

# Preparing for the 2nd DNS SSR Symposium

Dear Conference Participant,

In advance of the symposium, we would like to ask for a few minutes of your time so that we can start working towards the best possible outcome. We look forward to meeting with you in person in Kyoto next week.

*Hiroki Takakura, Co-chair*  
*Yurie Ito, Co-chair*

## Symposium Goal

The symposium will work along two main avenues of approach:

1. Understanding the meaning of “health” as it pertains to the DNS system and reviewing the **current state of the art** of measuring its health.
2. Identifying gaps in existing **techniques, mechanisms** and **metrics** for measuring the state of DNS system health.
3. Developing recommendations for improvements in how to monitor the system’s condition.

The main objective of the symposium will be to define what we need to understand in order to assess whether the DNS system is healthy and to determine what we need to measure in order to make such an assessment.

The definitions at which we arrive, and our understanding of them, can be used by a wide variety of people involved in the DNS system, from DNS operators to researchers, to not only measure the health of their own systems, but also to refine and study these measurements over time.

The symposium can touch upon issues such as instrumentation, visualization, data sharing and other contemporary problems that the participants identify.

## Method (and an assignment)

***Interaction, Improvisation, Creativity, and Involvement***  
*are key to successful outcome of the meeting.*

This symposium has been organized as a relatively free-format brainstorming session. Within the track session, there are no prepared lectures or presentations. Instead, led by session chairs, *participants are invited to share their perspective* to arrive at a common understanding of the current thoughts about DNS health, as well as the methodologies to measure that health.

The format of the symposium specifically allows divergence from these topics if the participants feel other venues will lead to the symposium’s goal more easily.

The agenda is separated into 4 main sections divided over the final two days. We would like to ask you to spend some time, *e.g.* during your flight, thinking about the issues for each session and try to be prepared to share that vision in a casual 3-to-5 minute presentation (no slides, please).

The four themes on the agenda are:

1. Global Internet-wide view:
  - a. What is the “Health” of the DNS in the context of Internet functionality?
  - b. How do failures in one part of the Internet impact other parts?

2. Taking a more local perspective:
  - a. Is “Health” a uniform and global property? Are there lessons we can learn about “health” of the DNS from the context of medical practice?
  - b. What are the key performance indicators that can be extracted from raw measurements?
3. Measuring the DNS:

Is Loose Coherency a disappearing norm?

  - What constitutes an “accurate” and “timely” answer?
  - How much does perspective influence the measurement?
  - Does the industry have common metrics?
4. Gap analysis
  - a. What is it that we want to learn and do we have the means to learn it?
  - b. Are all actors involved?
  - c. Recommendations?

In each session, the chairs will ask for volunteers to share their perceptive first. However, if that does not result in any discussion, the session chairs may ask any of the participants to share their ideas and perspectives.

# Brainstorms

*In the remainder of this memo you will find a number of brainstorms, thoughts and questions that may help you to structure your own thoughts on system health and the measurement of it. There may be significant overlap in these approaches.*

## **1st Brainstorm: What is the health of the DNS in relation to Internet Functionality and how do we measure that?**

1. Is health a uniform and global property? Does “health” depend on the point of the observer? Are there lessons that we can learn from e.g. medicinal practice?
  2. Loose Coherency: a disappearing norm?
    - a. What is an accurate and timely answer?
    - b. What is universally acceptable?
    - c. How does local divergence influence the system?
      - (1) How does coherency in the Root impact coherency of TLDs
      - (2) How does coherency in the TLD impact coherency at lower level
  3. How do failures in one part of the system impact other parts?
    - a. Within the DNS system itself
    - b. What is the impact on the Internet?
- 

## **2nd Brainstorm: Recognizing patterns**

4. What are you measuring today?
  5. Do you have an empirical model of the behavior from which you can detect changes? What changes are considered unhealthy?
  6. Do you have a qualitative model of the behavior?
- 

## **3rd Brainstorm: Phenomena in the DNS that have an impact on Health**

As an approach to assess the healthiness a system one could look at the phenomena that impact the system, and the phenomena that a system could cause.

1. What are the various DNS related phenomena one can look at to study if there can be a health impact?

Can an issue/phenomena have impact on the health of the DNS system? If so:

Is there a hypothesis on the effect on health in qualitative terms?

    - If so: Is there a quantitative understanding of the extent in which this phenomena/topic/issue occurs?
    - If not: What and where does one need to measure to reach quantitative understanding of the phenomenon
  2. Is there possible coupling between the qualitative behavior of identified phenomenon with other phenomenon?
  3. Is there a hypothesis on how that interaction works qualitatively?
  4. What and where does one need to measure to reach quantitative understanding of the phenomenon?
-

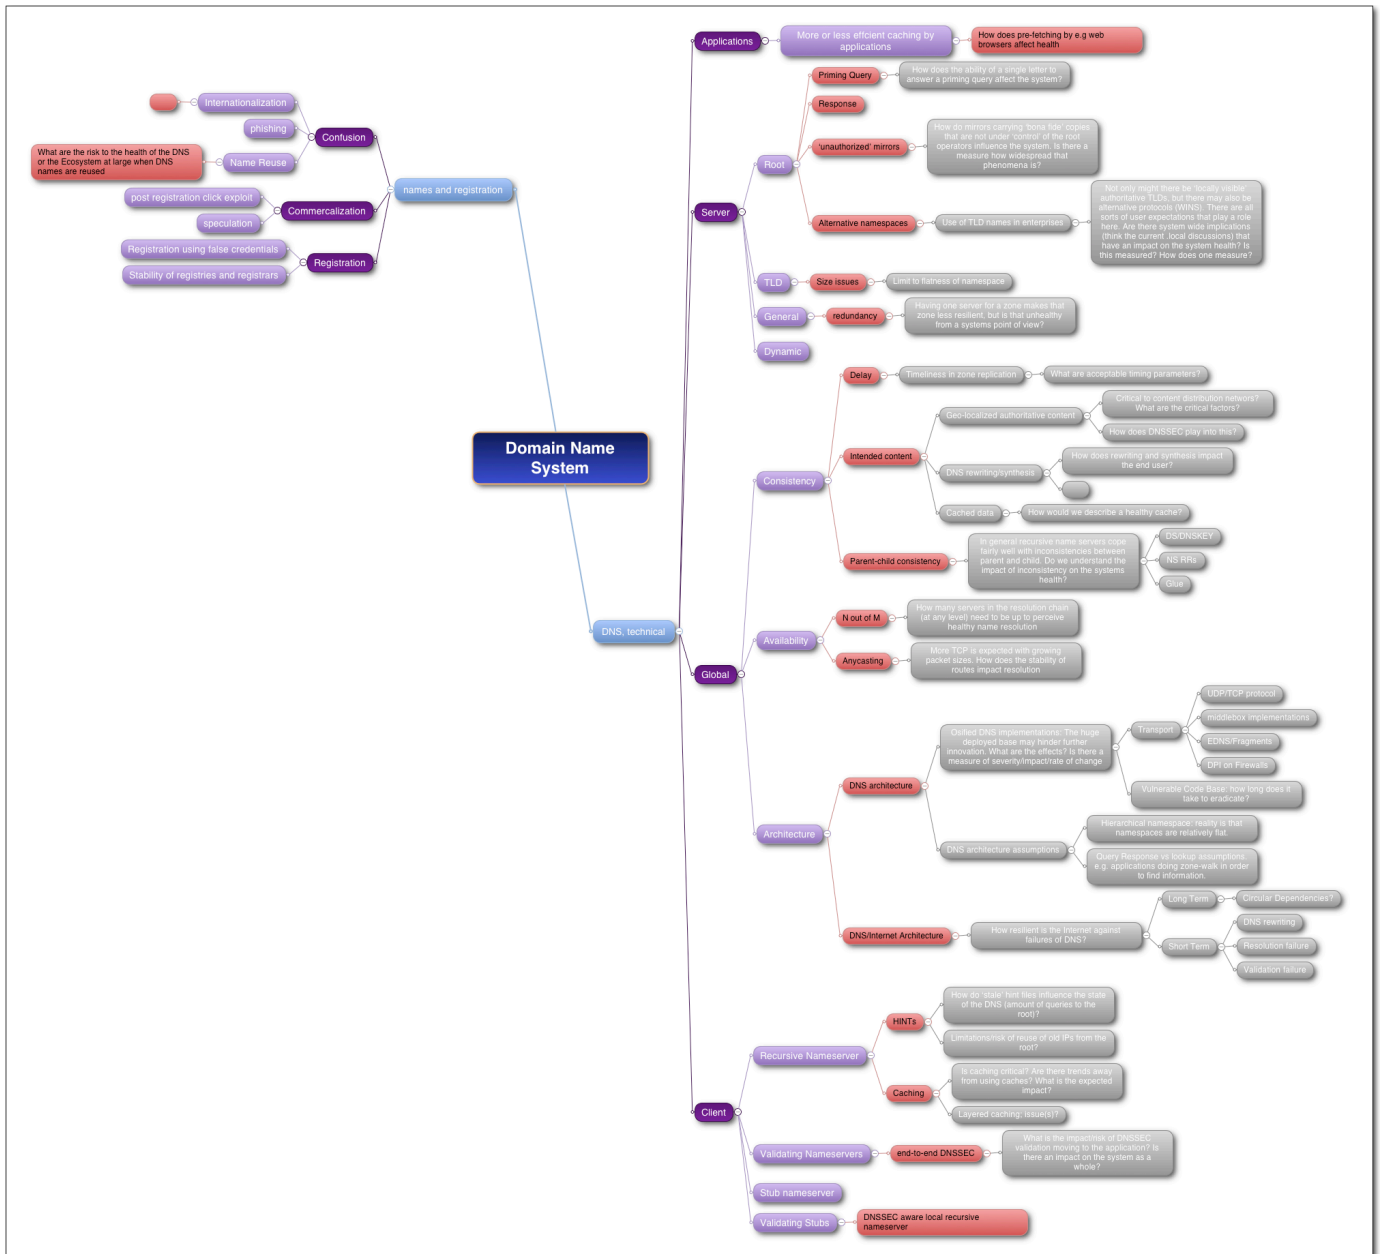
#### **4th Brainstorm: Understanding and measuring parts of the system**

1. We have to understand all the parts of the DNS system, value them for what role they play in the system. Understand how these subsystems interact. To be complete, a system can be a DNS server, a resolver, or the service provided by a TLD, or the root-system, or the ISP's resolver setup, or Google's recursive DNS service, etc. We can classify all these systems, so that they can be compared within their own class. Part of the symposium can be a brainstorm on what parts constitutes a system.
  2. If all the parts of the system are providing the services required (i.e. non of the parts are failing) we can call the system balanced. Which doesn't mean they perform well. Most parts can be measured, be it in CPU load, queries per second, records per zone, zones per host, etc. This part of measurement all relates to performance. Performances can only be compared when using the same units of course. A better performance of a subsystem might actually deteriorate overall performance. Again, this is a notion of balance. Part of the symposium can be a discussion on what can be measured, how it can be measured.
  3. When subsystem performances are measured using a common unit, against a reference baseline, we can compare and speak of a healthy system, or an unhealthy system. This part is really about setting that reference baseline.
-

# 5th Brainstorm: Questions related to parts of the DNS system

This is a longer brainstorm that looks at various parts of the DNS system. The text below is roughly based on the mind-map represented in the picture.

*It is not necessary to understand any of the detail in the mind-map; the goal of the picture is to get across that there are many different ways to look at the Symposium's topic.*



## 1. DNS, technical

The approach we take is looking at Applications that use the DNS, the Global operation, Recursive Caching Name servers and the client side of the system

### 1.1. Applications behavior

Some applications optimize their response time caching or prefetching data. This increases the load on the system.

1.1.1. How does pre-fetching by e.g web browsers affect health?

1.1.2. Walking the DNS tree (looking for zone cuts or specific authoritative data)

### 1.2. Server

#### 1.2.1. Root

What are (specific) root-zone phenomena?

##### 1.2.1.1. Priming Query

1.2.1.1.1. How does the ability of a single letter to answer a priming query affect the system?

##### 1.2.1.2. Response

Issues specific to responses to the root

1.2.1.2.1. [nothing identified yet]

##### 1.2.1.3. 'unauthorized' mirrors

1.2.1.3.1. How do mirrors carrying 'bona fide' copies that are not under 'control' of the root operators influence the system.

1.2.1.3.2. Is there a measure how widespread that phenomena is?

##### 1.2.1.4. Alternative namespaces

Use of TLD names in enterprises

1.2.1.4.1. Not only might there be 'locally visible' authoritative TLDs, but there may also be alternative protocols (WINS). There are all sorts of user expectations that play a role here.

1.2.1.4.2. Are there system wide implications (think the current .local discussions) that have an impact on the system health? How does one measure?

#### 1.2.2. TLD

##### 1.2.2.1. Size issues

###### 1.2.2.1.1. Limit to flatness of namespace

1.2.2.1.1.1. Are there any issues with respect to growing size of zone files?

1.2.2.1.1.2. Are there any issues with respect to moving to longer domain names?

#### 1.2.3. General DNS

Questions that are relevant with respect to operations at all levels

##### 1.2.3.1. redundancy

1.2.3.1.1. Having one server for a zone makes that zone less resilient, but is that unhealthy from a systems point of view?

1.2.3.1.2. The effects of having many servers serving the same zone.

##### 1.2.3.2. Dynamic

1.2.3.2.1. Does dynamic update of zones have a health impact?

### 1.3. Global

What are the global effects of operational practices

#### 1.3.1. Consistency

Phenomena that have an impact on the consistency/coherency of data obtained from the DNS

1.3.1.1. Delay

- 1.3.1.1.1. Timeliness in zone replication  
such as AXFR, IXFR replication
- 1.3.1.1.2. What are acceptable timing parameters?  
SOA values
- 1.3.1.2. 'Intended' content
  - 1.3.1.2.1. Geo-localized authoritative content  
Critical to content distribution networks? What are the critical factors? How does DNSSEC play into this?
  - 1.3.1.2.2. DNS rewriting/synthesis  
How does rewriting and synthesis impact the end user?  
Viva NXDOMAIN !?
    - 1.3.1.2.2.1. Wildcards at authoritative servers
    - 1.3.1.2.2.2. Recursive nameservers
  - 1.3.1.2.3. Cached data  
How would we describe a healthy cache?
- 1.3.1.3. Parent-child consistency  
In general recursive name servers cope fairly well with inconsistencies between parent and child. Do we understand the impact of inconsistency on the systems health for the following cases:
  - 1.3.1.3.1. DS/DNSKEY
  - 1.3.1.3.2. NS RRs
  - 1.3.1.3.3. Glue
- 1.3.2. Availability
  - 1.3.2.1. N out of M  
How many servers in the resolution chain (at any level) need to be up to perceive healthy name resolution
  - 1.3.2.2. Anycasting  
More TCP is expected with growing packet sizes. How does the stability of routes impact resolution
- 1.4. Architecture
  - Issues that have to do with the Internet/DNS architecture
  - 1.4.1. DNS architecture
    - 1.4.1.1. Osified DNS implementations: The huge deployed base may hinder further innovation. What are the effects? Is there a measure of severity/impact/rate of change
      - 1.4.1.1.1. Transport  
There are a number of issues that have to do with transport. Do we consider these health issues?
        - 1.4.1.1.1.1. UDP/TCP protocol
        - 1.4.1.1.1.2. Middleboxes with limited capabilities
        - 1.4.1.1.1.3. EDNS/Fragments
        - 1.4.1.1.1.4. DPI on Firewalls
      - 1.4.1.1.2. Vulnerable Code Base: how long does it take to eradicate?
    - 1.4.1.2. DNS architecture assumptions
      - 1.4.1.2.1. Hierarchical namespace: reality is that namespaces are relatively flat.
      - 1.4.1.2.2. Query Response vs lookup assumptions. e.g. applications doing zone-walk in order to find information.
  - 1.4.2. Internet Architecture
    - 1.4.2.1. How resilient is the Internet against failures of DNS?

- 1.4.2.1.1. Long Term
  - 1.4.2.1.1.1. Circular Dependencies?
- 1.4.2.1.2. Short Term
  - What is an acceptable level of issues of the following kind occurring?
  - 1.4.2.1.2.1. DNS rewriting
  - 1.4.2.1.2.2. Resolution failure
  - 1.4.2.1.2.3. Validation failure

## 1.5. Client

Specific perspective from the client side of the DNS architecture in its various appearances

### 1.5.1. Recursive Nameserver

#### 1.5.1.1. HINT

- 1.5.1.1.1. How do 'stale' hint files influence the state of the DNS (amount of queries to the root)?
- 1.5.1.1.2. Limitations/risk of reuse of port 53 on servers that used to be in the HINTS file?

#### 1.5.1.2. Caching

- 1.5.1.2.1. Is caching critical? Are there trends away from using caches? What is the expected impact?
- 1.5.1.2.2. Layered caching; are there issues when a cache is obtaining data from other caches, does DNSSEC change that picture?

#### 1.5.1.3. Validating Nameservers

End-to-end DNSSEC is a requirement

- 1.5.1.3.1. What is the impact/risk of DNSSEC validation moving to the application? Is there an impact on the system as a whole?

#### 1.5.1.4. Stub nameservers

- 1.5.1.4.1. Validating Stubs
- 1.5.1.4.2. DNSSEC aware local recursive nameserver

## 2. Names and registration

Looking at the issues that have little to do with the technical limitations of the DNS but issues that have to do with the use of domain names as and in identifiers.

### 2.1. Confusion

There are various phenomena that can cause confusion

#### 2.1.1. Internationalization

Mixing of scripts

#### 2.1.2. Phishing

#### 2.1.3. Name Reuse

What are the risk to the health of the DNS or the Ecosystem at large when DNS names are reused

##### 2.1.3.1. Speculation

##### 2.1.3.2. Reuse of names that have expired.

What are the potential issues?

##### 2.1.3.2.1. Running services to harvest data

#### 2.1.4. Registration issues

##### 2.1.4.1. Registration using false credentials

##### 2.1.4.2. Stability of registries and registrars

