

Scaling the Root

A study of the impact on the DNS root system
of increasing the size and volatility of the root zone

Jaap Akkerhuis
Patrik Fältström
Lars-Johan Liman

Lyman Chapin
Glenn Kowack
Bill Manning

Summary of Findings

Root Scaling Study Results session
ICANN meeting, Seoul, Korea, 28 October 2009

Root Scaling Study

- Requested by ICANN Board
- Steering group formed early 2009
 - representatives from SSAC, RSSAC, Staff
- Study team commissioned May 2009
 - Lyman Chapin (lead)
 - Jaap Akkerhuis, Glenn Kowack, Patrik Fältström, Lars-Johan Liman, Bill Manning
- TNO root scaling model
 - Developed by TNO under direction of and input from the Root Scaling Study Team

Root Zone Expansion

- New resource records for DNSSEC
- Internationalized top-level domain names
- New address records and glue for IPv6
- New gTLDs
 - Root zone becomes larger (size)
 - Change rate increases (volatility)

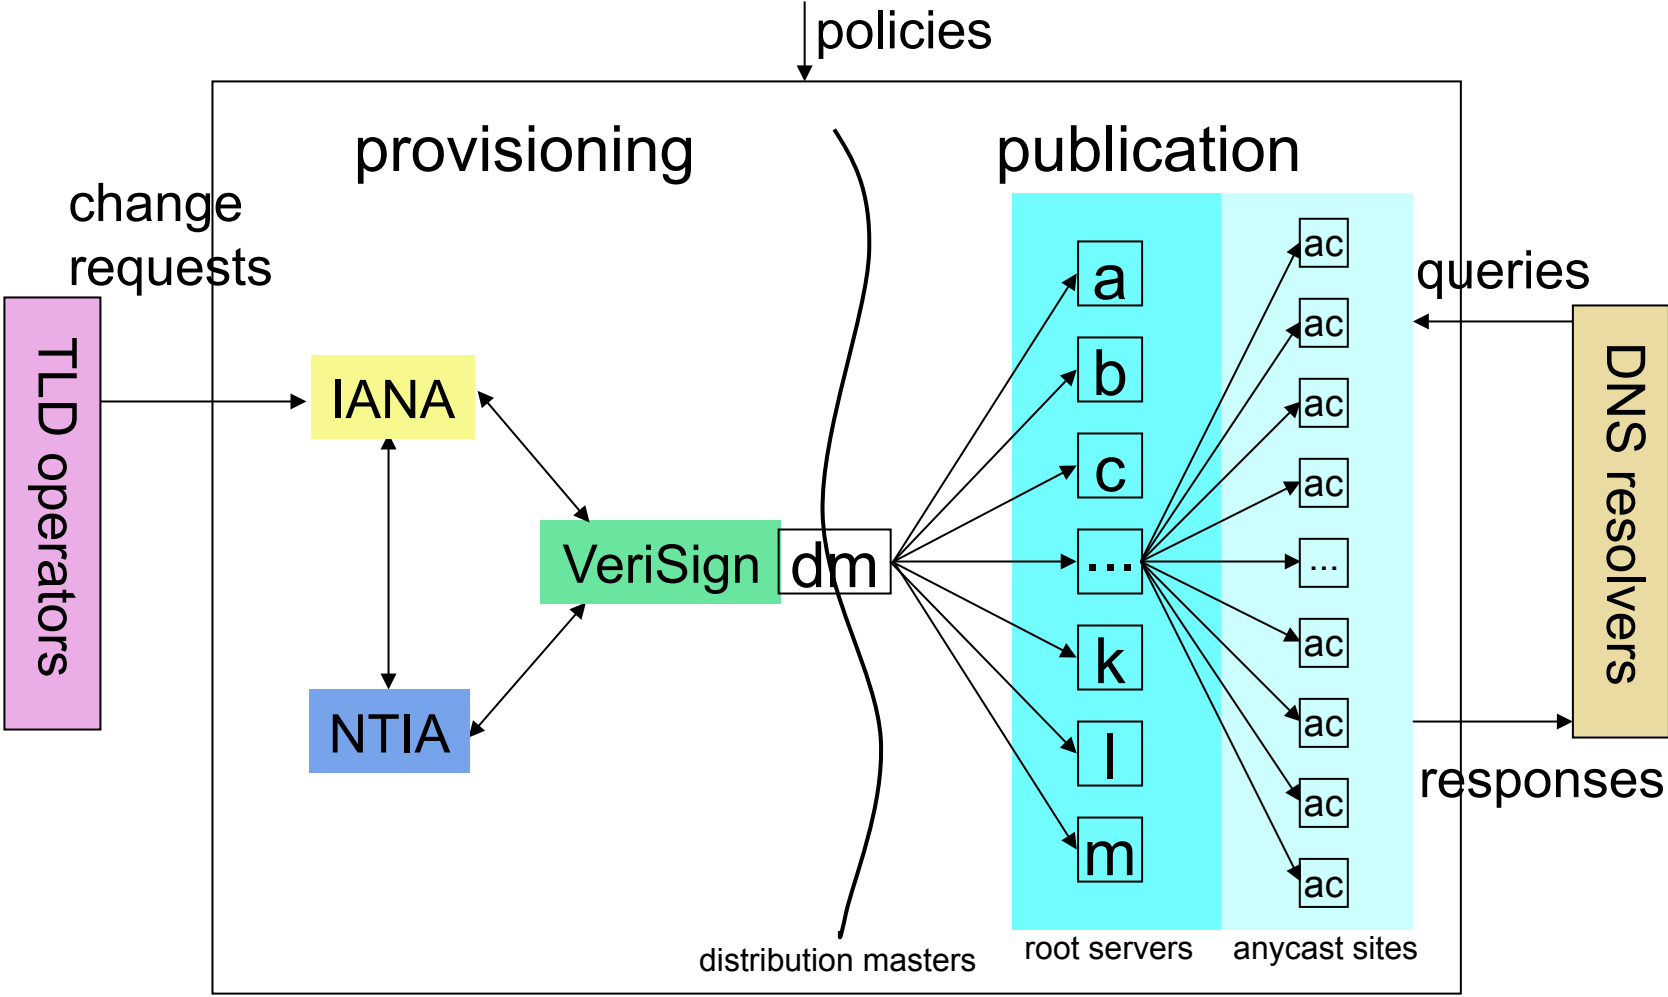
Information Gathering

- Interviewed all 12 root server operators
 - 2–4 hour telephone interviews
 - System designs
 - Management structures
- US Dept. of Commerce
- IANA function (ICANN)
- Root zone editor (Verisign)

Information Gathering

- Outreach to big DNS user organizations
 - Large corporations
 - Internet Service Providers
 - TLD operators – registries
 - TLD name server operators
- Presence at conferences
 - Announcements
 - Solicitated input

Root System Model



The Root Server System

- The root is a highly decentralized and dynamic system.
 - Many small moving parts, with engineers greasing cogwheels, replacing parts, rebuilding, and improving the system all the time.
 - The operators will ensure that the system adapts as the requirements on the systems evolve.
- There is no known sharp limit beyond which the system will suddenly “explode”.
- There is some instrumentation in parts of the system that will display warning signs when the system is under stress.

Introducing Changes

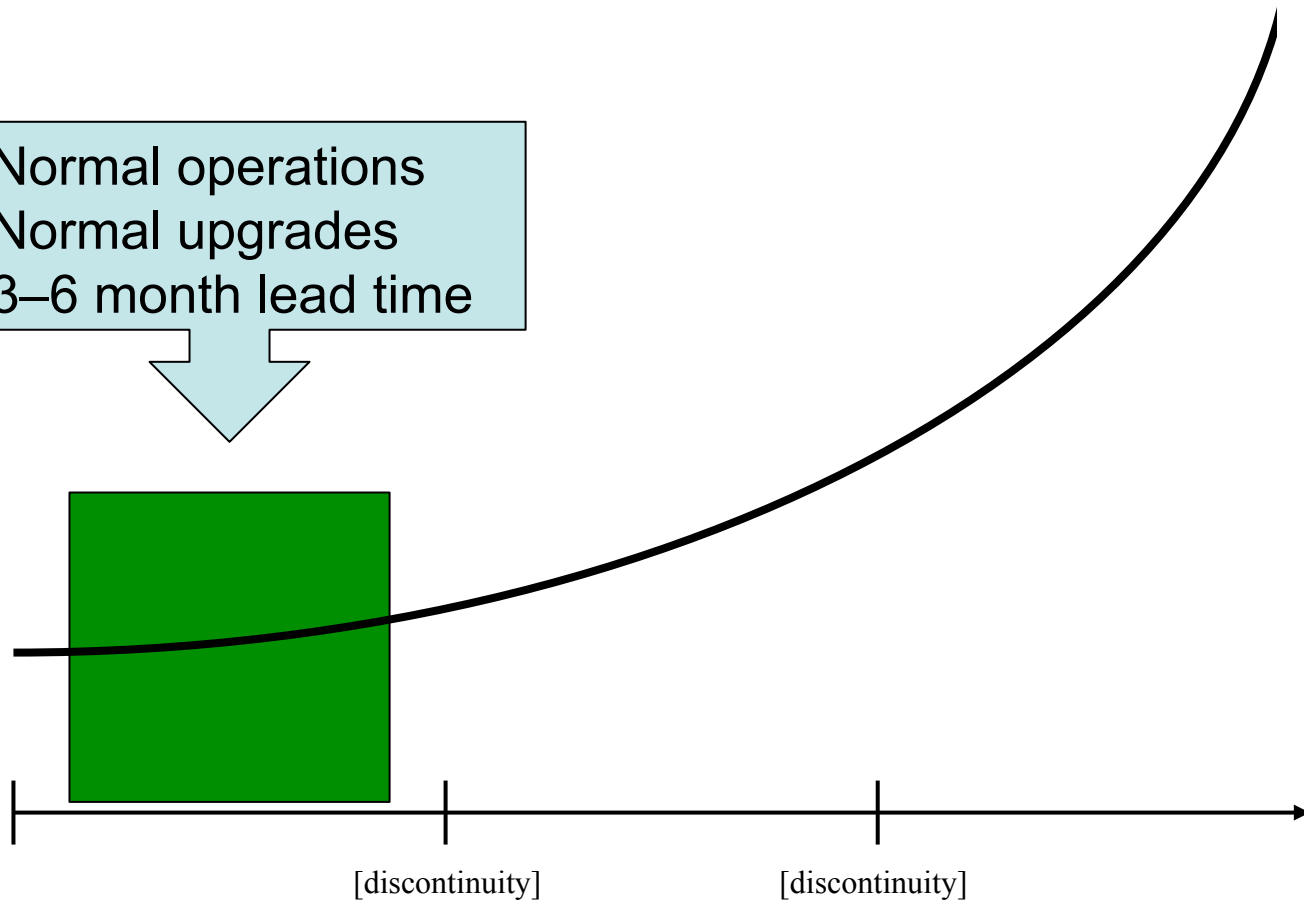
- The root system is flexible and adaptable!
 - But not instantly so.
 - If the operational requirements change at a controlled pace, the system can adapt and maintain operational headroom.
- Any increase in the size or volatility of the root zone involves risk.
 - It's all about managing the risks.

“Early Warning Systems”

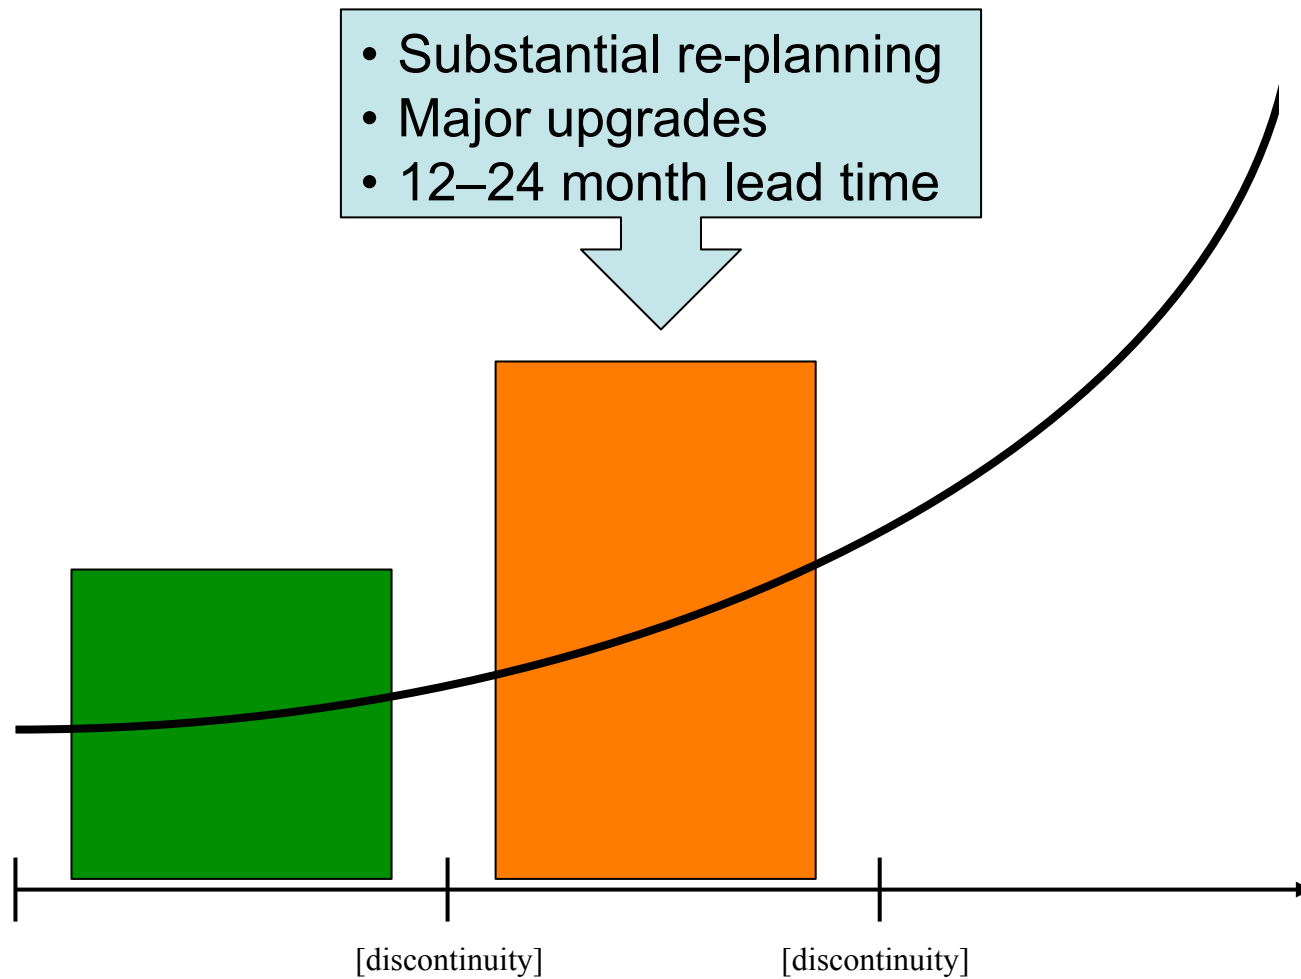
- Root system oversight should focus on “early warning” rather than threshold prediction.
- In order for “early warning” to be effective, changes to the root must be at a controlled pace.
- Additional instrumentation is needed to provide better early warning.
- Beware of “critical sections”!
 - Additions to the root that, once started, cannot be stopped (e.g. famous trademarks).

Dynamic Operating Regions

- Normal operations
- Normal upgrades
- 3–6 month lead time

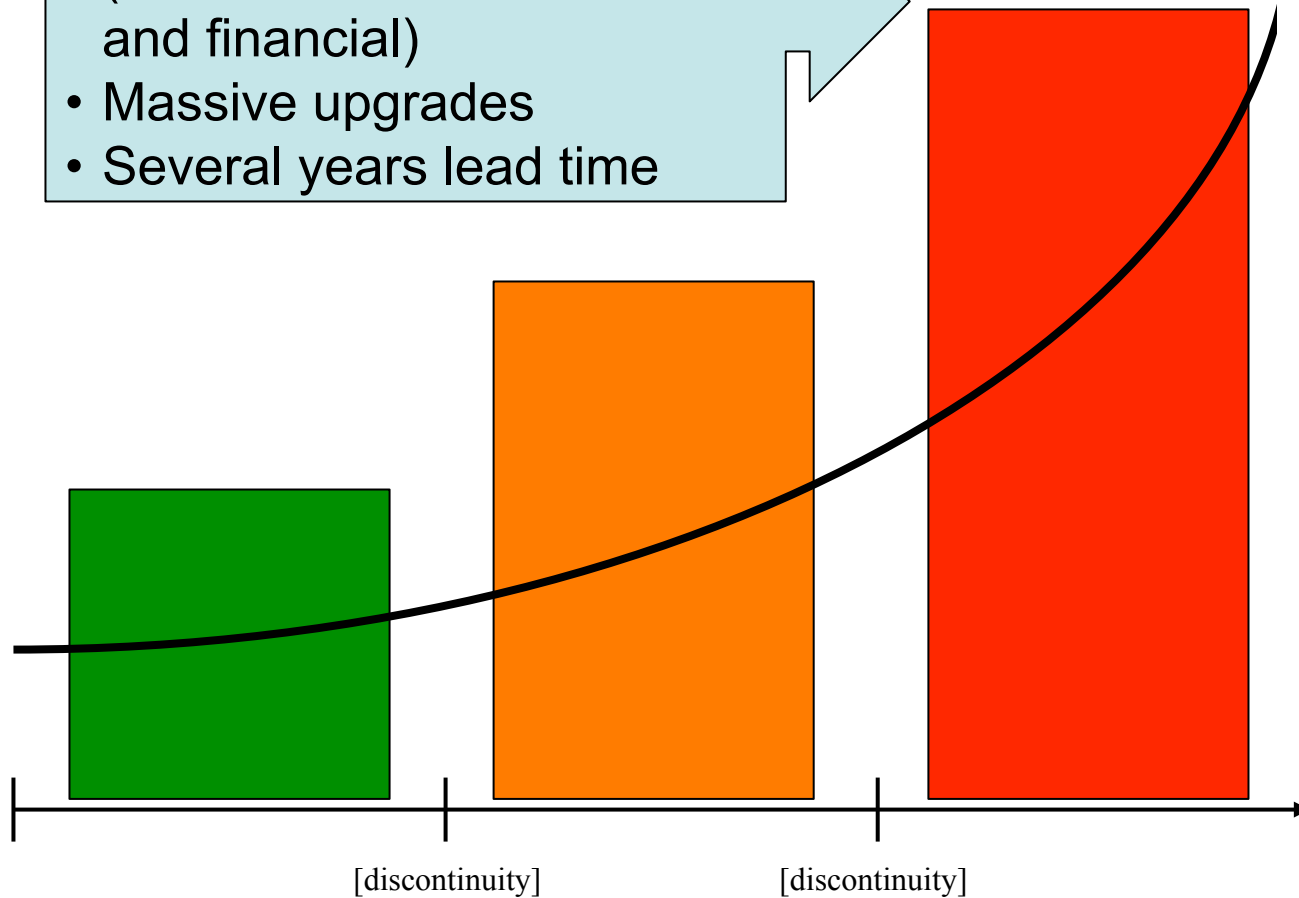


Dynamic Operating Regions



Dynamic Operating Regions

- Major redesign of entire root data distribution model, (administrative, technical and financial)
- Massive upgrades
- Several years lead time



Types of Stress

- Zone size
 - In itself not much of a problem ...
 - ... but a bigger zone leads to ...
- Zone volatility
 - Affects the volume of data per unit of time that needs to be transferred from the distribution masters to the individual root servers.

Other Issues?

- Increased response size
 - Will generate additional traffic.
- Query rate?
 - Well met by current amount of hardware.
- Malformed queries?
 - Yeah, there are plenty of those ...
 - Handled by software, and close cooperation with software vendors.
- Neither of these pertain to the content of the database (the root zone).

Sudden Step Changes

- Change the behaviour of the servers.
- *Introducing DNSSEC*
 - Handle DNSSEC record types
 - Significant increase in zone size and volatility.
- *Introducing AAAA (IPv6) records*
- *Introducing DNAME records*

Gradual Changes

- Adding more TLDs
 - More of “the same”.
- Adding IDN names
 - To the system, all TLDs are equal
gTLD = ccTLD = IDN TLD
 - Using NS records = more of “the same”
 - Using DNAME records = transient!
- Adding more AAAA (IPv6) records
 - More of what we already have

Findings


- The following findings should be read in the context that coordination needs to happen between all the involved parties – on both the provisioning side, and on the publication side.
- It's crucial to coordinate, plan, and watch all parts of the systems carefully.

Findings

- On the **provisioning side**, the ability to scale the root is completely dominated by the steps that involve **human intervention**.
- On the **publication side**, scaling the root primarily **affects poorly-connected** Internet locations.
- The risks associated addition of a **few hundred TLDs can be managed** without changing any actor's current arrangement.
- The risks associated with an **annual increase in the size of the root zone** on the order of thousands of new entries can be managed **only with changes to the current arrangements** of one or more actors.

Root Server Operator Headroom

Impact as factor of zone size



Effects of Root Zone Changes

	New TLDs	DNSSEC	IDNs	IPv6 addresses
Increases number of TLD entries in the root zone	X		X	
Increases size of the root zone file	X	X	X	X
Increases amount of data per TLD		X	X	X
Increases number of variables per TLD		X		X
Increases number of changes per TLD per year		X		X