TIR-20260225 Email Spoofing - Why Internal Emails Are Not Always Safe

2/25/2026

Prepared for:

Aspire Technology Partners
25 James Way
Eatontown, NJ 07724

NOTICE:

This document is marked TLP: CLEAR, allowing recipients to share this information globally without any disclosure limitations.

This report is governed by the terms outlined in the Master Services Agreement (MSA) or any other master agreement between Aspire Technology Partners and the client receiving this report, along with the Statement of Work (SOW) or Order detailing the services that led to its creation.

Contributor(s)

Portia S. Cole
CTI Threat Researcher
Aspire Technology Partners
pcole@aspiretransforms.com

TABLE OF CONTENTS

Executive Summary

Most people assume internal email is safe. If it comes from your own domain, it must be legitimate. That's how employees think, and in many environments, security tools reinforce that assumption. Internal traffic is often handled differently. It may bypass certain filters, skip warning banners, or be treated with a higher baseline of trust. Over time, that becomes invisible. No one consciously decides to trust internal email more, they just do and that's where the exposure starts.

Attackers have adapted to that reality. Instead of relying solely on obvious phishing campaigns from newly registered domains, they look for ways to inject messages into trusted mail flow paths. Misconfigured connectors, permissive Direct Send settings, and aging hybrid Exchange configurations can all create opportunities. If a malicious email lands in a mailbox appearing to originate internally, most users won't hesitate. The message doesn't feel like an attack, it feels routine.

The risk has shifted because of that. Security awareness programs have trained employees to be cautious of

TIR SUMMARY

ASPIRE

TLDR;

- **The Threat** - Attackers are shifting away from obvious external phishing and are instead delivering emails that appear to come from inside the organization. When a message looks internal, suspicion drops and response speed increases.
- **TTPs** - Abuse of Microsoft 365 routing and trusted mail paths, internal-looking subject lines tied to routine business activity, QR code–based credential harvesting, multi-stage redirect chains, and use of legitimate cloud-hosted infrastructure to reduce detection friction.
- **Aspire Case Study** - An internal-looking email containing a malicious QR code led a user to a credential harvesting page. The SOC detected abnormal activity quickly, reset the affected account, invalidated sessions, and prevented confirmed data loss.
- **Lessons Learned** - Internal does not equal safe. Mail flow configuration is part of the attack surface. Trusted routing paths must be reviewed regularly, and employees should verify sensitive requests through a separate communication channel – even when the message appears to come from inside the organization.

external senders, but internal-looking messages often bypass that instinct. Organizations now have to examine their own routing logic, trust relationships, and Microsoft 365 configurations like they do external threats. The real question isn't just whether phishing emails are being blocked, it's whether internal trust assumptions can be weaponized.

WHAT IS EMAIL SPOOFING

Email spoofing is the act of forging sender information so that a message appears to originate from someone it did not come from. The attacker manipulates visible header fields such as the "From" address or display name in order to impersonate a trusted sender.

SMTP, the protocol used to send email, was built for delivery, not identity validation. Authentication mechanisms such as SPF, DKIM, and DMARC were added later to improve domain verification. When properly configured and enforced, these controls reduce the likelihood of successful domain spoofing. When misconfigured or loosely enforced, spoofed messages may still be delivered.

Spoofing takes several forms

- **Display Name Spoofing** – The attacker alters the visible sender name but uses a different domain.
- **Domain Spoofing** – The attacker forges the sending domain in the header.
- **Internal Domain Spoofing** – The attacker attempts to make the email appear to originate from the organization's own domain.
- **Reply-To Manipulation** – The attacker changes where responses are sent, even if the visible sender looks legitimate.

The technique itself is not new. What has evolved is how it is combined with routing and infrastructure decisions to increase credibility.

Spoofing in Microsoft 365 Environments

Microsoft 365 evaluates incoming mail using multiple signals. These include SPF alignment, DKIM signatures, DMARC policy enforcement, IP reputation, and connector

associations. How these signals are interpreted can affect classification and delivery outcomes.

Mail flow connectors define how messages move between Microsoft 365 and external or on-premises systems. In hybrid environments, mail may transit between on-prem Exchange servers and Exchange Online before reaching the recipient. Application relays and device-based senders may also be permitted to route mail through defined pathways.

The presence of these components introduces complexity. Classification decisions may depend on how a message entered the tenant and whether it matched a defined connector. In some configurations, messages may be delivered without prominent external indicators if routing logic treats them as trusted.

This does not automatically indicate a vulnerability. However, it means that spoofing risk is not limited to open-internet delivery attempts. Configuration accuracy directly influences how spoofed messages are evaluated.

Microsoft 365 Factors That Influence Spoofing Risk:

- Connector IP scope and restrictions
- Hybrid Exchange routing paths
- SPF, DKIM, and DMARC enforcement level
- External tagging policies
- Transport rule configuration
- Application relay permissions

How Email Systems Are Built

Enterprise email environments are layered systems by design. A message may pass through a secure email gateway, Microsoft's filtering stack, hybrid Exchange infrastructure, and internal transport rules before reaching the inbox. Each layer exists to keep communication flowing and enforce policy. Connectors are authorized to move mail between systems and certain IP ranges are designated as trusted. Hybrid servers

are permitted to relay traffic between cloud and on-prem environments, while applications and devices are allowed to send automated notifications.

The risk does not come from trust itself; it comes from how that trust is maintained. Mail flow logic does not evaluate intent; it enforces rules. If a message meets the conditions of a trusted path, it is processed accordingly. Over time, connector scopes expand, legacy routing remains in place, and complexity increases. Internal classification then adds another layer. Messages **treated as internal may not display external warnings**, and users are less likely to question them.

Common Trust Assumptions Embedded in Email Architecture:

- Mail from defined IP ranges is legitimate
- Connector-authorized traffic is safe
- Hybrid routing paths are secure by default
- Internal domains require less scrutiny
- Authenticated relay traffic does not need additional inspection
- System-generated messages are inherently trustworthy

How Attackers are Abusing Microsoft 365 Routing and Trusted Email Paths

As organizations train employees to recognize security threats, threat actors adapt their tactics in response. Employees are trained to be cautious of external senders, and security gateways reinforce that caution by adding warning banners to incoming messages. As a result, only external phishing attempts are viewed as suspicious.

Attackers have shifted toward making their messages look internal, knowing that when an email appears to come from inside the organization, people respond faster and question it less. Abuse may involve:

- Leveraging misconfigured Direct Send
- Exploiting permissive mail connectors
- Injecting mail through hybrid Exchange routing
- Using legitimate cloud infrastructure to host phishing content

The objective is not to bypass authentication of a user account. It is to bypass suspicion. In several documented cases across the industry, Microsoft 365 environments with complex routing configurations allowed attackers to deliver spoofed internal messages.

Legitimate infrastructure further complicates detection. Hosting services such as serverless platforms and marketing email subdomains are used as redirect points. Reputation-based filtering becomes less effective when widely used domains are involved.

Observed Abuse Patterns

- Internal-looking subject lines referencing routine business activity
- QR codes replacing clickable links
- Multi-stage redirect chains
- Google or cloud-hosted sender infrastructure
- Minimal malware; focus on credential harvesting

Why Alarm Bells Often Do Not Trigger

When a message appears to originate from inside the organization, employees tend to lower their guard. There is no external warning banner and the sender name follows internal naming conventions. The request also fits within normal business activity and nothing visually seems out of place.

Security systems can be influenced by the same assumptions. Messages that traverse trusted connectors or internal routing paths may not be labeled as external. Filtering logic may treat them differently than internet-originated email. When the infrastructure involved is legitimate and commonly used, automated reputation checks may not immediately flag the activity.

Trust, in this scenario, is doing exactly what it was configured to do. The issue is not that the system failed. It is that the attacker positioned their message within expected boundaries. When the message looks internal and behaves like internal traffic, both human and technical caution decrease.

Common Factors That Reduce Suspicion

- No external sender banner
- Familiar internal display names

- Requests aligned with normal workflow
- Delivery through trusted connectors
- No obvious domain misspellings or spoofed lookalikes

An Aspire Case Study

In early 2026, a user received an email that appeared to originate from within their organization. The sender name aligned with internal formatting, and there was no external banner present to indicate that the message came from outside the company. Because the email appeared internal, the user did not question it before interacting with it.

The message contained a QR code. The user scanned it, believing it to be part of a legitimate business process. However, that action redirected the user to a credential harvesting page designed to capture login information.

Shortly after the interaction, abnormal activity was detected. Aspire's Security Operations Center (SOC) identified indicators consistent with credential compromise. The compromised account was immediately contained by the SOC. Sessions were invalidated, passwords were reset, and multi-factor authentication tokens were reissued. Fortunately, there was no data stolen and containment occurred before lateral movement was observed.

Direct Send Similarities

Microsoft 365 includes a feature known as Direct Send. It allows devices and applications to send email to internal recipients without authentication, provided specific routing conditions are met. While designed for operational purposes, the feature has been documented in industry reporting as a potential abuse vector when misconfigured.

Aspire's Cyber Threat Intelligence team recently published a Threat Intelligence Report analyzing how threat actors have exploited trusted Microsoft 365 routing paths in broader campaigns. Although this particular case was not confirmed to involve Direct Send, there are structural similarities worth noting.

Similarities

- The message appeared internal to the recipient

- No external warning banner was present
- Trust was derived from internal presentation

Differences

- There was no confirmed evidence of Direct Send exploitation
- Routing path abuse was not definitively identified
- The primary vector involved social engineering and credential harvesting via QR interaction

The similarity lies in the trust boundary, not the mechanism. In both scenarios, the attacker's objective is the same - position malicious content inside internal-looking traffic.

What Could Have Been Done Differently

On the user side, scanning a QR code without verifying the source is not safe. A secondary validation step, such as confirming the request through a separate communication channel, may have prevented interaction.

From the company side, this is a visibility problem. Internal-looking messages should not automatically blend in with routine traffic. Reviewing mail flow connectors, relay permissions, and routing logic on a regular basis makes it harder for attackers to hide inside trusted paths. In this case, the SOC caught the abnormal activity quickly, locked down the account, and stopped it before there was confirmed data loss.

The Risk for Organizations

When employees equate "internal" with "safe," attackers adjust their strategy. If they can make a message look like it came from inside the company, the chances of someone acting on it increase. Once credentials are handed over, the situation can move quickly. An attacker with access to a mailbox can read sensitive conversations, set up inbox rules to hide activity, or use that account to target others.

There is also financial risk. Business email compromise continues to result in real monetary loss across industries. A single internal-looking message requesting payment approval or updated banking details can be enough to trigger a costly mistake.

The bigger issue is overconfidence in perimeter defenses. Many organizations focus heavily on blocking threats from the outside while assuming internal routing is inherently trustworthy. But attackers do not always force their way in. Sometimes they step into existing trust relationships and use them.

Recent Attacks

Microsoft 365 Direct Send Abuse – July 2025

In July 2025, Proofpoint researchers reported an active campaign exploiting Microsoft 365's Direct Send feature. Threat actors used unsecured third-party SMTP relays to inject spoofed messages that appeared to originate from within targeted organizations. Although Microsoft flagged authentication failures, the messages were still delivered because they traveled through trusted routing paths. The emails blended into internal traffic, reducing both system friction and user suspicion.

The activity was not publicly attributed to a named threat group, but the tactics aligned with financially motivated phishing operations. The attack centered on inherited trust, not a software exploit.

Hybrid Exchange Routing Abuse – March 2025

In March 2025, researchers documented attacks against organizations running hybrid Exchange environments. Internal connectors configured to allow trusted relay between on-premises servers and Exchange Online were leveraged to inject internal-looking payment requests. In several cases, fraudulent transfers occurred before detection.

These incidents shared a common thread - attackers did not force entry through perimeter defenses. They analyzed mail flow architecture and positioned malicious messages inside trusted pathways.

THE RED FLAGS

Internal does not automatically mean safe. Messages that appear to come from inside the organization should still be evaluated with the same level of caution as external email. Attackers understand how employees are trained, and they adjust their tactics to fit within normal business communication.

Image 1: Spoofed Email Example

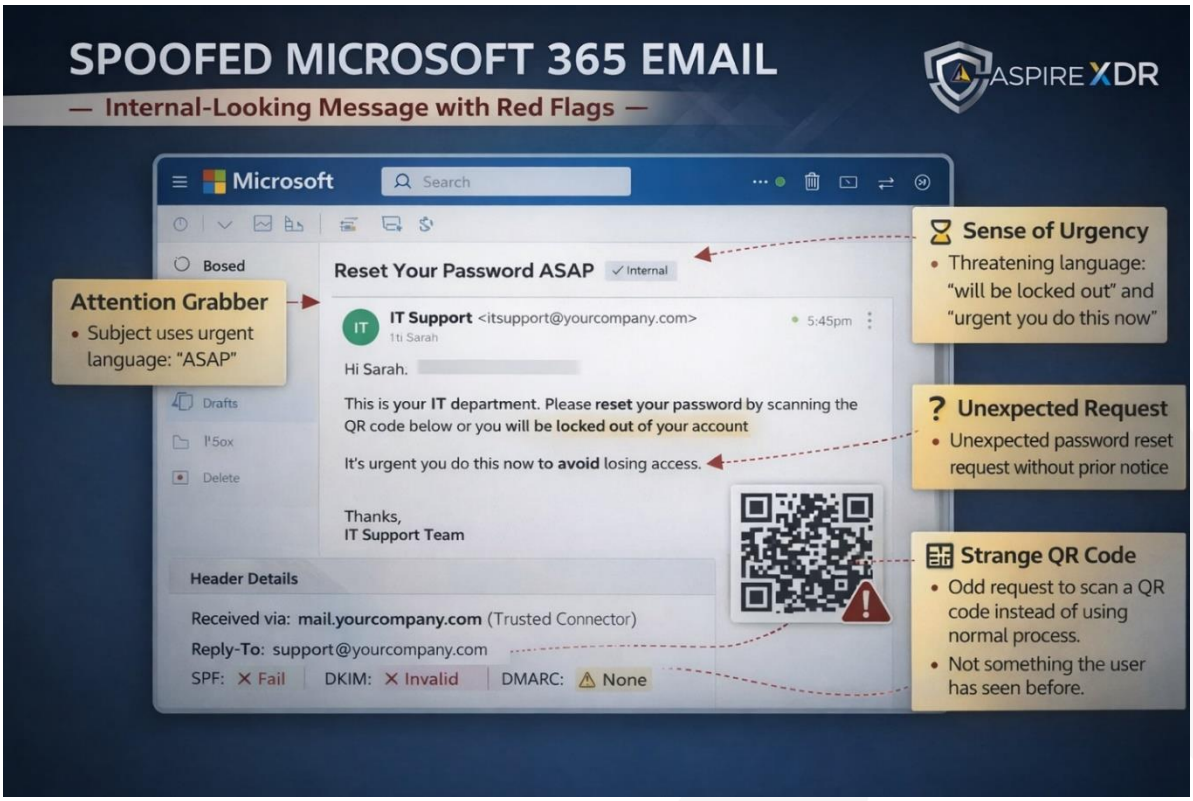

Before taking action on any request involving credentials, financial changes, or sensitive information, employees should pause and assess the context. Employees should pause and question:

- Why am I being asked to scan a QR code instead of using our normal login process?
- Did this person ever ask me for passwords or MFA codes before?
- Does this request match how we normally handle payments or approvals?

- Why does this need to be done right now?
- Is the tone slightly off, even if I can't immediately explain why?
- Am I being asked for access or information that isn't normally part of my role?
- If I reply or call the sender directly, would this request still make sense?

Verification through a separate, trusted communication channel remains one of the strongest defenses against internal spoofing and business email compromise.

Conclusion

Trust-based email architecture is necessary for business operations. However, implicit trust must be balanced with validation. Internal classification should not reduce scrutiny and routing paths should be treated as potential attack surfaces.

Organizations that regularly audit mail flow, restrict connectors, and train employees to question internal urgency will reduce exposure. The mailbox is no longer just a communication tool; it's also an attack surface.

Aspire's Recommendations

Modern email environments are more complex than they were a decade ago. As routing paths expand and hybrid configurations become common, defensive controls must keep pace with that complexity. Protection cannot rely on a single filter or banner; it requires layered awareness from both employees and the organization.

Employees should

- Verify sensitive requests through an independent communication channel
- Avoid scanning work-related QR codes from personal devices
- Report unexpected MFA prompts immediately
- Review sender details carefully
- Pause before acting on urgency

Technical controls matter just as much as individual vigilance. Organizations need to treat mail flow configuration as a security surface, not just an operational setting. That

means hardening routing paths and reducing the amount of inherited trust available to attackers.

Organizations should implement the following protections

- Disable Direct Send unless operationally required
- Restrict connector IP ranges to known infrastructure
- Enforce SPF, DKIM, and DMARC alignment
- Monitor for anomalous inbox rule creation
- Shorten refresh token lifetime
- Apply conditional access policies with token protection
- Conduct periodic mail flow configuration reviews

MITRE MAP

Aspire SOC TTPs

Initial Access	T1566 – Phishing
Credential Access	T1566.002 – Phishing: Spearphishing Link
Defense Evasion	T1583.002 – Acquire Infrastructure: Web Services
Command and Control	T1102 – Web Services

ASPIRE PROTECTS

- **Aspire Managed XDR (MXDR)**
  - Aspire Managed XDR (MXDR) combines next-generation eXtended Detection and Response (XDR) technology, a 24x7 security operations center (SOC), and the expertise of a US-based team of security professionals to identify and respond to threats across a broader attack surface.
  - Building on the existing capabilities of Managed Detection and Response (MDR), Aspire Managed XDR integrates and correlates telemetry from endpoints, network, cloud, email, identity, and more. This advanced platform creates valuable context enabling end-to-end visibility across all

threat vectors. Experienced security analysts and incident responders deliver 24x7 threat detection, analysis, and investigations – with both automated and human-led response actions to quickly mitigate cyberattacks.

- **Aspire Managed Detection and Response (MDR)**
  - Accelerate the maturity of your security operations and reduce risk with faster threat detection and response capabilities. Aspire MDR delivers around-the-clock protection across cloud, network, and endpoints in one integrated solution.
  - Our team of experts act as an extension of your IT operations to quickly identify and mitigate security threats before they impact your business.
- **Aspire Incident Response**
  - The ability to respond effectively to data breaches is critical to every organization. When attacks occur, IT leaders need a process that will ultimately maintain business continuity, protect the company reputation, and avoid fines, legal fees, and remediation costs.
  - Aspire offers incident response services to help you prepare for, manage, and recover from data breaches and network attacks. An experienced team uses industry-leading threat intelligence and the most current security technology to quickly respond to attacks, reduce damage, and minimize exposure.

INDICATORS OF COMPROMISE (IOCS)

Aspire SOC IoCs

IP Addresses

- 34[.]182[.]49[.]94

Domains

- n8ulc8hbb[.]cc[.]rs6[.]net
- cyan-reason-834476[.]framer[.]app
- accounts-g0033le-com-ca8f[.]1mustabd[.]workers[.]dev
- pq8gf4cab[.]cc[.]rs6[.]net
- cc[.]rs6[.]net

- framer[.]app
- workers[.]dev

SUPPORTING DOCUMENTATION

What Is Email Spoofing and How to Prevent it? - Spambrella

Phishing actors exploit complex routing and misconfigurations to spoof domains | Microsoft Security Blog

Internal phishing risk tied to complex email routing, Microsoft says

Phishing attacks exploit misconfigured emails to target Microsoft 365 - Infosecurity Magazine

Complex Routing, Misconfigurations Exploited for Domain Spoofing in Phishing Attacks - SecurityWeek

Microsoft Warns Misconfigured Email Routing Can Enable Internal Domain Phishing

EchoSpoofing is Back—and It's Even Easier for Attackers to Reach Inboxes | SECURITY.COM

PREVENTING EMAIL SPOOFING/PHISHING - Clearwater, St Petersburg, Tampa Fl.

Configure mail flow using connectors in Exchange Online | Microsoft Learn

Email Spoofing Fixes & Mail Exchange Tips | Dubex

Phishing attacks exploit misconfigured emails to target Microsoft 365 - Infosecurity Magazine

Microsoft Warns Misconfigured Email Routing Can Enable Internal Domain Phishing

Anti-phishing policies in Microsoft 365 - Microsoft Defender for Office 365 | Microsoft Learn

Microsoft 365 Has an Impersonation Problem

Sophisticated Microsoft Spoof Targets Financial Departments

Detecting Phishing and Account Compromise in Office 365 - Securonix

Attackers Exploit M365 for Internal Phishing | Proofpoint US

Forward thinking: How adversaries abuse Office 365 email rules

What Is Email Spoofing? Definition & Examples | Proofpoint US

What is email spoofing? | How it works & prevention | Cloudflare

What is Email Spoofing? | Barracuda Networks

APPENDIX II: DISCLAIMER

This document and its contents are not intended to serve as legal advice and should not be used as a substitute for it. The results of a Security Risk Assessment are intended to guide the implementation of diligent measures to reduce the risk of potential vulnerabilities being exploited to compromise data.

While the Services and this report may offer information that the Client can use to support its compliance efforts, the responsibility for evaluating and fulfilling compliance obligations rests solely with the Client, not Aspire. This report does not guarantee or assure the Client's compliance with any law, regulation, or standard.