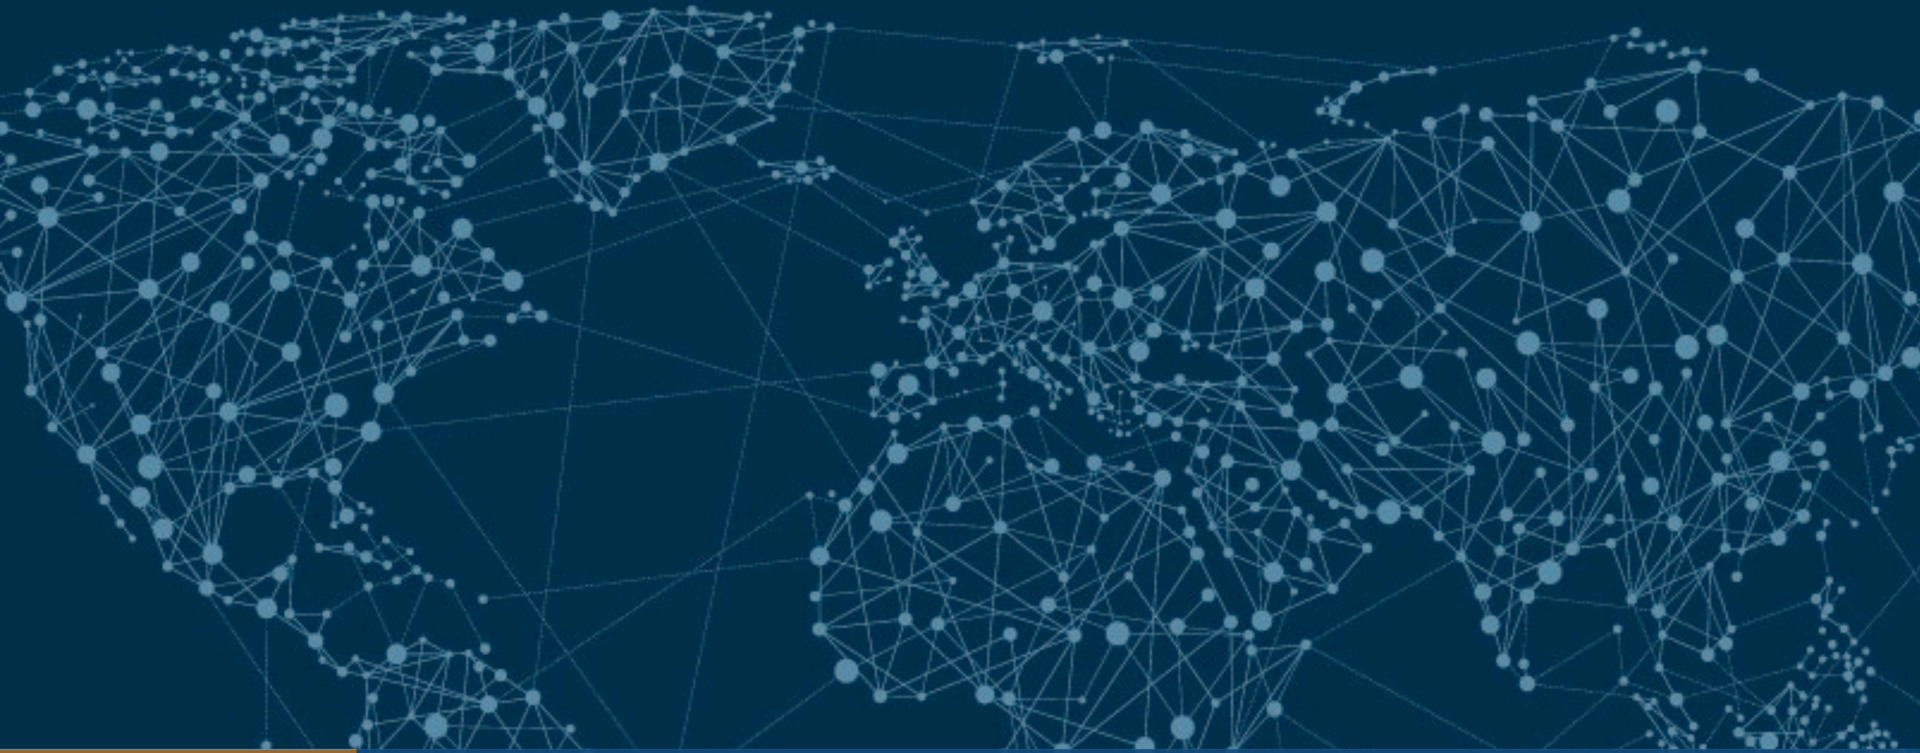

ICANN

Tutorial on Root Server System

Root Server System Advisory Committee | November 2016

Outline

1. Overview of Domain Name System
2. History of Root Server System
3. Root Server System Today & Its Features
4. Explanation of Anycast
5. RSSAC Publications and Current Activities

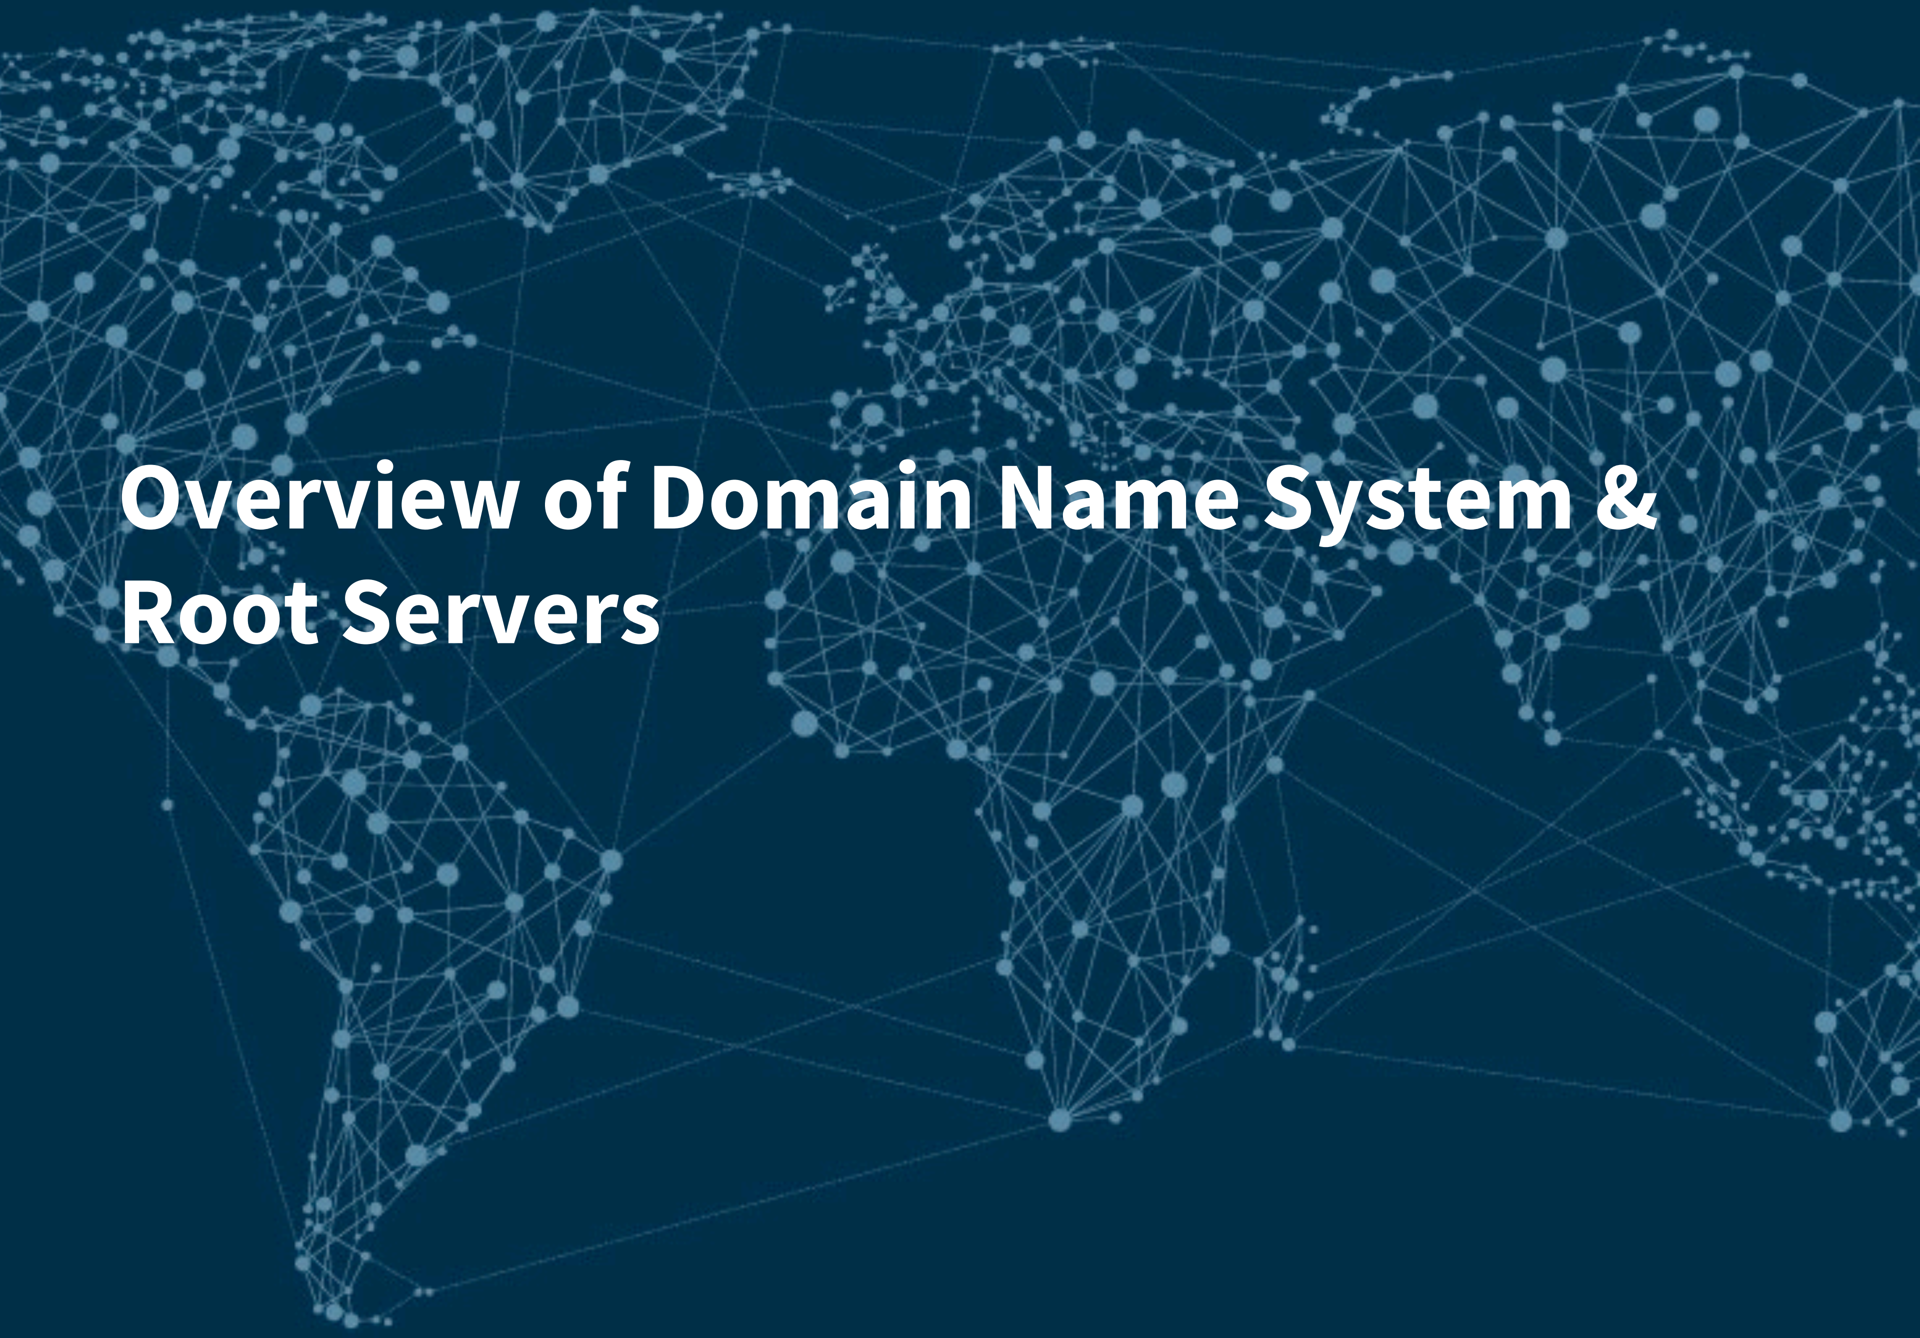

Overview of Domain Name System & Root Servers

Recap: Identifiers on the Internet

- The fundamental identifier on the Internet is an IP address.
- Each host (or sometimes group of hosts) connected to the Internet has a unique IP address.
- IPv4 or IPv6 (192.0.2.1, 2001:db8::1)
- Uniqueness guaranteed through allocation from a single pool (IANA-RIR system) and careful management within a network

Why DNS?

- ORIGINAL PROBLEM: IP addresses are hard to remember, and often change
- MODERN PROBLEM: IP addresses may also be shared, or multiple IP addresses may serve as entry points to a particular service; which IP address to use?

The Domain Name System

- A look up mechanism for translating objects into other objects:
    - Name to IP address www.example.org = 198.51.100.52
    - And many other mappings (mail servers, IPv6, reverse…)
- Globally distributed, loosely coherent, scalable, dynamic database

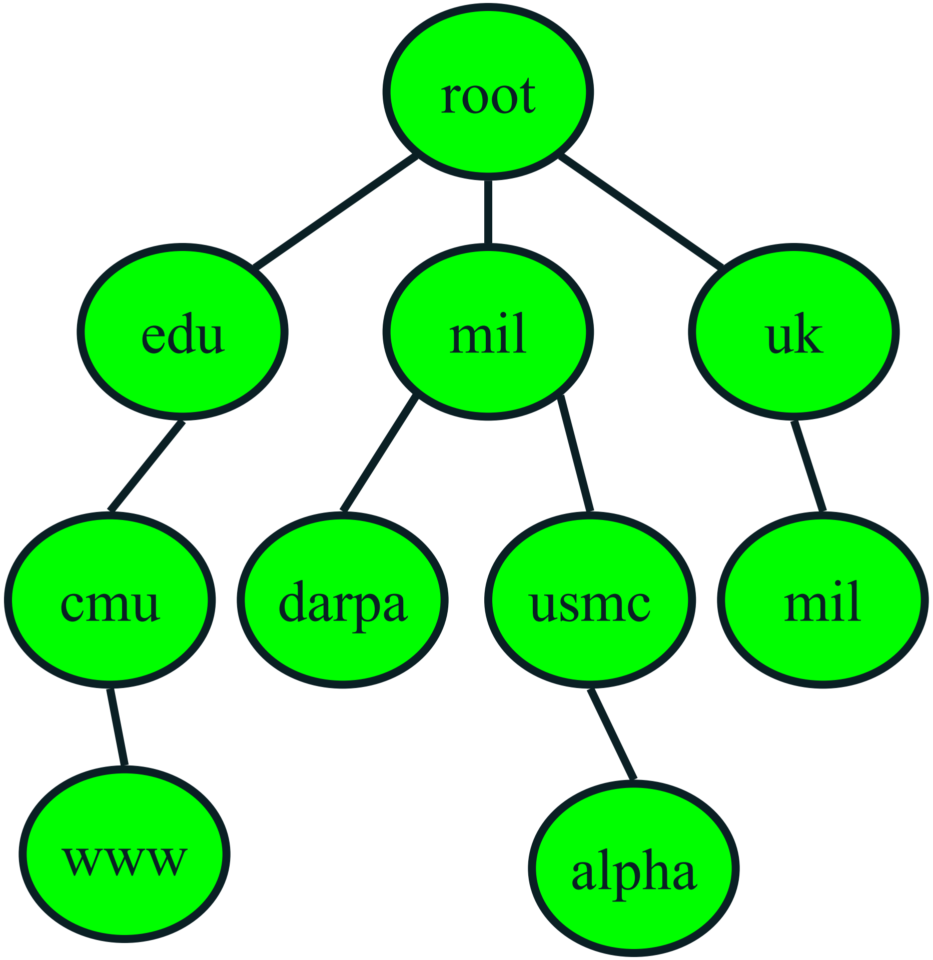

Domain Name Resolution Process

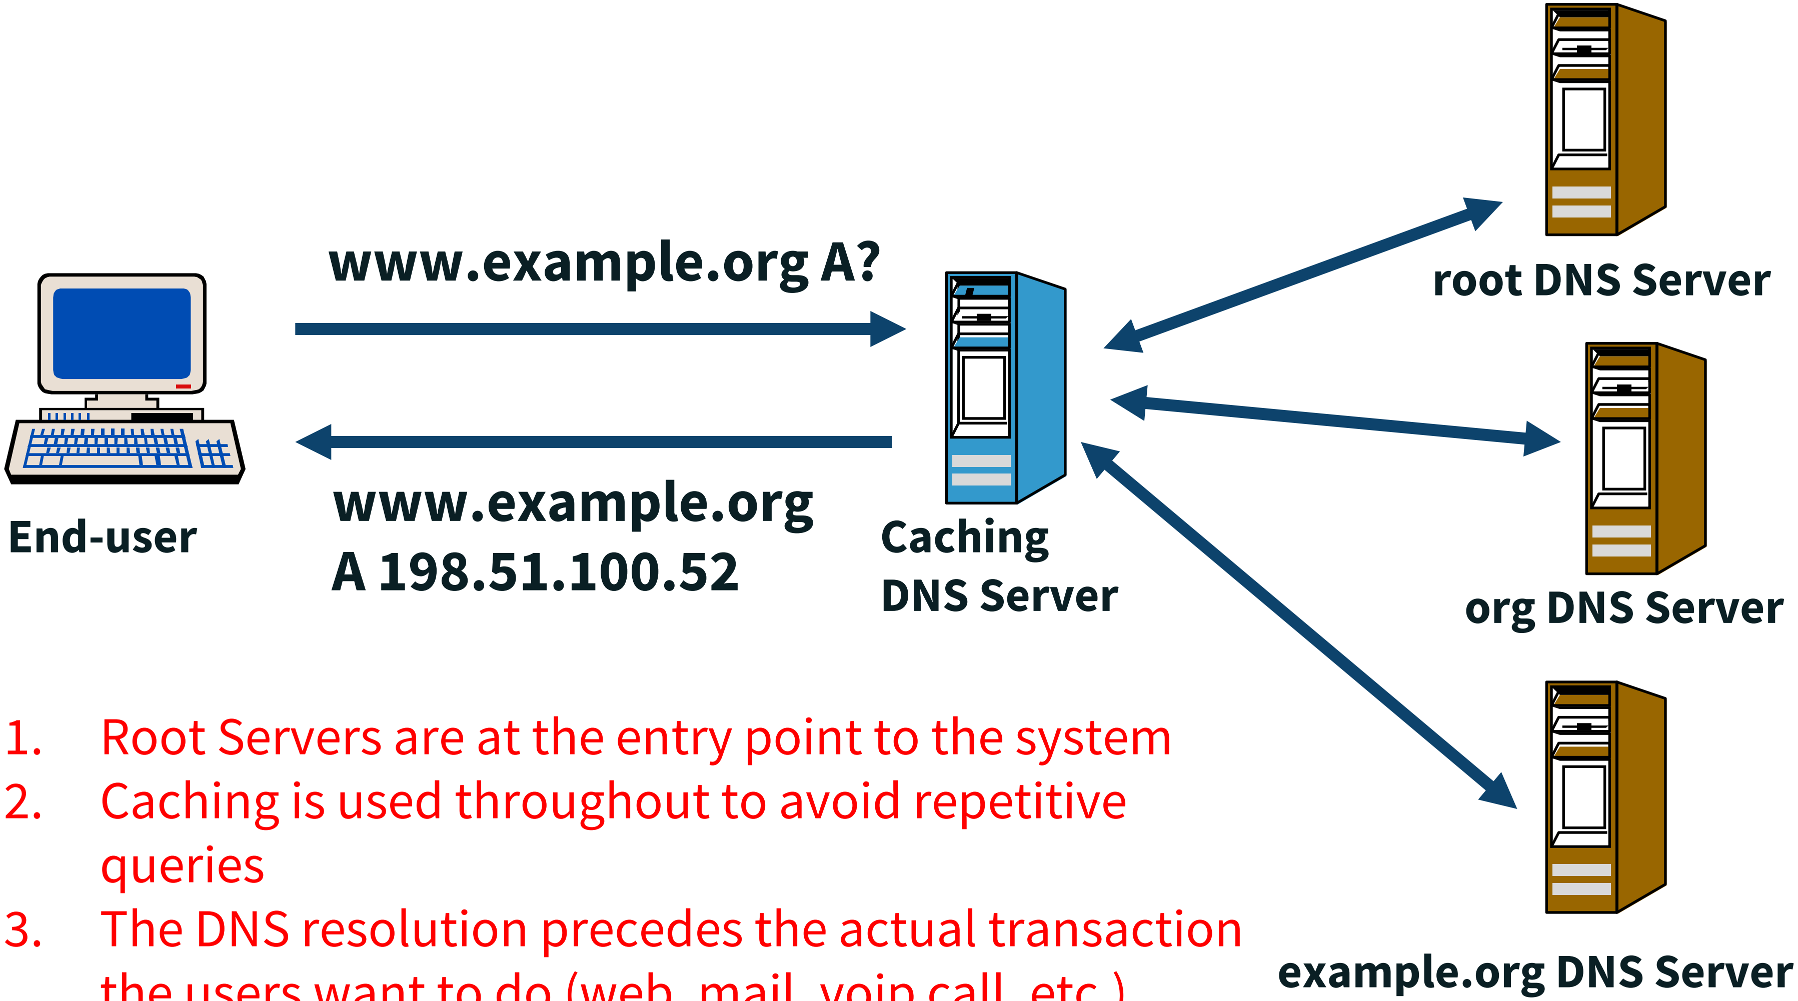

1. Root Servers are at the entry point to the system
2. Caching is used throughout to avoid repetitive queries
3. The DNS resolution precedes the actual transaction the users want to do (web, mail, voip call, etc.)

Domain Name Resolution Process

- Root servers only know who you need to ask next.
    - .com=>list of .com servers
    - .net => list of .net servers
    - .org => list of .org servers
    - ……
- Caching of previous answers means there is less need to query the root servers after the first question

Some Modern Refinements to DNS

- DNSSEC (Security extensions)
  - Cryptographic signatures on DNS data
  - Reduces risk of “spoofing”
  - Resolver has to validate
- Privacy enhancements
  - Queries can leak information
  - Standards being created to reduce this
- Anycast
  - Multiple servers share a single IP address
  - Improves latency and resilience
  - Protects against DDOS attacks

Root Zone vs. Root Servers

- Root zone
  - The starting point: the list of TLDs and nameservers
  - Managed by ICANN, per community policy
  - Compiled & distributed by the Root Zone Maintainer to all root server operators
- Root servers
  - Respond with data from the root zone
  - Currently distributed from 13 identities
    - [a-m].root-servers.net
    - Purely technical role = serve the root zone
    - Responsibility of the root server operators

The Root Server Operators

- 12 different professional engineering groups focused on
  - Reliability and stability of the service
  - Accessibility for all Internet users
  - Technical cooperation
  - Professionalism
- Diverse organizations and operations
  - Technically
  - Organizationally
  - Geographically

The Root Server Operators (2)

- The operators are not involved in:
  - Policy making
  - Data modification
    - Publishers, not authors or editors
- The operators are involved in:
  - Careful operational evolution of service (expansion as the Internet expands)
  - Evaluating and deploying suggested technical modifications
  - Making every effort to ensure stability and robustness

History of Root Server System

First Root Servers (1983-1986)

Name	IP Address	Software	Organization
SRI-NIC	10.0.0.51 26.0.0.73	JEEVES	SRI International
ISIB	10.3.0.52	JEEVES	Information Sciences Institute, University of Southern California
ISIC	10.0.0.52	JEEVES	Information Sciences Institute, University of Southern California
BRL-AOS	192.5.25.82 128.20.1.2	BIND	Ballistic Research Laboratory, US Army

Additional Root Servers - 1987

Name	IP Address	Software	Organization
SRI-NIC.ARPA	10.0.0.51 26.0.0.73	JEEVES	SRI International
A.ISI.EDU	26.2.0.103	JEEVES	Information Sciences Institute, University of Southern California
BRL-AOS.ARPA	192.5.25.82 128.20.1.2	BIND	Ballistic Research Laboratory, US Army
C.NYSER.NET	128.213.5.17	BIND	RPI
TERP.UMD.EDU	10.1.0.17 128.8.10.90	BIND	University Of Maryland
GUNTER-ADAM.ARPA	26.1.0.13	JEEVES	U.S. Air Force Networking Group
NS.NASA.GOV	128.102.16.10	BIND	NASA Ames

Expanding Root Service Outside US (1991)

Original Name	New Name	IP Address	Software	Organization
SRI-NIC.ARPA	NS.NIC.DDN.MIL	192.67.67.53	JEEVES	SRI International
A.ISI.EDU	A.ISI.EDU	26.2.0.103 128.9.0.107	JEEVES	ISI
BRL-AOS.ARPA	AOS.BRL.MIL	192.5.25.82 128.20.1.2	BIND	BRL, US Army
C.NYSER.NET	C.NYSER.NET	192.33.4.12	BIND	RPI
TERP.UMD.EDU	TERP.UMD.EDU	10.1.0.17 128.8.10.90	BIND	University Of Maryland
GUNTER-ADAM.ARPA	GUNTER-ADAM.AF.MIL	26.1.0.13	JEEVES	U.S. Air Force Networking Group
NS.NASA.GOV	NS.NASA.GOV	128.102.16.10	BIND	NASA Ames
NIC.NORDU.NET	NIC.NORDU.NET	192.36.148.17	BIND	NORDUnet

Renaming Root Servers to root-servers.net (1994-1995)

- By April 1993, the size of root hints response was approaching the 512 byte limit
- Bill Manning, Mark Kosters and Paul Vixie devised a plan to rename all the root servers from individual names to [a-i].root-servers.net
- IANA approved the plan and renaming was done in phases at the end of 1995
- Moving root servers to root-servers.net allowed for DNS label compression, thus four new root servers were added in 1997 to serve exclusively the root zone

Renaming Root Servers to root-servers.net

Original Name	New Name	Organization
NS.INTERNIC.NET	a.root-servers.net	Internic (operated by NSI)
NS1.ISI.EDU	b.root-servers.net	ISI
C.PSI.NET	c.root-servers.net	PSInet
TERP.UMD.EDU	d.root-servers.net	University of Maryland
NS.NASA.GOV	e.root-servers.net	NASA
NS.ISC.ORG	f.root-servers.net	Internet System Consortium (ISC)
NS.NIC.DDN.MIL	g.root-servers.net	DISA
AOS.ARL.ARMY.MIL	h.root-servers.net	Army Research Lab (ARL)
NIC.NORDU.NET	i.root-servers.net	NORDUnet

Adding Four Additional Root Servers (1996 – 1998)

- Jon Postel used a set of criteria to select new root server operators
  - Need (Europe, Asia)
  - Connectivity (both internal and external)
  - Commitment to send and respond to traffic without filtering
  - Community consensus: The potential operator should demonstrate the widest possible support from the community being served

Adding Four Additional Root Servers (1996 – 1998)

- In Europe, RIPE was chosen to run k.root-servers.net
- In Asia, WIDE was chosen to run m.root-servers.net
- j.root-servers.net stayed at NSI
- l.root-servers.net was transferred to ICANN as part of founding of ICANN

Root Server Planning After Postel's Death

- The root server operators all met in person and jointly agreed on the following principles
  - Operate for the common good of Internet reliability
  - The IANA as the source of the root data
  - Sufficient investment to operate responsibly
  - Proper notice and facilitate transition when needed
  - Recognition of the other operators

Root Server System Today & Features

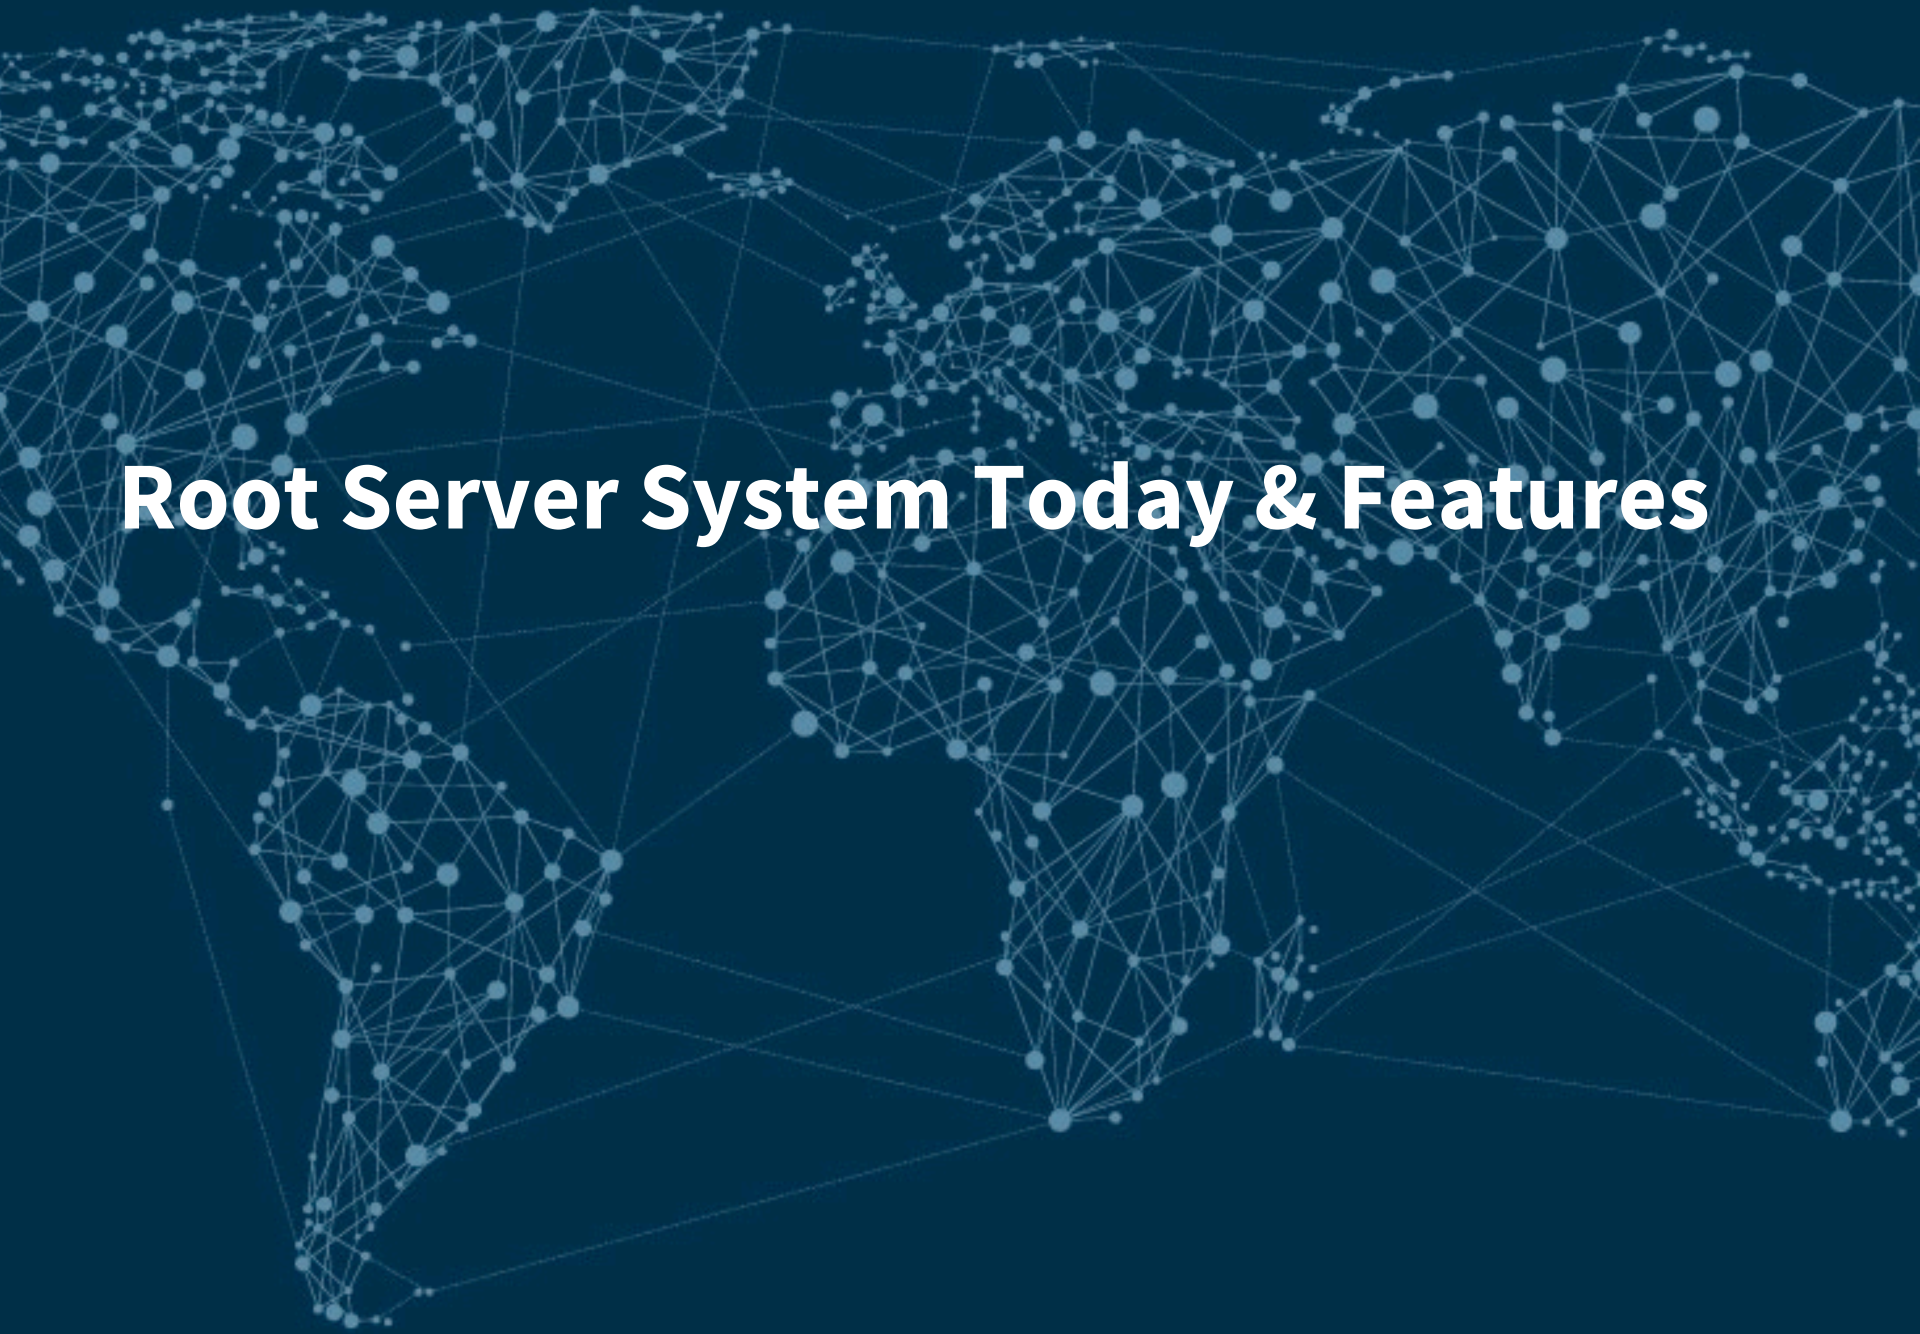

Root Servers Today - 2016

Hostname	IP Addresses	Manager
a.root-servers.net	198.41.0.4, 2001:503:ba3e::2:30	VeriSign, Inc.
b.root-servers.net	192.228.79.201, 2001:500:84::b	University of Southern California (ISI)
c.root-servers.net	192.33.4.12, 2001:500:2::c	Cogent Communications
d.root-servers.net	199.7.91.13, 2001:500:2d::d	University of Maryland
e.root-servers.net	192.203.230.10, 2001:500:a8::e	NASA (Ames Research Center)
f.root-servers.net	192.5.5.241, 2001:500:2f::f	Internet Systems Consortium, Inc.
g.root-servers.net	192.112.36.4, 2001:500:12::d0d	US Department of Defence (NIC)
h.root-servers.net	198.97.190.53, 2001:500:1::53.	US Army (Research Lab)
i.root-servers.net	192.36.148.17, 2001:7fe::53	Netnod
j.root-servers.net	192.58.128.30, 2001:503:c27::2:30	VeriSign, Inc.
k.root-servers.net	193.0.14.129, 2001:7fd::1	RIPE NCC
l.root-servers.net	199.7.83.42, 2001:500:9f::42	ICANN
m.root-servers.net	202.12.27.33, 2001:dc3::35	WIDE Project

Root Servers Today - 2016

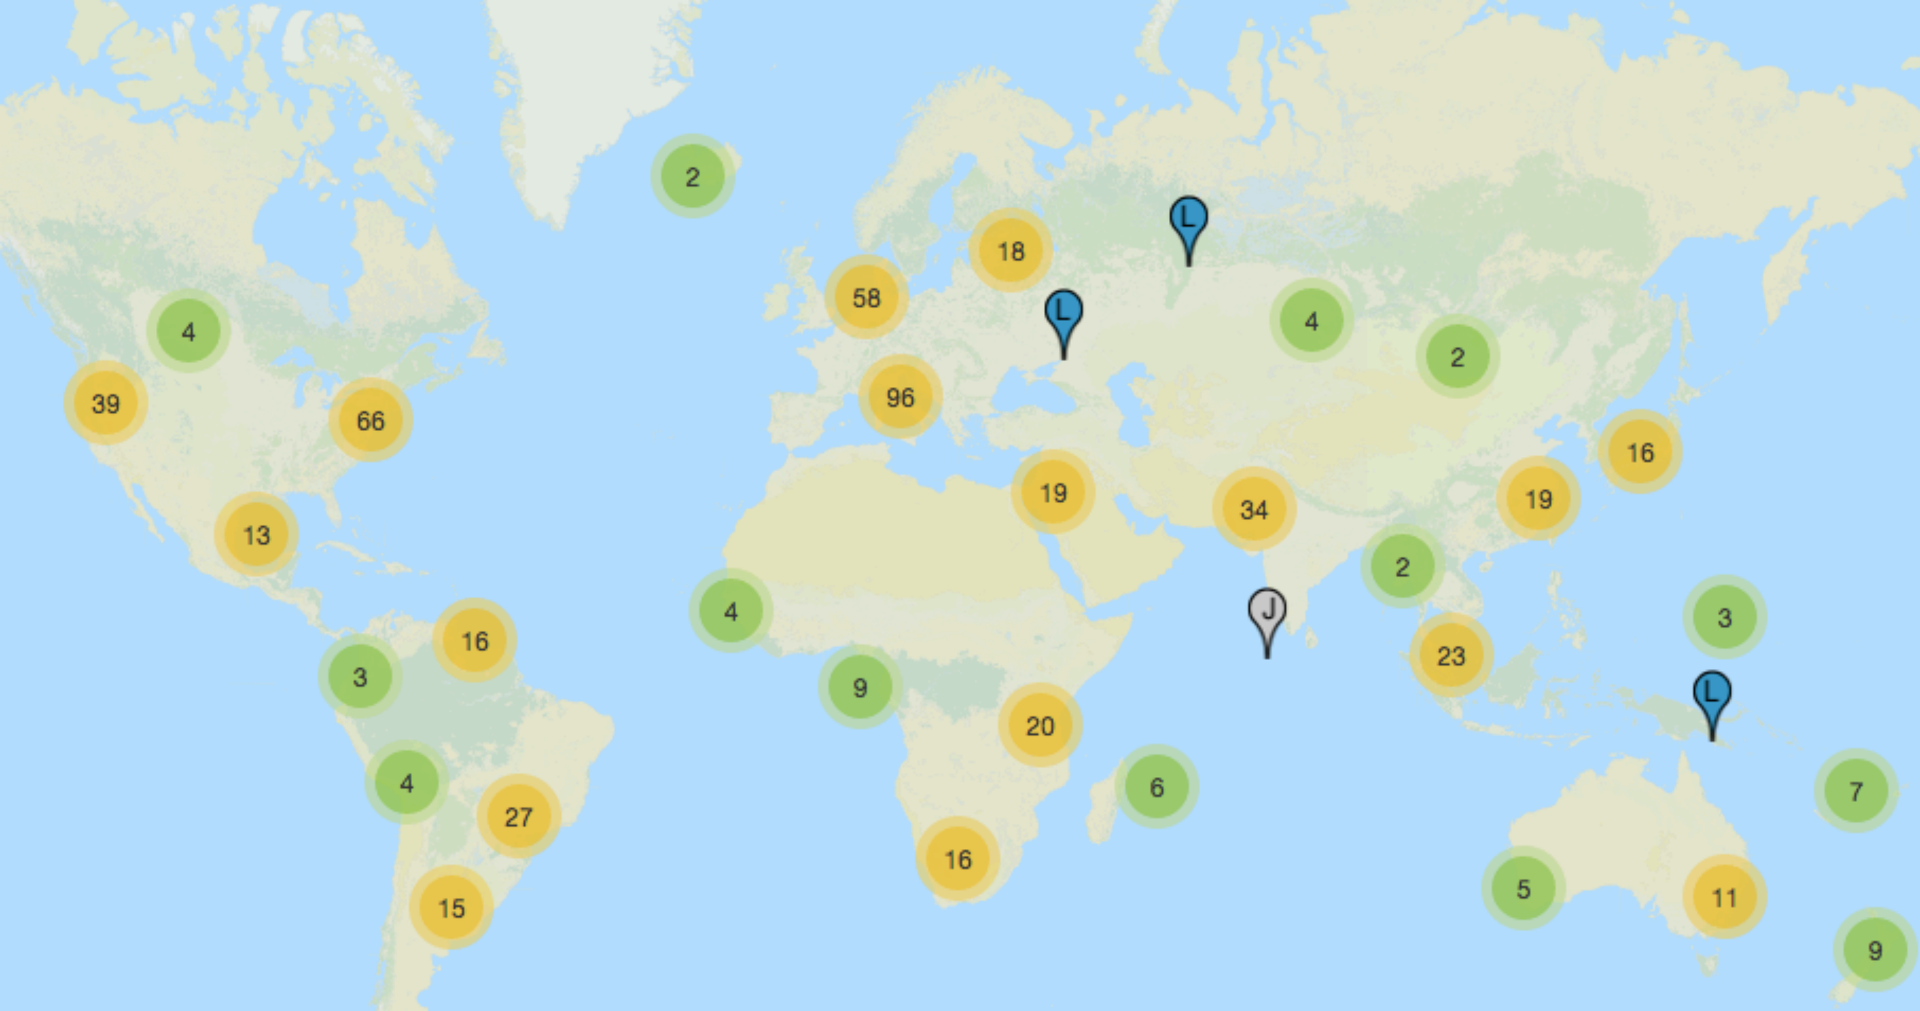

12 operators, 13 letters, over 600 instances around the world

Root Zone Management – Prior to IANA Stewardship Transition

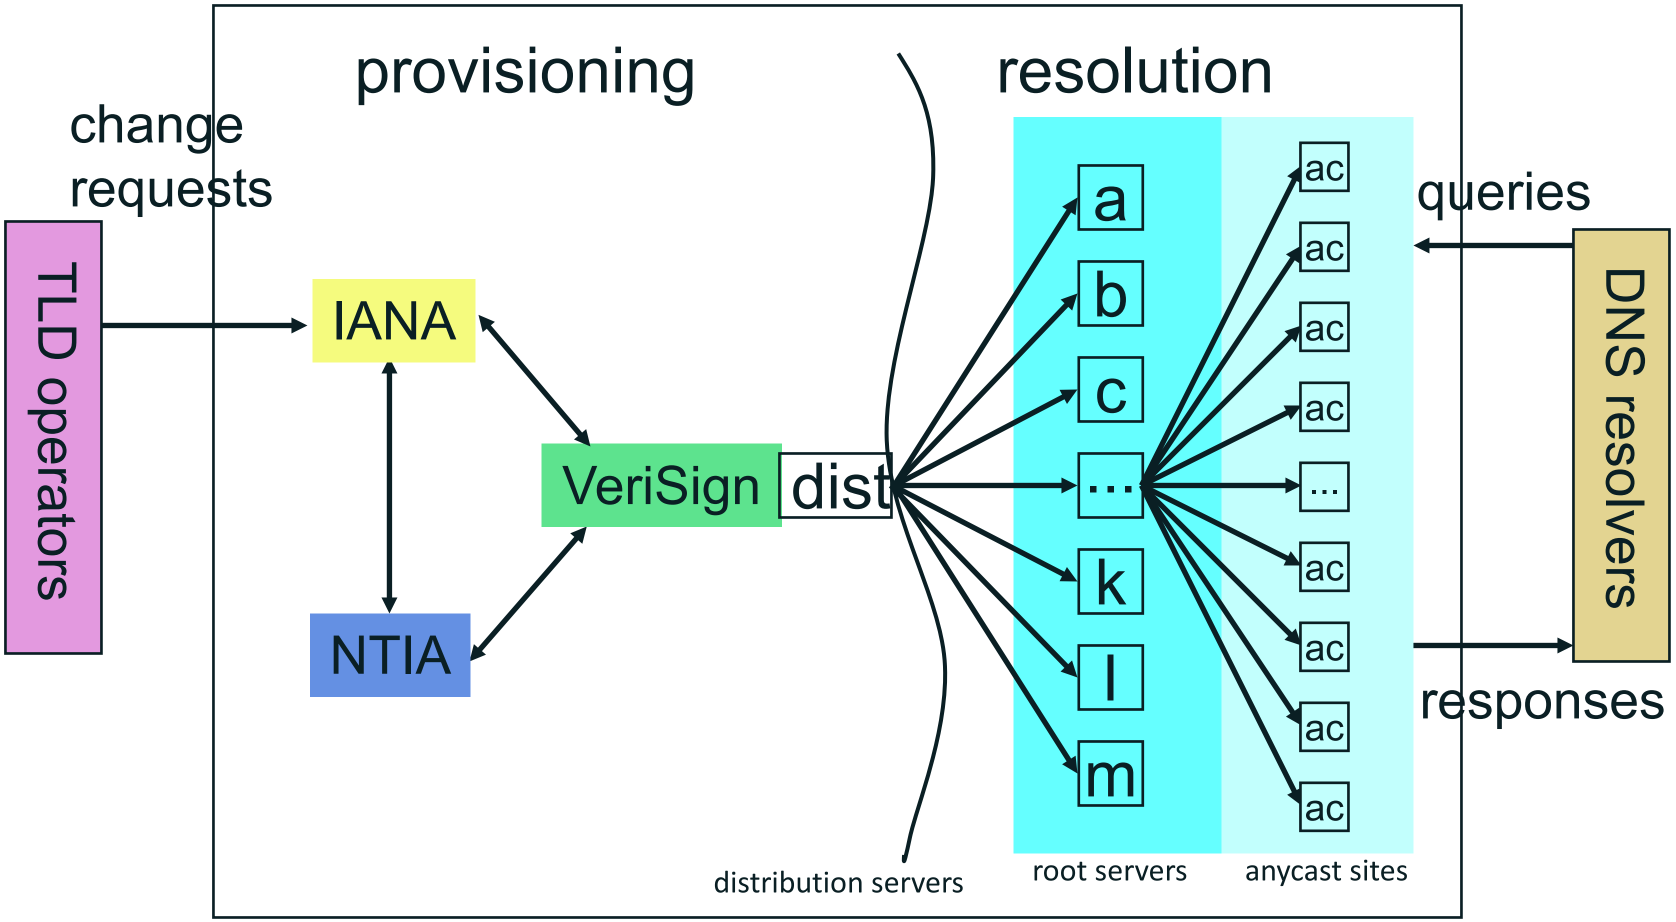

Root Zone Management – Post IANA Stewardship Transition

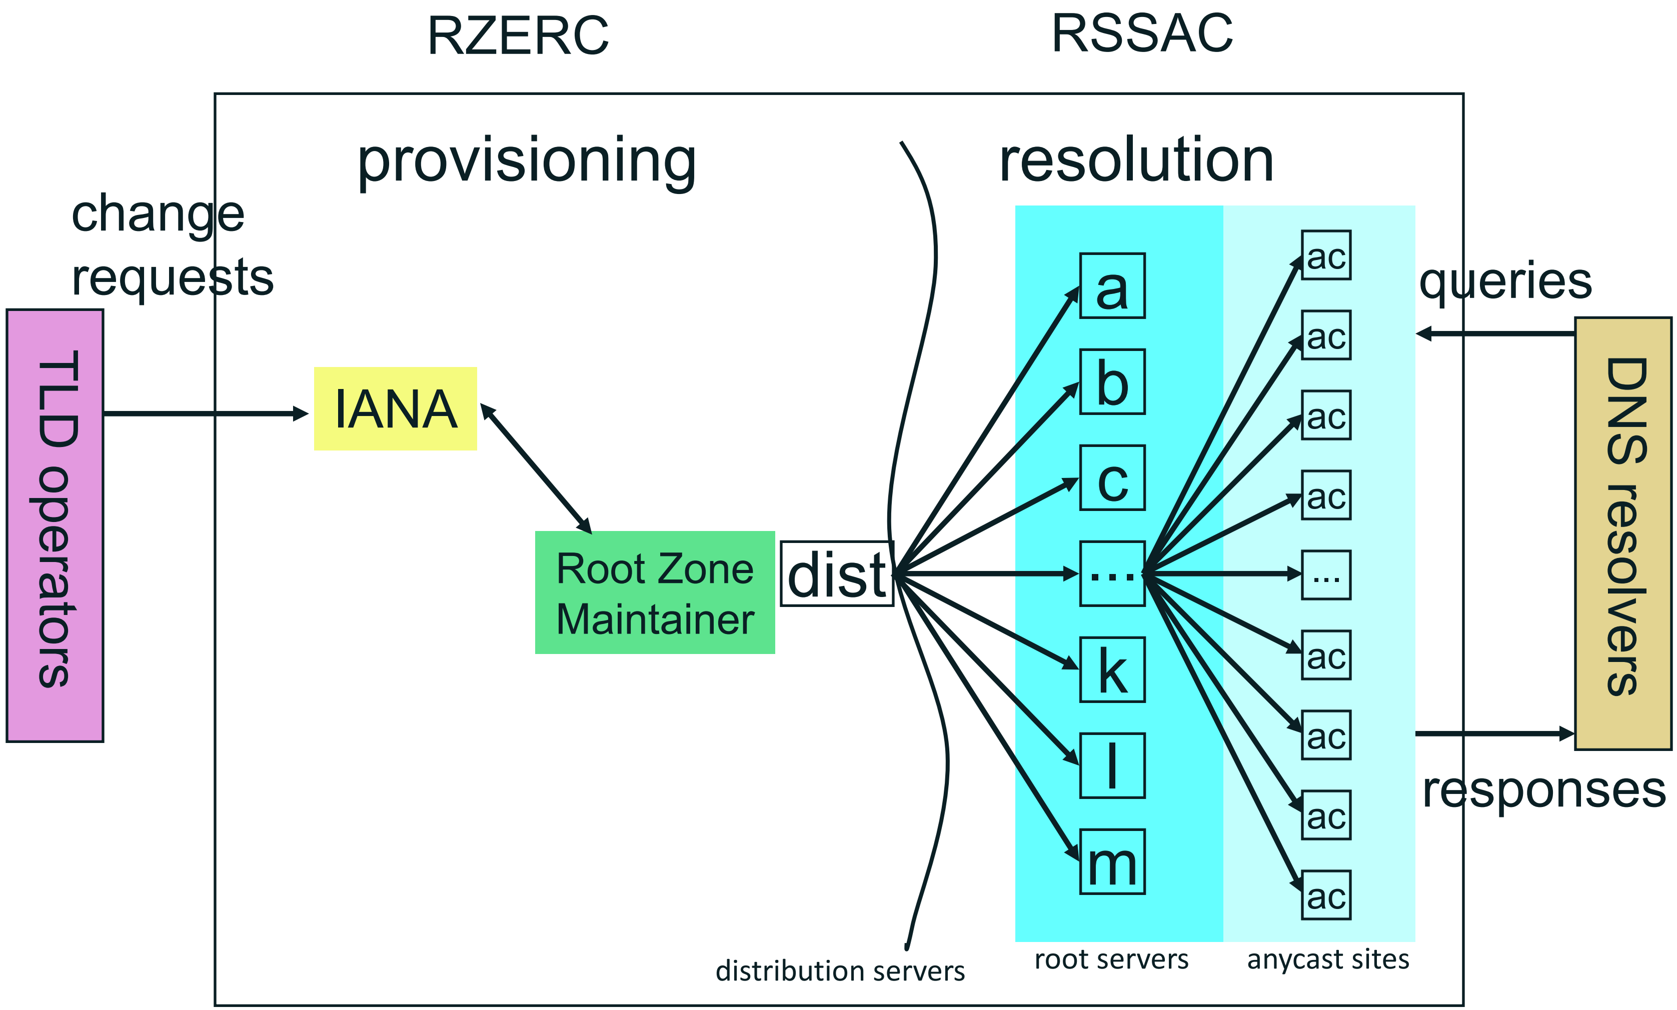

Features of Root Server Operators

- Diversity
  - Diversity of organizational structure (government labs, Universities, for profit companies, not for profit service)
  - Diversity of operational history
  - Diversity of hardware and software in use
  - Common best practices refer to minimum levels of
    - Physical system security
    - Overprovisioning of capacity
    - Professional and trusted staff

Features of Root Server Operators

- Cooperation and coordination
  - Cooperation takes place at industry meetings (ICANN, IETF, RIPE, NANOG, DNS-OARC, APNIC, ARIN, AFNOG, ..) and use of the Internet itself
  - Permanent infrastructure to respond to possible emergencies (telephone bridges, mailing lists, exchange of secure credentials)
  - Periodic activities to support emergency response capabilities
  - Coordination within established Internet bodies (RSSAC within ICANN, participation in evolving the DNS standard through IETF, data-sharing through DNS-OARC)

Responses to an Evolving Internet

- As the Internet evolves new requirements are put on the DNS system
  - Root server operators analyze the impact of and adopt new uses and protocol extensions on the service
    - IDNs, DNSSEC, IPv6, …
  - Increasing robustness, responsiveness and resilience
    - Wide deployment of distributed anycasts (over 600 instances around the world)

Myths Corrected

- Root servers do not control where Internet traffic goes, routers do
- Most DNS queries are not handled by a root server
- Administration of the root zone is separate from service provision
- None of the root server letters are special
- Root server operators are not hobbyists
- More than 13 servers. Only 13 technical identities
- The collective root server operators coordinate operation as a whole

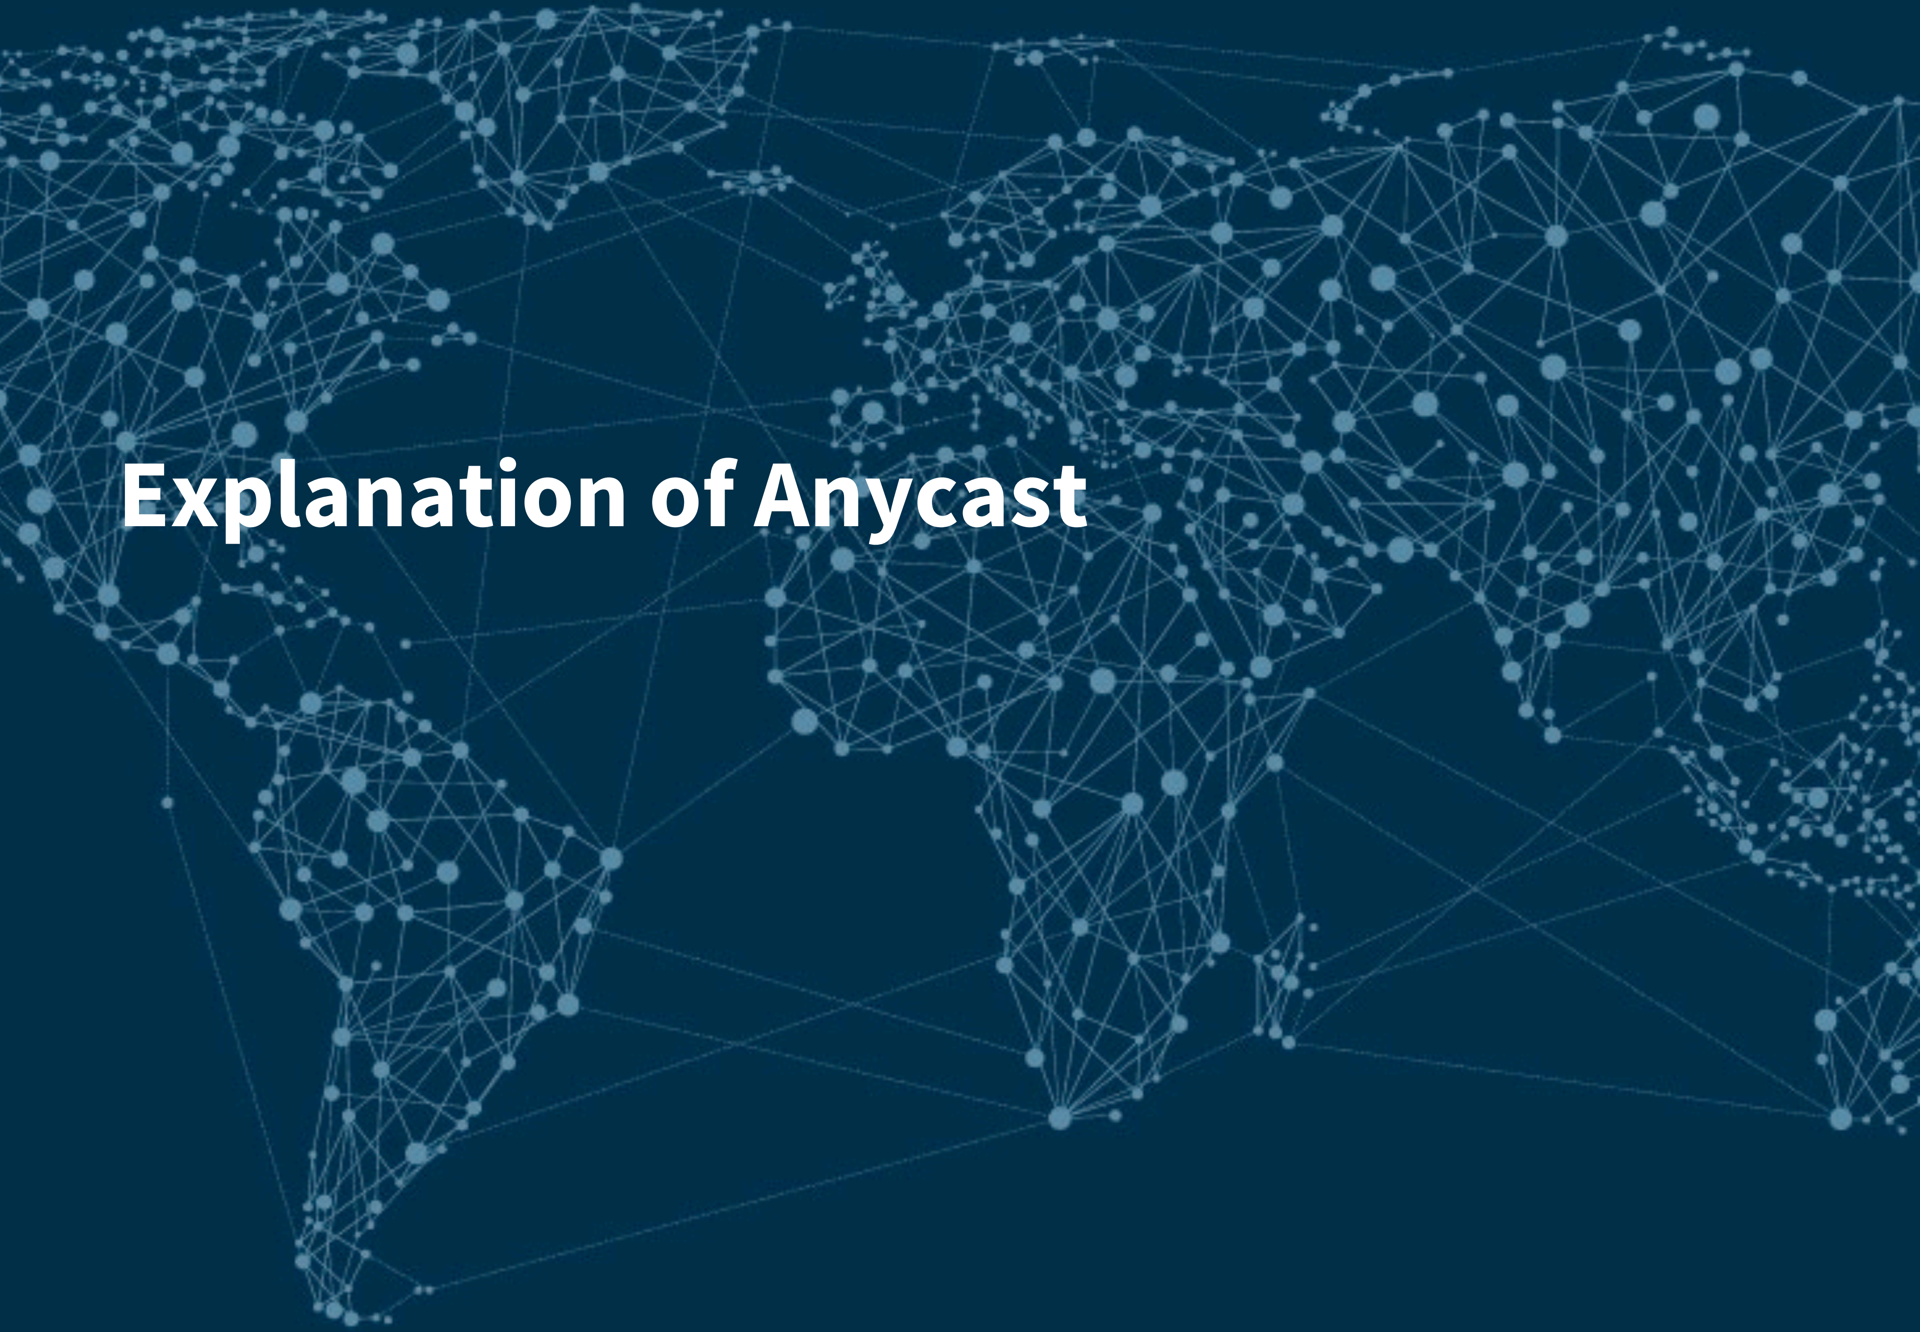

Explanation of Anycast

Unicast vs. Anycast

- Unicast
  - Packets from sources all go to the same destination
  - A single instance serves all sources
  - DDOS attack traffic all goes to single instance

- Anycast
  - Sources use destination based on intermediate routing policies
  - Multiple instances serve the same data to all sources
  - Sources get the data faster
  - DDOS attack traffic is sent to the closest instance
    - Frees other instances to service genuine traffic
    - Distributes the overall impact of the attack

Unicast

- Traffic takes shortest route to single destination

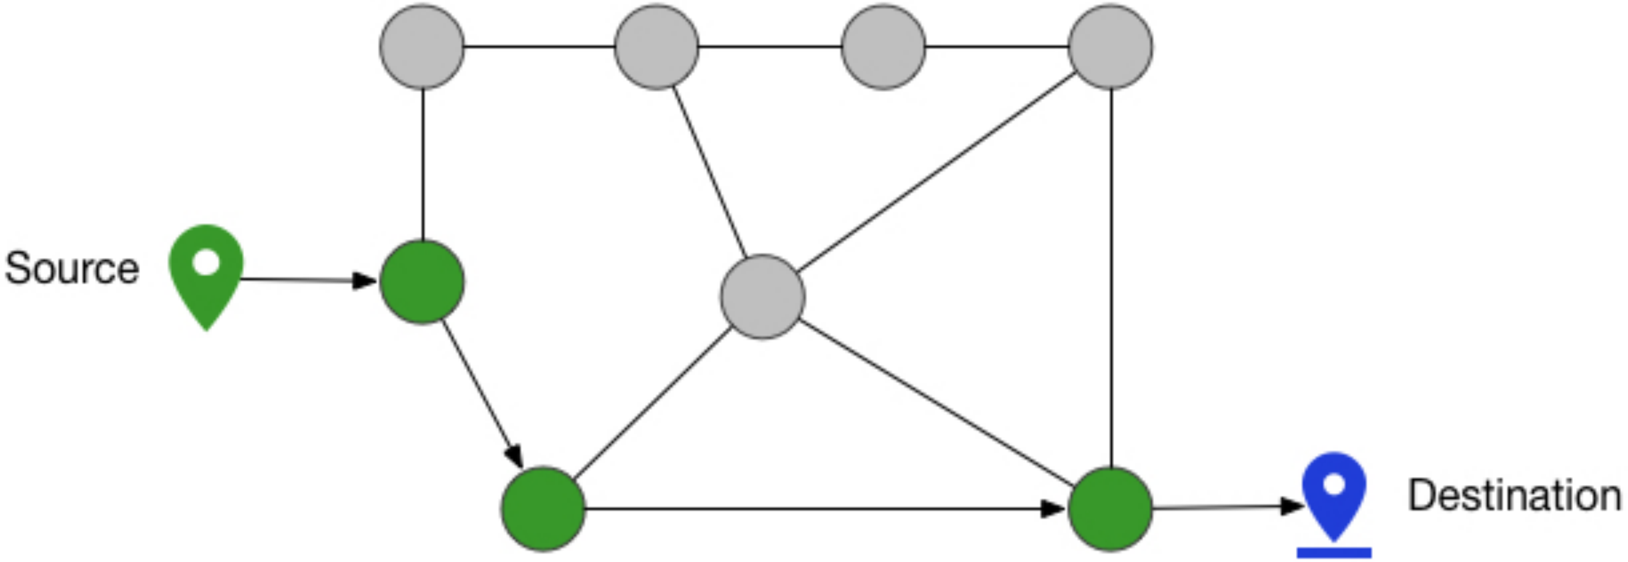

Anycast

- Traffic takes shortest route to closest destination
- Intermediate routing policies determine the destination for a source
- Path is shortened
- Data is delivered faster

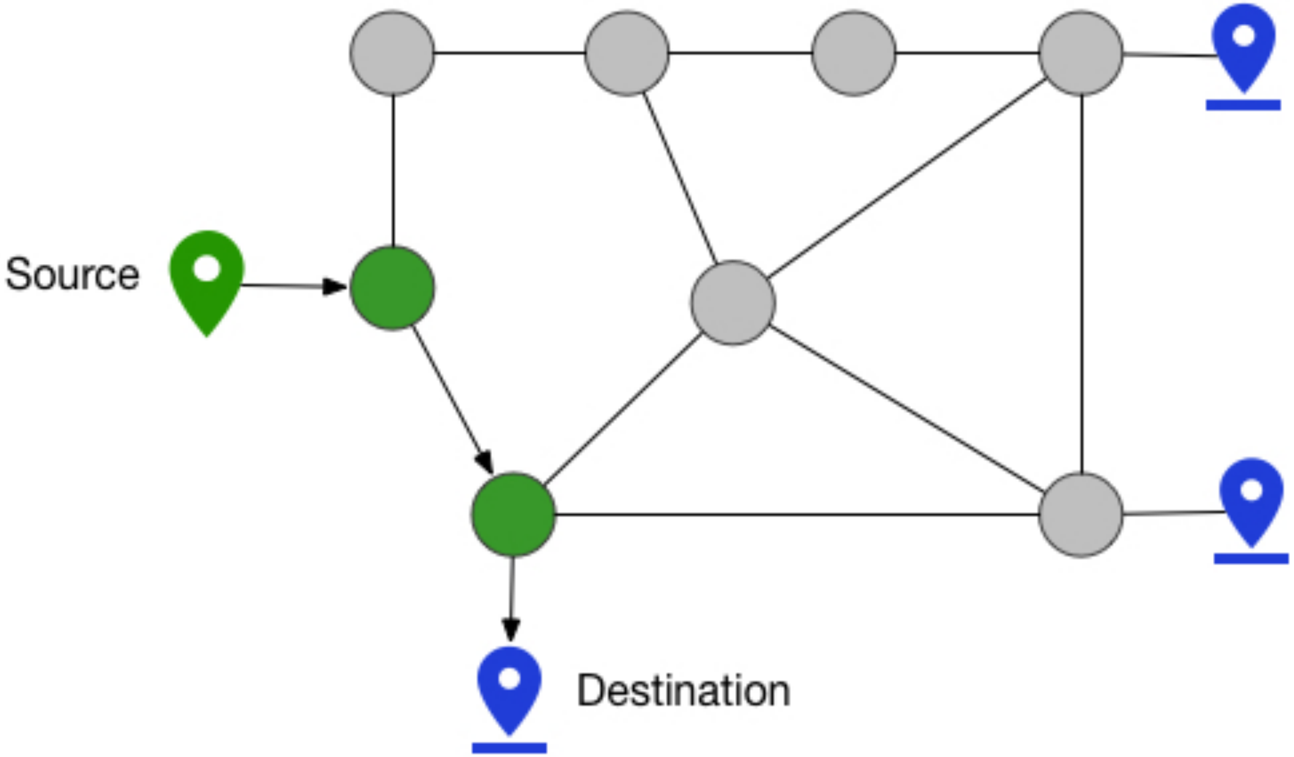

Anycast Under DDOS Attack

- DDOS attack traffic also takes shortest route to closest destination
- DDOS Traffic gets distributed across all destinations

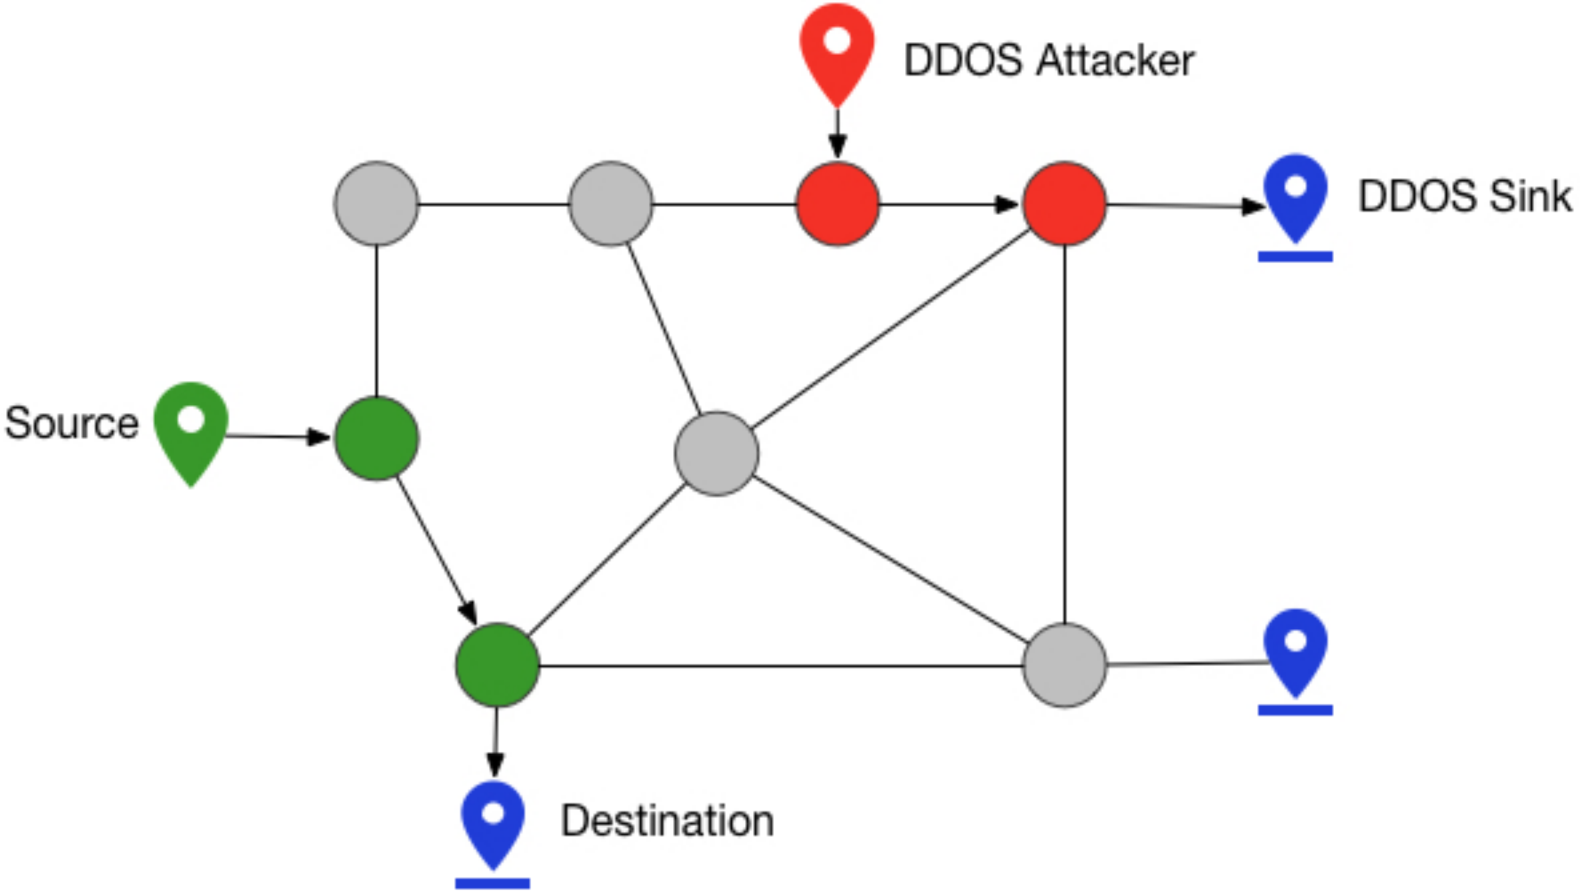

RSSAC and Recent Activities

What is RSSAC?

- The role of the Root Server System Advisory Committee ("RSSAC") is to advise the ICANN community and Board on matters relating to the operation, administration, security, and integrity of the Internet's Root Server System.
- (This is a very narrow scope!)

What RSSAC Does and Does Not Do

- RSSAC is a committee that produces advice – primarily to the Board but also to other ICANN bodies and other organizations involved in the overall DNS business.
- Root Server Operators are represented inside RSSAC, but RSSAC does not involve itself in operational matters.

RSSAC is here …

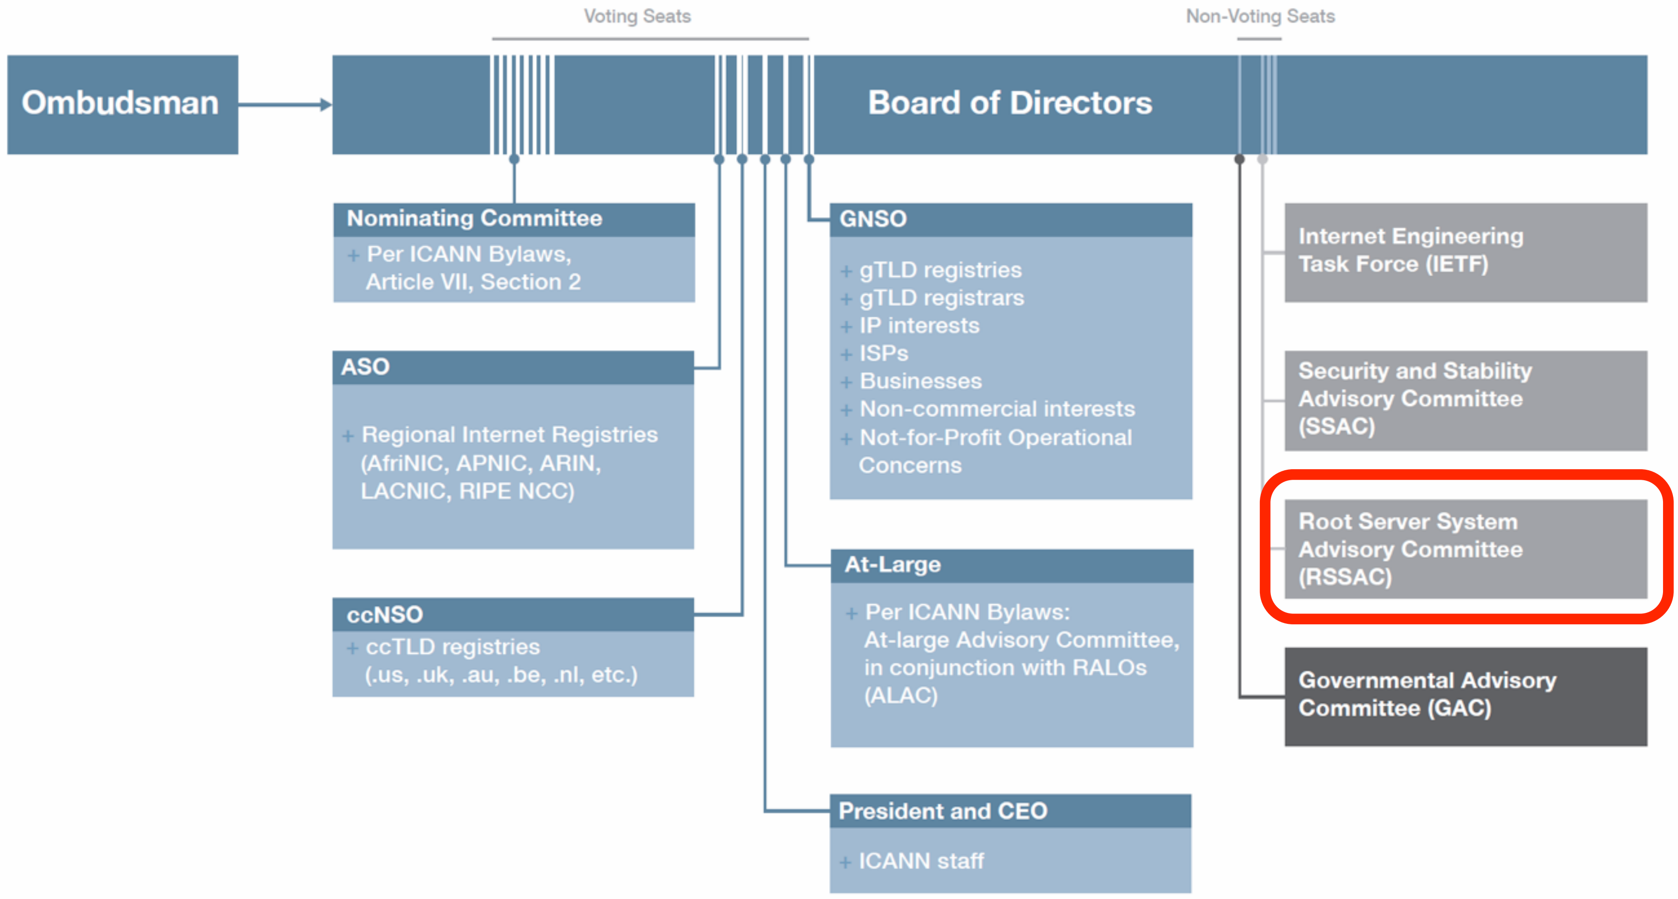
Voting Seats
Non-Voting Seats
Ombudsman
Board of Directors
Nominating Committee
+ Per ICANN Bylaws, Article VII, Section 2
ASO
+ Regional Internet Registries (AfriNIC, APNIC, ARIN, LACNIC, RIPE NCC)
ccNSO
+ ccTLD registries (.us, .uk, .au, .be, .nl, etc.)
GNSO
+ gTLD registries
+ gTLD registrars
+ IP interests
+ ISPs
+ Businesses
+ Non-commercial interests
+ Not-for-Profit Operational Concerns
At-Large
+ Per ICANN Bylaws: At-large Advisory Committee, in conjunction with RALOs (ALAC)
President and CEO
+ ICANN staff
Internet Engineering Task Force (IETF)
Security and Stability Advisory Committee (SSAC)
Root Server System Advisory Committee (RSSAC)
Governmental Advisory Committee (GAC)

RSSAC Organization

- RSSAC – composed of
  - – Appointed representatives of the root server operators
  - – Alternates to these
  - – Liaisons
- RSSAC Caucus
  - – Body of volunteer subject matter experts
  - – Members confirmed by RSSAC based on statement of interest

RSSAC Co-chairs

Brad Verd
Verisign
A/J-root

Tripti Sinha
University of Maryland
D-root

RSSAC Liaisons

- IANA Functions Operator (PTI)*
- Root Zone Maintainer (Verisign)*
- Internet Architecture Board*
- Security and Stability Advisory Committee*
- ICANN Board**
- ICANN Nominating Committee**
- Customer Standing Committee**
- Root Zone Evolution Review Committee**

* Inward Facing Liaison **Outward Facing Liaison

https://www.icann.org/groups/rssac

RSSAC Caucus

- Members
  - 77 Technical Experts (70% do not work in Root Server Operations)
  - Public statements of interest
  - Public credit for individual work
- Purpose
  - DNS experts who bring diverse expertise to publications
    - Expertise, critical mass, broad spectrum
  - Transparency of who does the work
    - Who, what expertise, which other hats
  - Framework for getting work done
    - Results, leaders, deadlines
- To apply, email rssac-membership@icann.org.

RSSAC Recent Publications

- Reports
  - RSSAC Workshop 2 Report [26 June 2016]
- Statements
  - RSSAC Response to GNSO Policy Development Process Working Group on the new gTLD Subsequent Procedures [6 October 2016]
  - RSSAC Statement Concerning the Impact of the Unavailability of a Single Root Server [8 September 2016]
  - RSSAC Statement on Client Side Reliability of Root DNS Data [28 June 2016]

RSSAC Workshop 2 Report

The RSSAC met for their 2nd workshop May 11-12 2016, in Reston Virginia, USA.

- Focused on three themes: Architecture, Evolution and Reinventing RSSAC

- RSSAC agreed to be the “front door” to the global DNS root service, and to the root server operators

- Three statements were conceived at the workshop
  - RSSAC Statement Concerning the Impact of the Unavailability of a Single Root Server (published on 8 September)
  - RSSAC Statement on Client Reliability of Root DNS Data (published on 28 June 2016)
  - Key Technical Elements of Potential Root Operators (forthcoming)

RSSAC Statement on Client Reliability of Root Data

On June 28 2016, the RSSAC published **“RSSAC Statement on the Client Side Reliability of Root DNS Data”**

- Reiterates that the operators of the root servers are committed to serving the IANA global root DNS namespace
- All root servers provide DNS answers containing complete and unmodified DNS data signed with DNSSEC
- The same cryptographically verifiable data is provided worldwide from all instances of these root servers to allow clients to detect tampering and ensure the integrity of the data

Statement on Unavailability of a Single Root Server

On September 8 2016, the RSSAC published **“RSSAC Statement Concerning the Impact of the Unavailability of a Single Root Server”**

- The loss of a single root server would not cause immediate stability issues for the root server system and the Internet that depends upon it.
  - High redundancy of the root server system guarantees availability and resiliency of the delivery service.
  - Caching (based on advertised TTL values) reduces the query load to root servers and limits the effects of an outage

- Root server system has experienced several real, large scale attacks.
  - None of these attacks resulted in any end-user visible error conditions.

Response to GNSO PDP on new gTLDs

On October 6 2016, the RSSAC published **“Response to the GNSO Policy Development Process (PDP) Working Group on the new Generic Top Level Domains (gTLDs) Subsequent Procedures”**

- If future plans for more top level domains are consistent with the past expansion program, the RSSAC does not foresee any technical issues.

- Recommends root zone management partners and root server operators to implement coordination procedures so that root server operators can notify ICANN in the event of stress on the root name service.

RSSAC Current Work

- History of the Root Server System
- Root Server Naming Scheme
- Key Technical Elements of Potential Root Operators
- Distribution of Anycast Instances

Current Work: History of Root Server System

In collaboration with root server operators, the RSSAC has produced a report to inform the community on the current root server system, and its history from beginnings to present day. The report:

1. contains a chronological history of the root server system from its origin to its current structure, divided into historical periods.

2. contains a description the current operators, and their histories in operating the root service, provided by each operator organization.

Current Work: Root Server Naming Scheme

On 9 July 2015, the RSSAC established a Caucus work party to produce **“History and Technical Analysis of the Naming Scheme Used for Individual Root Servers”** with the following scope to:

- Document the technical history of the names assigned to individual root servers;
- Consider changes to the current naming scheme, in particular whether the names assigned to individual root servers should be moved into the root zone from the root-servers.net zone;
- Consider the impact on the priming response of including DNSSEC signatures over root server address records;
- Perform a risk analysis; and
- Make a recommendation to root server operators, root zone management partners, and ICANN on whether changes should be made, and what those changes should be.

Current Work: Key Technical Elements

As an outcome of the RSSAC Workshop 2, held in May 2016, the RSSAC established a work party to produce **“Key Technical Elements of Potential Root Operators”**.

- Lists important technical elements for potential new root operators that would be a critical part of any potential root server operator designation process.
- Uses RSSAC001 and RFC 7720 as starting points, expands on them
- Multiple types of elements; Design, Experience & Networking, Diversity, Documentation, Data & Measurement

Current Work: Distribution of Anycast Instances

On 6 October 2016, the RSSAC established a Caucus work party to produce **“Best Practices for the Distribution of Anycast Instances of the Root Name Service”** with the following research questions:

- Given the state of current internet technology, what is the maximum latency a relying party should experience when transacting with the DNS root service as opposed to with a single “root server?”
- Will adding more instances in more topologically diverse locations make the system more resilient to Denial Of Service (DOS) attacks?
- If root operators were to coordinate their deployments of anycast instances, what considerations should be contemplated?
- Are there any regional or global technological risks (or benefits) if only a subset of operators (versus all or the majority of root operators) deploy anycast instances?

Questions?

- **For more information on the RSSAC**
- Main webpage:

  https://www.icann.org/groups/rssac
- RSSAC Publications:

  https://www.icann.org/groups/rssac/documents

- **For more information on the RSSAC Caucus**
- Caucus webpage:

  https://www.icann.org/groups/rssac-caucus
- To join send email to:

  rssac-membership@icann.org