A *MailChimp* GUIDE

EMAIL DELIVERY

FOR IT PROFESSIONALS

Introduction

So you've been tasked with putting together an *email marketing blaster thingy.* Maybe your boss wants to save some money, the marketing department needs better email tracking, or you just love a good challenge. Whatever the case, you're the poor person who has to build the email marketing machine.

Here at MailChimp, we send a few billion emails per month, so we know how to set up an email deliverability infrastructure. We've even created our own transactional (one-to-one) email system with Mandrill. Of course, you can sign up for free MailChimp and Mandrill accounts and let us handle the tough stuff. But if you're inclined to build your own deliverability engine you've come to the right place.

In this guide, we cover the complexities and hurdles involved with setting up a delivery infrastructure. Most of the information targets the technical and social requirements. And with everything in place, you'll have a better chance of hitting that inbox.

Now, let's get started.

What Is Deliverability?

Deliverability is an art and a science which ensures an email reaches its intended recipient. It's more complex than it seems at first glance, though. Deliverability is navigated by learning the expectations of ISPs, monitoring statistics, building a solid infrastructure, and a whole lot of trial, error, and patience.

When deliverability goes wrong, IP addresses end up on blacklists, emails get bulked, and you might find yourself explaining to the boss why the *email marketing blaster thingy* isn't working. ISPs are constantly adding new systems to stop your emails, and failing to monitor and secure your delivery infrastructure is a silent killer of email marketing. With that responsibility on your shoulders, you want the best possible infrastructure in place.

The bad, the bad, and the ugly

So what's the worst that could happen if you ignore this stuff?

- **ISPs will block you.** Not a 24-hour block or a "fix issue A and we'll allow you to send again" temporary block—they'll block you indefinitely or give you a "come back in nine months when you've cleaned up your act" response.
- **You will land on blacklist after blacklist**, filling out delisting forms (if the blacklist even offers delisting) and working overtime to fix the problems.
- **People will complain directly and report your email as spam.** Trust us —it's not pleasant.

Without the technology and manpower in place, you won't know there's a problem until it's too late. Once administrators start sending notifications you're blacklisted (if they're nice enough), not much can be done.

Terminology

As we progress through this guide, we'll use terms which might be new and scary. Become familiar with these deliverability terms before proceeding:

Abuse complaints

When a subscriber clicks a "Spam" button in their email client. For ISPs using feedback loops (FBLs), you can record these abuse complaints and unsubscribe the complainer. It's important to remove complaining subscribers because failure to do so is a common reason for getting blacklisted.

Blacklist

Lists maintained by companies revealing IPs or domains which are sending unsolicited email, engaging in bad email practices, or associating with a website engaging in bad practices. There are multiple blacklists, and some are more widely used than others. ISPs often maintain their own blacklists, which consist of public blacklists and internal blacklists based on engagement or direct complaints.

Bulking

When an email is routed to the spam folder instead of the inbox. Bulking can occur because of improper authentication, content resembling spam, or something in your infrastructure's history causing concern.

Dedicated IP

The use of an IP address for one client or department's traffic. No other traffic is sent over a dedicated IP.

Direct complaint

When a subscriber or ISP complains they no longer want to receive email from you by emailing the reply-to or abuse@ addresses. It's generally best to unsubscribe them and let them know they're unsubscribed. Don't ask them to do anything, beg them to stay, etc.

Engagement

How subscribers respond to the content you're sending. Are they regularly opening your email and clicking your links? If they're engaging with the content, your reputation and deliverability will improve. If they're not engaging and are deleting, marking as spam, or unsubscribing then your reputation will drop and affect your overall deliverability.

Feedback loop (FBL)

A process beginning when a recipient clicks "Spam" in their email client. If you're registered with the ISP's FBL process, they'll send you a specially formatted email which includes who complained so you can unsubscribe them from the list.

ISP

Any corporation or entity receiving email. Don't just assume the major email providers are the only ones capable of implementing technology to filter your email, provide FBL reports, etc.

MTA

An application responsible for transferring email from point A to B.

Reputation

How an ISP views your content and infrastructure. It's a grade tied to your IP or domain determined by algorithms which differ with each ISP. Generally, reputation is calculated by a formula of engagement, abuse complaints, bounces, and send volume.

Seed list

A list of email addresses with the major ISPs injected into a campaign to see whether the email is delivered. You don't touch these emails in any way; you just want to see where the email is placed or if it's delivered at all.

Shared IP

A group of IPs used for multiple clients. All of the traffic is spread across multiple IPs because the sending volume is not appropriate for a dedicated IP.

Spam trap

Stale email addresses which ISPs have turned into honeypots for catching senders engaging without permission. Generally, you'll see spam traps in old lists, purchased lists, or improperly collected lists. The best way to get rid of spam traps is using reactivation or pruning techniques on a regular basis. (And, of course, using strictly permission-based lists.)

Subscriber

A person who has given tangible and confirmable proof to engage in sending them commercial email.

Transactional email

A one-to-one email, usually sent to individuals or small groups instead of large lists. Transactional emails include receipts, confirmations, reminders, or any email personalized specifically to the recipient. We recommend delivering transactional emails separately from marketing emails to protect the reputation of your transactional infrastructure. Existing API-based transactional systems, such as Mandrill, can be separated from your marketing infrastructure while still connected with your application.

Unsubscribe

When a subscriber requests to be removed from a list. You're required to honor all unsubscribes, and failure to process an unsubscribe is a violation of CAN-SPAM.

Virtual MTA

Multiple MTAs running on one machine, all using different IPs and domains. Instead of having five machines for five sending domains in your infrastructure, you can have one machine configured with five (or more) virtual MTAs.

Whitelist

A list or registry of approved senders. Being on a whitelist doesn't mean you can violate terms or build a poor reputation. Doing so will get you removed from the whitelist.

Deliverability team

Team? What do you mean? It's just me!

Working on your own is fine, but we use "team" because you'll have to wear several hats. (If you plan on doing this delivery thing on a large scale, you'll probably want some help.)

The first order of business is planning out the responsibilities of the team. Every delivery infrastructure will have different needs, so we can't tell you exactly how to divide the roles. You might find team members gravitate toward specialities, and there are benefits to encouraging expertise. However, each member should have a general knowledge of deliverability and be able to handle any issue.

If you don't have a team, map things out so you're not stretched too thinly. Make to-do lists for tracking logs and reputation scores. Everything we discuss in this guide will need to be done on a regular basis or monitored frequently. Having a checklist can help track what's coming up next on your plate.

Sending emails means receiving emails—lots of emails

ISPs, email administrators, subscribers, and other contacts need to communicate with you. Generally, ISPs, email administrators, and anti-spam authorities will assume you've set up the typical email accounts, such as abuse@, postmaster@ and fbl@. It's important to create these addresses up front and get them right because changing them later can be close to impossible.

Remember to think about growth. For instance, if you want one easy-to-check address then set up each address separately and forward them to one inbox. Don't go down the tempting road of setting up a catch-all for everything. When the team grows, you'll want to easily move the responsibility of one or all of the inboxes to a new hire.

It's also vital these email accounts have spam filtering. As they're common addresses, they'll receive lots of unwanted email.

If using multiple domains or aliases, it's a good idea to ensure each domain and alias has the following email addresses properly forwarded to the main domain:

postmaster@

This is a common address for domains. It can be a catch-all for various types of email, but will be required for some registrations. Someone receiving unwanted email might try to send a complaint or a question to this address.

abuse@

A must-have address. It's used for handling direct complaints from subscribers and ISPs, or other permission-related issues. Sometimes your hosting facility will use this address if they see issues with your content or receive complaints directly.

fbl@

In order to process abuse complaints, you need to set up your feedback loop process. The ISPs will deliver to your fbl@ address. This address should only receive FBL emails, but the FBL process should also be smart enough to only process FBL requests in case it receives spam.

alerts@

Be sure to have an inbox which collects various alerts such as triggered events, MTA errors, and monitoring alarms.

And just a reminder, do not use forwarders and role addresses here. Additionally, prepare for failover by setting up redundant MX entries. These are important emails—don't tempt fate that something gets lost. Each of these emails require dedicated inboxes.

Email filters

Whichever email program you use, set up filters to help organize emails automatically into folders. In addition to the addresses above, you will also set up alerts which could be triggered by the infrastructure. Some of these will be very important and require immediate attention, while others might be useful for archiving immediately. Good email filters will help sort emails by priority and keep you from getting overwhelmed.

As part of maintaining email filters, keep a list of filters or an importable file for when new members join the team. Setting up filters can take a while, and experience is often the best way to know what's priority.

Hotmail, Yahoo, Gmail, etc.

Set up some email accounts with each major ISP or email system. We recommend one account for testing purposes and another for ISP-related support, tools, and postmaster issues. Hotmail's SNDS requires a Hotmail account, and Yahoo's FBL registration requires a Yahoo account.

Remember, don't use your personal accounts to avoid ownership changes if you leave your post.

IP/DNS

So you have the email addresses ready to go and now need a domain to attach to them. If sending commercial email for your company or organization, it's best to attach the email to your domain. If sending email for another company or if you have no idea where to start with IPs and domains, then this section should help:

Already have a domain?

Let's say you have the "exampledomain.com" domain, and you want to send a bunch of marketing email to opted-in subscribers. The first thing to understand is the effect sending commercial email can have on your domain.

If you're not sending to a legit list then you're putting your domain at risk for a poor reputation with the ISPs. Engaging in commercial sending is risky, so plan carefully. We recommend separating marketing and transactional emails across different domains. You could even use a separate transactional service, like Mandrill, side-by-side with your marketing infrastructure.

Because subscribers associate you with your domain, consider using a sub-domain. For instance, mail1.example.com. Choose your domain, sub-domain, and naming conventions wisely as this can have a significant effect on how ISPs and anti-spam authorities view you. As always, think about growth here, and use a letter or number scheme for your sub-domains. Do not use a numbering scheme for your root domain. It's best to associate one domain to one client if you're using a dedicated IP scheme.

I'm sending on behalf of someone else.

Do not tie someone else's commercial email to your domain. For example, say you're a design agency, and Joe's Bait & Tackle wants you to send email for them. We advise associating all email traffic to Joe's Bait and Tackle's domain or creating a new domain. (Using your domain could adversely affect your domain's reputation.)

When creating a new domain think about growth, and use a letter or number scheme for the sub-domain as described previously.

I'm sending on behalf of several companies.

Let's say you're sending email for several companies at once. The first determination you need to make is whether to use a shared or dedicated IP scheme. Later, we'll describe the differences and reasons to use one versus the other.

If using a pool of IPs for several customers, ensure the domain is consistent across the board. In other words, don't use different domain names in the shared IP pool, but use the xxxx1.exampledomain.com schema.

IP ranges

It's important to think about how much email you're going to send because the sending volume will determine how many IPs are needed. Generally speaking, you don't want too few IPs in case you experience more volume than expected. But too many IPs and volume gets spread thin, which looks suspicious. There has to be a balance of volume to IPs.

ISPs know who you are, and they know your IP blocks. They're probably smarter than you are, and if you think your little operation isn't being watched you're sorely mistaken. Sending too much volume from one IP, sending from too many IPs, or sending too little from a range of IPs can all lead to poor reputation. So what's the right number? Word to the Wise has a series of articles providing insight into the right balance of IPs for you.

Shared IPs

If sending emails for multiple clients or customers, it's sometimes better to use a shared pool of IP addresses. You can place several senders into the pool, maintaining a consistent send frequency. With a shared IP pool, you want to get the number of IPs right so you're sending the right amount over each one.

But there is a downside to a shared IP pool. One sender, or set of senders, can affect the reputation of the others. It's important to consider the ISPs are generally looking at the IP or domain to determine reputation. Most people might then say, "That's a good enough reason to put everyone on a dedicated IP!" But sending too little or infrequently can be just as damaging. If sending emails for multiple users, a good rule of thumb is to keep the small users on a limited number of IPs to maintain a consistent volume for them.

Make sure the code which hands the email off to the MTA is properly and evenly spreading your email over the pool of IPs. In other words, don't let one IP get 80% of the content and another IP get 20%. ISPs like to see an even flow of emails. No big jumps! If you're sending too much volume and deliverability drops, add another IP or set of IPs. Ensure the pool of IPs is segmented by domain. If you need several pools, use domains to denote the different pools and sub-domains to differentiate each MTA/VMTA.

Please note, if using DKIM authentication, which we cover a bit later on, make sure the Signing Domain Identifier (d=) parameter points to the sub-domain instead of the root domain. Doing so helps maintain separate reputations for each sub-domain.

Dedicated IPs

If you plan to send lots of volume (at least 50-75K twice a week,) send frequently (at least 10-15K daily,) or send transactional messages, you'll want to use a dedicated IP. In cases of even higher volume, you might need to use a pool of dedicated IPs for the traffic. In that case, make sure the IPs carry the same domain with different sub-domains.

Dedicated IPs are great because the traffic is isolated to a specific sender. But it's important your sending be consistent and the quality high. It's your traffic and *only* your traffic, so nobody else is at fault if things go badly.

Unlike shared IPs, dedicated IPs can be whitelisted with the ISPs. See the Word to the Wise article on ISP Information for some helpful tips. Additionally, dedicated IPs can use custom domains specific to the organization sending over each IP. If the custom domain matches the from address, the email will appear more authentic to the ISPs and the recipients.

Purchase IPs

It might seems like a great deal to lease IPs from your datacenter or hosting provider, but think longterm here. If you ever switch providers, and chances are you will want to eventually, you won't be able to take the IPs with you. So all that time spent building great reputation for the IPs will be lost, and you'll need to start over.

Instead, consider buying the IPs directly from ARIN, the American Registry of Internet Numbers. It's like buying IPs wholesale, and you can take them with you to any datacenter or hosting provider you might work with in the future.

Domain registrars

It's important to use a reputable domain registrar with high standards and, most importantly, takes abuse seriously. Don't associate your IPs with a registrar known to be used by spammers. Look into their abuse policies as your own reputation could suffer by proximity.

WHOIS contact information

Add *all* of your contact information to all of your IP and domain WHOIS records. Make sure a physical address is listed along with the organization name associated with the emails. Also ensure the email addresses for abuse and general requests are present. Check each and every record if someone else sets them up for you.

The WHOIS information is also important when registering for whitelisting, FBLs and other registration processes. If the WHOIS records don't match up you'll be unable to register. For example, Microsoft SNDS checks your WHOIS records match up with the rDNS of the IPs you're registering. Avoid domains by proxy and other privacy services to mask IP/domain ownership.

Most postmasters frown upon this.

Setting and changing WHOIS contact information should be handled through the Shared WHOIS Project (SWIP) process, which is another benefit to purchasing from ARIN. When leasing non-ARIN IPs, you might not be able to change the contact information. It's vital for ISPs to know exactly who is sending emails and who to contact about abuse, meaning the contact information for each IP should point to you and not your hosting provider or registrar.

Reverse DNS

Set up forward-confirmed reverse DNS for your IPs before registering with ISPs, providers, or any type of whitelist.

MX Records

Ensure your MX Records are properly configured. Generally, your hosting facility will help with the setup. Then check the records at MXToolbox or DNS Stuff.

Test IPs

Before you send over your IPs, visit Sender Score, Sender Base, and AOL to find out the history and reputation of your IPs. You don't want to purchase from your registrar or hosting company an IP with an existing bad reputation.

Test DNS

Constantly test the DNS to make sure it is set up correctly and everything resolves correctly. There are many free DNS test tools available, but the information they provide can be limited. For a quality paid tool with a good free option, try DNS Stuff.

Publish your IPs

Once the IPs are set up, make sure to publish them somewhere on your website. This can be used as a reference for ISPs who want to whitelist (or blacklist) your IPs.

Transactional messaging

We advise not sending transactional messages (signups, unsubscribes, online receipts, etc.) over the same IPs or domains as marketing-related email. Instead, separate these messages into their own "worlds." If you include transactional messages in the same domain/IPs, you won't be able to apply for whitelisting.

Also, if someone wants to unsubscribe from your weekly newsletter but still receive their online bill, you'd want that traffic to go over different domains. If they block the newsletter traffic or unsubscribe they would still receive their bill.

There are API-based and SMTP-based transactional message systems, such as Mandrill, which can incorporate with your application to handle the one-to-one emails while your delivery infrastructure tackles the heavy-duty marketing campaigns.

Authentication

If sending email in any commercial capacity, you have to use some form of authentication. We highly recommend using all four current forms and probably any new ones released in the future. Do *not* listen to someone who says authentication is optional. Some ISPs use authentication to determine whether an email is valid and will bulk your email if you don't send with their recommended authentication method. If you don't have authentication in place you might make it to the inbox, but if the content is questionable or the reputation of the IP/domain is poor then you could see higher levels of bulking.

These are the six primary forms of authentication to set up:

SPF

Sender Policy Framework (SPF) is easy to set up and makes it harder for spammers to spoof an email from your domain. You can even cheat with the SPF setup wizard.

Domain keys

Domain Keys are still used by some smaller ISPs, but DKIM is preferred by most. If you want to ensure delivery, you can sign emails with Domain Keys, but it's not an absolute requirement.

DKIM

It's important to fully understand DKIM and configure it properly. DKIM is an authentication method to prove an email originated at a specific domain and has not been changed during delivery. Most ISPs do not require DKIM, but it does help to boost your reputation. And if you use DKIM, you better get it right because a failing DKIM is a lot worse than no DKIM at all.

DKIM consists of salient values, such as the s=, d=, and i= parameters. Note that the Signing Domain Identifier (d=) parameter can be configured to point to a different domain. For instance, if your MTA's domain is mydomain.com, you can configure your DKIM signature to use another domain, such as your client's domain. This is becoming a highly recommended industry standard so the from address domain and DKIM signature match when sending traffic over a dedicated IP.

SenderID

Authentication developed by Microsoft and used by several big ISPs.

DMARC

Domain-based Message Authentication, Reporting and Conformance, or DMARC, is a specification using SPF, DKIM, and alignment to authenticate the email. The key to alignment is ensuring the domain of the from email address matches the domain, or sub-domain, or the SPF and DKIM record. In order to pass DMARC, either the SPF or DKIM must authenticate and align with the from address. Thus, only dedicated IPs with custom domains should use DMARC since shared IPs might use a domain which does not align with the from address.

Testing authentication

After setting up all the authentication, you need to test it and test frequently in case anything changes.

Start by testing with the Port25 verifier, which sends an email to check-auth@verifier.port25.com. If that works successfully, move on to real testing with the major ISPs.

Send an email to the ISPs using the various authentication types, and make sure the email passes properly. Even if it gets to the inbox, you should physically open the email, look at the headers, and make sure it's passing authentication. After going through all the details of setting up your delivery infrastructure, the last thing you want is an authentication error causing your email to get blocked.

MTA

We highly recommend using a commercial MTA product so you can manage multiple IPs, error correction/handling, and connection/sending to ISPs. We use PowerMTA from Port25, an incredible product which serves out lots of email. Why use a commercial product? Because things will break (or you will break something), and you'll want someone smart enough to fix the disaster when things don't come back up as they should. Commercial vendors also understand what you want to do with their product, and they'll help you with configurations and maintenance.

There are lots of commercial and open-source MTAs available. Some popular MTAs include PowerMTA, Message Systems, Cold Spark, Postfix, Gmail, Strongmail, and many new ones coming out all of the time. We recommend testing a bunch of MTAs until you find the perfect fit for your sending volume with all of the needed features at the right price. Commercial products generally provide benefits over open-source products, such as monitoring and configuration, general ease of use, and, most importantly, support.

So you've got all these domains, IPs, and servers set up. Now you have to match this IP to that MTA, send really slowly to Yahoo, faster at certain times to this other ISP, and then there's that Russian domain blocking your emails entirely....

This is the tough part—getting the configurations dialed in.

There are several configuration levels, and it's important all these settings are properly tuned. There are configurations at the server level, domain level, virtual MTA level, and even domain-specific per ISP. We recommend starting with the default configurations included with your MTA to see what works best, and then tweak as you go along. You'll find what works for another sender may not apply to your infrastructure. This is one of the reasons we recommend a commercial MTA—so you can get proper support during the setup process.

And when all is said and done, you need to test, test, and test some more.

Hosting and hardware

Choosing the proper hosting facility is a big part of designing a delivery infrastructure.

Cloud environments like Amazon or RackSpace Cloud are not suited for sending email. MTAs need bare metal and fast drives. An MTA's activity is bound mostly to disk and CPU, for which cloud and virtual environments aren't well-suited. The cloud is also home to spammers, so it's best to keep your sending IPs off these networks to avoid backlash from ISPs. We've seen the effects of tying sending infrastructure to the cloud, and the results are poor performance and reputation.

Some delivery applications build the email and hands it off to the MTA. While other applications package the data and let the MTA handle the email construction. Both methods have pros and cons. So if your application is in the cloud and the MTA lives elsewhere, it's important to pick the option with the lowest bandwidth usage. The more data sent between application and MTA, the longer the transfer time and higher the bill.

When choosing a hosting provider or datacenter, make sure they have a good reputation within the anti-abuse community. Their abuse teams need to be responsive and well-trusted. Some ISPs will send abuse complaints directly to your provider, who should then forward those complaints to your abuse@ address for review. Additionally, your provider needs to investigate and remove any users in their system sending spam, or your reputation could be affected by proximity.

Security

Your delivery infrastructure is precious and vital, and security is *extremely* important. It's the type of system which can be used maliciously in the wrong hands, and there's plenty of valuable email data to be stolen. Don't do anything silly like allowing open relay. Secure your infrastructure behind a firewall, and, better yet, require VPN access to get to your systems. That includes the monitoring and dashboard utilities. Secure anything and everything related to your delivery infrastructure.

Rate limiting

Another critical step in the MTA configuration is rate limiting (AKA: throttling.) Rate limiting allows ISPs proper time to process and filter spam, ensuring transactional emails don't get backed up. Without rate limiting in place, ISPs would be even more overwhelmed than they already are. The ISPs all have different sending limits on a per hour, per day basis. ISPs can throttle your sending volume when it's too high or too low. Consult with your MTA vendor on proper rate limiting configuration. And don't think you can get away with simply sending as fast as possible—you'll fast-track yourself to getting bulked or blocked.

Error correction/handling

Once you hit rate-limit thresholds, send too much spam, or have any number of other issues, the ISPs might start returning error messages. A good commercial MTA will allow you to handle these errors and adjust accordingly.

Some ISPs will want you to slow down the sending, stop sending for a period of time, or change your habits due to bad engagement, bad reputation, etc. Most commercial MTAs have recommended settings and allow further customization for when errors occur. Take this aspect of your setup seriously because failure to do so will get you in serious trouble. Understand each ISP is different, ISP preferences change throughout the year, and ISPs develop new error codes which will require tweaking settings and configurations. You'll need resources to keep an eye on error handling and attend to tweaking the configurations to ensure successful deliverability.

Bounce Handling

Bounce handling is a critical factor in your delivery infrastructure. When sending to a non-existent email address or an email account which is full, the receiving server will return a message with the bounce reason. The bounce message will contain a Diagnostic-Code header, which will give the bounce code, reason, and any other information pertinent to the bounce.

Clean bounces from subscriber lists, but you'll want to handle soft bounces differently from hard bounces. And in some cases, such as spam bounces, handle the bounce as though it didn't bounce at all. Unless you want to manually clean bounces, install a bounce processor in your application which takes in bounces and handles them accordingly. This can get complicated because bounce reasons aren't fixed and can be customized by the ISP. So a 521 error code with Joe's mail server might be a hard bounce, but a 521 with Jane's mail server might mean something totally different.

VERP

The sending application which pass emails to the MTA needs to include a variable envelope return path (VERP) in the header. A VERP'd address looks something like this:

```
    some_unique_identifier(s)-
therecipient=theirdomain.com@mail1.yourdomain.net
```

Notice the recipient's email address in the string. When the bounce occurs, you can get the recipient's email address out of the bounce record and handle the bounce accordingly. Your unique identifier(s) might be a Customer ID, Account ID, etc. Your bounce processor has to parse the VERP'd address and record the bounce accordingly, so include whatever's necessary. However, VERP addresses should remain within 72 characters because some ISPs will disregard the VERP'd address or, in some cases, reject the entire email.

Hard bounce

Hard bounces are email messages returned to the sender because the recipient's address is no longer valid. A hard bounce could occur because the domain doesn't exist or because the recipient is unknown. Here's an example of a hard bounce diagnostic code:

```
    Diagnostic-Code: smtp;550 5.1.1 - Invalid mailbox: 
xxxxxx
```

Process hard bounces as soon as possible, but you might want to build a policy which allows two consecutive hard bounces before the recipient's address is cleaned. You don't want to keep sending email to a bounced email address because ISPs include this as part of their reputation formulas. If sending to a bunch of bad addresses it could mean you have an old list, which ISPs factor when determining whether or not your email is spam.

Soft bounce

Soft bounces are email messages returned to the sender because the mailbox is full or for other reasons, usually temporary. Here's an example soft bounce diagnostic code:

```
    Diagnostic-Code: smtp; 552 xxxxxxxx MAILBOX FULL
```

Process soft bounces after a certain number of occurrences or sequential occurrences. So if an address soft bounces X times in a row, remove them from your list. We recommend making that "X" number configurable, and you can determine whether to use consecutive bounces or bounces over the life of the list. Just remember, a soft bounce can be a temporary issue, and it's possible the address is valid.

Bounce categories

Bounce Categories can be anything the ISP wants them to be, which can be frustrating. Administrators might use specific bounce categories for just your

IPs, if you're blacklisted, or for reporting bulked email. And some administrators even mix soft and hard bounces. In other words, you should be able to distinguish bounces, but it's a constant challenge to process them correctly.

Categorization

Since bounce categories are all over the place, what do you do?

Categorize bounces based on codes or diagnostic information. Set up a way to store regular expressions, and if a match is found re-categorize the bounce. If you know Joe's email administrator throws a 521 with the message "your IP XX.XX.XXX.XXX is listed on uribl" when you're on a blacklist, you can categorize it as a soft or spam-related bounce. When the block is removed, you can send to Joe again. Build something into your bounce processing code which will take a bounce, look at the diagnostics provided, and categorize it appropriately if there's a matching regular expression.

Using professional, paid tools, such as BoogieTools, can save time and effort. Bounce handling can be set up quickly and reliably with tested settings and systems. And if you do run into any configuration problems, you'll be able to get support from knowledgeable experts.

Feedback Loops

Feedback loops allow you to receive reports from most of the major ISPs when someone reports your email as spam. Essentially, you register with the ISP to become part of the FBL process, and when someone clicks "Send to Spam," the ISP will send you a specially formatted request in ARF format.

Process the report as soon as possible by unsubscribing or removing the recipient from the list. Unless you want to deal with each report manually, make sure your application is set to handle FBLs automatically.

Email headers play a large role in tracking down who sent and received the so-called spam. But not every ISP supplies a standard ARF report, including VERP addresses and custom headers denoting account, recipient, etc. They don't have to include the recipient in the returned email, and they can and will munge your headers. Make sure your FBL process has the capability to look at multiple headers and, if it fails to process them, alerts someone to handle the FBL report manually.

Keep in mind the recipient's email address can be fully redacted from the headers, so embed an email ID or encrypted email address you can easily look up. An example could be something as simple as a header X-FBL:

```
campaignID.listID.emailID
```

If you fail to remove a person who reported spam, it can lead to serious problems. If you continue sending to them, they can very easily report you to blacklists, ISPs, your registrar, and your hosting facility. All of those will affect your deliverability and reputation.

You should also keep a history of the abuse complaints. It's fine to keep the FBLs in an email account, but it's extremely helpful to store the data so it can be retrieved for statistics later. This information should be available at your fingertips.

FBL registration resource

Word to the Wise has a great list of ISPs which allow you to participate in FBL programs. It's critical to register all of your IPs. If you have a lot of IPs this can take a while, so leave yourself enough time to work through each registration process. When registering, don't use a personal email address—use the role-based emails set up earlier.

Some ISPs will require an account with their domain. Yahoo also requires DKIM in order to get their FBL data. But most ISPs use Return Path for FBL processing, which uses Microsoft's JMR program. It does take some time and requires your IP registrar confirm you're the owner of the IP ranges being registered. Leave yourself plenty of time to get this set up before sending.

Corporations and larger technology companies might send ARF reports as well. If sending a lot of emails to a specific domain, contact their abuse department or email administrator to see if they offer FBL reporting.

FBL maintenance

Register any new IPs with FBLs prior to sending email from them. It's also a good idea to check once or twice a year everything is still in place with the ISPs.

Email Headers

If sending commercial email, you should include some industry-standard header information to help with tracking abuse.

Below are the most commonly found headers in commercial email. Your MTA can use merge features to munge this data into your headers.

Reporting abuse

Remember that abuse@ address you set up? You can do two things to allow people to report abuse:

The nicest way is including a link in the email header, which will take recipients to a page with some details about who you are and a form to fill out about the incident. That report would be sent to your abuse@ address for processing.

The other method is including a message in the email, such as, "To report abuse, send an email to abuse@exampledomain.com." The recipient can send an email directly from their email program, including whatever information they choose to provide.

Either way, include an X-Report-Abuse header in the emails for the ISPs.

```
    X-Report-Abuse: Please report abuse here:
http://www.exampledomain.com/abuse?
u=account_id&id=campaign_id&e=subscriber_id
```

List-unsubscribe

Some ISPs are now integrating with the List-Unsubscribe header instead of reporting spam. The header should include a link to an unsubscribe form.

The presence of this header can sometimes turn on images or add a special icon in the inbox/email to denote a safe sender.

```
    List-Unsubscribe: <mailto:unsubscribe-account_id-
campaign_id-subscriber_id@exampledomain.com?
subject=unsubscribe>,
<http://www.exampledomain.com/unsubscribe?
u=account_id&id=campaign_id&e=subscriber_id>
```

Unique identifiers

Include some unique identifiers in your headers. These would help identify the sender and recipient if an ISP were to send a message with redacted header information.

Depending on how your system is engineered, you might need to include an Account ID, Campaign ID, and Subscriber ID. Combine this into one header, such as an X-Data header in the format of:

```
    X-Data:
company_name.account_id.campaign_id.subscriber_id
```

It's important to obfuscate each of those values for security reasons. Don't use unencrypted data of any sort.

Precedence: Bulk

Gmail requests any marketing or mass email include the Precedence: Bulk header to improve delivery. There are some theories including this header could lower the email's priority over other messages waiting to be delivered. Additionally, some ISPs will stop the sending of out-of-office replies to emails with the Precedence: Bulk header. But if you're having trouble delivering to Gmail, adding the header could score some brownie points with them.

Getting Started With Sending

Warming up IPs

So you got the IPs, domains, servers, emails, and all of the other fun stuff set up. Time to open the flood gates and just send all the emails, right? Unfortunately not because you don't have a reputation with the ISPs yet, and they don't like sudden introductions.

It's important to warm up IPs to slowly build a good reputation. First, check your IP's reputation with Sender Score, Sender Base, and AOL. If the IP starts with a low reputation, perhaps from a previous owner, try to exchange for a new IP or warm it up very slowly. At this point, most emails will bulk until the IP reputation improves. Even if the reputation is good or neutral, you still can't send a lot of email from a fresh IP. Send 100 the first day, 200 the next, and so on. Slowly increase the volume over 24-hour periods. Some MTAs have a warm-up capability built in and will gradually increase volume and handle all this for you.

The ISPs are getting to know you and learn your content and traffic patterns, making the warm-up phase critical. Give it a few days and allow the ISPs time to learn who you are. If problems arise, give it some more time before contacting the ISPs. Warming up is a process, and there could be some dark secrets in the IP's past. If your IP reputation isn't good, warm up much more slowly and work with each ISP to repair the reputation.

Any time you contact an ISP, first review their requirements, fully investigate the issue, and fix any problems which need to be addressed.

Whitelisting IPs

When dealing with dedicated IPs, go through all major whitelist registrations. You can register shared IPs, but some ISPs won't allow them (and don't lie—they know what you're up to.) Some whitelists have terms of service which allow them to boot you for certain behavior. So once you're listed, it doesn't mean you'll remain listed. If you get in bad standing, you might have to re-register or perform certain steps to get back in good standing.

Whitelisting is tedious, so leave plenty of time. Generally, you'll need some sending under your belt before you can register. Usually 15 days is enough, but sometimes you can get away with less. Take the time to register with all the whitelists and just like FBL registrations make sure you maintain these. If denied whitelisting, try again in thirty days *only* if your sending quality improves.

Abuse.net

It's a big lookup database for abuse contact information tied to IP/domain. Register with abuse.net so people can get in touch about unwanted email.

DNSLW

Register with DNSLW, a whitelist commonly used by SpamAssassin.

Certification

A few companies provide IP certification, but this isn't something you pay for and then send whatever you want. Providers have a vetting process, and they'll want to analyze your sending history, content, etc. before bringing you on as a customer. It's generally worth it if you can spend the money, but each certification has its own caveats.

Certification.EQ

Return Path provides sender certification, which is essentially high-end whitelisting covering a large ISP footprint. Certification will improve deliverability almost immediately once it kicks in. It can also turn on images automatically and provide other benefits with ISPs. The downside to certification is you're required to maintain a high level of deliverability. Go outside the acceptable boundaries and you can temporarily or permanently lose certification. For instance, if you attempt to send shared traffic over a certified IP.

Monitoring and Exception Reporting

Boom! Your infrastructure is set up and emails should flow smoothly to the ISPs. But how do you know? It's important to have good visibility into your deliverability. Here are some products and tools which can help:

Engagement monitoring

One critical monitor is the use of seed lists, which include email addresses set up at the major ISPs and used solely for the purpose of seeing whether emails land in the inbox, bulk, or disappear into the ether. If sending from one dedicated IP, simply place a seed list into your subscriber list. If using several IPs, randomly place the seed lists onto email and watch inbox placement in some automated fashion. Don't move, open, or touch emails when they arrive, if they arrive, because ISPs use predictive algorithms to deliver future emails based on past engagement. Remember not to use the seed lists too often, as they can have a negative effect on list performance because those emails are not opened or responded to.

Also set up monitoring to determine how many emails bounce, how many recipients open and click the email, how many subscribers unsubscribe, and how much FBL activity occurs per IP. ISPs and email administrators keep a close eye on this information, and if you see high abuse complaints, unsubscribes, or bounces then the ISPs know something is wrong with your list-collection techniques. If experiencing low opens and clicks, it could be due to poor inbox placement, unengaging content, or lack of proper segmentation.

Scraping MTA logs

Set up some form of monitoring on your MTA logs. Most MTAs will have ways of handling exceptions, but they won't be able to handle all exceptions. For instance, some ISPs will report back if your IP is on a blacklist. That's something you want to be aware of in real time. You should set up some scripts to actively monitor your IPs and domains to check the major and minor blacklists. But you also want to actively mine this data in your MTA logs. Some ISPs and blacklists don't have a way to look up your IP's status, so by scanning logs you can catch the exceptions which might occur. There are other scenarios where you want to scrape MTA logs for spam trap addresses, fatal errors, etc. The MTAs are full of useful information, and being able to search the logs will prove handy when troubleshooting problems.

Static/Dynamic error handling

We touched briefly on using your MTA to handle certain errors and scraping your MTA logs to find out when errors occur—but there are some errors which require different or additional actions.

The first type is a static error. An ISP, such as Comcast, might throw a static error which looks something like this:

Comcast block for spam. Please see http://help.comcast.net/content/faq/BL000000

If you weren't scraping logs or using your MTA's error handling, you'd never know this error took place. So we recommend using both the MTA error handling *and* log scraping to look for errors and send alerts. In this instance, the MTA should back off sending to allow enough time to unblock the IP with Comcast. You can even switch all Comcast traffic to another MTA, which the MTAs should be able to handle.

The next step would be to find out what's causing the issue. Having searchable logs is key here. Find out the who, what, where, and when of the incident. Fix the issue. If it's a specific sender, stop their sending. If it's a specific recipient, block them from being sent any emails. Then manually fill out the Comcast unblock form. After you receive an unblock notification, turn the traffic back on if the problem will not persist. Failure to fix the issue will cause emails to bounce continually until the underlying problem is corrected.

Other errors are dynamic errors. Some ISPs will throw a dynamic error which requires slowing down or completely stopping sends for a few hours. Configure your MTA and log scraping to send an alert so your infrastructure can respond properly. If dynamic errors continue for an IP, investigate the cause and remediate. That might involve speaking with the ISP to find out the issue, but investigate thoroughly first. There are tons of codes and resolutions, and ISPs add new ones each year. Work with your MTA vendor and development team to build an application and infrastructure that's fully aware of these errors and can handle them cleanly.

Blacklist monitoring

There are several monitors and tools which will show if you're blacklisted. Unfortunately, no, they don't include all of the blacklists, and most monitors will not be "real-time" enough for your needs.

Here's a list of some major and minor blacklists:
AHBL, ANT, Backscatter.org, BARRACUDA, BURNT-TECH, CASA-CBL, CASA-CBL+, CASA-CDL, CBL, CYBERLOGIC, DEADBEEF, DNSBLINFO, DULRU, EMAILBASURA, FABELSOURCES, FIVETEN, GIRL, GRIP, HIL, HIL, HILLI, ICMFORBIDDEN, IMP-SPAM, IMP-WORM, INTERSIL, ivmSIP, ivmSIP/24, KEMPTBL, KUNDENSERVER, LASHBACK, LNSGBLOCK, LNSGBULK, LNSGDUL, LNSGMULTI, LNSGOR, LNSGSRC, MSRBL-Combined, MSRBL-Images, MSRBL-Phising, MSRBL-Spam, MSRBL-Viruses, NERD, NETHERRELAYS, NETHERUNSURE, NIXSPAM, NJABL, NJABLDUL, NJABLFORMMAIL, NJABLMULTI, NJABLPROXIES, NJABLSOURCES, NLKUNBLACKLIST, NLKUNWHITELIST, NOFALSEPOSITIVE, NOMOREFUNN, ORID, OSPAM, PDL, PSBL, RANGERSBL, RATS-Dyna, RATS-NoPtr, RATS-Spam, REDHAWK, RRBL, SCHULTE, SDERB, SENDERBASE, SERVICESNET, SOLID, SORBS-BLOCK, SORBS-DUHL, SORBS-HTTP, SORBS-MISC, SORBS-SMTP, SORBS-SOCKS, SORBS-SPAM, SORBS-WEB, SORBS-ZOMBIE, SPAMCANNIBAL, SPAMCOP, Spamhaus-ZEN, SPAMSOURCES, SPEWS1, SPEWS2, SWINOG, TECHNOVISION, TRIUMF,

UCEPROTECTL1, UCEPROTECTL2, UCEPROTECTL3, VIRBL, WPBL, WSFF, ZONEEDIT, CSMA, DUINV, ORVEDB, RSBL, SPAMRBL

And that's not even all of them!

People can run blacklists out of their mom's basement, and corporations can have their own blacklist. Also keep in mind your IPs and domains can be blacklisted, redlisted, or graylisted. Red listed senders are related to senders currently in the blacklist, and gray listed senders are generally associated with email marketing. It's usually not possible to delist IPs and domains from red and gray lists, but most ISPs won't block senders on red or gray lists unless there are other extenuating circumstances, like content blocks.

It's pretty easy to set up monitoring with most blacklists, and there are many scripts publicly available on the web to help. Additionally, Spamhaus DBL is a realtime queriable database of domains found in spam messages. Some blacklists require registering IPs in order to use their lookups, and failure to do so could get the IP you're checking from listed as well.

Spam filter monitoring

Spam filters are used by most major and minor ISPs to locate spam based on email content. If they're not using Spam Assassin it's probably a commercial-grade spam filter. These filters will catch most, if not all, spam, but sometimes they can be aggressive. Check your content, and test as much as possible.

Return Path and Litmus offer scanning through a few major spam filters, but you can also do some scanning internally prior to sending campaigns. It's easy to install Spam Assassin on a test box and run content through it. Run the content through other spam filters and test accounts with the major ISPs as well, just to cover your bases. The difficult part is weeding out false positives and ensuring good content isn't flagged as spam. Try tweaking the Spam Assassin rules to suit your needs. Just keep in mind some people run default installs of Spam Assassin while others run custom rules.

If you can't get your hands on the filters used by major ISPs, it's at least good to know the products used. Return Path provides some information about spam filters used by major ISPs: Yahoo uses a proprietary spam filter called SpamGuard; Hotmail, MSN, and Live use BrightMail. Other commonly used filters at the enterprise level are MessageLabs, Barracuda, and Forefront. These filters would affect B2B email communication, so it's important to keep them in mind.

Content fingerprinting is becoming widely used by ISPs and enterprise/corporate email administrators, as well. Cloudmark offers a high-quality product for detecting spam through fingerprinting.

Familiarize yourself with the spam filters used, and, if possible, get these filters into your infrastructure to ensure you're not sending out content which will get blocked.

Monitoring tools

Return Path

Return Path offers several products and tools which give insight into how good your delivery is (or isn't.) They offer tools for reputation monitoring, seed list monitoring, blacklist monitoring, and previewing emails in over 30 email clients with spam filtering analysis.

Seed list and campaign monitoring allow you to attach a list of monitored addresses to emails. Return Path then collects the delivery stats from the major ISPs, determining whether the message ended up in the inbox, bulked or went missing.

Reputation monitor provides helpful data for determining IP issues, but it should just be one data point as not all major ISPs contribute.

Mail Monitor

Mail Monitor uses seed lists to provide deliverability monitoring and spam filter testing. You can see which of the major ISPs delivered, bulked, or blocked the emails. Additional information about how well emails fare against spam filters, such as Cloudmark, SpamAssassin, Barracuda, and Symantec.

The MTA can be set to automatically append seed lists to each email campaign, campaigns for specific users, or random campaigns. One advantage to seeding random campaigns is a good overview of how the entire infrastructure is delivering and not just specific users or IP addresses.

IBM Enterprise Marketing Management

IBM Enterprise Marketing Management recently purchased Unica and Pivotal Veracity, delivery monitoring tools with a good reputation. We can't speak directly to them, but we know a few ESPs use them heavily.

Microsoft SNDS

The SDNS tool should be used as a secondary tool to troubleshoot issues and gather information. Register your IPs and SNDS provides information such as number of spam traps hit, abuse complaint ratio, and volume per IP. It's a great tool to check when a client is having delivery issues to Hotmail, MSN, Live.com, or Outlook.com.

Deliverability troubleshooting

It's not me, it's them.

Let's say a big ISP is blocking your email. The first thing you do is email them about it, right? Nope!

Don't email your buddies at AOL or talk to someone who knows someone. That's not what deliverability is about. First, start with your infrastructure, your logs, your lists, your content, etc. Not them—*you*. Generally, most issues are on the sending side, and the issue stems from poor engagement, stale lists, or spammy content. If, and only if, you have checked and fixed everything on your end should you engage with an ISP.

I'm blacklisted.

Generally, getting blacklisted boils down to bad list etiquette. You can't send commercial email because the recipients verbally told you it's okay. You can't send commercial email because you bought a list, rented a list, or a vendor said it was cool. There are all sorts of blacklists and different ways to land on them, but all the ways we know of start with bad list etiquette.

Chances are, you know you're doing something wrong, or your client knows they're doing it. Just do the right thing, unless you enjoy filling out forms and emailing people about how great your list is.

When you get blacklisted, find out as much about the incident as you can from the listing company. Usually, they'll at least provide a date and subject line. If you're scraping your MTA logs, you likely caught the incident when it occurred or just after it occurred and can work backwards—similar to working an ISP block. Figure out what caused the issue, fix it, and then, and only then, do you contact the listing company.

Corporate domains/business-to-business

Corporate domains, small ISPs, and international ISPs can employ similar technology to the major ISPs. Corporate domains are sometimes more stringent on rate limiting and spam policies because their servers can't handle large volumes or they want to reduce workers' time deleting junk emails. If you're getting blocked at Bigco's domain or a small ISP, treat it just like a major ISP. Do some research, and read up on any public information prior to contacting them.

If marketing to government or large non-profit organizations, they'll likely be using a Barracuda device. We recommend getting one as well to check your content. In some cases, these organizations might reach out for different issues or policy violations. Generally, they'll ask you to comply with their standards, and it's important you comply as quickly as possible.

Some corporations and small businesses have such strict policies they might decide to blacklist you if anything looks suspicious. Good thing you set up all those email addresses so companies can alert you to potential issues and possible ways of resolving those issues.

Application Development Considerations

Delivery infrastructure isn't just about servers configured with DNS and running an MTA. Software has to be written to handle the construction of the email, creating headers, handling bounces, etc. We've touched on several elements needed in your software, but we'll summarize and add a few more for your consideration:

Agent

Most ESPs have an agent which constructs the pieces of data and content to pass off to the MTA. Some opt to have the agent handle merging the data and passing to the MTA for sending. Others supply the data and the content to the MTA, which handles the merge process.

Bounce processing

Ensure you have some code, script, or an intern with lots of free time to process bounces. VERP'd headers will allow for easier processing, but remember to include enough data in the VERP'd address to find the recipient and remove them from the list.

Click tracking

Allowing your marketing team or customers to know who's clicking links and which link they're clicking is vital. Tracking clicks helps understand how recipients engage with the email. Those who do not engage are more likely to unsubscribe, bounce, or complain in the future, which can lead to poor deliverability. Periodically removing non-engaged recipients will improve deliverability and reduce strain on the infrastructure.

Email headers

The application or MTA should allow you to pass in custom headers. Include a VERP address, unsubscribe link, abuse contact information, and other unique identifiers. Most commercial monitoring tools will require a unique identifier in the headers. Consult with the monitoring tool company and your development team to determine the best header option. You can provide multiple headers to uniquely identify your email, but returned bounce and FBL emails might not contain the same custom headers as there are no set standards.

FBL

You'll also need code, script, or an intern to process FBL requests, but keep in mind the returned data can be munged, and some headers might not be intact. Set up notifications for when an abuse complaint fails to process. Don't go live without FBL processing, and register your domains and IPs with all ISPs.

Merge data

Most MTAs have merge capability, which uses a specialized syntax to merge the data and content. For instance, to start the email with "Dear John," where the name would be specific to each recipient, you could use a merge tag for the name data and allow the MTA to handle the processing. This is a widely used feature in commercial email to be considered when designing your application and sending infrastructure.

Also, your marketing department or customers will expect to use conditional content to send specific text, links, images, etc. to each recipient depending on their data, such as zip code or age. All of this can be done with merge tags, application code, and commercial MTAs.

Open tracking

Similar to click tracking, offer open tracking to your marketing team or customers. The industry standard is generally a 1x1 pixel image embedded in the body of the email. When the user has images turned on, the image is downloaded from your server and the open is recorded. This is important, just like click tracking, because you want stats on list-engagement metrics.

Transactional emails

In addition to sending marketing emails in bulk, the application will also need to handle transactional emails, such as signups, unsubscribes, confirmations, etc. Since it's best to keep marketing and transactional emails separate for reputation purposes, and also because you would not be able to whitelist IPs sending both bulk and transactional emails, we recommend creating separate systems or using third-party transactional services, such as Mandrill.

Wrap Up

Sounds like fun, right? *Email marketing blaster thingy* sure didn't seem like such a monumental effort at first. We've been at this for years and currently send billions of emails per month. Guess what? We're still learning and adapting to the constantly changing field. New technologies, practices, ISPs, and forms of abuse appear seemingly over night. This guide makes for an excellent beginning to a life-long devotion to email marketing.

So what's the future got in store for your delivery infrastructure? Being able to send both marketing and transactional emails? How about analytics for subscriber engagement? Setting up email marketing doesn't just end with requisitioning a server and installing some software. Even following every step of this guide, you still need to know what to do tomorrow and the day after and the day after that.

We raise a glass to those who try it. To everyone else, remember, we offer free MailChimp and Mandrill accounts if you'd rather leave the fun stuff to us. Because we love what we do.