Zoho ZeptoMail

Email infrastructure and how email sending works

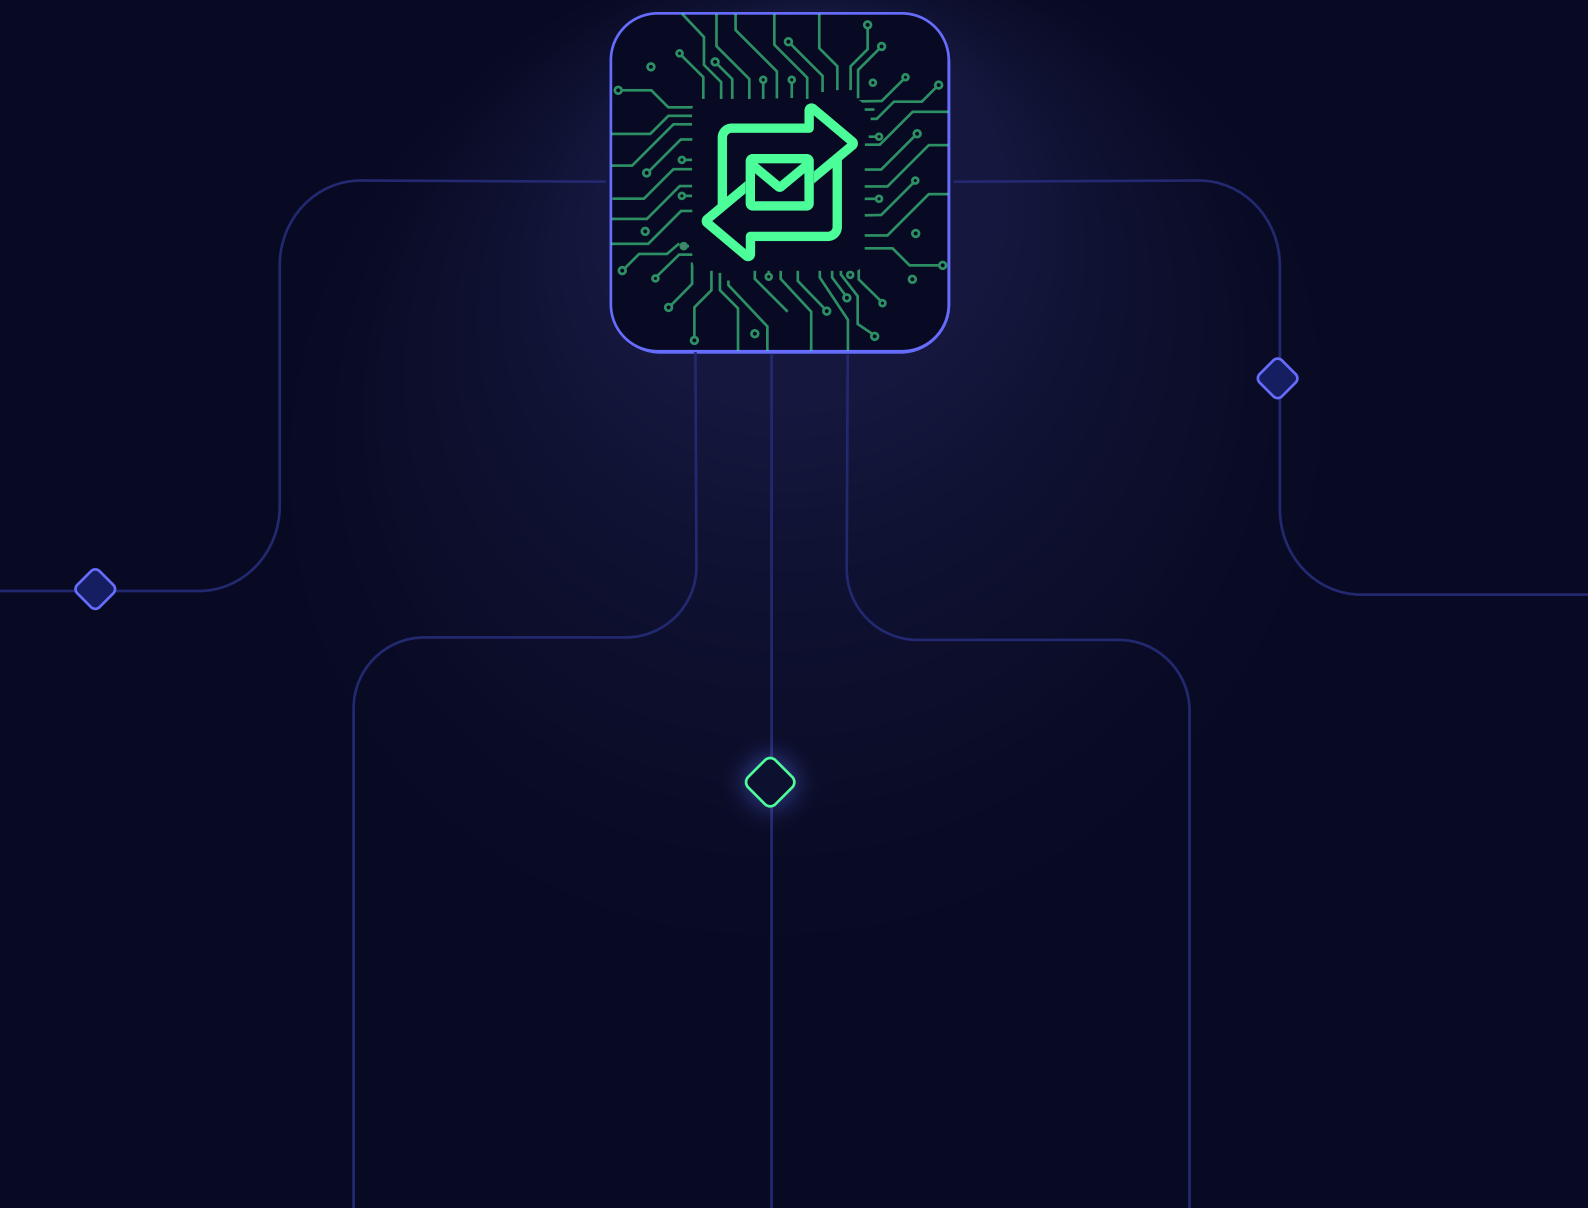

How is an email sent? You open your email client or application, type out your email, and hit send. And magically, the email that you typed is dropped into the inbox of the intended recipient. This is what every user sees. But the successful delivery of the email has less to do with magic and more to do with a series of well-designed protocols and processes. Once you hit the send button, multiple gears turn to make the delivery happen, and if you're curious to know how this works, this article has everything you need.

Table of Contents

1.A brief history of email

Email is short for electronic mail. It is the internet-equivalent of a physical mail that goes through the postal system. Instead of the postal system, an email goes through electronic devices and an email infrastructure.

Let's begin with a brief history of how today's email came to be. The origin of email pre-dates the invention of the internet itself. It could even be considered the fueling force in the invention of the internet.

1965: Noel Morris and Tom Van Vleck from Massachusetts Institute of Technology (MIT) created a mail command that could leave a message to users that were sharing a single computer. Back then, computers were often shared by big groups of people, and their mail command allowed communication between different users of the same computer. While this is a long way off from what email means today, it's still the first recorded concept of an electronic message.

1971: Ray Tomlinson invented an email protocol closest to current day email that could send messages between different computers on the same system. In his case, his program would communicate between computers on the ARPANET system. ARPANET then eventually evolved into the internet we know today.

1976: Email attracted attention when Queen Elizabeth II tried out ARPANET's electronic mail program during a visit to Malvern, England. Her username was HME2.

1978: Where there's email, there's spam. Gary Theurk sent the first known spam email when he sent marketing messages to hundreds of ARPANET users. He's sometimes referred to as the father of spam.

During this same time, V.A Shiva Ayyadurai laid claim to the title of "inventor of email," as he's said to have created an interoffice software that digitized physical inboxes and outboxes and called it "email." He claims that earlier versions were more similar to text messages, and his program is the first version of what we know as email. Irrespective of whether he really was the inventor of email, he's most likely the person that coined the term "email."

1982: Simple Mail Transfer Protocol (SMTP) came into the picture to advance the way servers received and sent out emails.

1988: Microsoft launched its first email product, MSMail, that was available for the public to purchase and use

Early 1990s: Somewhere between 1993 and 1996, the first version of webmail was created. Previously, you had to use a specific program to send and receive emails. The first web-based email program was launched in 1995.

In the following years, email has been upgraded repeatedly. The email we know and use today is an amalgamation of all these improvements.

The modern email is used by different entities for different purposes. Personal emails are used to communicate with family and friends, and business emails to communicate with customers and colleagues. Even with many new and fancy modes of communication today, email is still considered to be the king of official communication for most businesses. Whether it's one-to-one business emails, bulk marketing emails to promote their business, or transactional emails that convey acknowledgement to a customer, email plays a key role in the success of a business. Understanding how an email is sent will allow businesses to leverage emails better for their business.

2. Elements of an email

An email contains three parts—the envelope, header, and body. Let's look at an example of an actual physical letter to explain these elements.

Envelope: This is equivalent to the actual envelope that carries the letter. It has metadata that tells the system how and where to route the email. This is often invisible to the users because it's used purely to route the email to the right server.

Header: The header is like the "From" portion inside a letter telling you who wrote it. In an email, the header contains the from address and how it was sent for the recipient's benefit. This header information isn't used for email delivery.

Body: The body is the actual letter. In the case of an email, it's the content of the email. It includes the text, images, videos, and any other data-like attachments.

3. Email infrastructure

From the outside, how an email is sent seems simple. We click on "send" and it appears at the recipient's end. But to make this possible, a meticulous email infrastructure is built and maintained.

Email infrastructure is the complex system that comprises all of the software and hardware required to deliver an email. If you run a business, you're bound to be sending out emails for various reasons like marketing and transactional emails. Each of these emails goes through this system to be delivered. So if you're a business owner or associated with sending email for a business, understanding email infrastructure at least at a basic level goes beyond just satisfying your curiosity. This understanding can help troubleshoot any future issues or even improve your business's email infrastructure.

4. Components of an email infrastructure

To begin understanding how an email is sent, we'll first learn about some prominent components in the email infrastructure.

4.1 Mail user agent (MUA)

An email client or webmail is also known as a mail user agent. It's a web-based or computer application that can be used to send and view emails. An MUA will download emails from a user's remote storage mailbox to display them to the user. It's the component that users generally interact with. Some examples of MUAs are Outlook, Apple Mail, Zoho Mail, and Gmail.

4.2 Mail submission agent (MSA)

An MSA acts as an intermediary between the MUA that initiates the email and the mail transfer agent (MTA) that delivers the email. An MSA receives the outgoing email from the email client (MUA) and performs various checks to authenticate the email and eliminate spam. It checks for sender credentials, formatting mistakes, spam content, and compliance with email protocols, among other things. If the email checks out, it hands the email over to the MTA to be delivered.

While an MTA can perform the action of an MSA, having a separate MSA can be beneficial in many ways. Having an MSA will notify the sender of potential mistakes before delivery attempt, whereas in the case of an MTA, the sender is notified of mistakes only after delivery is attempted. It prevents spam, phishing, and malicious content to a certain extent. This, in turn, decreases the chance of delivery issues or rejections.

4.3 Mail transfer agent (MTA)

The mail transfer agent is responsible for the actual relaying of the email. It receives the email from the MSA and transmits it to the destination. The MTA first accepts the email from the MSA. Then the MTA looks up the destination of the domain where the email has to be delivered, often using mail exchange (MX) records. The MTA uses SMTP relay to deliver the email to the next step.

4.4 Message delivery agent (MDA)

The message delivery agent receives the email from the MTA and stores it to deliver it to the MUA (email client) of the recipient. The MDA is the component of the email infrastructure that's also referred to as the recipient email server.

4.5 Email protocols

Email protocols play an important role in the email infrastructure. They're the components that enable the transport of email across these different agents. There are three prominent protocols:

- **Simple Mail Transfer Protocol (SMTP):** Where there's email, there's spam. Gary Theurk sent the first known spam email when he sent marketing messages to hundreds of ARPANET users. He's sometimes referred to as the father of spam.

- **Post Office Protocol (POP):** POP is the protocol responsible for the accessibility of the delivered email. Once the email is delivered to the recipient server, POP will retrieve the email to be viewed. After downloading the email from the server, it deletes the message from the server.

- **Internet Message Access Protocol (IMAP)**: IMAP is similar to the POP protocol in theory. It's also used to retrieve emails from the server and display them, but it's a more complex protocol and different from POP in some ways. For example, unlike POP, IMAP leaves a copy of the downloaded email on the server, and it also has the ability to maintain folder structures.

4.6 Email domains

Even the most infrequent email user would have come across email domains. If your email address is rachel@zylker.com, zylker.com is the email domain. It's equivalent to the address we write down when addressing a physical mail. It denotes where the email needs to be delivered.

Email domains can be broadly classified into two groups: shared domains and custom domains.

- **Shared domain:** These are often owned by major ESPs. They will allow users to sign up for a mailbox, and the domain will be their common domain. For example, gmail.com, outlook.com, or zohomail.com.

- **Custom domain:** Custom domains are often owned by the sender alone. For example, if you own a business called Zylker Mobiles, your custom domain will be zylkermob.com.

Whichever email sending domain you should opt for completely depends on your purpose. If you're looking for an email address for personal purposes like writing to your friends and family, then a shared domain provided by an ESP is sufficient. But if you're a business owner or using the email address for some professional use, it's best to have your own custom domain.

4.7 IP address

An Internet Protocol (IP) is your address on the internet. It contains a string of multiple numbers separated by periods, like 112.3.43.235, 222.22.4.54, or 54.2.34.023. In the context of sending email, a unique IP address is given to your domain by your Domain Name Server (DNS), which will be used to locate where your email needs to be delivered. IPs are important for data to be shared or communicated over the internet.

Much like IP addresses, there are two types of IP addresses—shared and dedicated IPs.

Shared IP: ESPs generally maintain IP addresses that are shared by a pool of users for email sending. Users with similar sending patterns are often assigned to the same pool. The email activities of each user in the shared IP will affect the other users, good or bad.

- **Dedicated IP:** These are IP addresses that aren't shared. If you're assigned a dedicated IP, only you will be sending emails from the IP. While you have complete control over your sender reputation with a dedicated IP, if you're not a high-volume sender with a well-established reputation, dedicated IPs might be detrimental.

Take a look at our article on IP addresses to see which one is a good fit for your email sending.

4.8 Domain Name Servers (DNS)

Name servers are the basic unit of a DNS system. They hold information about all of the domains they host. DNS is responsible for assigning IP addresses to domains, authentication records, and much more. During the process of email sending, DNS is constantly referred to. First, it's to navigate the email to the right recipient using the information hosted in the DNS, like MX records. Second, DNS contains all of the authentication records like SPF and DKIM that are used to check the authenticity of the email sender while the email is being delivered.

4.9 Feedback loop

A feedback loop (FBL) is an important component of email infrastructure that comes into play after the email has been delivered. FBL will notify email senders if their emails have been marked as spam. This gives the senders an opportunity to fix issues in their emails or, sometimes, even remove the recipients from their recipient list. While FBL is an important component, not every sender uses this. While a registration process is involved, it's crucial to stay informed of your email's performance so you can improve your email sending practices.

5. Types of email infrastructure

An email infrastructure is necessary for any email sending, and there are multiple types of email infrastructure to choose from. These types can be classified based on two factors—who manages and how it's managed.

Classification based on who handles the infrastructure includes:

- **Managed email infrastructure:** The infrastructure in this category is often handled by a third-party provider. The software, hardware, analytics, data, and email sending is all handled by the third party. Businesses simply subscribe to their services. This is beneficial because it saves you the time and resources that go into handling an email infrastructure.

- **In-house email infrastructure:** In contrast to managed infrastructure, the entire system is built in-house by the business sending the emails itself. There is no third-party involvement. While this gives businesses a certain level of control and customization, the strain it puts on time, resources, and budgets can be heavy. It also requires a large, well-equipped team to build and maintain the infrastructure. You can attain the control and customization required without the development and maintenance strain simply by picking a suitable third-party provider.

Classification based on how the infrastructure is managed includes:

- **Open-source:** This email infrastructure is completely free and allows change and customization to fit the user's choices. This type of system has no payment plan or subscription, which means businesses are free to scale up without having to worry about their budget. But along with the flexibility comes some strain on your resources and time. The service can also be abused by spammers, and this can affect your deliverability.

- **On-premise:** This email infrastructure hosts its data on hardware at their physical premises or another location of their choice. All software and other components run from this hardware. On-premise infrastructure provides complete

control to the business and offers the ability to customize functionalities for their specific needs. Some advantages of on-premise infrastructure include:

- **System control**
- **Functionality customization**
- **Data security**
- **Easier integration with other organizational systems**

- **Cloud-based:** The infrastructure and its associated components are hosted on a cloud service owned by a third-party provider. While this isn't complete control, it does come with its own benefits:

  - **No hardware setup or installation cost**
  - **Easy and quick to get started**
  - **Upscaling is effortless and swift**
  - **Minimal downtime with a comprehensive backup system**
  - **Less strain on time and resources for maintenance**

Zoho ZeptoMail is one such transactional email service that's hosted on the cloud and comes with reliable email infrastructure.

- **Hybrid:** This is a combination of an on-premise and cloud-based infrastructure. If there are aspects of either of these systems that you cannot give up, a hybrid email infrastructure will work best for your business. You can choose the hardware you think you'll spend time maintaining while leaving the rest to the third-party vendor. Similarly, you can keep sensitive data on-premise and the rest on the cloud. This will help you optimize your time, costs, and resources.

6. How is an email delivered?

6.1 Hit "send" on the MUA

This step is the one that everyone is familiar with. MUAs, like email clients or webmail

interfaces, are where the email sending process begins. You log into your inbox, open the mail composer, and draft your email. Once the email is ready, you hit the send button. This initiates the email delivery process.

6.2 Checks by the MSA

Once the send button is triggered, the email is transferred to the MSA. The MSA is a middleman between the MUA and MTA/SMTP server. It receives the email from the MUA and performs certain checks to see if the email is in the right has minor errors, and more.

6.3 Transfer to an MTA

Once the check is done by MSA and the email is found to be compliant, the MSA hands over the email to the MTA on the outgoing SMTP server. This email is in the MIME format (Multipurpose Internet Mail Extension) that accommodates the transfer of media files like audio, video, images, and more. After the email is received by the outgoing server, it again performs certain checks to validate the email's authenticity. The sender details, associated data, and compliance with existing rules of the email account are verified.

This transfer of email from the client to the server and additional transfers are performed using SMTP, which is a set of rules for email transfer between email clients and servers.

6.4 Pinpoint recipient location

This step can vary depending on where the recipient is located. The SMTP server checks to see if the recipient is also part of the same location as the sender; that is, the recipient address and the sender address belongs to the same domain. If this is the case, the email is kept on the same server to be accessed by the recipient's client later.

The more likely scenario is that the recipient address isn't part of the same domain. In this case, once the SMTP server is done running validating the email, it contacts the DNS to retrieve the location of the recipient in a format that's understandable by the server. To do this, the MX lookup is employed. Every email address hosted in a domain will have a MX record added to the associated DNS. This MX record contains information about the host and IP address of the email address server. Using MX lookup, the

outgoing SMTP server will fetch the details and use it to route the email in the right direction.

6.5 Server-to-server connection

Once the server identifies the location of the recipient, the sender server attempts to establish an SMTP connection with the recipient server. This is referred to as the "SMTP handshake." During this connection, the sender and recipient information is exchanged to verify authentication.

In cases where the recipient server isn't available to establish the connection directly, the email message is routed through multiple servers to facilitate the email's delivery to its final destination. This is called "SMTP relay."

6.6 Spam checks by the recipient server

Once the email reaches the recipient server or message delivery agent (MDA), it runs a series of checks to authenticate the email. MDA checks for spam, phishing, malicious content, and other security breaches using email protocols like SPF, DKIM, and DMARC.

- **Sender Policy Framework (SPF):** This determines whether the sending server has permission to send the email. costs, and resources.
- **DomainKeys Identified Mail (DKIM):** This email protocol uses a digital signing process to ensure that the email hasn't been tampered with after being sent.
- **Domain-based Message Authentication Reporting and Conformance (DMARC):** This email protocol informs the recipient server of the SPF and DKIM status that can be expected and what to do when the expected standard isn't met by an email.

6.7 Delivery to the recipient client

Once all of the checks have been performed on the email and the email has passed the verification, the email is finally accepted by the recipient server. The recipient MUA then uses the IMAP or POP protocols to download the email to the email client where it can be displayed to the recipient.

7. What happens if the delivery attempt fails?

The reasons for which the email delivery can fail are broadly classified into two groups—permanent reasons and temporary reasons.

7.1 Permanent reasons

If the email cannot be delivered due to permanent reasons, it's considered a hard bounce. The email is bounced back to the sending server with a notification of the bounce reason. Hard bounces are often characterized by 5xx series error codes.

Some possible reasons for hard bounces are:

- **Invalid email adress:** When the email has been sent to an email address that doesn't exist either due to an error in the email address or because the address was deleted.
- **Invalid domain:** When the email is sent to an email address whose domain doesn't exist.
- **Missing MX records:** MX records are used to identify the location of the recipient and route the email. If the MX record is missing, there's no way the sending server can deliver the email.

7.2 Temporary reasons

If the email can't be delivered because of temporary reasons, it's called a soft bounce. This email is added to the retry queue, and delivery is periodically tried in the hopes that the bounce reason has been resolved. If the delivery attempt fails after a certain number of times, the email will then be returned to the sending server. Soft bounces are often characterized by 4xx series error codes. Some possible reasons for soft bounce are:

- **Mailbox full:** The recipient's mailbox storage is full, so the email cannot be delivered at that instant. If the mailbox has storage left during any of the subsequent attempts, the email will be delivered.

- **Server occupied:** If the recipient server is too busy to receive the email, the email will be set aside to retry delivery.

- **Greylisted:** Sometimes the recipient server will greylist senders for temporary reasons on suspicion. This often resolves itself, and the email will be delivered in later attempts.

Some of these reasons can be placed in one category by certain email providers and another category by other providers. While in most cases these reasons fall under the categories mentioned here, it's best to check the standard set by the recipient's email provider in case of bounces.

8. Wrapping up

While how an email is sent and the component or processes involved are too vast to be understood in a single read, this article gives you an overview and a basic idea of how an email works and the key components at play. Understanding how email infrastructure works is helpful for picking the right infrastructure for your email sending that can ensure good delivery for your emails.