

Enabling Experimentation and Research for the DNS


Wes Hardaker
<hardaker@isi.edu>
22 Jun 2020

John Heidemann
<johnh@isi.edu>

Robert Story
<rstory@isi.edu>

Michael Elkins
<elkins@isi.edu>

Yuri Pradkin
<yuri@isi.edu>

 Copyright © 2020 by Wes Hardaker
Release terms: CC-BY-NC 4.0 international

These servers are waiting for you



What will you do with them?

Why Experiment with the DNS Now?

- Originally, DNS was an experimental replacement to static hosts
- Now, the world's economy relies on its “perfect” performance
- This makes innovation in **naming and identification** difficult
- Academic research limited to small-scale lab tests
- “DNS 2.0 will never happen”

But Is The DNS Fully Cooked?


- How do new versions of software compare?
- Does running software on a line-card really work?
- Are there DDoS defenses still to be studied? Which are best?
- What if all DNS went over IPv6? Or over TCP? Or QUIC?
- Are there better load-sharing techniques specific to DNS?
- What if clients asked multiple questions at once?
- What if authoritative servers needed to handle 100% DOH?
- What effect does TTL caching have on authoritative servers?
- How does the roll out of caching effect operational traffic?
- What are the regional differences in traffic patterns?
- How is Internet usage shaped over days, weeks, months and years?
- How are DNS services affected by cloud deployments? IoT deployments? New mobile networks?
- ...

Reality: the Internet's landscape keeps changing – DNS will always have to keep up with current trends

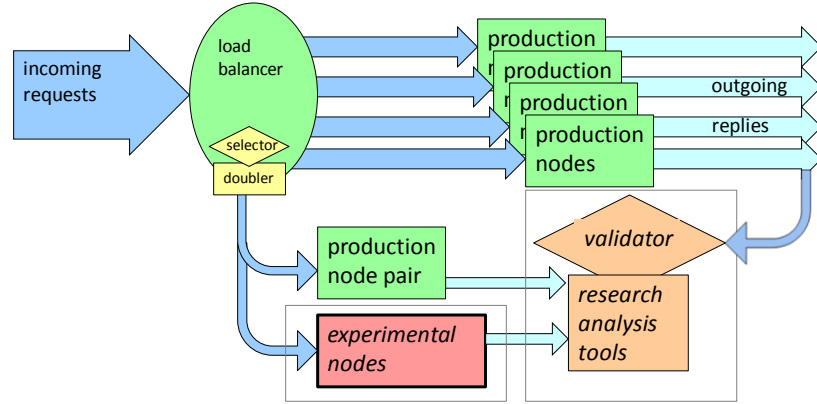
Our (new) experimentation platform

- ISI's NSF-funded DIINER project goal:
 - To support innovative naming research with real-ish experiments
- Available today:
 - Production-equivalent hardware
 - Live and replayed data streams
 - Curated datasets of “interesting” traffic events
- Future plans:
 - Additional hardware, web portal, data comparison engines

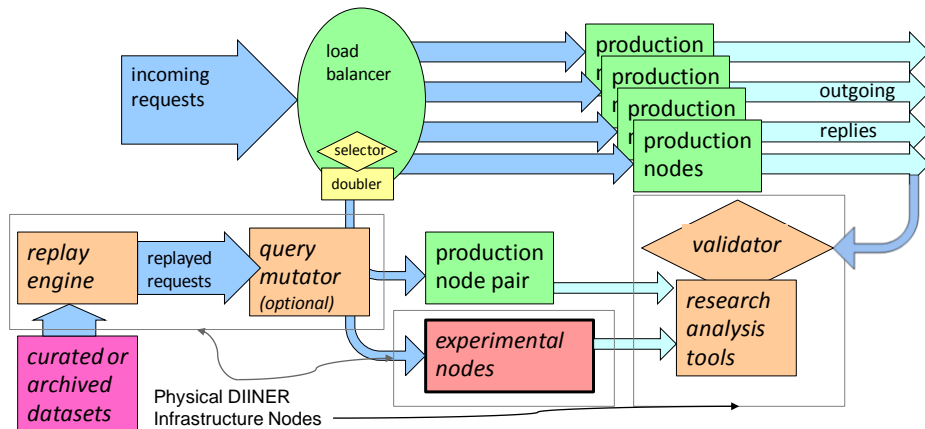
b.root-servers.net (and others) today



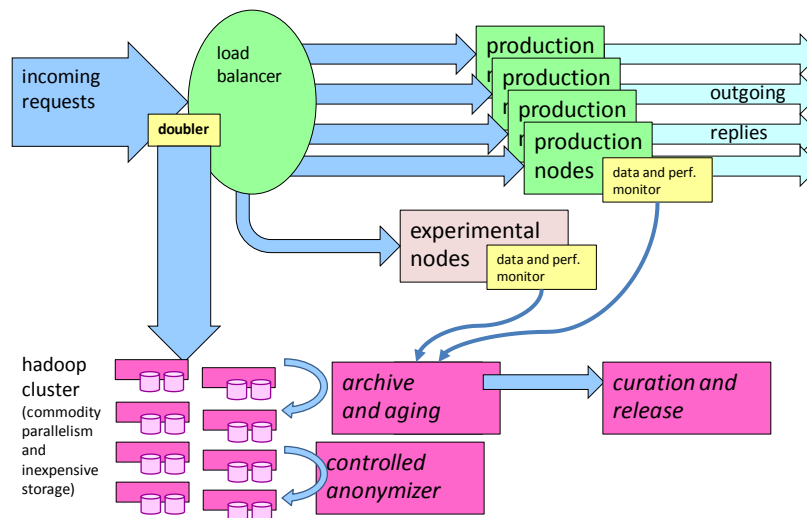
Planned Research Experimental Platform



Enabling Repeatability



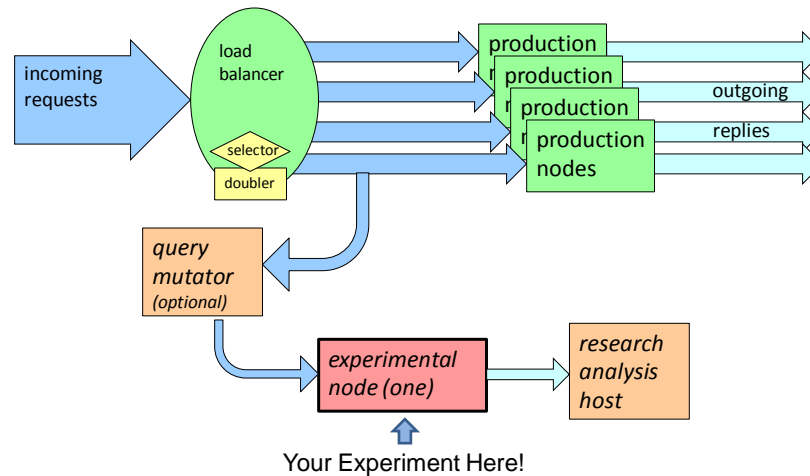
Data Capture and Curation



Plans and Timeline: Infrastructure

- Near-term:
 - Clone a B-Root production stream as a parallel test system ✓
 - Simultaneous monitoring of production vs test ✓
 - Installed stream replay and performance tools ✓
- Long-term:
 - Automatic result curation and packaging
 - Query mutating
 - Bring your own DNS (anycast, zones, resolvers, etc)
 - Real-vs-production comparison engine
 - Hardware expansion
 - Web-based console for project and experimentation management

Our Current Early-Setup



Plans and Timeline: Collaboration

- Near-term:
 - Hold a (virtual) workshop to discuss future needs and plans (Today!)
 - Work closely with some early-adopters
- Long-term:
 - Increase the number of participants we can support
 - Coordinate and collaborate with external research organizations (*DNS-OARC!*)
 - Hold annual workshops for collaboration

Our First Testbed Experiment

- Goal: compare operational bind with an experimental knot deployment
- Production systems: ISC's Bind
- DIINER testbed: Knot
- Traffic used:
 - mirrored B-Root DNS/UDP



Come Run an Experiment

- Initial testbed is operational today
- **We need you for beta testing**
 - Seeking early adopters
 - Mail: Wes Hardaker <hardaker@isi.edu> or John Heidemann johnh@isi.edu
 - <https://ant.isi.edu/diiner>
- **Future goals**
 - Partnerships with organizations and researchers (you!)
 - Host additional DINR collaboration events
 - Testbed expansion plans of both hardware and software