BEC Active Defense Overview

What is active defense?

Active defense is the process of interacting with an attack or an attacker for the purpose of collecting more robust intelligence about the threat. Specifically, active defense helps better understand what happens after an attack is successful. Examples of active defense tactics include engaging with BEC actors or seeding credential phishing sites with controlled personas.

How does the Threat Intelligence team identify attacks that feed our BEC active defense engagements?

A query is used to identify attacks targeting Abnormal customers that are labelled as BEC attacks or contain characteristics that are indicative of BEC attacks. The query used to identify these messages can be found here.

The results of this query are deduplicated and added to a review queue. Before an active defense engagement is initiated an analyst reviews each message to determine whether it is, in fact, a malicious, response-based BEC attack that is suitable for an engagement.

What outcomes of BEC active defense engagements are/can be used by other Abnormal teams?

There are numerous byproducts of active defense engagements that can be used by other teams at Abnormal for various purposes. Below are some examples of unique outputs of the active defense process that can be used by other teams.

- **Message Classification Corrections:** Correcting attack type/goal classifications of messages that have already been reviewed by the product and/or CSI based on the outcome of active defense engagements.
  - **Consumer(s):** Data Science, EPD
  - **How It Can Be Used:** Provide customers with a more accurate picture of threats Abnormal has stopped and helps retrain existing models to more accurately classify messages in the future.
- **Malicious BEC/ATO IP Addresses**: IP addresses known to be associated with malicious actors collected during the course of active defense engagements.
  - **Consumer(s):** EPD
  - **How It Can Be Used:** Supplement detection with known-bad IP signals.
- **Post-Attack Attacker Communication & Behavior:** Content of attacker communications received after an initial message collected during an active defense engagement.
  - **Consumer(s):** EPD
  - **How It Can Be Used:** Improve FNs for second-stage BEC attacks. Identify “red alert” messages that indicate a potential successful attack. Support POVs by identifying successful BEC attacks based on second-stage email content. Provide customers with additional attack context to better understand the potential impact of a blocked attack.
- **Financial Impact:** Identification of specific attack-level financial impact based on the results of active defense engagements.
  - **Consumer(s):** Data Science, EPD
  - **How It Can Be Used:** Provide customers with more specific insights into potential financial impact of email-based attacks and tailored ROI of Abnormal products. Update global business value assessments based on internal, real-world data instead of external, third-party data.
- **Threat Group Identification:** Because we have more complete information about BEC attacks, we’re able to use that information to more effectively attribute attacks to specific threat groups based on cluster analysis.
  - **Consumer(s):** EPD
  - **How It Can Be Used:** Provide customers with additional context about the actors associated with BEC we detect.

Data from active defense engagements can be found in a DataBricks dashboard here.

What other outputs are derived from BEC active defense engagements?

In addition to providing a unique data source to aid detection and provide additional attack context to customers, active defense engagements also produce other outputs that are used for non-product purposes.

One of these byproducts is confirmed BEC mule accounts that actors are requesting BEC funds be sent to. While these accounts may have limited usefulness for detection purposes, they are instrumental in building strategic relationships with third parties, (individual financial institutions, ISACs, law enforcement) based on mule account intelligence sharing.

Additionally, because active defense engagements are a core source of our intelligence for BEC attacks, it feeds most of the external content (blogs, intelligence briefs, whitepapers, webinars, etc.) the Threat Intel team develops for marketing purposes.

What information is used to create an BEC active defense engagement?

The only artifacts from an original BEC attack that are used to create an active defense engagement are the attacker’s email address and email subject (unless the subject contains personalized information). All other information used in an engagement is fictitious and is not associated with the original attack. The display name (e.g., the impersonated executive) for the attacker is randomly generated. Body content from the original attack is not used when initiating an engagement.

The Threat Intel team has created a pool of 200+ fictitious personas that are used to communicate with BEC actors. When initiating an engagement, the first “response” from a persona is created to look like the persona was the original target of the BEC campaign. If an email subject is personalized (i.e., referencing the sender/receiver name or company name), the subject for the engagement is modified to reference the persona information.

Is any customer information used in an active defense engagement?

No. No customer information is used when creating an active defense engagement. As discussed above, we aren’t impersonating the original recipient of a BEC attack, engaging with a BEC actor on behalf of an organization, or referencing the original message in an active defense engagement. Rather, an engagement is constructed to look like one of our personas was targeted by the actor as part of the BEC campaign.

Are active defense engagements a manual process?

While some parts of the active defense process are manual, a majority of the process will be automated once the entire tool capability has been built out. The review of pending messages prior to the initiation of an engagement is manual, but features have been implemented to increase the efficiency of the review process. Once a pending message has been approved for an engagement, the initial message is constructed automatically and, depending on the type of attack, follow up responses are also automated. For some attacks, such as gift card BEC attacks, nearly the entire engagement is automated from beginning to end.