Sinch E-Mail 365 – Deliverability Best Practices Guide

Version 15 – October 2020

sinch.com

Revision History

Version	Date	Description
1.0	2016-11 -31	V1.0
2.0	2017-02-28	SAP Marketing Cloud specific highlight
3.0	2017-03-10	Updates to deliverability watch list.
4.0	2017-05-06	Branding updates
5.0	2019-03-09	IP Warmup schedule updates
6.0	2019-04-23	MX record updates
7.0	2019-05-16	IP Types and the differentiations within sending high-email volumes
8.0	2019-07-07	Updated ‘Plan for a validated ramp-up”section
9.0	2019-10-30	Feedback Loops
10.0	2019-12-04	Suppressions
11.0	2020-01-22	Added list of roles based email addresses that are suppressed
12.0	2020-02-24	Yahoo Feedback Loops. Bounce retry frequency
13.0	2020-03-03	Update to bounce retry frequency

14.0	2020-03-23	IP allocation guidelines based on email volume, transactional email guidelines, updated information on Gmail feedback loops
15.0	2020-10-22	Warmup process for Microsoft. Additional information on MX records

Table of Tables

Table of Contents

The purpose of this document is to provide Sinch E-Mail 365 and platform customers an understanding of email deliverability and the associated best practices with it.

1 Deliverability and why is it important

Email deliverability measures the ability to deliver email messages in recipients' inboxes for email marketing campaigns and to avoid messages being lost, blocked or driven to the spam folder.

By extension, email deliverability also encompasses the secondary folder issue. Marketers should be in the inbox but also to be in the good folders and to avoid junk mail or if possible the Google's promotion tab.

Email deliverability is an issue due to address quality and spam filtering. Issues associated with address quality are easily managed through good address verification and list hygiene practices. Deliverability issues due to spam filtering and blocking are more complex to manage.

There are two main indicators or metrics for measuring deliverability: "classic" deliverability rate and inbox placement rate. Each is a based on a different measurement methodology and inbox placement is often seen as more accurate since it measures emails delivered in an inbox as a percentage of all emails sent.

Deliverability is a crucial issue for brands as the single most important factor in improving response and conversion rates. However, improving deliverability can be a complex issue to handle. This is because there are many factors in play—content, authentication and infrastructure, list quality, subscriber engagement, spam filters, and more.

2 Deliverability watchlist

Based on sending experience; we have compiled a list of things that are important in achieving high deliverability. We have also attempted to detail out why these are important.

2.1 Authentication: Both DKIM and SPF based authentication are important

In our view, the # 1 factor for getting caught in SPAM filters shall be a lack of either DKIM (Domain Keys Identified Mail) or SPF (Sender Policy Framework) based authentication. The authentication assures a receiving mail server that the email service provider is sending in an authorized capacity.

ISP spam filters have different levels of emphasis on DKIM or SPF; for Sinch E-Mail 365; both these authentication methods must be used.

For more details on DKIM and SPF based authentication; please read Annexure 1 of this document.

2.2 Sender reputation: Both Domain and IP reputations impact deliverability

In general; a lot of attention is focused on IP reputation or sending score and only a few email providers (for example, AWS, sparkpost) discuss sending domain reputation. This tends to lend (in our view) to the mistaken belief that if a shared IP range has good sender score (or IP reputation); a brand can quickly establish high volume email marketing campaigns expecting high deliverability.

This notion is not completely accurate in that:

IP reputation has large role; but not the only role in deliverability and

It ignores the past sending behavior of the sending domain (“from address“).

In our view; it is important to establish high reputation for both the sending domain and IP addresses together. In particular, if the sending domain is expected to be a high-volume sender; it is preferred to assign a dedicated IP address.

We recommend sending different types of email-content (promotional/ transactional) via separate domains. Sending promotional emails over the same domain as transactional email can affect the reputation of the domain.
Customers should consider that also a domain with very low traffic could be blacklisted Therefore, merge domains with very little traffic to others with more.

2.3 MX records: MX records are important and mandatory for deliverability

A MX record is mandatory to avoid failing anti-virus checkers against the registered domain/sub domain for the email service and also to avoid a bad sender reputation. The MX record is also used for “reply” tracking. The “reply to” address allows a recipient to respond back to an email. When a marketing / transactional email is replied to, the mail transfer agent looks up DNS entries to determine the MX record of the recipient’s mail server.
MX record is required for the sender domain (it is assumed that both the sender and reply-to address are from the sender domain address)

2.4 IP Types and the differentiations within sending high-email volumes

The right IP-strategy has an impact on: reputation, high delivery rate, inbox placement
For the usage of e-mail delivery, there is 3 kinds of IPs available

1.Shared IP
Several senders share one IP address for sending e-mails

2.Dedicated IP
An IP address is only used by one sender

3.CSA compliant IP (shared or dedicated) available for select EU Senders (subject to availability)

Specially certified IP addresses, have certain requirements and the criteria are based on strict European requirements that basically cover internationally applicable law in e-mail marketing.

Advantages and better inbox placement through:

a. Whitelisting with some ISP´s
b. Less throttling with some ISP´s
c. Prioritized management of mails at ISP´s
d. Less SPAM-marking
e. Protects the IP-Reputation

2.5 Which IP to use?

Shared IP	Dedicated IP
You are sending emails: • Not consistently • Low volume (not over 100,000/per month) • With peaks (e.g. campaigns delivered not over some days)	**You are sending emails:** • Consistently • High volume (more than 100,000 / per month) • At constant high volume (i.e. your email traffic is not intermittent with high peaks)
Advantage: • Reputation is generated by sharing with other senders to generate consistent volume on the shared IP address • Senders benefit from existing reputation • With existing infrastructure and an already good reputation a faster warm up is possible	**Advantage:** • Independent of the sending behavior of other senders • Very scalable and flexible • Sender has fullcontrol • No throttling due to other senders • Splitting or separating of business processes with dedicated IP´s (marketing / transactional)

Challenge: • Bad senders can influence the reputation of the IP,addresses however mail providers today are very good at differentiating also based on the sending domain and then block just abusive sender.	**Challenge:** • By using a dedicated IP you are responsibility for your reputation. The warm up period should be well planned and reputation should be continuously reviewed. Measures like address maintenance and campaign planning is necessary.

Table 1: Shared IP vs. Dedicated IP

2.6 CSA IPs Pros and Cons

CSA IPs - Pros	CSA IPs - Cons
• CSA compliant IP is the IP address or block of IP addresses that are aligned with the legal, technical standards and policies of Certified Sender Alliance. • CSA's certificate of quality shows professionalism in e-mail marketing. • The criteria are based on strict European requirements and basically cover internationally applicable law in e-mail marketing. • Protects the IP reputation	• The Sending criteria is stricter hence customers are more likely to have more sender reputation issues if they don't follow the CSA criteria. • CSA compliant IPs (shared or dedicated) are available only for select EU Senders (subject to availability)

Table 2: CSA IPs Pros and Cons

2.7 IP Allocation

Recommendation of dedicated IP's related to email sending volume

Guideline/No. of IPs	Daily Volume
1/2	400,000
2	800,000
3	1,600,000
3	2,500,000
4	3,500,000
5	5,000,000
7	7,500,000
10	10,000,000
11	12,500,000
12	15,000,000
13	20,000,000
15	30,000,000
17	50,000,000
20	80,000,000

Table 3: IP Allocation Guide

2.8 Plan for a validated ramp-up

A validated ramp-up Scheduler across 2-6 weeks helps establish sending history for both sending domain and associated IP address(es). An illustrative ramp up schedule is provided below, but you can find our Warm Up Template **HERE**

Execution Timeframe		Ramp up factor
Day	**Volume**	**overall ramp up factor**
1	50	
2	60	1.20
3	72	1.20
4	86	1.20
5	104	1.20
	End of Week Checkpoint	
6	145	1.40
7	203	1.40
8	284	1.40
9	398	1.40
10	558	1.40
	End of Week Checkpoint	
11	892	1.60
12	1,427	1.60
13	2,284	1.60
14	3,654	1.60
15	5,847	1.60
	End of Week Checkpoint	
16	10,525	1.80
17	18,944	1.80
18	34,100	1.80

19	61,380	1.80
20	110,484	1.80
	End of Week Checkpoint	
21	198,870	1.80
22	357,967	1.80
23	644,340	1.80
24	1,159,812	1.80
25	2,087,662	1.80
	End of Week Checkpoint	
26	3,757,791	1.80
27	6,764,024	1.80
28	12,175,243	1.80
29 **30**	21,915,438 39,447,788	1.80 1.80
	End of Week Checkpoint	
31	71,006,019	1.80
32	127,810,833	1.80
33	230,059,500	1.80
34	414,107,100	1.80
35	745,392,781	1.80
	End of Week Checkpoint	
36	1,341,707,005	1.80
37	2,415,072,610	1.80
38	4,347,130,698	1.80
39	7,824,835,256	1.80
40	14,084,703,461	1.80

Table 4: IP Warmup

As you will now be moving to a new infrastructure environment using new IP's then IP warming becomes mandatory and is not optional. Please refer to our Template for Dedicated and Shared IP recommendations

Some important warmup rules are as below:
It is important to mail to the most recent & active subscriber list during the warm up period.
Maintain a regular rhythm and distribution between various ISP's in the recipient list (i.e. not to send to 1 yahoo account in week 1 and 50 in week 2).
Always start with 50 email recipients / day and grow the recipient list as per our recommended schedule. It's important to attempt to keep to our proposed % increases per ISP e.g. no more than 18% / Day for Microsoft Domains
No Warm Up plan is a guaranteed success therefore it's important to review and adjust warm up based on Bounce / Complaint / SPAM statistics (as above).

- We recommend sending different types of email-content (promotional/ transactional) via separate IPs i.e. do not mix traffic from marketing and transactional on the same IP.
- IP warmup also applies to transactional email use cases. The IP should be warmed up slowly by continuous increase. If a customer migrates existing traffic from another provider, the customer should split the volume at the beginning and start with a small number of messages daily and slowly increase.
- Warmup process for Microsoft
    - Customers to understand, that there is no generic IP "Unblocking" possible. Start with warm up process and the request for delisting must happen/go hand-in-hand.
    - De-listing request only makes sense, as soon as there is production traffic (warm up). Otherwise MS is seeing no reason for a blocking.
    - Expectation of the customer must be set accordingly that this can happen, but also that this is quite normal in the warm up phase.

- The warm up process must not be stopped when IP blocking happens and the plan should be followed. Otherwise, MS Support may classifies this as erratic sending behavior and use this as reason to deny any mitigation.

2.9 Ensure double opt-in & maintain list hygiene

What makes an email list “double opt-in” is that any person who subscribes must confirm their request twice.

The first time is when the user submits their email address to the web-based form.

After the initial request is received by the email list software a special confirmation email is sent to the address the person input into the form. This is the second opt-in. The email contains a link which the recipient must click to confirm their subscription request. Once they have done this they have “double opted-in”.

If the target subscriber list is old; one approach to ensure double opt-in is to send the target recipients an email asking for their confirmation to be a part of the campaign list.

Constant clean-up on the subscriber list based on “double op-in“ ensures list hygiene.

2.10 Avoid SPAM traps (sending emails to hard bounced addresses)

A spam trap is an email address traditionally used to expose illegitimate senders who add email addresses to their lists without permission. But they are also set up to identify email marketers with poor permission and list management practices.

Sending to a spam trap can be very damaging to your sender reputation as well as causing deliverability problems

Ensuring that the subscriber list is cleared of email addresses that hard bounce helps avoid SPAM traps.

2.11 Valid reply-to address. Do not use of DO-NOT-Reply@ or noreply@

ISP’s process the reply-to header differently; but here are important considerations:

Do not use Yahoo and AOL based reply-to addresses. AOL and Yahoo have made changes to their DMARC policy, which tells other mail systems that they should reject any mail with a "from" address at the Yahoo or AOL domains, if the message did not actually come from Yahoo or AOL.

Avoid using addresses such as “noreply@..“;“do-not-reply@..“. These are seen as non-engaging and liable to increase the propensity to be filtered as SPAM’s.

It is recommended that a verified email address be used as a “reply-to“ address.

2.12 Maintain consistent traffic pattern

Maintaining a consistent email send pattern that includes:

- **Rhythm** - no spikes in email traffic and no more than 2X times previous day’s send
- **Distribution** – ensuring that the distribution of subscribers is consistent during regular campaigns from the same sender.
- **Sending times** – Consistency in send-time behavior helps reduce ISP flagging email as “potential” spam and throttling the traffic.

ISP’s reward predictability and penalize extreme variations by aggressively defining variances as “potential” spam.

2.13 Do not use too many or shorted URL’s. Use Text & HTML both. Do not use encoding.

ISP’s are especially sensitive of shorted URL’s such as using bit.ly and/or encoded emails; flagging such emails as SPAM.

2.13.1 Improving email content: Things to consider

When a message is received by a remote mail server, it first considers aspects such as reputation, blacklists, and information from the headers received during the SMTP conversation. This allows the receiving software to decide whether to allow, filter, or block the incoming message based on domain or IP address reputation and authentication.

If the receiving server allows the message to move through the next stage, then the full content of the message is received and evaluated. Based on the information within the body of the message and the subject line, it determines whether to mark the message as spam or deliver it to the inbox. Anti-spam software is constantly learning by gathering information over time from user feedback (retrieving legitimate message from the junk folder or marking messages as spam). So, ensuring consistent inbox placement from a content perspective requires regular fieldwork and adaptation for senders.

2.14 How to improve content?

There is no ‘perfect’ email template, so it is important to be weary of the guidelines provided. There is a long list of criteria used in determining whether a message should be considered spam, and every single mailbox provider and anti-spam software has its own ‘secret sauce’ when it comes to reading and understanding the content of incoming messages. The following are some steps you can take to improve the chances your messages will pass content filtering.

2.14.1 Balance text and imagery

Don’t create messages as a single large image, as this is a common spammer technique used in attempt to bypass spam filters. Embedding large images in emails or using a lot of graphics can also slow the email server’s ability to process mail. As a result, content spam filters will often flag such emails and stop delivery. Keep in mind also that some mailbox providers turn images off by default, so it's likely images won't be seen anyway. As a rule of thumb, we advise a good balance of text and images. The overall goal is to have enough text in the body of the message so subscribers will understand what is being conveyed whether images are on or off.

2.14.2 Check your HTML

Most emails today are created in HTML, so having a nicely formatted HTML message is a good start. Broken HTML can lead to a poorly rendered message and generate complaints if recipients believe it's a phishing attempt. Make sure your HTML is free of syntax errors and formatting errors.

2.14.3 Check your Spam Score

The likelihood of email being classified as spam depends on various factors. Spam alerts can be triggered by unusual HTML formatting or table constructions, excessive links, or dubious wording in the subject line and email body.

2.14.4 Test, test, test

Testing message content in a pre-deployment tool such as Return Path's Inbox Preview can help to identify potential spam filter issues before you send. Once you identify content that is being flagged by spam filters, continue testing to isolate what is causing the issues (subject lines, URLs/links, text, and/or images). Content testing can be a time-consuming process, but well worth the effort.

2.14.5 Avoid base64

Messages that have a base64 encoded body or subject line are more likely to be flagged as spam by anti-spam software, mainly because this is a known tactic used by spammers to hide the content from anti-spam software.

2.15 Fingerprinting

Most senders know that the content of their email is scrutinized for “spammy” content, but it’s interesting to understand the methods used to examine content. One well-known method of analyzing content is called “fingerprinting.” Some technology providers are known for creating fingerprints of email content.

Fingerprinting in and of itself is not a filter, but it is a technology that helps mailbox providers make decisions about email content.

Fingerprints are hashes or checksums of content. These hashes are many times smaller (64 bytes) than the content that they’re generated from, which makes them easier to store. Once the fingerprints are created and stored, they can be compared to other fingerprints. The result of the comparison helps filters decide whether email is spam by scoring the similarity of fingerprints, meaning if your fingerprint is highly similar to a fingerprint belonging to email that has been confirmed as spam, then your mail will likely be flagged as unwanted mail.

At the risk of repeating some of the points highlighted above; here is a quick check list around email content:

- ✓ Offer a safe and secure unsubscribe option (Note:
- ✓ Clearly communicate your privacy policy
- ✓ Validate you are adhering to applicable anti-spam and privacy laws and policies
- ✓ "Opt-out" - Consider moving your subscription management options to the top of the email vs. the bottom
- ✓ Make sure you are communicating a unique value proposition clearly, even with images turned off

- ✓ Change the Subject lines of your messages to be more relevant
- ✓ Have a clear call to action and make sure it is above the fold; top left and 2 - 4 inch section of the page
- ✓ Choose content wisely and verify URLs look normal and point to valid domains
- ✓ Format a reply header to ensure subscribers see your “friendly” e-mail address
- ✓ Sell the offer, not the product
- ✓ There should be one offer and one call to action
- ✓ Repeat the call to action in the email. Don’t only have an image with the call to action, include text and image links if possible
- ✓ Define alt tags for images that include your subject line or headline
- ✓ Keep your message simple and to the point (avoiding fluffy marketing speak whenever possible)
- ✓ Avoid large blocks of text (break into paragraphs)
- ✓ Don't use too many images; rule of thumb is image-to text ratio is 30/70 or less. i.e., Hugo Boss suit vs. amount of content
- ✓ If possible, employ A/B testing to see what layouts, subject lines, time of day and approaches work best
- ✓ Does the subject line include the topic and benefit?
- ✓ Keep subject lines to 40 characters or less
- ✓ Offers should be specific and tangible
- ✓ Messages over 100k in size get blocked more often. Aim for 35 to 75k
- ✓ Add text reminding subscribers where they opted-in to receive your e-mail.

3 Appendix A – Using DKIM and SPF for authentication

3.1 Domain Key Identified Mail – DKIM

DKIM provides a method for validating a domain name identity that is associated with a message through cryptographic (public-private keys) authentication.

Identification of an email's association with a trusted domain enhances deliverability.

3.1.1 What is Domain Key Identified Mail?

A DKIM signature is associated with the email body and selected parts of header

The DKIM Signature (Private Key) is transmitted in the email header.

The DKIM Signature (Public Key) is stored in the DNS.

This DKIM signature public key is maintained in the DNS records as _domainkey.yourdomain o The text record associated with the DKIM public key is stored in the DNS server

To support multiple concurrent public keys per sending domain, the DNS namespace is further subdivided with "selectors".

Selectors are arbitrary names below the _domainkey. namespace.

For example, selectors may indicate the names of your server locations (e.g. mta1 or mta2), the signing date (e.g. january2005, february2005, etc…), or even the individual user.

The most important thing is: selector indicates your DomainKeys/DKIM public key location. For example:

- if your domain selector is s1024, your public key record will be s1024._domainkey.yourdomain;

- if your domain selector is mta1, your public key record will be mta1._domainkey.yourdomain.

The selectors can also be used by a third-party email service to create signatures associated with the customer domain.

Unlike S/MIME and OpenPGP based authentication mechanisms, DKIM does not modify the body.

Instead, it places its parametric information into header fields that are typically not shown to the recipient.

Therefore, DKIM's can be entirely invisible to recipients. This is a significant advantage.

3.1.2 What does DKIM not do?

DKIM does not offer any assertions about the behaviors of the identity doing the signing.

DKIM does not prescribe any specific actions for receivers to take upon successful (or unsuccessful) signature validation.

DKIM does not provide protection after message delivery.

DKIM does not protect against re-sending (replay of) a message that already has a valid signature. Therefore, a transit intermediary or a recipient can re-post the message in such a way that the signature would remain valid, although the new recipient(s) would not have been specified by the originator.

3.1.3 How does DKIM work?

3.1.3.1 Sending servers (this could be a 3rd party ESP’s server):

There are two steps to signing an email with DKIM:

The domain owner generates a public/private key pair to be used for signing outgoing messages (multiple key pairs are allowed).

The public key is published in a DNS TXT record, and the private key is made available to the DKIM-enabled outbound email server.

When an email is sent by an authorized user of the email server, the server uses the stored private key to generate a digital signature of the message (which is inserted in the message as a header) and the email is sent as normal.

3.1.3.2 Receiving Servers

1. The DKIM-enabled receiving email server extracts the signature and claimed From: domain from the email headers.
2. The public key is retrieved from the DNS system for the claimed From: domain.
3. The public key is used by the receiving mail system to verify that the signature was generated by the matching private key. A match effectively proves that the email was truly sent from, and with the permission of, the claimed domain and that the message headers and content have not been altered during transit.
4. The receiving email system applies local policies based on the results of the signature test. For example, the message might be deleted if the signature does not match.

3.1.4 A quick guide on making DKIM work for you

Typical Scenario: Customer manages their DNS administration for consuming applications. If the customer does not manage and an 3rd party agency is the webmaster; these tasks will need to be performed by the 3rd party webmaster.

All upstream applications consuming Sinch email services will be supported by DKIM to enhance delivery reliability. Here are simple steps that every customer needs to undertake:

1. In the email account set up provisioning form (provided by the Sinch sales executive), please highlight the necessary sub-domains that will be used for sending emails. Illustrative sub-domains are engg.company1.com or marketing.company1.com. Sinch will set up the email service accounts to map to

these sub-domains and pass back the necessary credentials which shall include the DKIM parameters such as the selector that will be used by an Sinch partner of high reputation to sign emails on your company’s behalf.

a) Public key: Typically, it will look as below:

```
v=DKIM1; k=rsa;

p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCrLHiExVd55zd/IQ/
J/mRwSRMAocV/hM3jXwaHH36d9NaVynQFYV8NaWi69c1veUtRzGt7yAioX
qLj7Z4TeEUoOLgrKsn8YnckGs9i3B3tVFB+Ch/4mPhXWiNfNdynHWBcPcbJ8
kjEQ2U8y78dHZj1YeRXXVvWob2OaKynO8/lQIDAQAB;
```

b) TXT Record / Selector information.

2. DKIM keys will be provided in 3 formats to support various DNS server configurations. As an example, a DKIM key pair was created for an illustrative engineering.customer.com as follows:

a) Bind9 format:

```
selector._domainkey.engineering.customer.com IN TXT (
v=DKIM1;t=s;p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDQ35
G5S9EWxbBFgKWUzwIrBTeZ6YmvcwWuJ2GC8FAIdAi6wOGQYDCRPrLM2
WaCwDAQP+NYDlU7flJOAUuD7G4HGrEEmZfm57wl7TV5IoqPBorlv55BhLS
ONtWM1fqk1CSrgqQrmhWQcPxeZ19dxWV2wjKLmIRuzUI6vMPtF58J
KwIDAQAB)
```

b) TinyDNS format:

```
‘selector._domainkey.engineering.customer.com:v=DKIM1;p=MIGfMA0GCSq
GSIb3DQEBAQUAA4GNADCBiQKBgQDQ35G5S9EWxbBFgKWUzwIrBTeZ6
YmvcwWuJ2GC8FAIdAi6wOGQYDCRPrLM2WaCwDAQP+NYDlU7flJOAUu
D7G4HGrEEmZfm57wl7TV5IoqPBorlv55BhLSONtWM1fqk1CSrgqQrmhWQc
PxeZ19dxWV2wjKLmIRuzUI6vMPtF58JKwIDAQAB:3600::
```

c) Raw format (if needed).

```
-----BEGIN PUBLIC KEY-----

MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDQ35G5S9EWxbBFg
KWUzwIrBTeZ6YmvcwWuJ2GC8FAIdAi6wOGQYDCRPrLM2WaCwDAQP+N
```

YDlU7flJOAUuD7G4HGrEEmZfm57wl7TV5IoqPBorlv55BhLSONtWM1fqk1C
SrgqQrmhWQcPxeZ19dxWV2wjKLmIRuzUI6vMPtF58JKwIDAQAB

-----END PUBLIC KEY-----

3. Have the Webmaster update the DNS records by doing the following:
    a) Associate the easylink selector with the subdomain, by updating the TXT record. The selector name and TXT details will change with each new sub domain and will be provided by Sinch
    b) Map the DKIM public key with the sub-domain.
    c) Update the policy record value (optional and only available in some DNS servers). If the policy record value is “-“, all emails will be sent with the DKIM signature.
    d) Have the public key parameter “t” set to “y”, to enable testing.
4. Send a confirmation to Sinch email address on DKIM set up; so Sinch can test the set up and confirm that the Hybris set up is good to go.

3.1.5 What does Sinch do?

Once, Sinch receives confirmation, the Sinch Sales engineer will do a quick test to confirm the DKIM key set ups are operational.

These tests will involve the following:

1. **Verify domain:** This will confirm that the domain/sub-domain set up is accurate in the ESP, the ownership is accurate and a basic reply to address is set up.
2. **High level test:** Use tools available at DKIMcore.org (http://dkimcore.org/c/keycheck) to set the DKIM key association with the subdomains and the selector. Please see illustrative screenshot below:

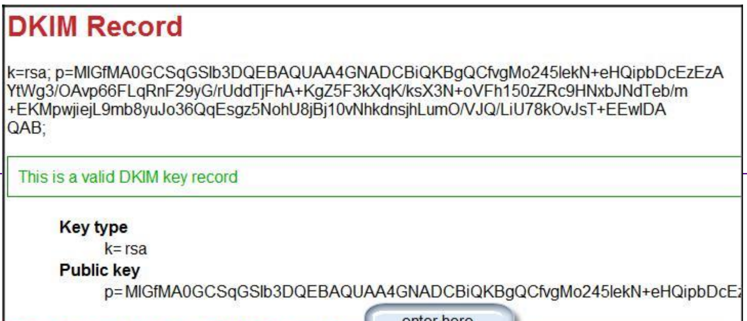

DKIM Record

k=rsa; p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCfvgMo245lekN+eHQipbDcEzEzA
YtWg3/OAvp66FLqRnF29yG/rUddTjFhA+KgZ5F3kXqK/ksX3N+oVFh150zZRc9HNxbJNdTeb/m
+EKMpwjiejL9mb8yuJo36QqEsgz5NohU8jBj10vNhkdnsjhLumO/VJQ/LiU78kOvJsT+EEwIDA
QAB;

This is a valid DKIM key record

Key type
k= rsa
Public key
p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCfvgMo245lekN+eHQipbDcEz

As a result, the DKIM key shall be returned:

k=rsa;

p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCfvgMo245lekN+eHQi
pbDcEzEzAYtWg3/OAvp66FLqRnF29yG/rUddTjFhA+KgZ5F3kXqK/ksX3N+oVF
h150zZRc9HNxbJNdTeb/m+EKMpwjiejL9mb8yuJo36QqEsgz5NohU8jBj10vNhk
dnsjhLumO/VJQ/LiU78kOvJsT+EEwIDAQAB;

3. **Detailed tests**: Send signed test emails to a known SMTP server to check deliverability (based on known spam filters rules).
4. **Confirmation**: Confirm set up is OK to customer via email. This is undertaken by a sales engineer.

3.2 Sender Policy Framework - SPF

3.2.1 What is SPF and how does it work?

Sender Policy Framework (SPF) is a simple email-validation system designed to detect email spoofing by providing a mechanism to allow receiving mail exchangers to check that incoming mail from a domain comes from a host authorized by that domain's administrators.

If a domain publishes an SPF record, spammers and phishers are less likely to forge e-mails pretending to be from that domain, because the forged e-mails are more likely to be caught in spam filters which check the SPF record. Therefore, an SPF-protected domain is less attractive to spammers and phishers.

Because an SPF-protected domain is less attractive as a spoofed address, it is less likely to be blacklisted by spam filters and so ultimately the legitimate e-mail from the domain is more likely to get through.

3.2.2 How to set up an SPF record for your domain?

To set up an SPF record for your domain, create a TXT record against the domain/ sub-domain and insert the following values:

<domain/ sub-domain> in TXT "v=spf1 ip4: 149.235.15.0/24~all".

4 Appendix B - Feedback Loops

A **feedback loop** (FBL), sometimes called a **complaint feedback loop** is a service offered by some ISPs that report back complaints (when a recipient hits the spam or junk button in on their email it is considered a “compliant”) to senders. It is a form of feedback by which an ISP forwards the complaints originating from their users to the sender. It’s provided to aid senders in keeping a clean list of recipients by removing such recipients from their lists and ensure they won’t continue to receive unwanted emails. This will ensure that recipients aren’t submitting multiple complaints, keeping the compliant rate low and improving deliverability because ISPs won’t be as likely to quarantine or reject their messages.

4.1 ISPs that provide feedback loops to the Sinch E-Mail 365 platform

- Outlook
- AOL
- BAE Systems
- Bluetie / Excite
- Comcast
- Cox
- Earthlink
- Fastmail
- Hotmail
- Italia Online / Liberio
- Laposte
- Locaweb
- Mail.ru
- Open SRS

- QQ.com
- Rackspace Roadrunner
- Synacor
- Telenor
- Telstra
- Terra
- Tucows
- Usa.net
- United Online / Juno / Netzero (UOL)
- Verizon
- XS4ALL
- Yahoo
- Yandex
- Zoho.com

4.2 Gmail Feedback Loops:

- Gmail does not provide standard FBL functionality to ESP's. In order to support Gmail feedback loop, we embed a Feedback-ID header consisting of identifiers that will uniquely identify sender campaigns. These identifiers will be used by Gmail to report aggregated complaints in the Gmail postmaster tools FBL dashboard. Feedback-ID header will consist of 4 identifiers that uniquely identify the individual campaigns, Feedback-ID: a : b : c : SenderId, where
  - a= campaignID (Campaign Identifier specific to Customer)
  - b= customer sender email address (customer Identifier)
  - c= systemID (Identifier for the type of mail)
  - SenderId= Customer Account ID (Sender's unique Identifier)

  SenderId is the only mandatory identifier. If the other identifiers are not set up then those identifiers will go unused

This header will only be used by Gmail and to get aggregated feedback loop information customers need to sign up on google postmaster portal https://postmaster.google.com/u/0/managedomains feedback loop information will be provided by Gmail to registered customers on this portal. However this this information from Gmail is only provided as aggregated data, actual spam complaints made by recipients are not available
Steps to sign up
https://support.google.com/mail/answer/6227174?hl=en&ref_topic=6259779

Please contact your Sinch Account Executive to check if this update is available on your account

4.3 Yahoo Feedback Loops

- Yahoo has completely changed their Feedback Loop system. They have now developed a new feedback loop program where customers need to register up on their portal for the yahoo feedback loop, and provide an email address where yahoo will send the customer complaint notifications. https://help.yahoo.com/kb/sign-manage-yahoo-complaint-feedback-loop-program-sln3438.html

5 Appendix C – Suppressions

5.1 Suppression Lists

Recipient email addresses after a bounce may be added to a suppression list based on the reason for the bounces Suppressions are not based on 'hard bounce' or 'soft bounce' but based on the reason or the bounces. The function of the Suppression list is to avoid a recurring send to mailbox / user that has bounced. This helps to avoid impact to IP reputation. After expiry of the period the address will be automatically removed. In case the address is still used and runs into error, the period will be extended accordingly. For more details on the reason and the time period of suppressions please refer to the table below. If an email is sent to an address on a suppression list then the sender will receive a notification with statusText = 'SUPPRESSED'. Please refer to the API guide for more details on these notifications

	Reason detail	**Suppression days**
1	over quota:	7 Days
2	quota exceeded:	7 Days
3	mailbox full:	7 Days
4	mailbox unavailable:	1 Day
5	unknown local:	1 Day
6	does not exist:	1 Day
7	user unknown:	1 Day
8	address rejected:	1 Day
9	no such user:	1 Day
10	Unknown recipient:	1 Day
11	Host or domain name not found:	1 Day
12	address is administratively disabled:	7 Days
13	nullmx:	7 Days
14	Syntactic incorrect address:	7 Days
15	no mailbox:	7 Days

16	invalid mailbox:	7 Days
17	unknown user:	7 Days
18	mailbox not found:	7 Days
19	Connection refused:	7 Days
20	Unknown or illegal alias:	7 Days
21	unrouteable address:	7 Days
22	recipient rejected:	1 Day
23	cannot be delivered:	1 Day
24	no such mailbox:	1 Day
25	is disabled:	1 Day
26	Loops back to myself:	1 Day
27	account has been disabled:	1 Day
28	address unknown:	1 Day
29	No information available:	1 Day
30	Malformed or unexpected name server reply:	1 Day
31	address invalid:	1 Day
32	user inactive:	1 Day
33	No route to host:	1 Day
34	recipient not found:	7 Days
35	exceeded storage:	7 Days
36	account.*exist:	45 Days
37	doesn't have.*account:	45 Days
38	address.*unknown:	45 Days
39	unknown.*alias:	45 Days
40	account.*disabled:	45 Days
41	address.*disabled:	45 Days
42	mailbox.*unknown:	45 Days
43	user.*not found:	45 Days
44	mailbox.*unavailable:	45 Days
45	Sender Policy Framework:	45 Days

Table 5: Suppression rules

5.2 Role based Suppressions

In regards to IP’s registered for Marketing email sending we apply a global suppression rule for roles based email addresses. The role based email addresses are normally generic email addresses which typically define a responsibility rather than a person and which go to several people i.e. admin@, info@, work@, etc. Therefore, they are unlikely to give their consent to receive marketing emails. Having roles addresses is considered bad practice in general

Sending to these email addresses mean that you won't get the full effect of the email marketing campaign as they are not really targeted, they could go to several people, most of which will not want to receive the emails because they have not opted in to receive marketing emails. Which then means these types of addresses are associated with high bounce rates and spam complaints. Additionally there have been studies conducted by some ESPs that reinforces the claim that role addresses on a list have a negative effect: even a single role address on a list is associated with more bounces, more unsubscribes, and decreased delivery and engagement rates.

- abuse@
- admin@
- billing@
- compliance@
- devnull@
- dns@
- ftp@
- hostmaster@
- inoc@
- ispfeedback@
- ispsupport@
- list-request@
- list@
- no-reply@
- noreply@
- null@
- phish@
- phishing@
- postmaster@
- privacy@
- registrar@
- root@
- security@
- spam@
- support@
- sysadmin@
- unsubscribe@
- usenet@
- uucp@
- webmaster@
- www@

maildaemon@ noc@	tech@ undisclosed-recipients@	

Table 6: List of roles based suppressions

6 Appendix D – Retry Frequency

All soft bounces are retried in a 48 hour window

- The first two retries after 30 minutes.
- Thereafter one retry after every ~ 65 minutes.
- After 48 hours it ends with final status if there was no successful delivery possible.
- Total number of retries should be around 48 (including the final status)

7 Appendix E – Marketing Vs. Transactional email

Marketing/Promotional e-mail

- Contains commercial content for a commercial purpose
- Promotes a brand/product/service
- Is sent to a large number of interested parties or customers
- Is campaign controlled
- Must comply with local laws

e.g. Advertising & sales campaigns, newsletter, shopping offers, discounts, product launches, 3rd party offers

Transactional e-mail

- One-to-one communication
- To individuals, not a large recipient list
- Event triggered, controlled by website or app
- Personal communication

e.g. Welcome mail, notification, reminder, confirmation, receipt, password change, request emails

*When sending transactional mails, it must be clearly defined WHO is sending the message (e.g. orders@sinch.com). WHAT is the purpose of a message and WHY the message is sent

	Marketing/Promotional Mail	Transactional Mail
Source	Sent through a campaign/marketing platform	Automated generated by an event via an application / platform
Comply with local laws	Yes	Yes

E-Mail Type	Coupons, advertising, info about new products, newsletter, sales campaigns, newsletter, shopping offers, discounts, product launches, daily deals, third party, offers	Welcome mails, orders, delivery notifications, receipts, notifications, reminders, confirmations, password changes, request mails
Ratio	One to many	Usually oone to one
Sender list	Sender list is generated through opt-in. Unsubscribe for opt-out necessary Large list of recipients	To individuals initiated by a purchase, with agreement thru Eg. Delivery confirmation by email
Volume	Large	Low
Unsubscribe link	Mandatory (CSA, local laws) Details on unsubscribe in API specification	Not mandatory (but look at local laws, often other way to get out of communication like buttons etc. must be given or other way offered for communication)
From Address	Brand / product related	Application type related
Reply to	This should be a real email contact. A 'no-reply' to address is not recommended for marketing emails	Should be a real email contact because a reply from recipient is expected. Trustful and responsible address absolutely necessary
IP & Domain Management	IP´s with focus on marketing use case	IP´s with focus on transaction use case
Domain Management	Use a separate domain for marketing emails so that you don't mix your domain reputation with transactional emails.	Use a separate domain for transactional emails so that you don't mix your domain reputation with marketing emails.

	Separate domains with a very high volume (>> 1,000,000 emails per month)	Too little traffic is also a blacklist risk. Therefore, merge domains with very little traffic to others with more.
ISP´s reputation checks	SPF and DKIM mandatory	SPF and DKIM mandatory
Update senderlist (Bounce Mgmt.)	Required	Required

Table 7: Marketing Vs. Transactional emails