

MUST Shuffle Resource Records

Why Shuffling DNS Answers Turns Out to Matter

DNS OARC
2024-10-26
Prague, Czechia

Shane Kerr <shane.kerr@ibm.com>
Back-end Engineer

IBM

The Problem

New mostly-secondary customer was using simple DNS-based load balancing to spread load out across their servers.

This means that they have multiple address records for a given name, and rely on clients connecting to them more-or-less at random.

IBM NS1 has never randomized (or "shuffled") secondary answers.

We only shuffle responses if the customer is using the SHUFFLE filter chain on a specific record.

Our customer saw huge bias in results hitting their servers. These were coming directly from *their* customers, who were using an unknown set of resolvers.

IBM NS1's Solution

We updated our DNS server to allow us to flag whether a given customer wants their answers shuffled.

We cannot just randomize for everyone. This might break things in the other direction!

That is, if all of a sudden a customer started getting results distributed at random, it might impact their services in unexpected ways.

Customer issue solved. 🎉

Is This a Real Problem?

We trust that this really impacted our customer!

But is this a generic problem or some artifact of a specific setup, or their customers' setups?

We have always returned consistent answers since NS1 was founded more than a decade ago.

We rely on resolvers to randomize answers.

If a resolver does not randomize, then it will skew results for any clients that use it.

Other authoritative servers return results in a deterministic order.

I believe NSD also works this way, and probably more.

Possibly the DNS standards should be more explicit?

We can use RIPE Atlas do a simple experiment to see what resolvers do.

1. Configure a zone with an A record with 5 answers.
2. Query the zone from 50 random RIPE Atlas probes using their local resolvers.
3. Download the results, parse the packets, figure out if the answers are in order.

Repeat this experiment, but add a SHUFFLE filter chain to the record. (This randomizes the order of records returned.)

Without randomization/shuffling on our side:

28 of 50 probes got the results back in-order.

18 of these seemed to be using Cloudflare 1.1.1.1 or similar.

With randomization/shuffling on our side:

0 of 50 probes got the results back in-order.

By chance, we would expect to see 1 in 720 results back in order so 28 of 50 is a strong signal that resolvers are not randomizing.

Note: Cloudflare randomizes answers from cache, not initial response.

Raw Data

<https://atlas.ripe.net/measurements/80339030>

<https://atlas.ripe.net/measurements/80339222>

Ask me for Python code used for analysis, if you want to see it. (It's only 32 lines of Python.)

Now What?

Consider changing our defaults so we randomize for new customers by default?

Consider approaching existing customers and offering to enable shuffling?

Make a strong recommendation to our customers to use SHUFFLE and SELECT 1 filters instead of multiple answers?

Contact public resolver operator and see if they want to change their behavior (in progress).

Present results to DNS community and/or blog it. (You are here!)

Possibly make an IETF Internet draft which says basically:

"SHUFFLE ALL THE THINGS!"

As an authority server, not randomizing cannot help you (since most resolvers do shuffle), but it can hurt you (since some resolvers do not shuffle).

Questions?

© 2024 International Business Machines Corporation

IBM and the IBM logo are trademarks of IBM Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

THIS DOCUMENT IS DISTRIBUTED “AS IS” WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT, SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

Client examples are presented as illustrations of how those clients have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

Not all offerings are available in every country in which IBM operates.

Any statements regarding IBM’s future direction, intent or product plans are subject to change or withdrawal without notice.