



# Using Relationship Graphs to Mitigate Email-based Threats

## Contents

Introduction .....	2
More on the problem.....	2
More on the solution .....	3
Building and maintaining graphs is challenging .....	4
In closing.....	5
References .....	5



## Introduction

Despite significant advances in cybersecurity, email remains a prime target for threat actors. Organizations worldwide remain vulnerable to increasingly sophisticated email-based attacks such as Business Email Compromise (BEC) that bypass traditional security measures. The persistent exploitation of email as an attack vector is not coincidental—it remains one of the most widely used communication channels in business environments, making it an attractive gateway for malicious actors.

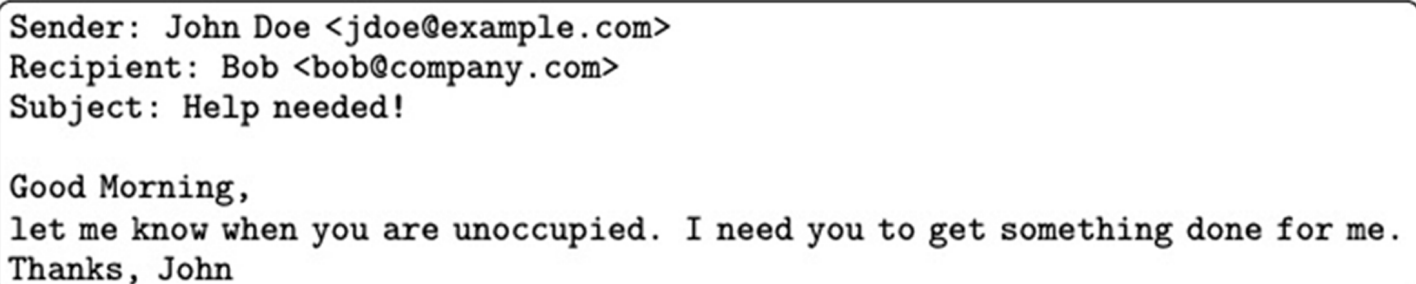
Email security systems employ multiple techniques to detect threats, with many relying heavily on content analysis. These systems scan for maliciousness links, malicious attachments, and language patterns typically found in phishing attempts. However, sophisticated attackers have adapted by crafting emails whose content is virtually indistinguishable from legitimate business communications.

How, then, can security systems detect malicious emails when their content appears legitimate? The answer lies in leveraging non-content contextual information derived from relationship graphs. These graphs map connections between senders and receivers, providing crucial context that content analysis alone cannot detect.

By analyzing the patterns and history of communications between various entities, email security systems can identify anomalies that signal potential threats—even when the email content itself appears innocent. This approach represents an important advancement in email security, allowing for the detection of sophisticated attacks that would otherwise bypass traditional content-based filters.

## More on the problem

Consider the following example (see Figure 1) of a malicious email that cannot be conclusively declared malicious or benign based solely on its content:



Sender: John Doe <jdoe@example.com>  
Recipient: Bob <bob@company.com>  
Subject: Help needed!

Good Morning,  
let me know when you are unoccupied. I need you to get something done for me.  
Thanks, John

Figure 1. An example of a Business Email Compromise email that can't be identified based on content alone.

This type of email presents a significant challenge for content-based detection systems for several reasons:

First, the content is entirely consistent with legitimate business emails. The language, formatting, and overall presentation match what one would expect from a professional business contact focused on an important task.

Second, if security systems were configured to flag all emails of this type, the result would be an overwhelming number of false positives. Security teams would be inundated with alerts about legitimate business communications, quickly leading to alert fatigue that may cause staff to ignore genuine threats.

This scenario illustrates the fundamental limitation of content-based detection: when (potentially) malicious emails are crafted to match legitimate communication, traditional security methods fall short. Organizations need additional context to differentiate between genuine business communications and sophisticated impersonation attempts.

## More on the solution

Email security systems can overcome the limitations of content-based analysis by leveraging information contained in email headers. One particularly effective technique involves the use of relationship graphs, which store comprehensive data about communication patterns between individual senders, receivers, and their respective domains.

A relationship graph (see Figure 2) maintains detailed records of communications, tracking the frequency and nature of interactions between various entities. This approach enables security systems to answer critical contextual questions, such as:

- Does a particular user in the organization receive frequent emails from a specific external user?
- Does an organization receive regular communications (across all users) from another organization?

- Has any user in the organization previously received an email from a particular external domain?




Figure 2. Senders and receivers (recipients) in a relationship graph.

Beyond individual relationships, these graphs also store global information about senders and receivers. This allows security systems to evaluate questions like: How many times has a specific domain been observed sending email messages across all monitored organizations?

Additionally, relationship graphs track information about senders whose messages have previously been identified as malicious. This historical data enables the system to quickly identify when a new email arrives from a known malicious source, even if the content of that email appears legitimate.

Implementing relationship graphs provides email security systems with additional signals that enhance those derived from content analysis. By combining these multiple data points, the overall detection performance of the system improves substantially, allowing it to identify sophisticated threats that would otherwise go undetected.

In real-world applications, such as Cisco® Secure [Email Threat Defense](#) (ETD), relationship graphs have proven highly effective. In many cases, multiple maliciousness indicators are identified, including sender information derived from the relationship graph.

Figure 3 shows a screenshot of an email from a service provider processed by ETD. The subject of this email is innocuous. The body of the email requests contact information—an occurrence that by itself is insufficient to raise an alarm. However, the email is flagged because the sender has previously been associated with malicious activity, and this sender rarely communicates with the receiver—determinations made possible through relationship graph analysis.

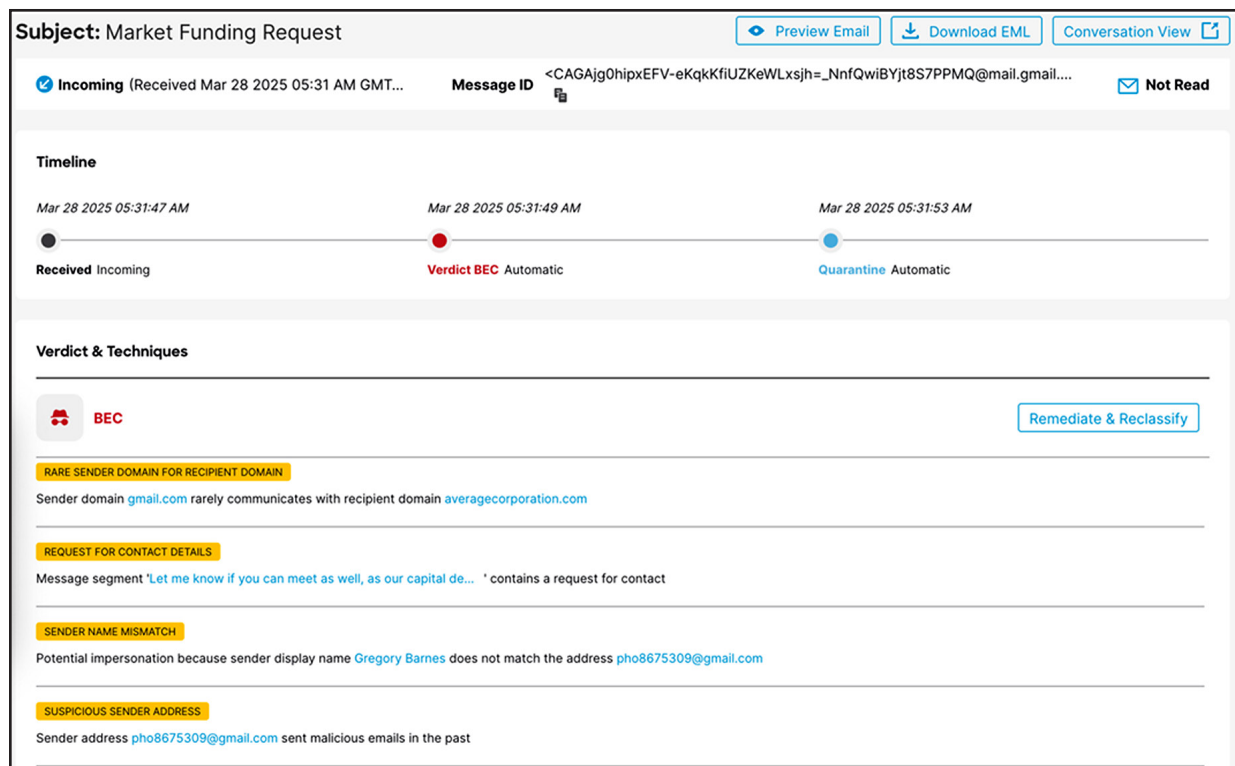


Figure 3. Cisco ETD processing a BEC email

## Building and maintaining graphs is challenging

While relationship graphs provide significant security benefits, implementing and maintaining them presents several technical challenges:

**Query and update latency** is a critical consideration. For effective threat detection, query operations must be extremely fast—ideally under 10 milliseconds—to ensure timely detection and maintain high email processing throughput. Updates to the relationship graph can be batched to reduce computational demands but must be frequent enough to keep the data current and relevant.

**Scale** presents another major challenge, particularly for large organizations or those with high email volumes. Processing and analyzing massive datasets require substantial computational power, memory, and storage. As organizations grow and communication volumes increase, relationship graphs must scale accordingly without compromising performance.

**Privacy** concerns must also be addressed when analyzing relationships between entities. Since email communication often contains sensitive information, access to relationship graph data must be carefully controlled to protect user privacy and comply with data protection regulations.

Cisco's technical team has successfully overcome these challenges to make relationship graphs an integral part of ETD:

- Specialized data structures have been developed to enable lightning-fast queries while supporting efficient batched updates. These custom structures are optimized for the specific patterns and requirements of email relationship analysis.
- The graph implementation is designed to use computing resources efficiently while meeting stringent performance criteria. This optimization ensures that relationship graphs remain cost-effective even at enterprise scale.
- To minimize communication latency, relationship graphs are maintained in multiple geographic locations. This distributed approach ensures that query responses are delivered quickly regardless of where the query originates.
- Strict security measures protect the data contained in relationship graphs. Access is limited to system administrators, and all graph data stored on disk is encrypted to prevent unauthorized access.

Further, the ETD technical team continues to improve the relationship graph by finding new computing resource efficiencies and algorithmic refinements.

Finally, the investment in maintaining relationship graphs delivers value beyond direct threat detection. Information derived from these graphs also improves the functioning of other components within ETD. In certain cases, relationship graph data is provided to large language models, helping the system better distinguish between malicious and legitimate emails through enhanced contextual understanding.

## In closing

With two decades of experience in email security, Cisco continues to lead the industry in developing innovative techniques to combat evolving threats. The large-scale implementation of relationship graphs in ETD exemplifies the company's commitment to advanced security solutions that bolster traditional content analysis.

By mapping the complex web of communication relationships between senders and receivers, ETD can detect sophisticated threats that would otherwise bypass conventional security measures. This approach has proven particularly effective against business email compromise and other advanced impersonation attacks.

The relationship graph technology represents just one component of Cisco's comprehensive approach to email security. For organizations seeking to protect themselves against today's most sophisticated email threats, solutions that incorporate relationship analysis provide a critical layer of defense.

For more detailed information on ETD's technical architecture, interested readers are encouraged to consult the references below, which provide in-depth explanations of the system's design and capabilities.

To experience these capabilities first-hand, start a [free trial](https://www.cisco.com/c/en/us/products/security/email-threat-defense-free-trial.html) by going to <https://www.cisco.com/c/en/us/products/security/email-threat-defense-free-trial.html>.

## References

1. "[How large language models enhance Cisco Secure Email Threat Defense](#)," May 2024, Cisco Systems Inc.
2. Brabec et al., "[A Modular and Adaptive System for Business Email Compromise Detection](#)," August 2023, arxiv.org.