

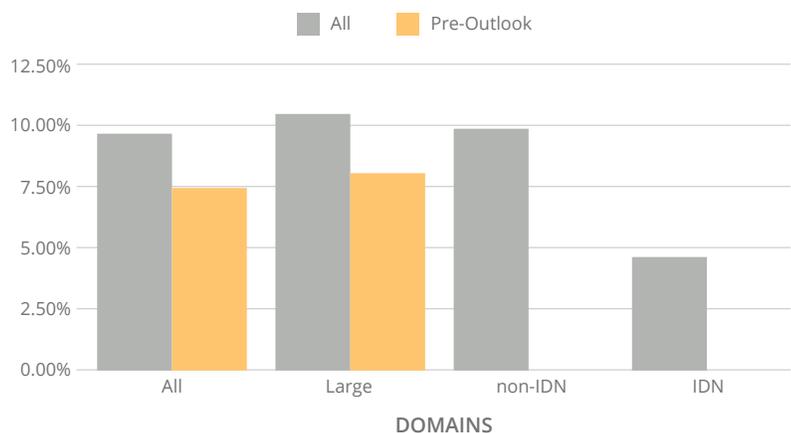
EAI Readiness in TLDs

Overview

Email Address Internationalization (EAI) has been slowly growing in acceptance. How widely accepted is it? We made an informal survey of mail servers across domains registered in hundreds of top-level domains (TLDs) to see whether they have mail servers that accept EAI mail.

TLDs ranged from the largest, .COM, down to tiny TLDs with a handful of names. Overall, 9.7% of the domains sampled were EAI ready. In the largest TLDs (over a million names), 10.5% of the domains sampled were EAI ready. Microsoft's Outlook.com recently became EAI ready; before that the numbers would have been 7.41% and 7.93%. IDN TLDs were significantly less ready than non-IDN, 4.7% rather than 9.8%, but most IDN TLDs are small and none of the ones sampled have many active mail servers.

All and Pre-Outlook %



The fraction of domains within a TLD that accepted EAI mail ranged from 100% down to 0%, with the extremes all being in domains with very few mail servers. In domains with over 50,000 registrations, it ranged from a high of 25% of sampled domains being EAI ready in .solutions to a low of none in .ren.

What did we check?

We started with the zone files for about 1250 TLDs. Most of them were gTLD zones from ICANN's CZDS (Centralized Zone Data Service.) We also have zone files of older TLDs such as .com, .org, and .info with whom we have access agreements, and about a dozen ccTLDs that either have access agreements or allow zone downloads from their DNS servers. The largest zone is .COM with over 100 million names, the smallest are some new gTLDs that haven't started registering names yet with one or two. Over 400 of the TLDs had no MX records at all due to very few registered names.

Testing methodology

Our tests used two programs. The first read through the zone files, trying to look up MX records for each domain. In small TLDs, we checked every name. In larger ones, we sampled every 50th name until we found 1000 MXes or ran out of names. For .com, which is much larger than any other TLD, we sampled until we found 10,000 MXes. The program creates a list of domain, MX, IPv4 address, and optionally IPv6 address. For example, here are some of the MXes from the .post TLD:

malawipost.post. webmail.malawipost.post. 193.247.58.16

correios.mo.post. mail.correios.mo.post. 202.175.50.70

postakenya.post. mail.postakenya.post. 139.162.254.245

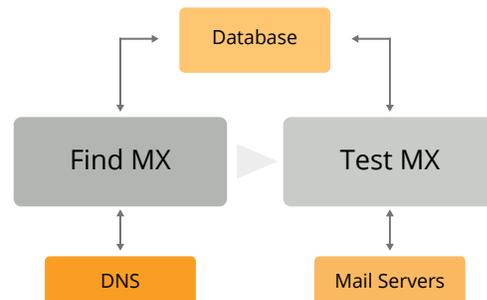
If a domain has more than one MX record, the program chose the highest priority record (the one with the lowest priority number) since that is the one that sending systems will try first.

The list of domains and MXes was passed to the second program, which tested each MX for EAI readiness by doing a partial SMTP session. First it connected to the server and received the SMTP banner. Then it issued an EHLO command, and received the EHLO response, then it issued a QUIT command and disconnected. The SMTP banner often contains the name of the mail server software, such as “gsmtp” for Google’s gmail, or “Postfix” or “Exim”, open source mail programs. The EHLO response contains a list of keywords identifying the optional SMTP features that the server supports. We check for the SMTPUTF8 keyword that means the server accepts EAI mail, and also for the 8BITMIME keyword that means that the server can handle mail with unencoded UTF-8 characters. For each server tested, we remember whether it supports 8BITMIME and SMTPUTF8 and the server software if the banner identified it. A surprising number of MX records point at non-existent or broken servers, so if the test failed, we remember how it failed, e.g., invalid address like 0.0.0.0, rejected connection, or timeout.

To make the testing process reasonably fast, we used a great deal of parallelism and caching. The MX lookups were done in parallel batches of 1000 at a time, and the mail server tests 100 at a time. To avoid overloading the servers we were testing and to avoid looking too much like a spambot, we only tested each server once no matter how many domains used it. Nearly every server we checked handled 8BITMIME whether or not it handled SMTPUTF8, other than old versions of MailEnable and sendmail.

Notes on the numbers

Our methodology over-represents small TLDs since we tried to get about the same number of samples from each TLD regardless of size. The estimate of 10.5% for TLDs over 300K names somewhat avoids the oversampling. The 9.7% number for all domains is clearly too high due to statistical quirks. For example, the .kred TLD has 6700 names, but only four of them have MX records. Those four MXes all happen to point to gmail, which is EAI ready, so kred is listed as 100% EAI ready even though most of its domains have no mail at all.



Mail Software and EAI

We identified 15 popular packages and systems. This list is far from complete since there is a great deal of other mail software, and many mail systems don't say what software they're running.

Mail Packages

These are a variety of open source and commercial packages.

Postfix: popular open source mail package for Unix and Linux systems originally from IBM. It has EAI support, but it's not turned on by default.

Sendmail: the original Unix mail program, now largely supplanted by Postfix. Both free and commercial versions now maintained by Proofpoint. It also has EAI support not turned on by default.

Exim: popular open source mail package for Unix and Linux systems from the University of Cambridge (UK). It also has EAI support not turned on by default.

MailSite: Commercial Windows mail package. No EAI support

Microsoft: Exchange servers, different from hosted outlook.com

Haraka: high performance open source mail package written in Javascript using the node.js framework. Recent versions support EAI.

MailEnable: Commercial Windows mail package. No EAI support.

MDaemon: Commercial Windows mail package. No EAI support.

CommuniGate Pro: Commercial multi-platform messaging package. Optional EAI support.

Hosted Mail

These are popular hosted systems that support many domains.

Gsmtp: Google's mail system. EAI ready.

Outlook.com: Microsoft's mail system also known as Hotmail and Live.com. EAI ready.

Nemesis: Mail system from 1&1/Schlund, the large German-based hosting provider. No EAI support.

Amazon SES: Hosted mail service for Amazon cloud customers. No EAI support.

Bizsmtp: Godaddy's hosted mail service. No EAI support.

About Email Address Internationalization (EAI)

Email Address Internationalization (EAI) is the protocol that allows email addresses with internationalized domain names (IDNs) in the domain part and/or Unicode (non-ASCII) characters in the mailbox name to function within the traditional email environment. Email software and services providers need to make specific changes to support EAI, which is an important component of Universal Acceptance. These changes include making sure that their applications support Unicode characters and that their application announce support of EAI when conversing with other email systems.

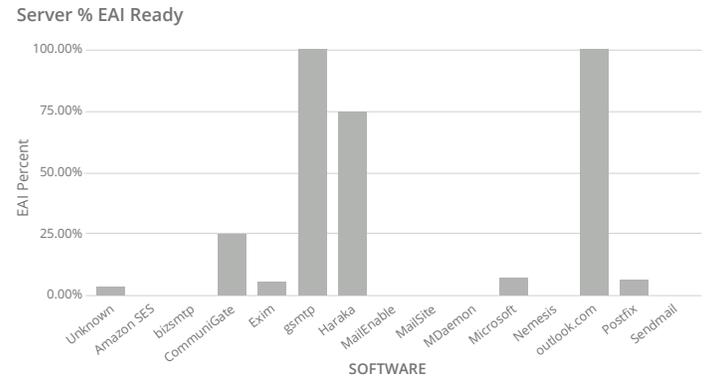
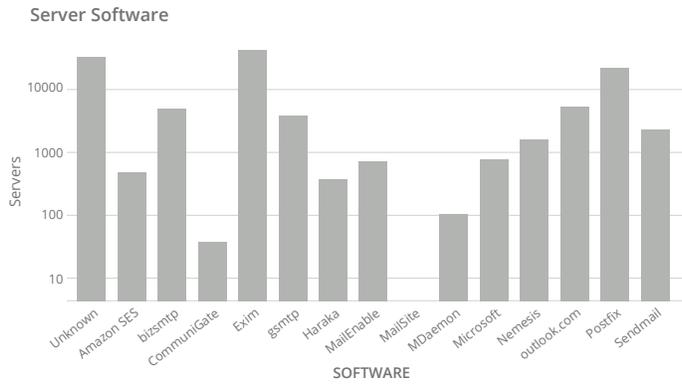
About Universal Acceptance Steering Group (UASG)

The Universal Acceptance Steering Group is an Internet community initiative that was founded in February 2015 and tasked with undertaking activities that will effectively promote the Universal Acceptance of all valid domain names and email addresses. The group is made up of members from more than 120 companies (including Apple, GoDaddy, Google, Microsoft and Verisign), governments and community groups. The UASG receives significant financial and administrative support from ICANN. For more information, visit <https://uasg.tech/>.



Report: EAI Readiness in TLDs

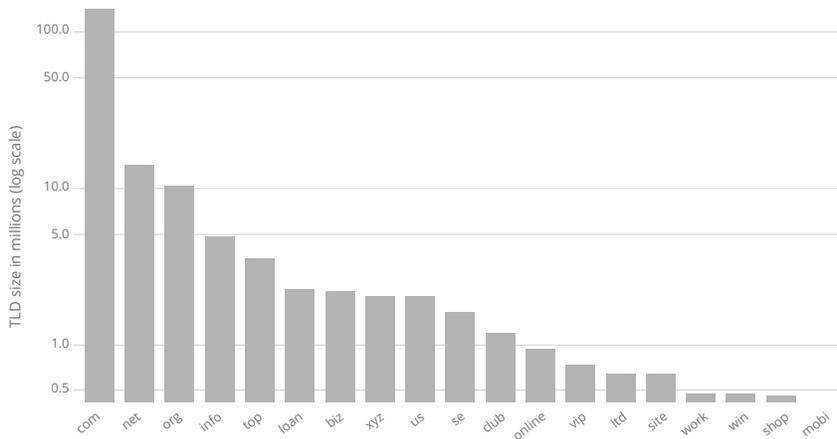
Summary numbers



Hosted mail

These are the largest domains we sampled, and their percentage of EAI ready domains.

Largest TLDs



EAI Ready Domains in Large TLDs

