

Load Balancing Email Using DNS MX Records

There are two ways to load balance email using standard DNS MX records. The first technique is to define multiple MX records with the same priority. The second is to define multiple A records with the same name and different IP addresses. In addition to providing load balancing, both methods also provide automatic failover.

DNS and MX Records

DNS (Domain Name System) provides a lookup of domain names to IP addresses. When given an email address that contains a domain name, DNS can determine what IP address the email is sent to. The sending email server performs a lookup of the MX record on the destination domain of the email address to find out where the email is sent. The MX record typically returns a machine name such as `mail.mydomain.com`. Next, a second DNS lookup is performed to translate that machine name into an IP address. The sending email server can then open communications with the destination email server and transfer the email via the SMTP protocol.

Defining Multiple MX Records with the Same Priority

The first method of load balancing using MX records is to define multiple MX records with the same priority:

Record *FQDN		Record Type	MX Preference	Record Value
mydomain.com	IN	MX	10	mail1.mydomain.com.
mydomain.com	IN	MX	10	mail2.mydomain.com.
mydomain.com	IN	MX	10	mail3.mydomain.com.
...				
mail1.mydomain.com.	IN	A		192.168.2.1
mail2.mydomain.com.	IN	A		192.168.2.2
mail3.mydomain.com.	IN	A		192.168.2.3

**Fully Qualified Domain Name*

The MX preference field is relative to any other MX record for the zone and can be any value between 0 and 65535 although low values are preferred and, by convention only, the preferred value is 10. Any number of MX records may be defined.

In the example above, the MX preference of 10 represents an equal weight amongst all three SMTP servers. When this method is used and a sending SMTP server looks up the MX record, the DNS server delivers all IP addresses associated with the MX records, indicated in the order defined by the `rrset-order`. The `rrset-order` defines the order in which multiple records of the same weight are returned in DNS. This statement is fully implemented in BIND versions 9.2.3 and higher. By default, the `rrset-order` returns the MX record information in random order followed by a cyclic order (round-robin). The sending email server then selects one IP address based on its own SMTP algorithm and uses it for delivery.

In some cases, the predetermined SMTP algorithm may cancel out the definition of the `rrset-order` statement. The latest versions of Sendmail, Exim, and Postfix clearly document that they randomly select from equal preference servers (Postfix allows control of this selection with the `smtp_randomize_addresses`). As of this writing, Qmail, Courier-mta, and Microsoft (Exchange and IIS SMTP) do not appear to have documentation that indicates how they handle MX records of equal priority.

Based on the SMTP algorithm of the sending email server, the load is balanced between the multiple destination email servers randomly returned by DNS. This method works best in protecting email flow due to the multiple MX options available in DNS. In the case of a non-reachable MX record, the sending SMTP server selects another MX record to send to, ensuring email is delivered.

RELEASE 1
SEPTEMBER 2005

Load Balancing

distributes communications and processing activity evenly across a computer network so that no single device is overloaded.

MX record (Mail Exchange record)

is an entry in a domain name database that identifies the mail server that is responsible for handling emails for that domain name.

DNS (Domain Name System)

is an internet service that translates domain names into IP addresses.

SMTP (Simple Mail Transfer Protocol)

is a protocol for sending email messages between servers.

Defining Multiple A Records With the Same Name and Different IP Addresses

The second method of load balancing using MX records is to define multiple A records with the same name and different IP addresses:

Record *FQDN		Record Type	MX Preference	Record Value
mydomain.com.	IN	MX	10	mail1.mydomain.com.
...				
mail.mydomain.com.	IN	A		192.168.2.1
	IN	A		192.168.2.2
	IN	A		192.168.2.3

*Fully Qualified Domain Name

When multiple A records are defined, load balancing is under the full control of the DNS Server. This technique eliminates the dependency on the SMTP algorithm of the sending email server. The DNS server follows the rrset-order record and delivers the IP addresses in random order. The sending email server then sends email to the first IP address.

In order to avoid complications when using multiple A records, make sure to add PTR records for all A records. If the receiving email system does reverse look-ups, the PTR records for all the A records must point to mail.mydomain.com. If all the A records do not point to mail.mydomain.com, the receiving email system can detect a mismatch and may reject the email if this is part of its anti-spam algorithm.

For questions about the Barracuda Load Balancer, please visit <http://www.barracuda.com/load> or call Barracuda Networks for a free 30-day evaluation at 1-888-ANTI-SPAM or +1 408-342-5400. For more information on our other security and productivity solutions, please visit <http://www.barracuda.com/products>.

About Barracuda Networks Inc.

Barracuda Networks Inc. combines premise-based gateways and software, cloud services, and sophisticated remote support to deliver comprehensive security, networking and storage solutions. The company's expansive product portfolio includes offerings for protection against email, Web and IM threats as well as products that improve application delivery and network access, message archiving, backup and data protection.

Coca-Cola, FedEx, Harvard University, IBM, L'Oreal, and Europcar are among the more than 100,000 organizations protecting their IT infrastructures with Barracuda Networks' range of affordable, easy-to-deploy and manage solutions. Barracuda Networks is privately held with its International headquarters in Campbell, Calif. For more information, please visit www.barracudanetworks.com.

PTR (Reverse Pointer Records) are used only in reverse (IN-ADDR.ARPA) domains. There must be exactly one PTR record for each Internet address.



Barracuda Networks
3175 S. Winchester Boulevard
Campbell, CA 95008
United States
+1 408.342.5400
www.barracuda.com
info@barracuda.com