator
  - Responsible for the part of the hierarchy

Example:

MMIX controls *.mm-ix.net
MMNOG controls *.mmnog.net

Naming Hierarchy

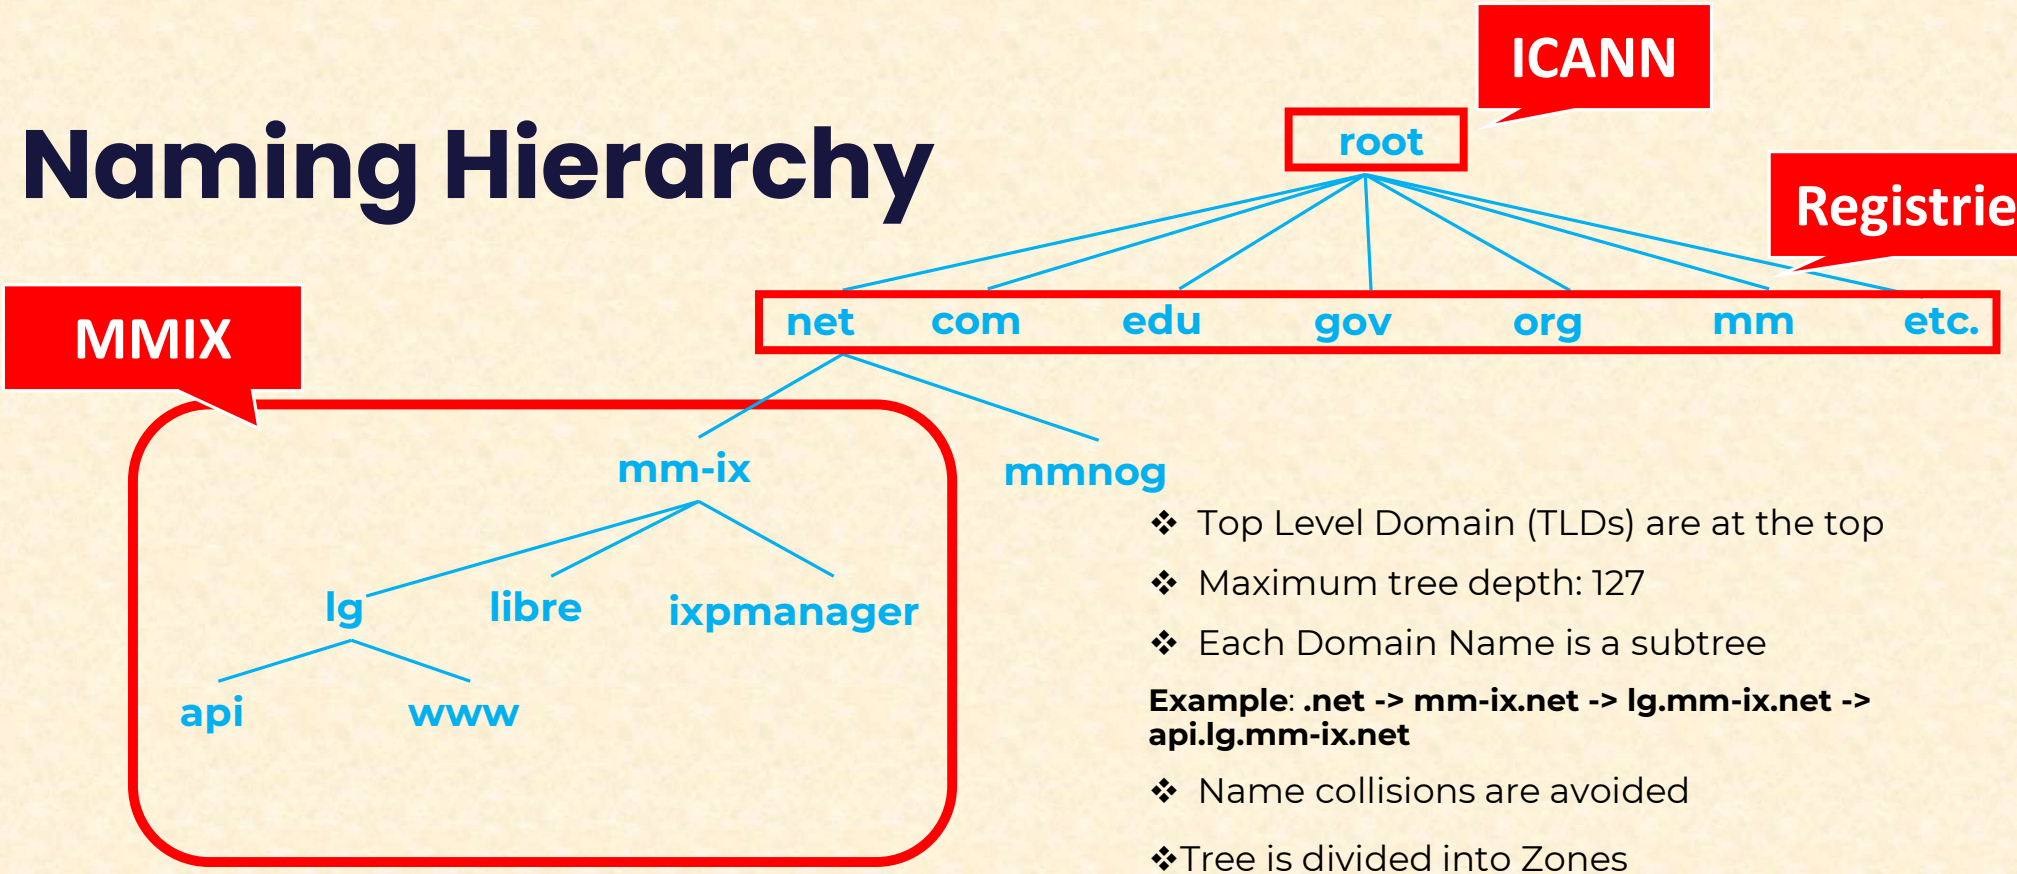

- Top Level Domain (TLDs) are at the top
- Maximum tree depth: 127
- Each Domain Name is a subtree

Example: **.net -> mm-ix.net -> lg.mm-ix.net -> api.lg.mm-ix.net**

- Name collisions are avoided
- Tree is divided into Zones
  - Each zone has an administrator
  - Responsible for the part of the hierarchy

Example:

MMIX controls *.mm-ix.net
MMNOG controls *.mmnog.net

Root Name Servers

- Responsible for the Root Zone file
  - Lists of the TLDs and who control them
  - ~272KB in size
  - 13 root servers, labled A -> M
  - All are anycasted, i.e. they are globally replicated

com.	172800	IN	NS	a.gtld-servers.net.
com.	172800	IN	NS	b.gtld-servers.net.
com.	172800	IN	NS	c.gtld-servers.net.

Authoritative Name Servers

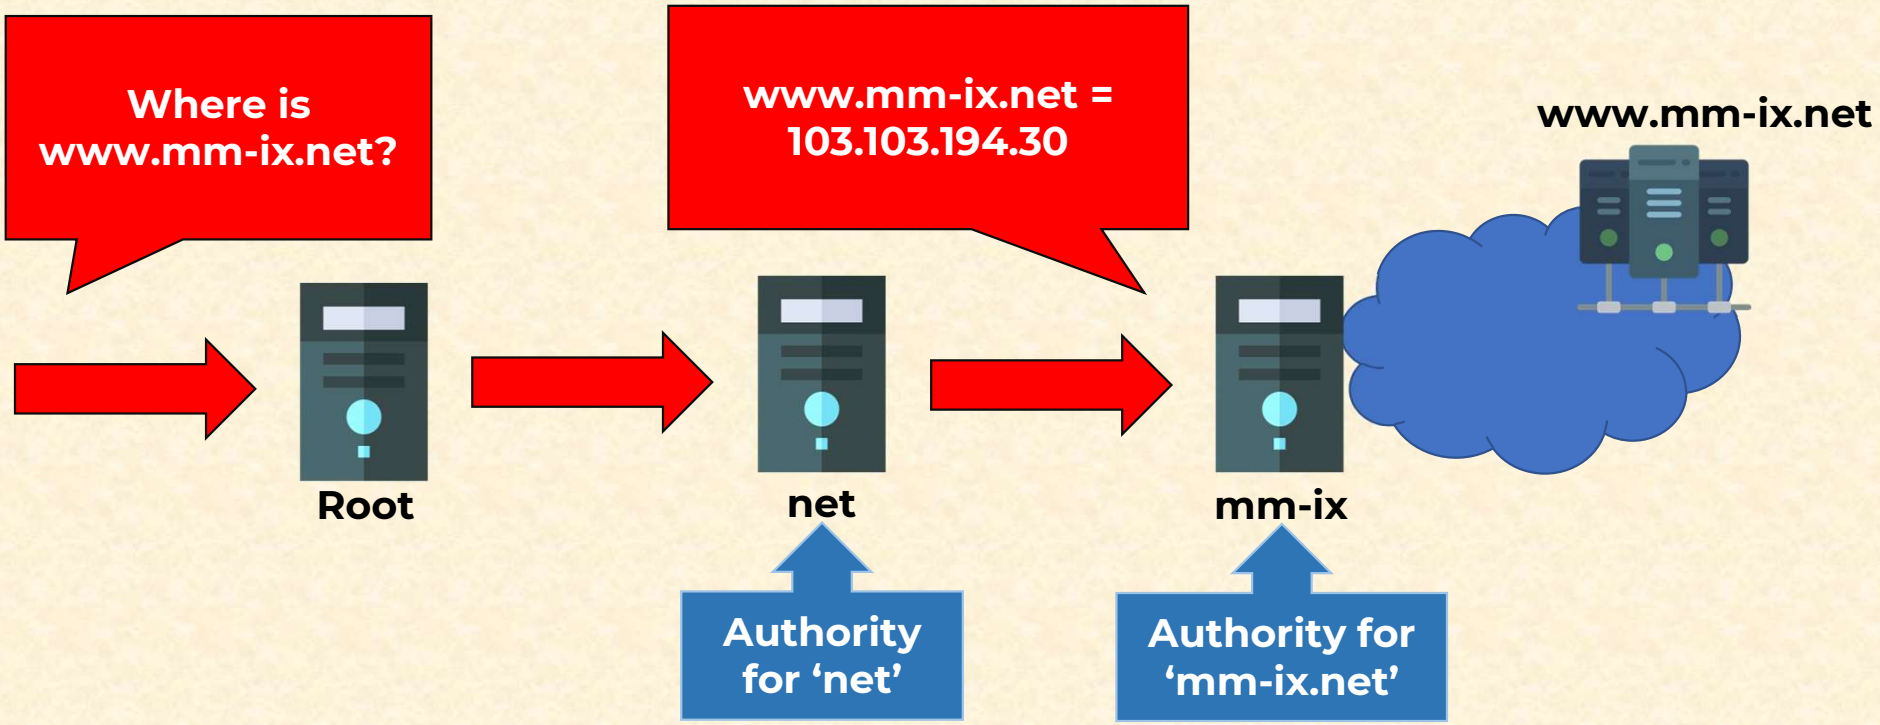

The authoritative name server will have the IP information

❑ Stores the name -> IP mapping for a given host

What is Anycast?

- Network routing method!
  - Multiple routing paths to a group of endpoints that are each assigned the same IP address.

- Route to one of several destinations / one-to-one-of-many association.
  - Routing is determined by one of two schemes:

    Network Layer Anycast scheme: the router selects a destination optimal for the user and provider, based on number of hops.

    Application Layer Anycast scheme: the router may also take into account additional calculations, such as server availability, time to response, number of connections, and so on.

Why use Anycast with DNS?

- With Anycast DNS, we can enable a group of DNS servers with single IP address, to respond to DNS queries based on the geographical location of a DNS client.

- Benefits of the DNS anycasting:
  - Enhancing DNS response time
  - Simplify DNS client settings
  - Extra layer of redundancy
  - Protect against DNS denial of service attacks

Anycast for Authoritative Name Servers

- Many registrars, enterprise providers, and hosting companies provide authoritative DNS services that host the DNS records for top, second, and third-level domains, as well as deeper subdomains.

- Using anycast, recursive lookup requests are resolved by the nearest authoritative DNS server, ensuring the lowest possible latency.

Anycasting the Authoritative Name Servers

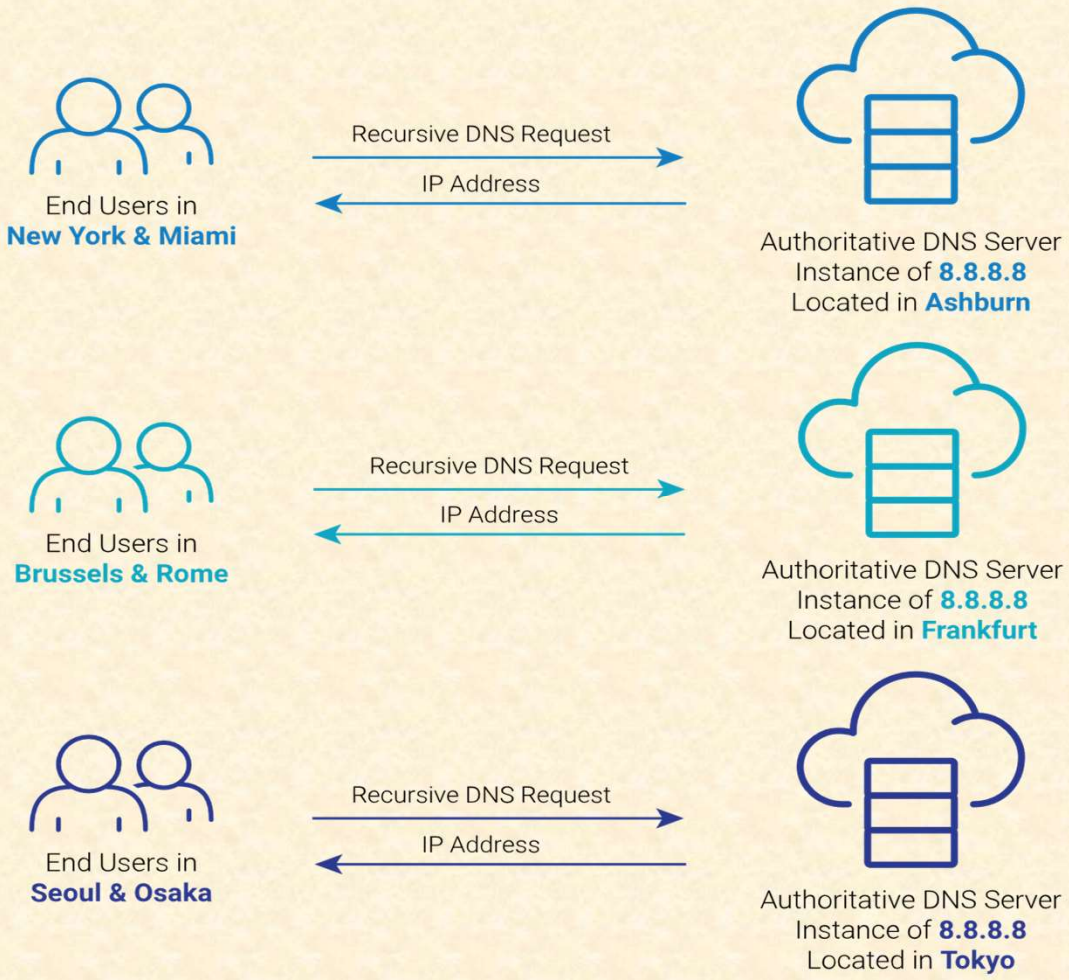

Anycast for Recursive DNS resolver

- ISPs can use anycast for recursive DNS resolver to improve the speed and reliability of their customers' Internet browsing experience.

- By anycasting, ISPs can allow all of their customers to configure a single IP address that will reach the nearest recursive DNS server, and seamlessly failover to the next closest server in the event of a location fails or is taken offline.

Anycasting of the well-known public recursive DNS service

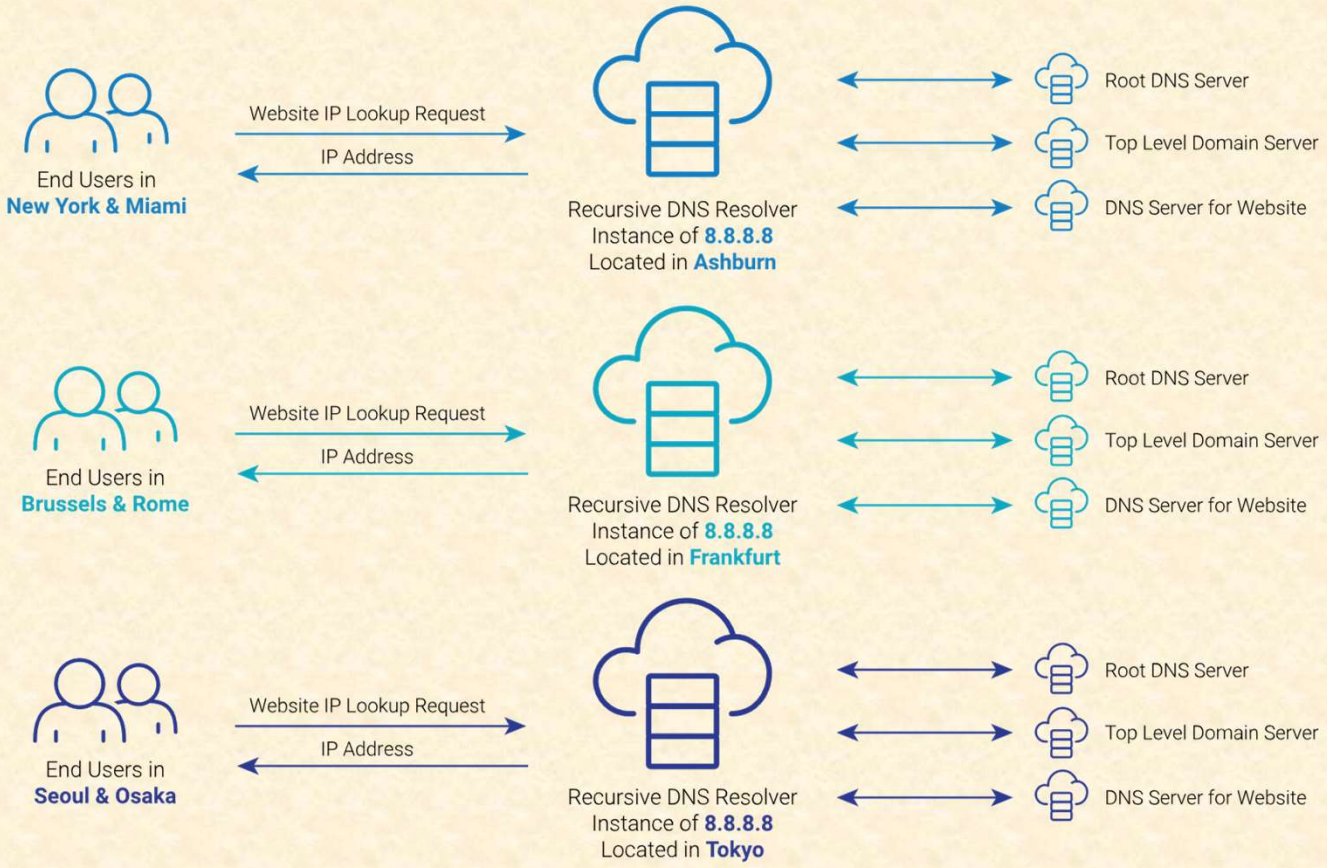

MMIX Anycast Authoritative DNS Service

Vision:

- **.mm** domains localization.

Type:

- Authoritative DNS

Commercial:

- Free hosting for .mm registrars and also individual registrants.

Topology:

- Anycast; distributed nodes around the world, more nodes in local.

AnyCast – Advertising All IPs to Users from Various Locations

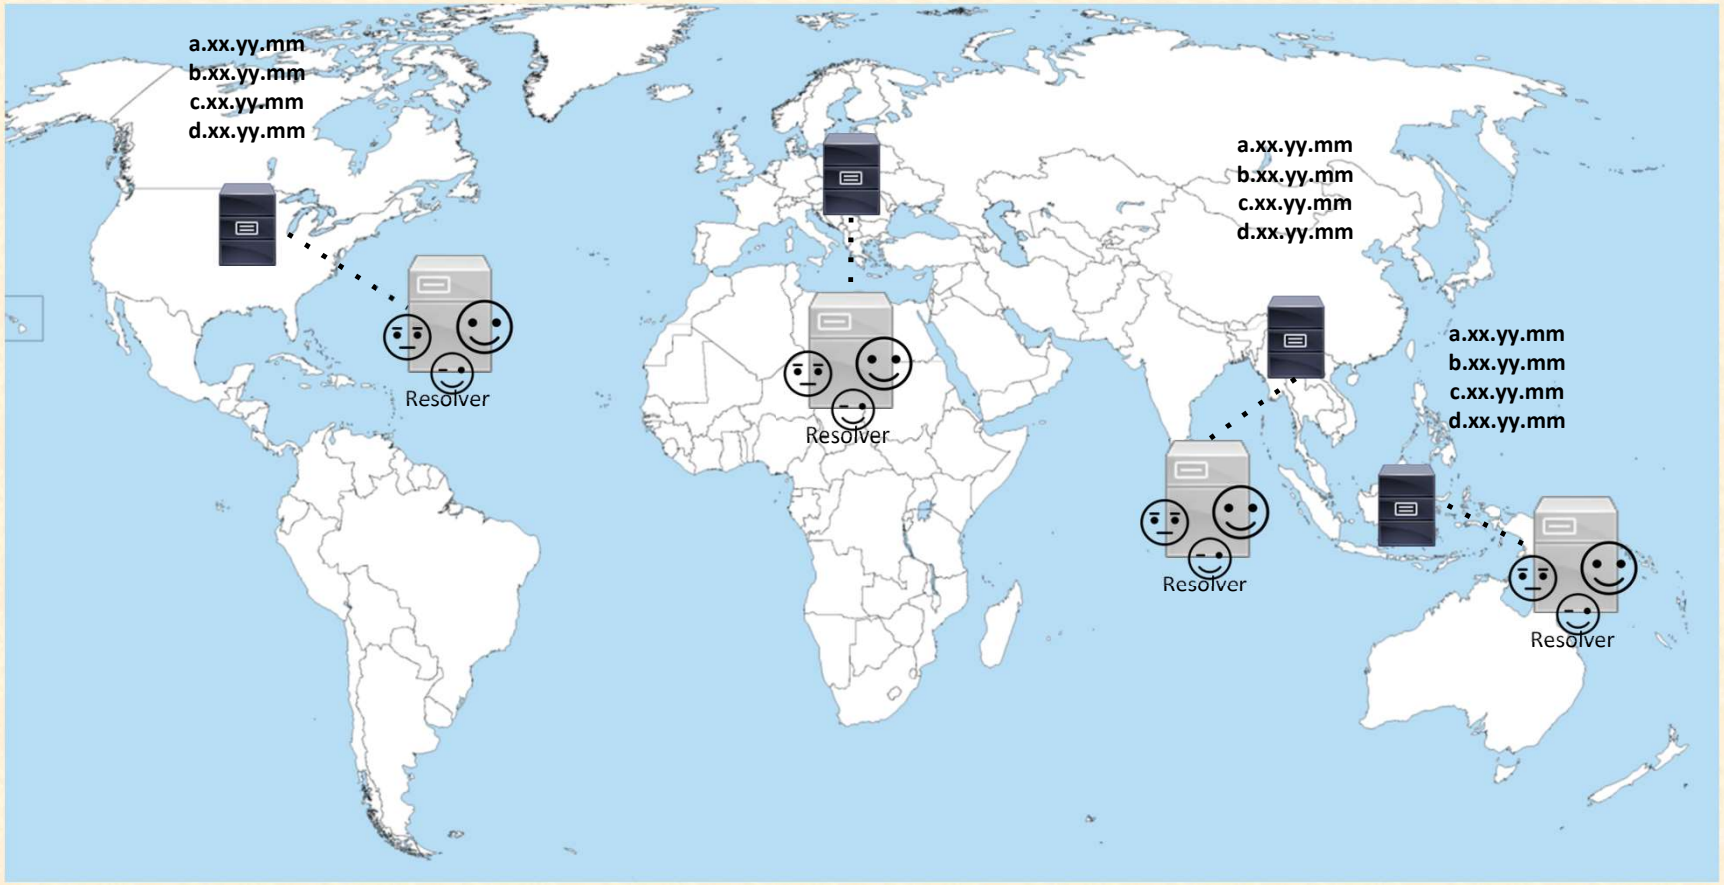

Phase I

Distributed nodes

1. One (1) node at MMIX, Yangon.
2. One (1) node at MMIX, Mandalay
3. At Least 2 nodes at local ISPs.

Service

Available for all **.mm** registry and registrars.

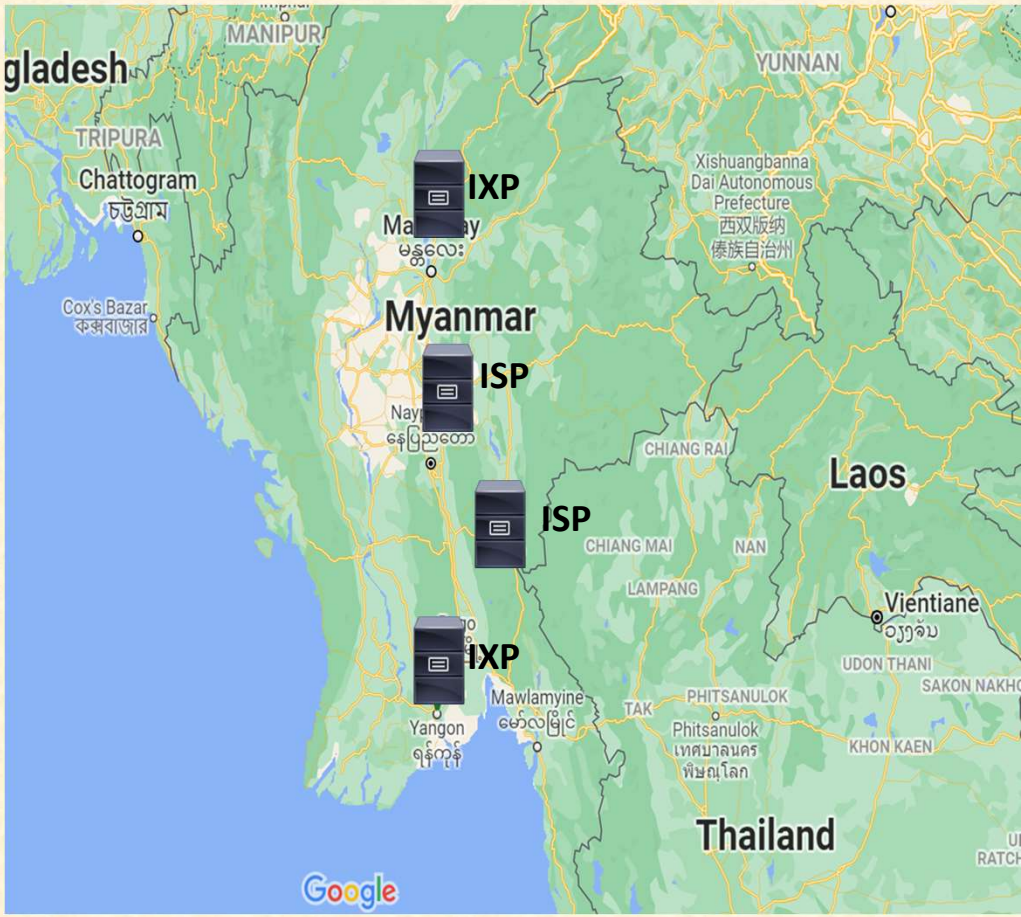

Phase II

Distributed nodes

- Same or more local nodes as Phase I
- Additional at least 2 more nodes on the cloud (Oversea)

Service

- Same services as Phase II
- DNS hosting available for Individual Registrants also.

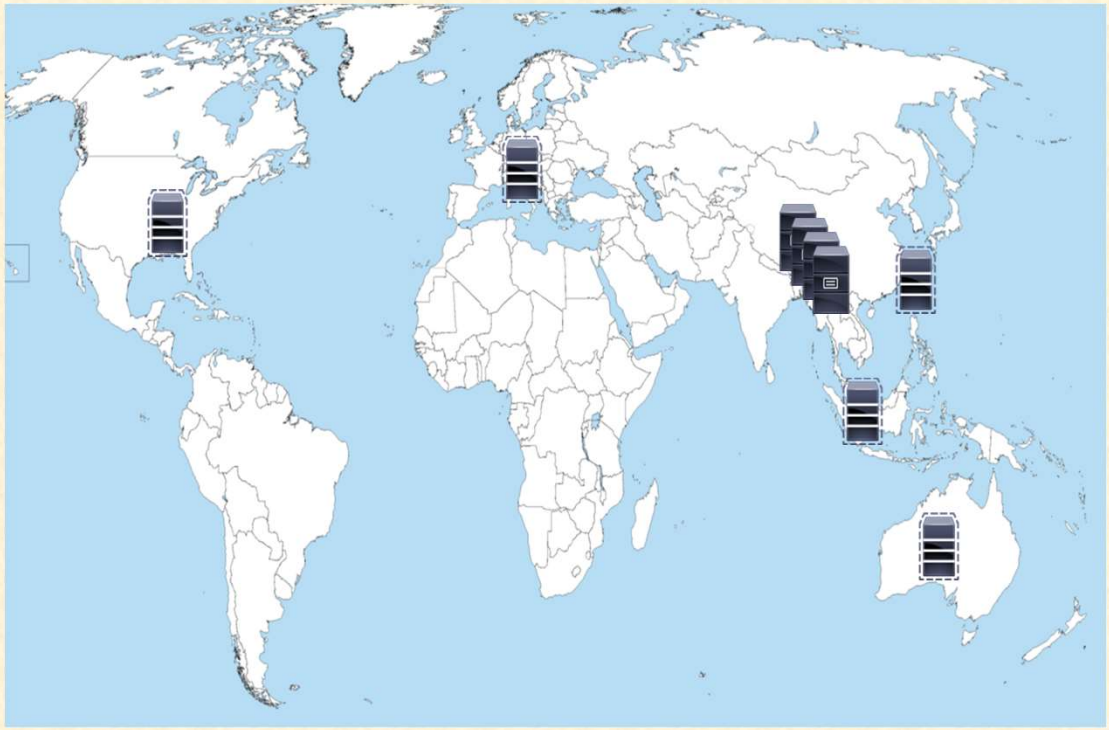

MMIX Anycast Structure

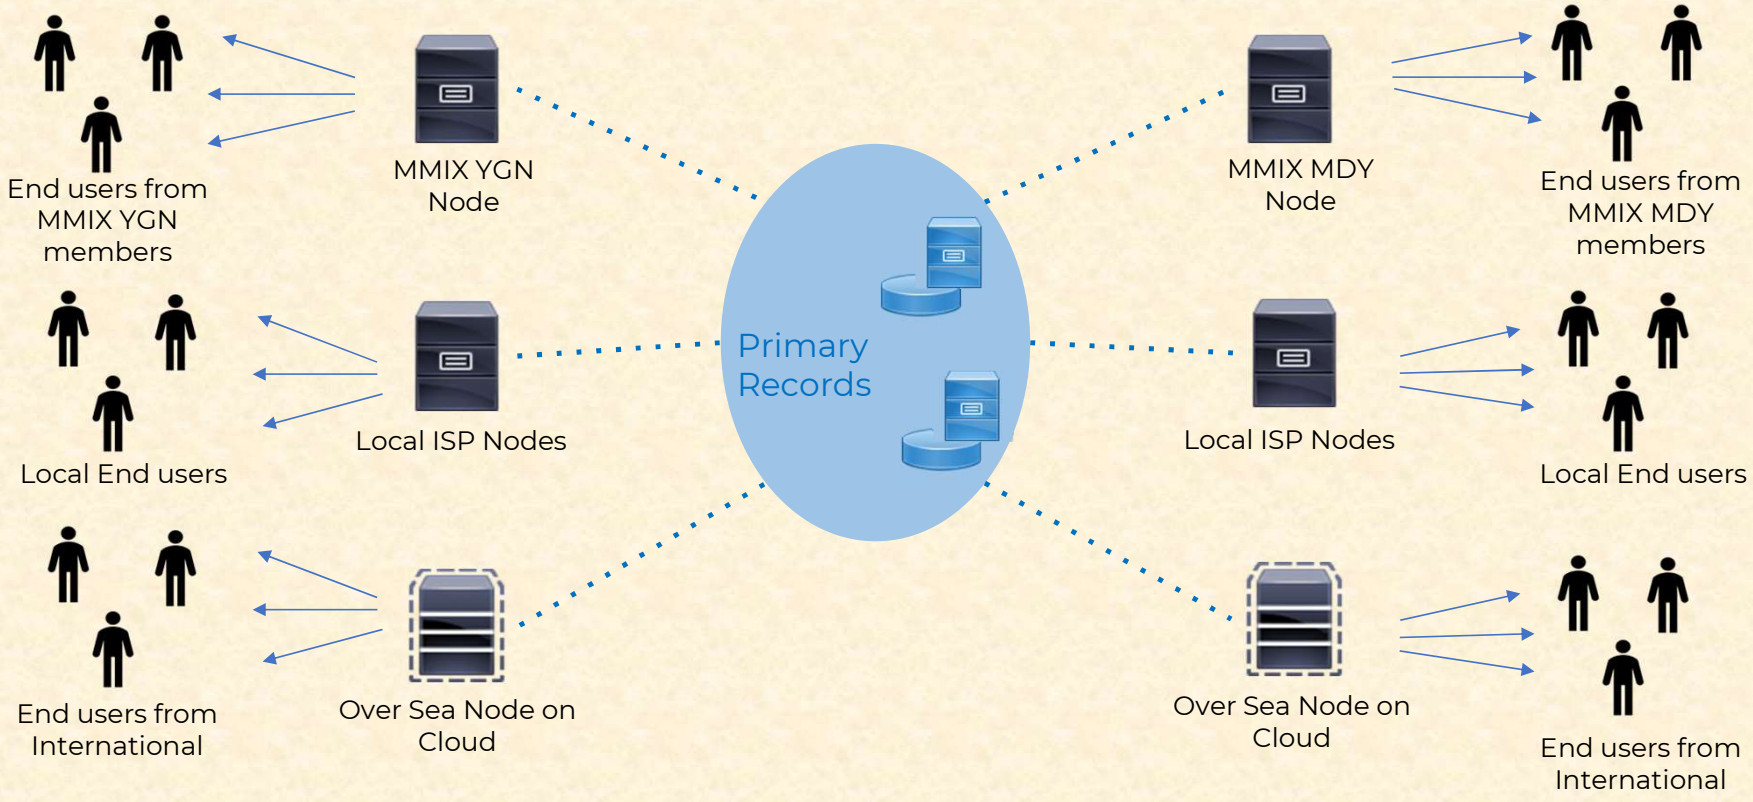

How it's work!

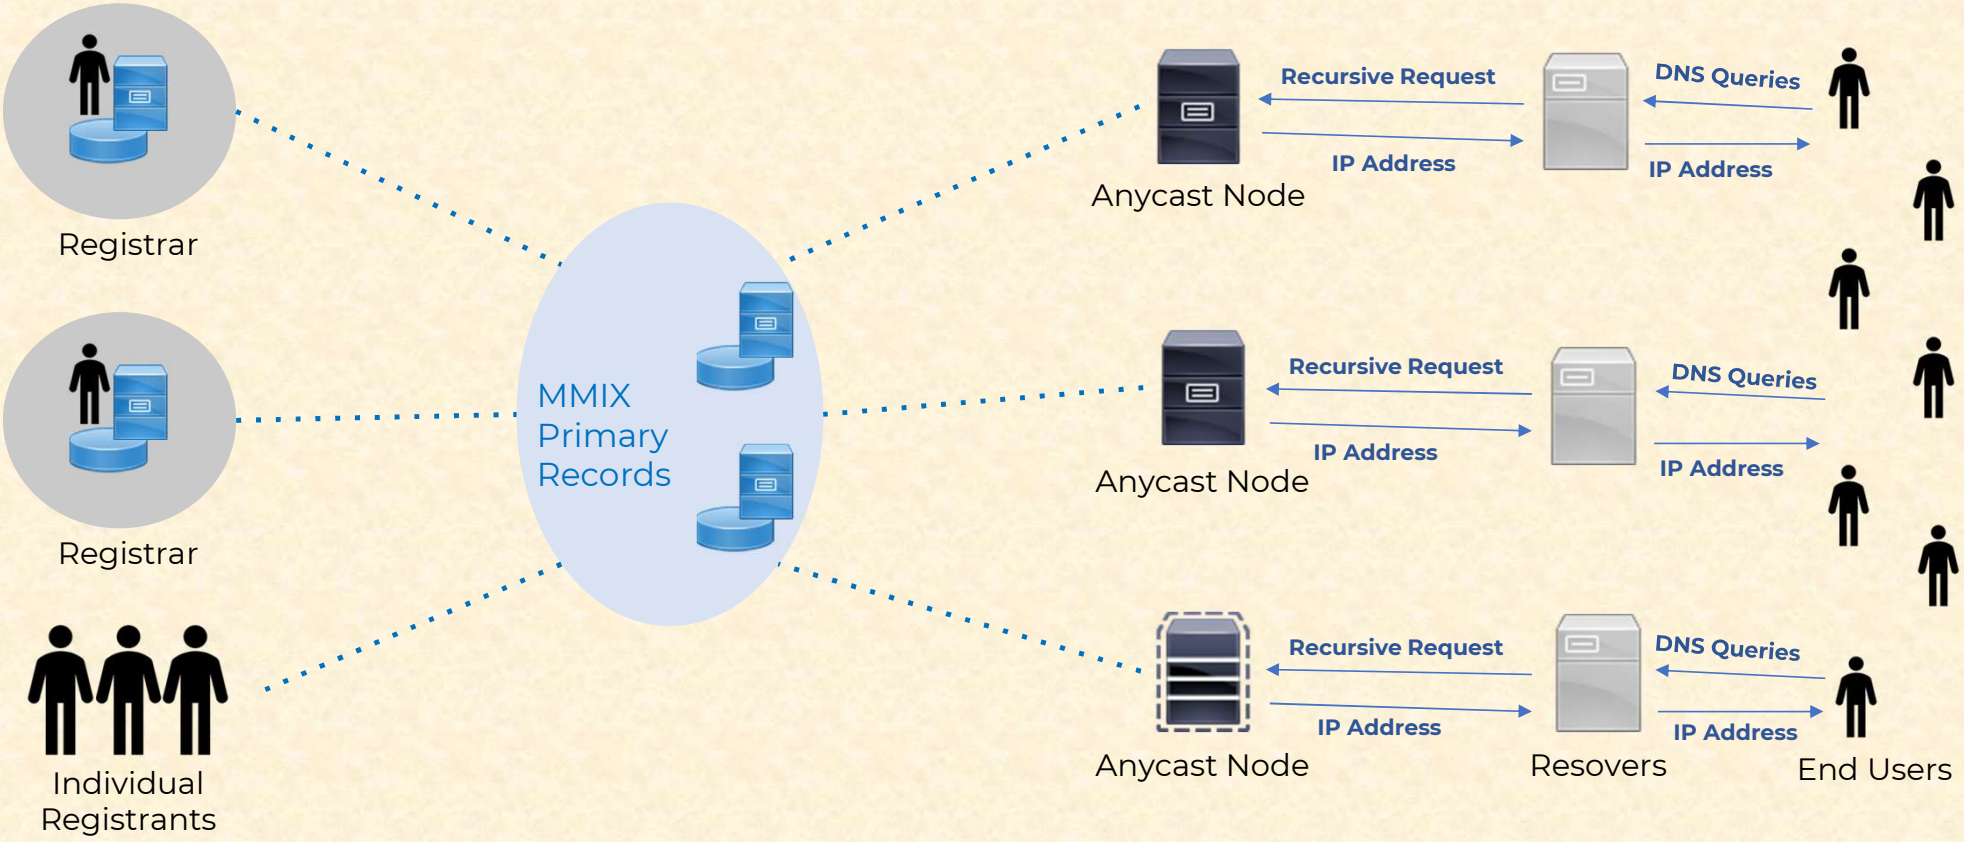
Registrar
Registrar
Individual Registrants
MMIX Primary Records
Anycast Node
Anycast Node
Anycast Node
Recursive Request
IP Address
Recursive Request
IP Address
Recursive Request
IP Address
DNS Queries
IP Address
DNS Queries
IP Address
DNS Queries
IP Address
Resovers
End Users

DNS@MMIX

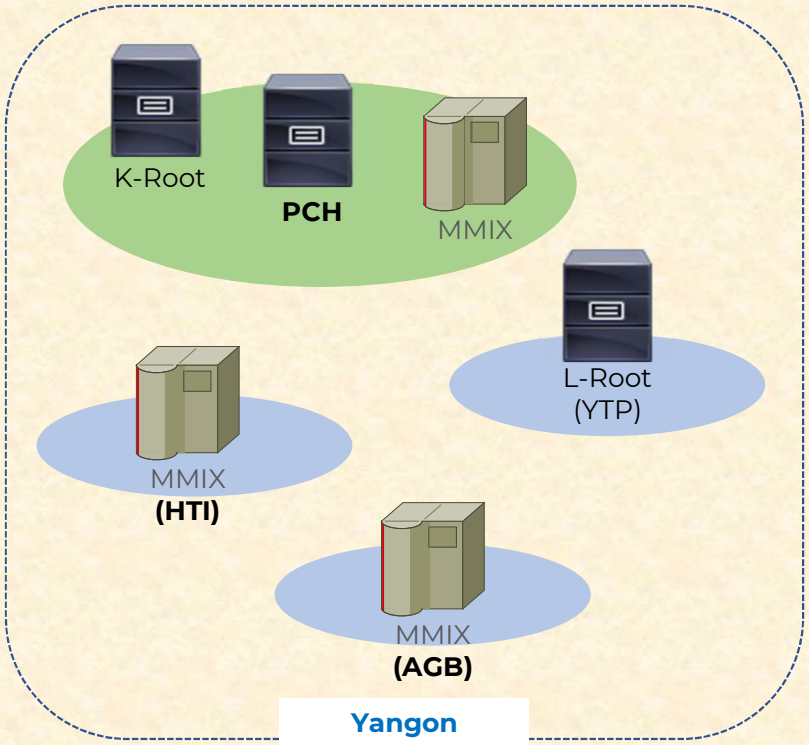
K-Root
PCH
MMIX
L-Root
(YTP)
MMIX
(HTI)
MMIX
(AGB)
Yangon

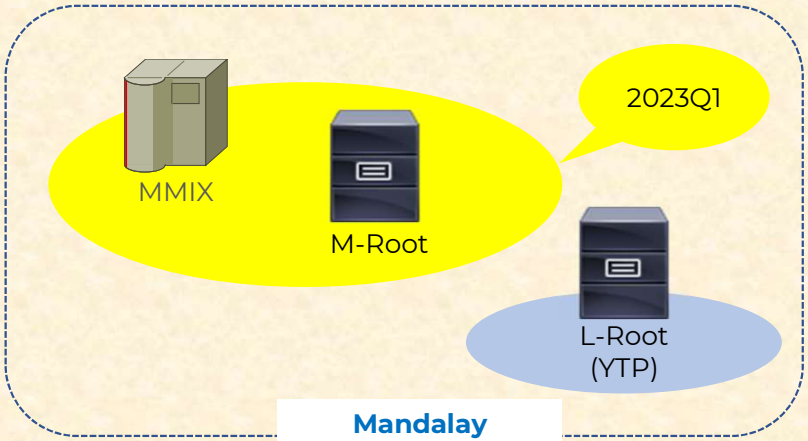
MMIX
M-Root
2023Q1
L-Root
(YTP)
Mandalay

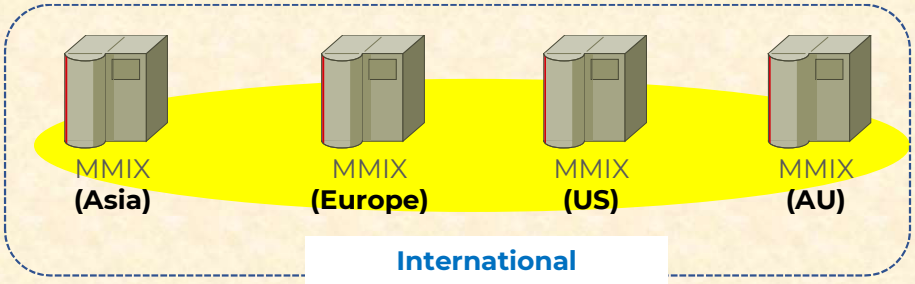
MMIX
(Asia)
MMIX
(Europe)
MMIX
(US)
MMIX
(AU)
International

Registrar with owned NS

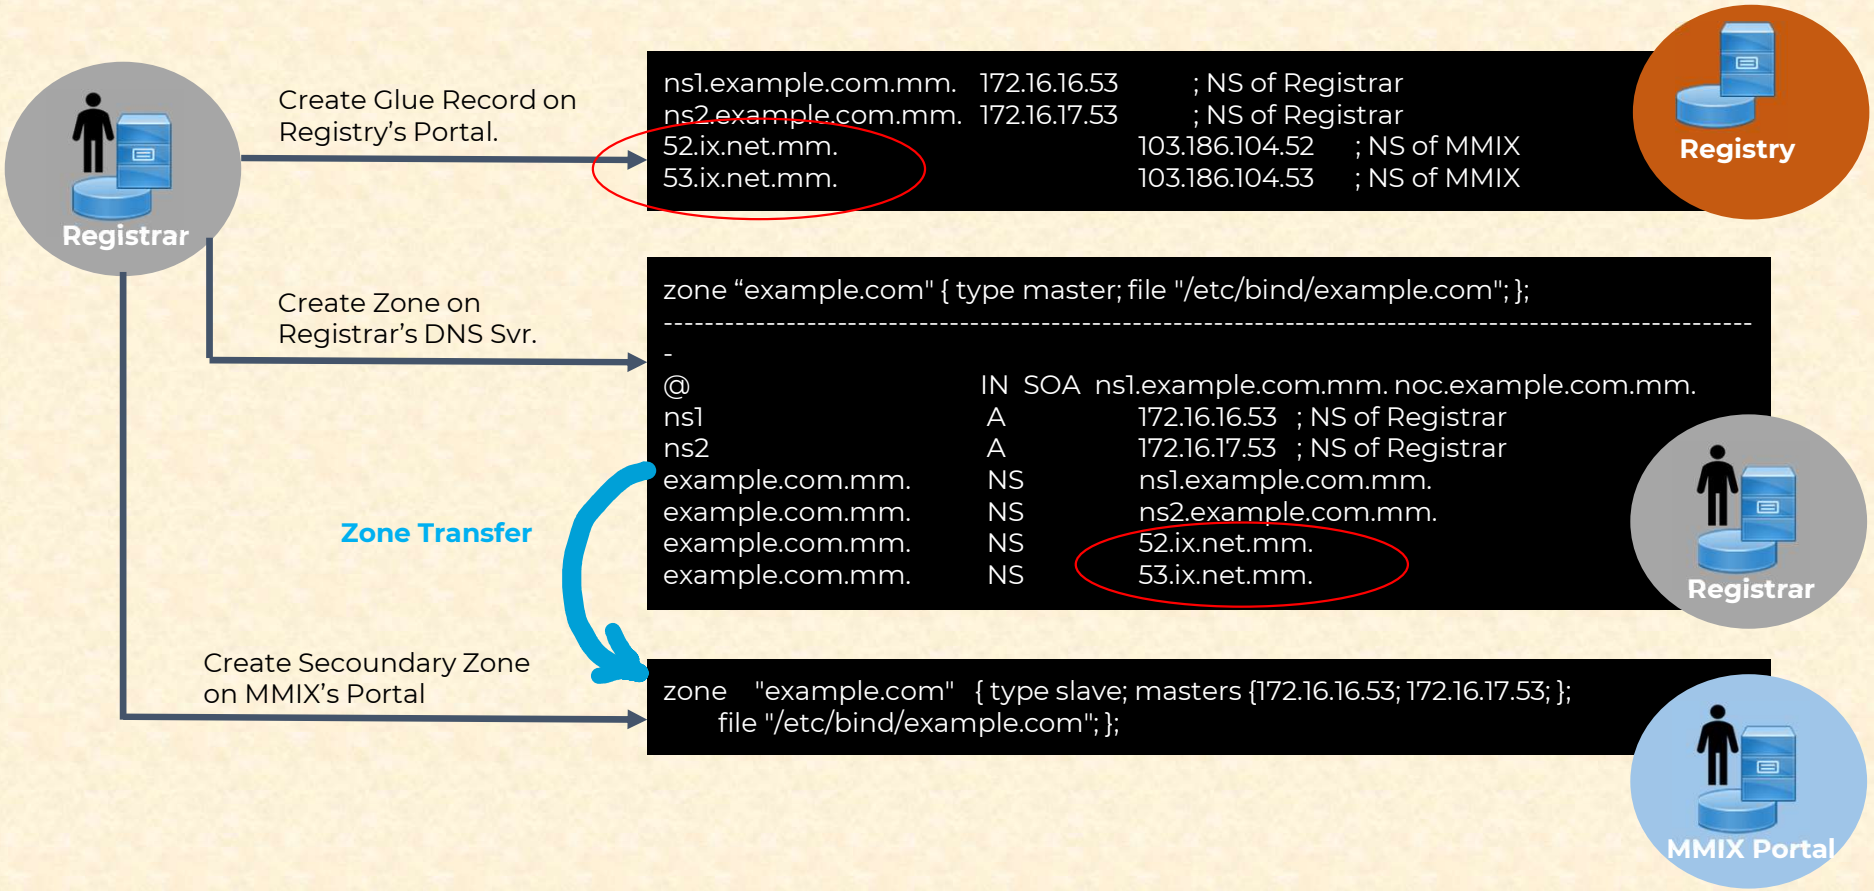

Registrar without owned NS

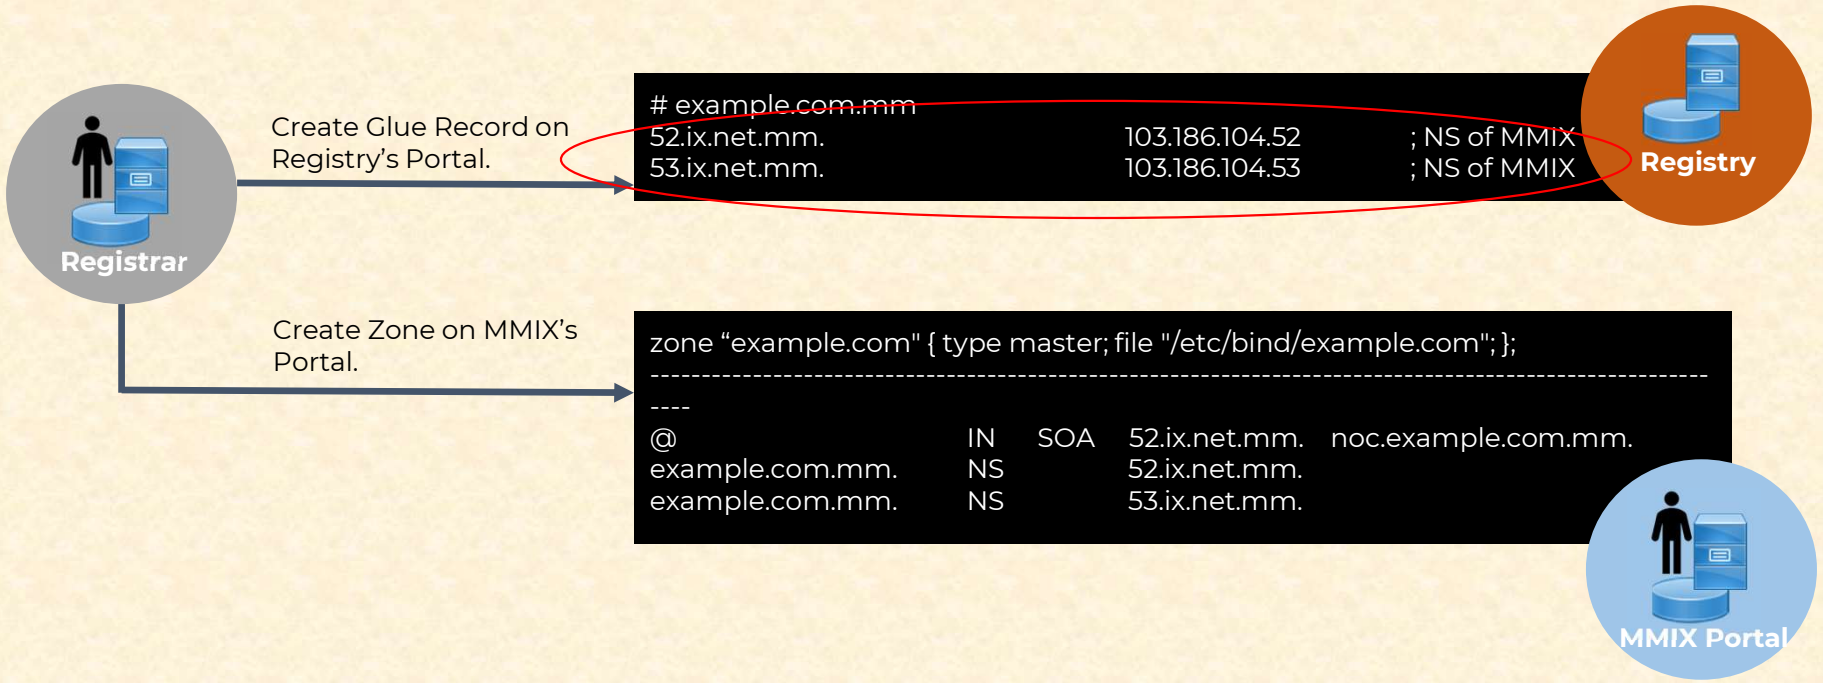

Registrant

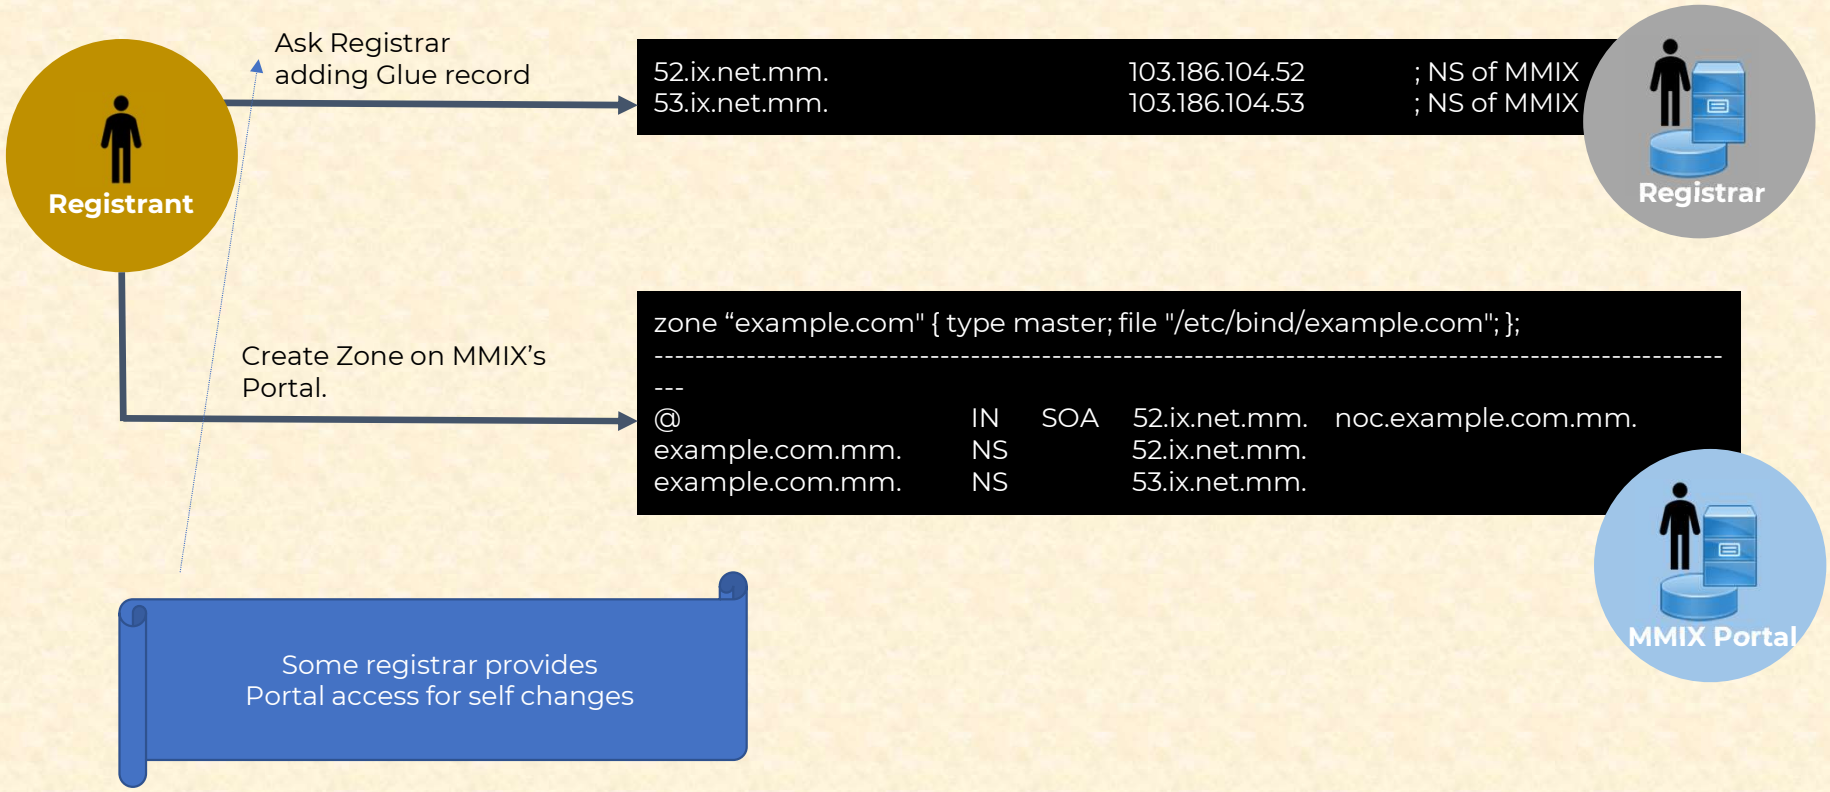
Registrant
Ask Registrar
adding Glue record
52.ix.net.mm. 103.186.104.52 ; NS of MMIX
53.ix.net.mm. 103.186.104.53 ; NS of MMIX
Registrar
Create Zone on MMIX's
Portal.
zone "example.com" { type master; file "/etc/bind/example.com"; };

@ IN SOA 52.ix.net.mm. noc.example.com.mm.
example.com.mm. NS 52.ix.net.mm.
example.com.mm. NS 53.ix.net.mm.
MMIX Portal
Some registrar provides
Portal access for self changes

Registrant: Create Reversed Zone

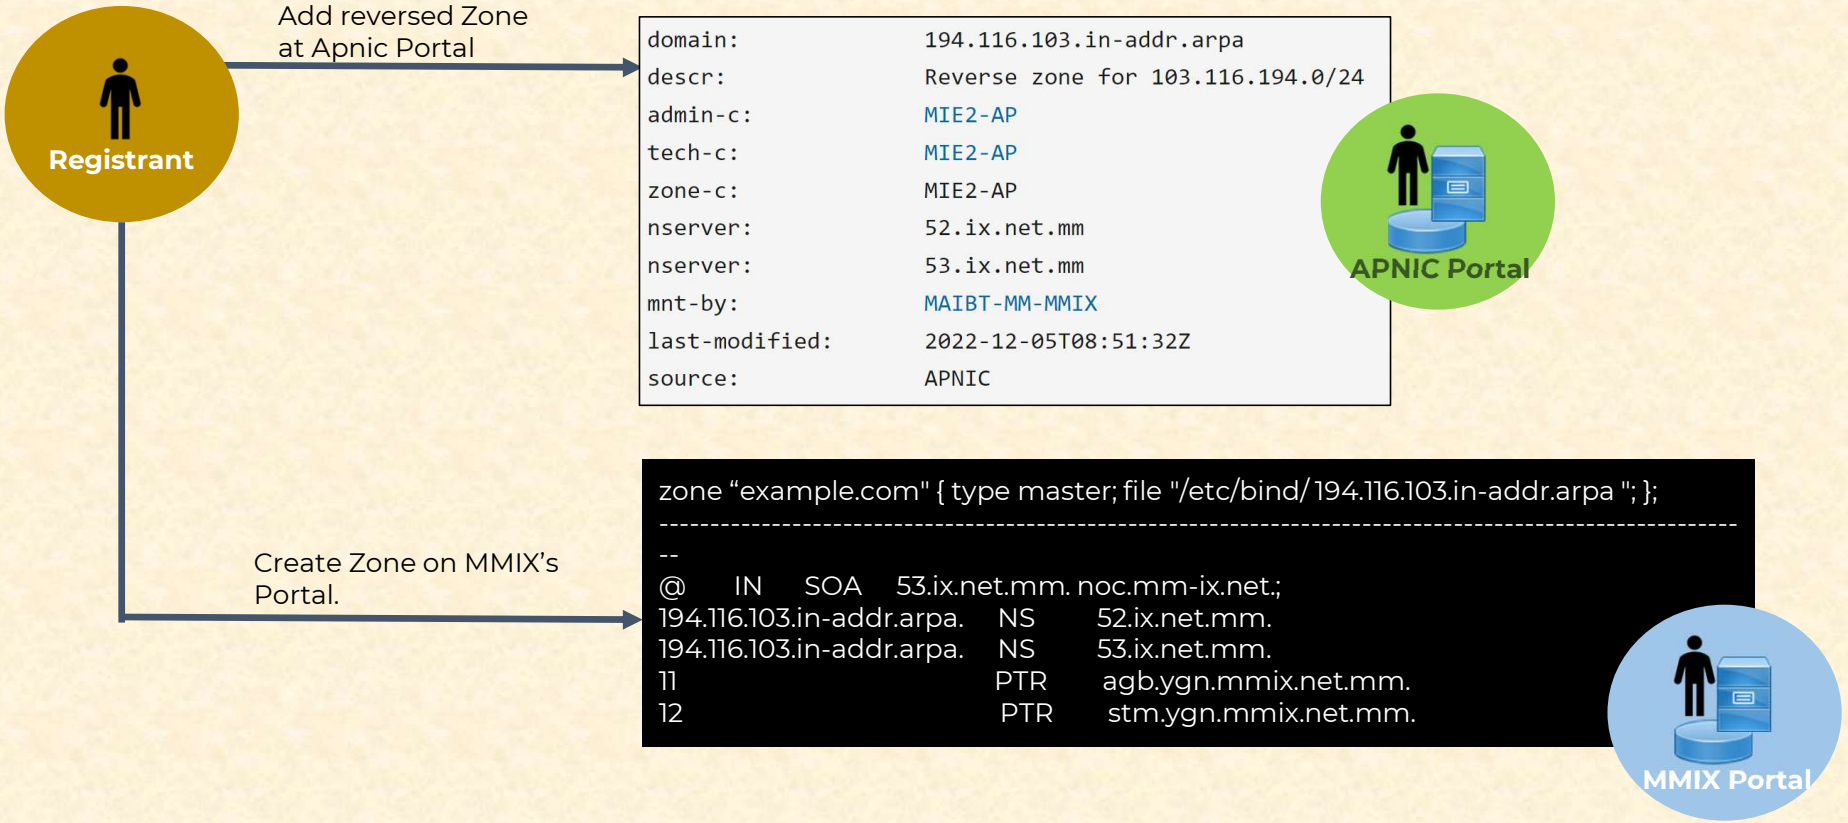

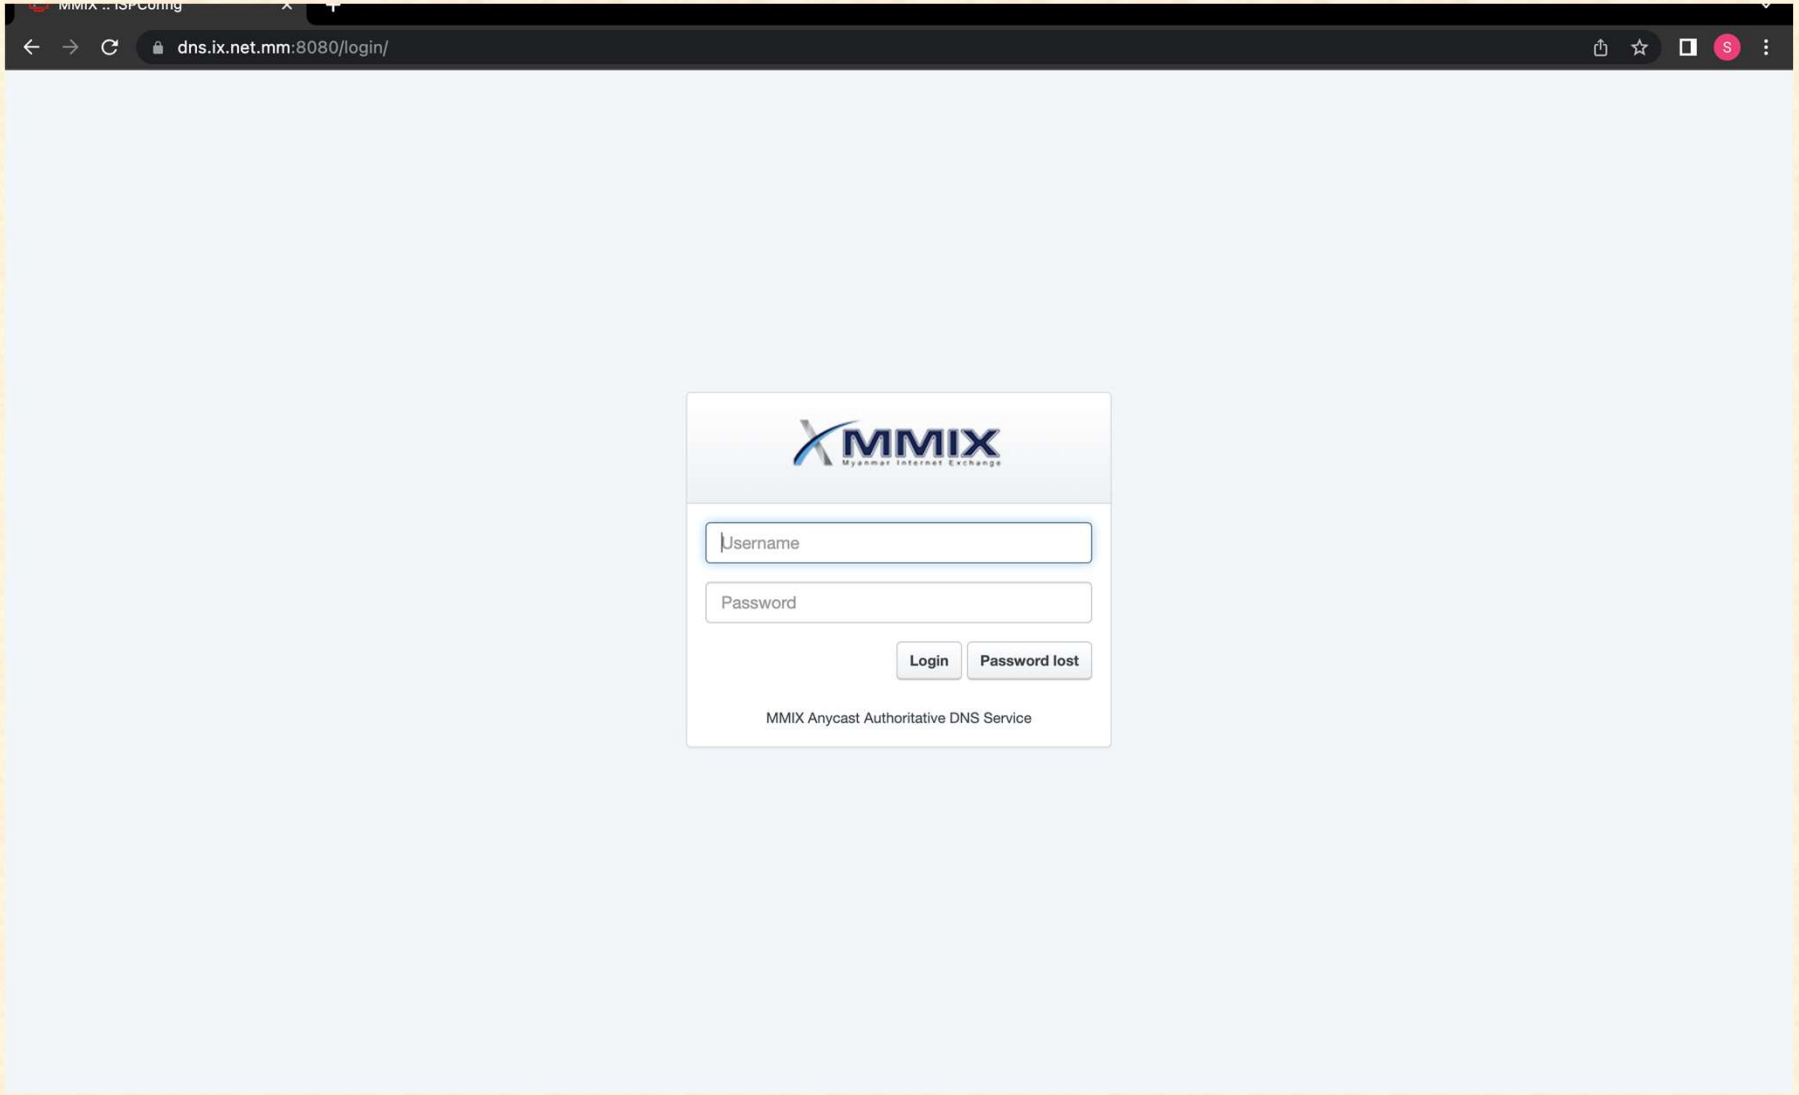
dns.ix.net.mm:8080/login/
MMIX
Myanmar Internet Exchange
Username
Password
Login
Password lost
MMIX Anycast Authoritative DNS Service

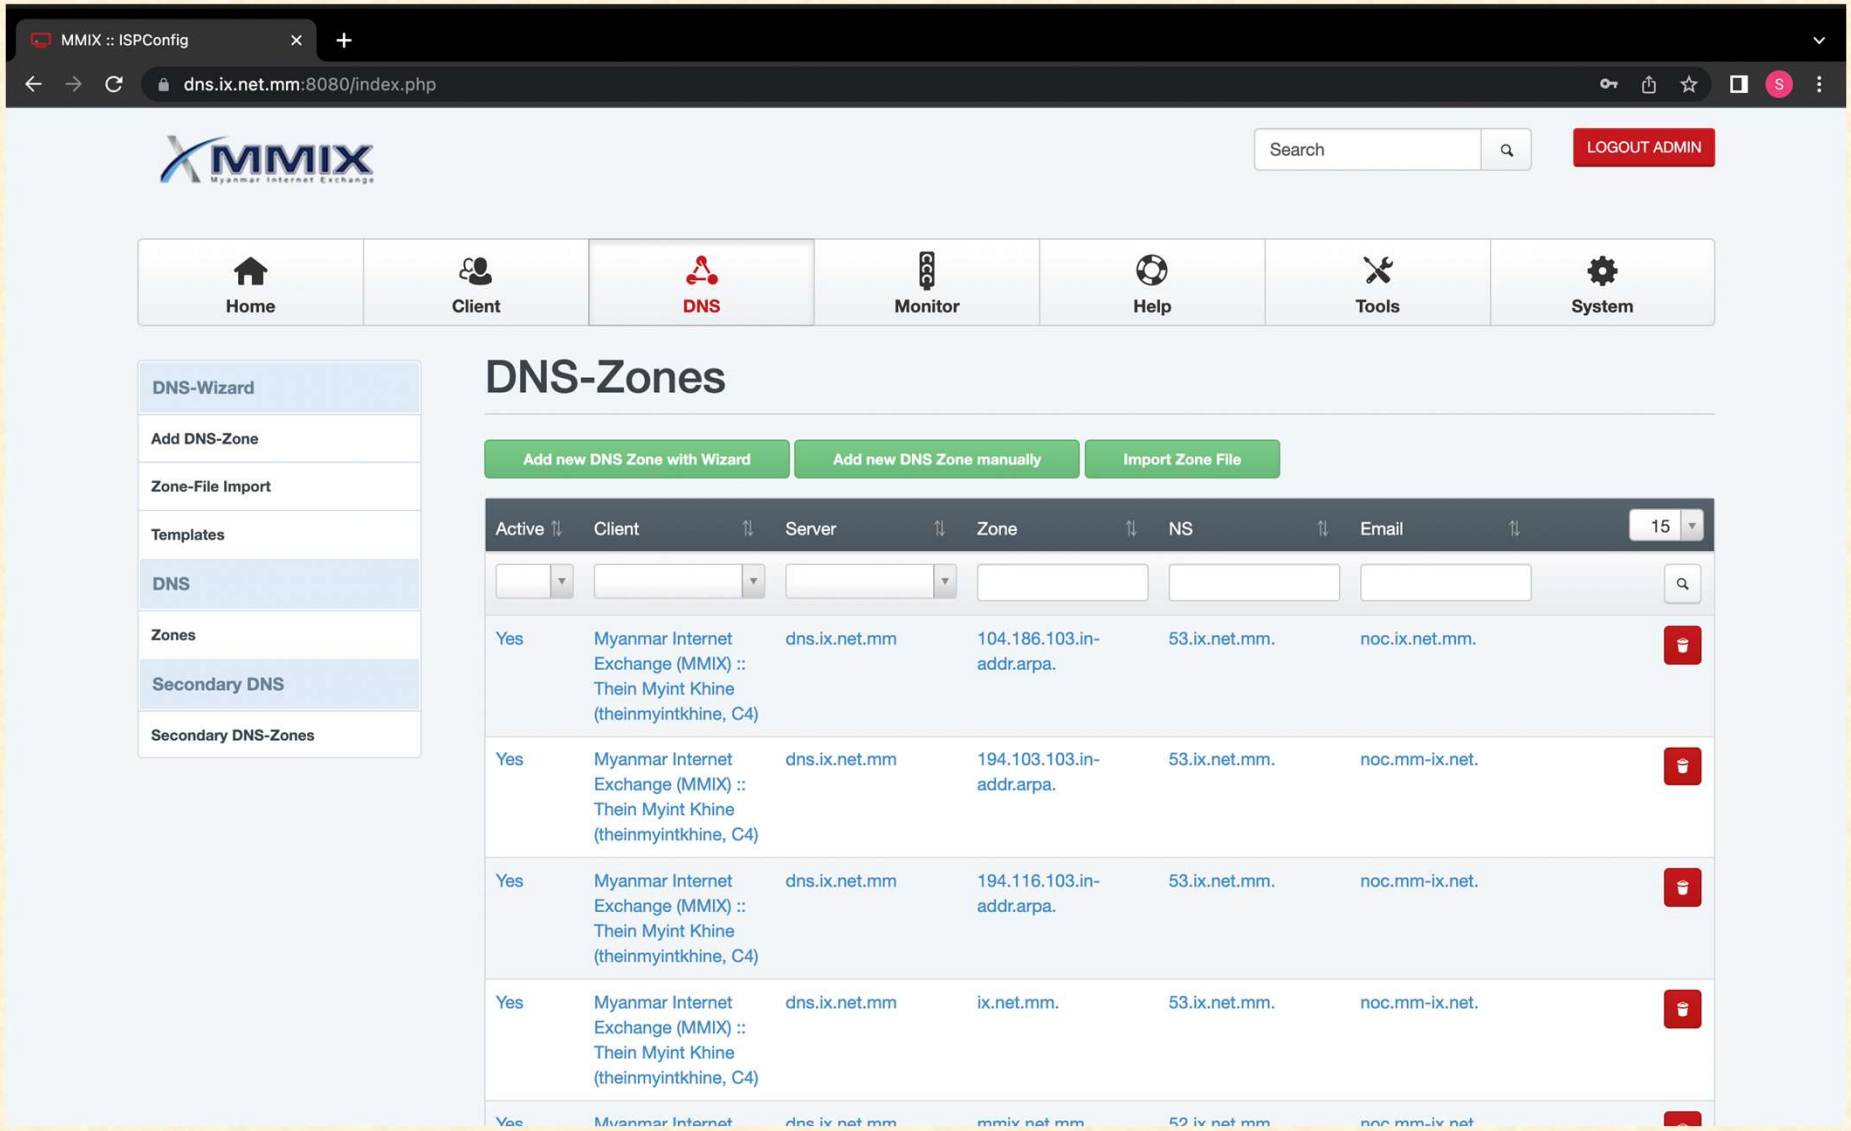
MMIX :: ISPConfig
dns.ix.net.mm:8080/index.php
MMIX
Myanmar Internet Exchange
Search
LOGOUT ADMIN
Home
Client
DNS
Monitor
Help
Tools
System
DNS-Wizard
Add DNS-Zone
Zone-File Import
Templates
DNS
Zones
Secondary DNS
Secondary DNS-Zones
DNS-Zones
Add new DNS Zone with Wizard
Add new DNS Zone manually
Import Zone File
Active
Client
Server
Zone
NS
Email
15
Yes
Myanmar Internet Exchange (MMIX) :: Thein Myint Khine (theinmyintkhine, C4)
dns.ix.net.mm
104.186.103.in-addr.arpa.
53.ix.net.mm.
noc.ix.net.mm.
Yes
Myanmar Internet Exchange (MMIX) :: Thein Myint Khine (theinmyintkhine, C4)
dns.ix.net.mm
194.103.103.in-addr.arpa.
53.ix.net.mm.
noc.mm-ix.net.
Yes
Myanmar Internet Exchange (MMIX) :: Thein Myint Khine (theinmyintkhine, C4)
dns.ix.net.mm
194.116.103.in-addr.arpa.
53.ix.net.mm.
noc.mm-ix.net.
Yes
Myanmar Internet Exchange (MMIX) :: Thein Myint Khine (theinmyintkhine, C4)
dns.ix.net.mm
ix.net.mm.
53.ix.net.mm.
noc.mm-ix.net.

Example: Anycast for the recursive DNS resolver

Software Resources	Hardware Resources
CentOS 7.5 64 bit/ Ubuntu 20.04 64 bit	CPU Core – 4 with 2 Socket
rpcbind-0.2.0-44.el7.x86_64	RAM – 8 GB
bind-chroot-9.9.4-61.el7.x86_64	HDD – Sata SAS 15k RPM
bind-license-9.9.4-61.el7.noarch	
bind-utils-9.9.4-61.el7.x86_64	
bind-9.9.4-61.el7.x86_64	
bind-libs-lite-9.9.4-61.el7.x86_64	
bind-libs-9.9.4-61.el7.x86_64	
iptables-1.4.7-16.el6.x86_64	
iptables-ipv6-1.4.7-16.el6.x86_64	
quagga-0.99.22.4-5.el7_4.x86_64	

Assigned anycast address

Anycast address as an additional loopbacks

```
[root@dc-anycast-dns network-scripts]# ifconfig lo:0

lo:0: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 179.100.0.254  netmask 255.255.255.255
        loop  txqueuelen 1  (Local Loopback)
```

Configure the Name Service

Configuring named service to listen on anycast address

```
[root@dc-anycast-dns etc]# vim /var/named/chroot/etc/named.conf
options {
        listen-on port 53 { 127.0.0.1; 179.100.0.254; };
        directory       "/var/named";
        dump-file       "/var/named/data/cache_dump.db";
        statistics-file "/var/named/data/named_stats.txt";
        memstatistics-file "/var/named/data/named_mem_stats.txt";
        allow-query     { localhost; 192.168.0.0/16; };
        allow-query-cache       { localhost; 192.168.0.0/16; };
        allow-recursion { localhost; 192.168.0.0/16;  };
        version "go to sleep" ;
        recursive-clients 100000;
};
```

Configure Quagga & BGP

Configuring zebra.conf

```
[root@dc-anycast-dns quagga]# # vim /etc/quagga/zebra.conf

hostname dc-anycast-dns.link3.net
!
enable password NothingToSay
!
interface eth0
ip address 192.168.0.226/30
!
interface lo:0
ip address 179.200.0.254/32
!
interface lo
!
line vty
!
```

Configure Quagga & BGP

Configuring bgpd.conf

```
[root@dc-anycast-dns quagga]# vim /etc/quagga/bgpd.conf
hostname dc-anycast-dns.link3.net
password NothingToSay
log stdout
!
router bgp 65430
 network 179.200.0.254/32
 neighbor 192.168.0.225 remote-as 23688
 neighbor 192.168.0.225 description BTS
 neighbor 192.168.0.225 activate
 neighbor 192.168.0.225 next-hop-self
 neighbor 192.168.0.225 remove-private-AS
 neighbor 192.168.0.225 soft-reconfiguration inbound
 neighbor 192.168.0.225 prefix-list anycast out
 neighbor 192.168.0.225 prefix-list default in
!
ip prefix-list default seq 15 permit 0.0.0.0/0
ip prefix-list anycast seq 5 permit  179.200.0.254/32
```

Configure MPE router

Configuring BGP from router

```
router bgp 23688
 network 192.168.0.224 mask 255.255.255.252
 neighbor 192.168.0.226 remote-as 65430
 neighbor 192.168.0.226 description DC-DNS_Anycast-SERVER
 neighbor 192.168.0.226 activate
 neighbor 192.168.0.226 next-hop-self
 neighbor 192.168.0.226 default-originate
 neighbor 192.168.0.226 remove-private-as
 neighbor 192.168.0.226 soft-reconfiguration inbound
 neighbor 192.168.0.226 prefix-list anycast-DNS-in in
 neighbor 192.168.0.226 prefix-list default out
ip prefix-list anycast-DNS-in seq 10 permit 179.200.0.254/32
ip prefix-list default seq 5 permit 0.0.0.0/0
```

Failover anycast nodes

```
#!/bin/bash
DNSUP=`/usr/bin/dig @179.100.0.254 localhost. A +short`
if [ "$DNSUP" != "127.0.0.1" ];
then
echo "Stopping Anycast...."
 /etc/init.d/bgpd stop
 /etc/init.d/zebra stop
 echo "Stopped: DC Anycast DNS has stopped working, BGP has already been shutdown, Please check the system right now."
| mailx -S smtp=smtp.notifcation.net:25 -s "Alert: Stopped - DC Anycast DNS has stooped working" nothing@notifcation.com
else
 echo "Everything's good... Do nothing..."
fi
```

Thanks You