Global Information Assurance Certification Paper

Assisted Security Investigations Using Cognitive Computing

GIAC (GSEC) Gold Certification

Author: Lori Cole, LoriSaysRAWR@gmail.com
Advisor: Mohammed Haron
Accepted: November 20, 2019

Abstract

The purpose of this research is to illustrate the application of cognitive computing and machine learning concepts through the building and training of a chatbot that simulates human conversation for cybersecurity investigation scenarios. The SOC chatbot will offer best-practice advisory dialogue to security analysts as they proceed through security incident investigations, thus simulating technical mentorship. As a security analyst progresses through various investigations, they will become more practiced in the recommended and appropriate workflows, gain investigative tool proficiency, and become more confident in handling standalone investigations. The SOC chatbot will serve as a training tool for less experienced analysts and afford more time to upper-tier analysts to respond to escalated security incidents, as they will no longer need to walk through incidents alongside junior analysts. Security analysts serving in a tier 1 SOC role are ideal end-users of the SOC chatbot. As the first line of defense, their primary function is to address SIEM events. They are familiar with basic security concepts, incident ticketing systems, and hold the appropriate level of access for data gathering and external research.

1. Introduction

Enterprises strive to keep up with the current cyber threat landscape and adequately defend their infrastructure. Most are reliant upon manual processes while struggling with a lack of resources, skills, and budgets. Organizations are pragmatically building and maturing security operations centers (SOCs), focused on threat detection and response, to better protect themselves and keep up with the overwhelming amount of cyber threats (Gartner, 2019). A SOC provides centralized cybersecurity event monitoring, detection, and response capabilities. One of the main hurdles that SOCs face is the sheer volume of security events; analysts are being overwhelmed by the number of alerts and the number of investigations that require their attention. The annual Cisco Cybersecurity Report states that gaps continue to exist between alerts generated and those that are investigated. An average of 44 percent of daily SIEM alerts is not investigated. There is a lack of trained personnel who can meet the demand to investigate all alerts (Cisco, 2018). According to a Fidelis Cybersecurity research study, most SOC analysts can handle 7 or 8 investigations in a day (Fidelis, 2019). In addition to an alert fielding capacity issue, SOCs are facing a training issue, as many organizations struggle to recruit, train, and retain qualified SOC analysts (MITRE, 2014). Automation is becoming increasingly important for SOCs, including automating investigative assistance and cybersecurity analyst training. (Cisco, 2018).

One way to strengthen and force multiply a security investigation capability is by incorporating cognitive computing. Cognitive computing is composed of computer science (the processing of information) and cognitive science (the understanding of human brain functionality) and can lend self-teaching algorithms, visual and speech recognition, natural language processing and generation, and informed reasoning logically to optimize a human-dependent process (IBM, 2017). Security investigations guided by cognitive computing can extend and supplement the security investigation capabilities of an SOC and serve as a training resource. In this research effort, cognitive computing resources were used to assess the feasibility of assisted security investigations in the form of an interactive chatbot.

Author Name, email@address

The IBM Watson Assistant service was utilized for the SOC chatbot dialogue build, trial, and implementation. End users will interact with the chatbot via its public link (hosted on IBM Cloud Service):

https://assistant-chat-us-south.watsonplatform.net/web/public/468015c2-c0eb-48a7-9487-20a0be429d60

The SOC chatbot will assist in the fundamentals of a security incident investigative process (SANs, 2017):

- Prioritize, assign, categorize the incident
- Identify and extract IOCs (e.g., IPs, hashes)
- Perform initial analysis (reputational lookups, threat level scoring)
- Get running processes from the machine
- Get network connections from the machine
- Contextualize with threat intelligence
- Confirm or refute the threat
- Escalate for IR or begin the remediation process

Expected benefits include:

- Reduction in the meantime to resolve security incidents
- Increased autonomy from tier 1 security analysts
- Standardized workflows for incident handling

The following Draw.io dialogue illustrates the logical flow of the SOC chatbot investigative assistance.

Author Name, email@address

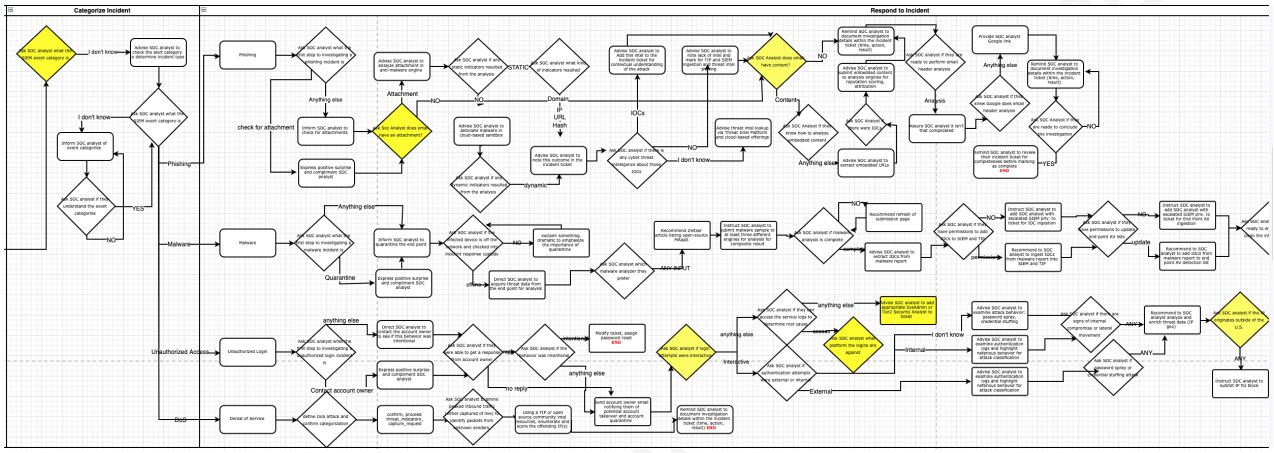

Figure 1. SOC chatbot draw.io dialogue, logical mapping: SIEM alert acknowledgment, event categorization, and event-specific response. Full Draw.io diagram and sourcing provided in Annex A.

Four common dialogue scenarios concerning investigative workflows specific to security incident types are represented within the SOC chatbot in addition to initial data gathering conversation (e.g., end user name, event categorization):

- Malware
- Phishing
- Denial of Service
- Unauthorized access

Cybersecurity incidents involve data points unique to the event type (e.g., network security logs for a denial of service attack, authentication logs for a privilege escalation attempt), and thus each incident type is handled via a specified investigative process. The SOC chatbot investigative dialogs lead security analysts through steps to gathering the information needed to satisfy basic questions about a security incident such as:

- Where did the attack originate from?
- What was the attack vector?
- When did the attack occur?
- Do we need to escalate this incident?

Author Name, email@address

The SOC chatbot also reminds security analysts to note important data findings in their incident ticket and bring in additional internal assistance as required.

2. Research Methodology

The SOC chatbot was created per the above-defined requirements using Watson Assistant, on IBM Cloud PaaS. Development steps included: identifying a real-world cybersecurity task that could be improved using artificial intelligence and machine learning, logical mapping of scenario dialogues, creating necessary data points for dialogue training (e.g., intents, entities), creating dialogue decision models within Watson Assistant, training Watson Assistant, and finally deploying the chatbot for user interaction.

The SOC chatbot dialogue tree model was partitioned into three conversation segments: user orientation and SIEM alert acknowledgment, event categorization, and event response workflow (Figure 2).

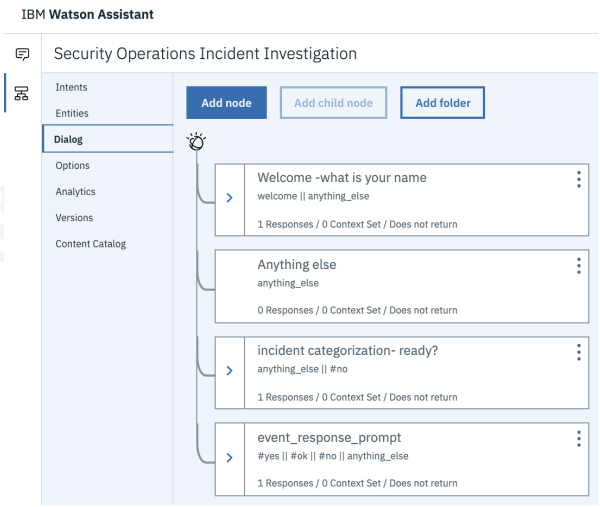

Figure 2. SOC chatbot Watson Dialogue mapping: SIEM alert acknowledgment, event categorization, and event-specific response workflow.

The chatbot dialogue uses intents, entities, and context variables to define the interaction occurring with the end-user, represented in a hierarchy of nodes, and executed from the

top node. Dialogue intents are the specific purpose of input during a dialogue. 30 unique intents representing the various dialogue segments (alert acknowledgment, event classification, and response) were created with over 500 user examples and their likely synonymization. Entities are used to gather additional information in the conversation. 21 entity values representing security incident characteristics were created ranging from high-level classifications (event IDs, or event type) to specific symptoms of targeted cyber-attacks (privilege escalation failure, backdoor detected), and respective synonyms. The application of context variables for end-user addressing offered a personalized feel to the chatbot conversation, while disambiguation, autocorrection, and irrelevance detection nodes were configured for error handling when responses were unrecognizable and a question needed to be rephrased or repeated (e.g., confirming incident event type by asking for event ID number, asking for incident characteristics to determine event type, if unrecognized, then rephrase). The overall dialogue flow of the SOC chatbot mimics security investigation basics: categorize alerts, conduct, and document the investigation, and respond prescriptively.

The "Try it out" panel within IBM Watson Assistant-enabled real-time testing of dialogue scenarios and their node configurations. The below figures display dialogue panels for the scenarios discussed above, showing chatbot responses and end-user input.

Author Name, email@address

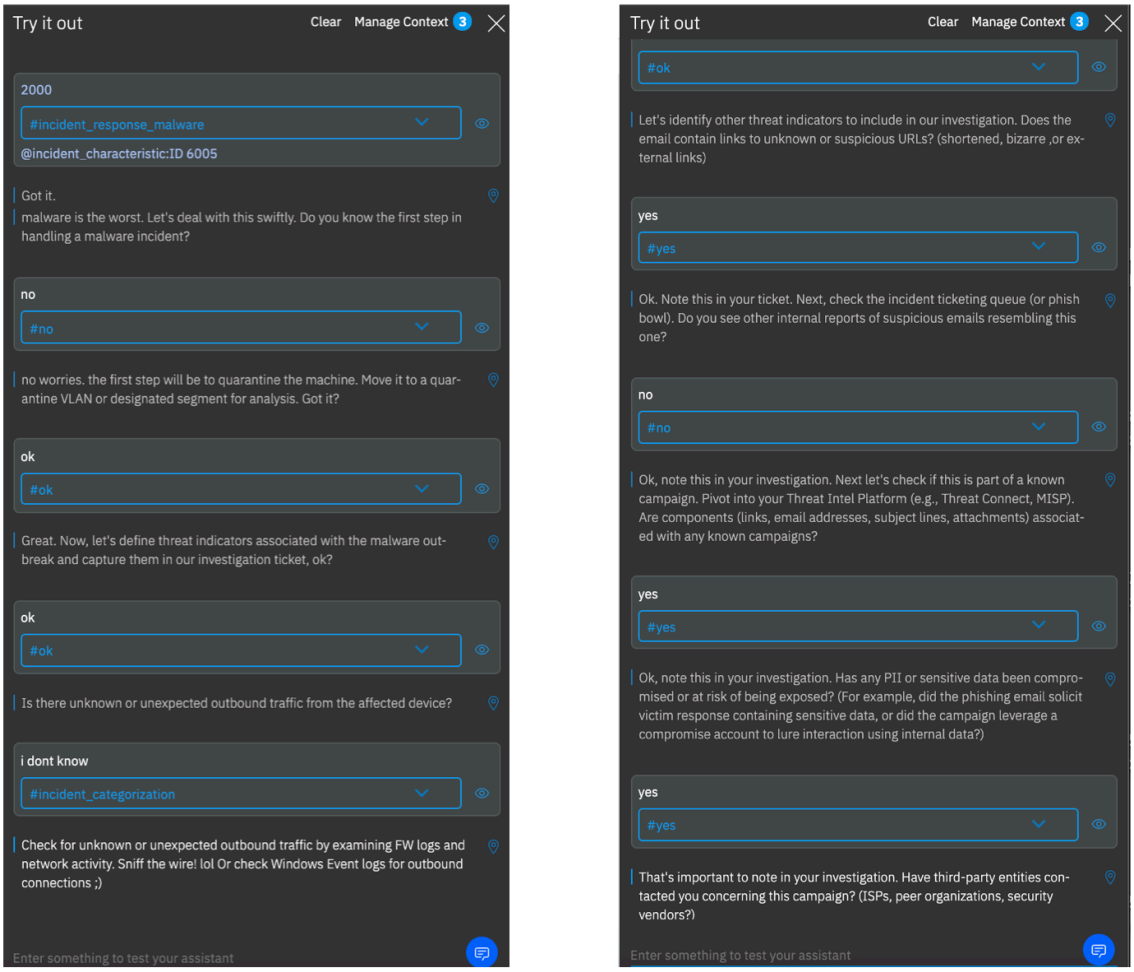

Figure 3. Watson Assistant "Try it out" panel displaying SOC chatbot handling of user input: uncertainty variable and decisive variable.

Figure 3 shows the SOC chatbot prompting the user for input concerning unknown or unexpected outbound traffic from the affected device when the user is uncertain; a categorization intent is leveraged to rephrase the question and guide the user in locating the solicited information. The SOC chatbot confirms the user to input specific to the event type, incorporating provided variables into the affirmation.

Despite thoughtful dialogue tree design and conditional responsive nodes, the SOC chatbot sometimes encountered unexpected user input and reacted with scenario reset. As seen in Figure 4 (below), the SOC chatbot classified some gibberish text as

Irrelevant, and some gibberish was handled with a scenario reset, identifying specific node improvement and indicating that further training is required.

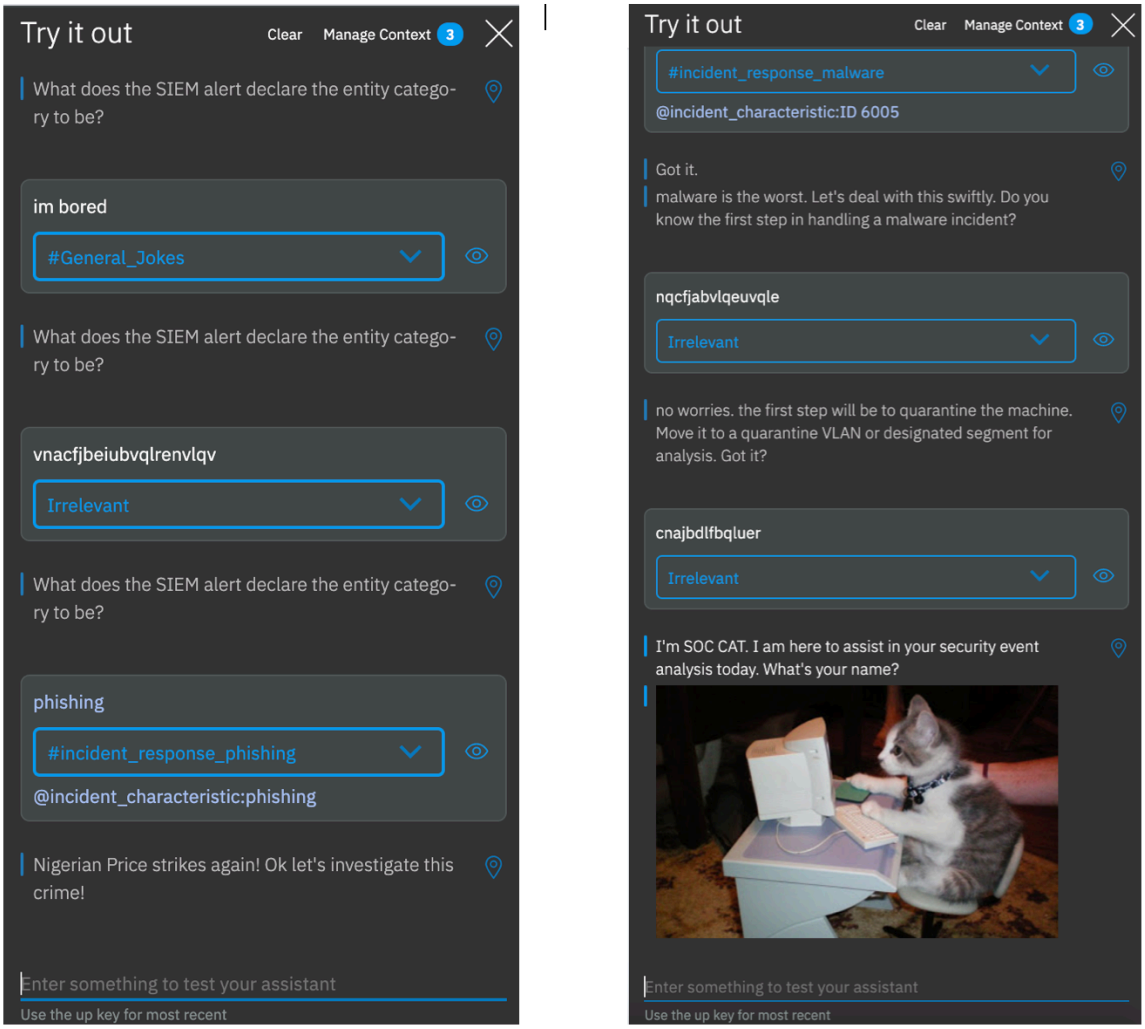

Figure 4. Watson Assistant "Try it out" panel displaying SOC chatbot handling of user input: unexpected variables.

2.1.1. Conversational Analytics Findings

The SOC chatbot Assistant was soft-deployed onto IBM Watson Cloud and trialed for two days via a preview link shared among several tier 1 security analysts. On day one, each participant was asked to participate in 3-4 conversations with the SOC chatbot.

Author Name, email@address

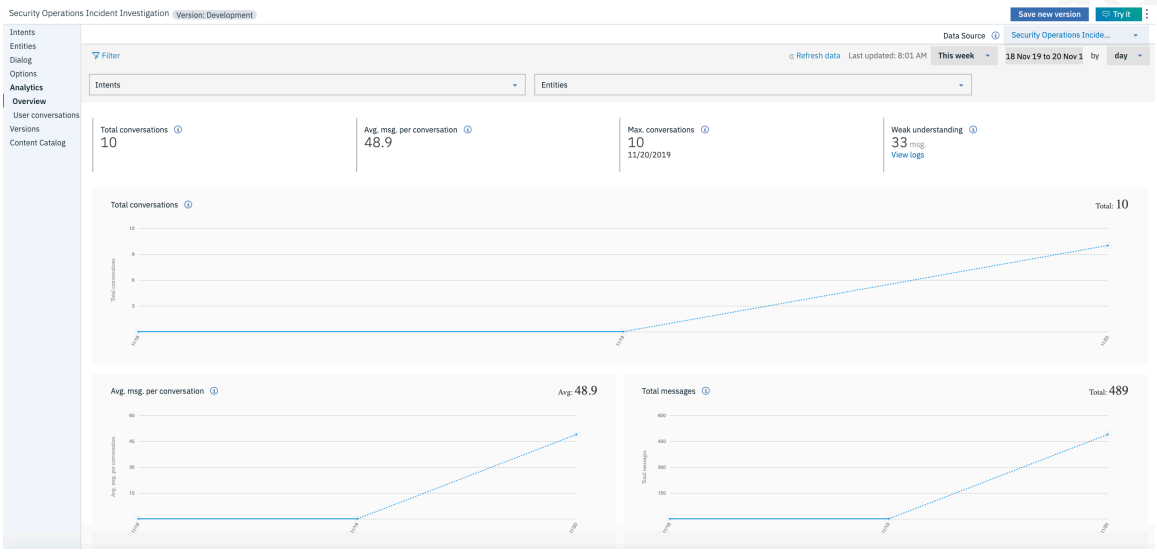

Figure 5. Watson Assistant Dialogue Analytics Overview panel and usage visualizations, Day 1

The Analytics page of Watson Assistant offers statistics representative of external traffic (from users) that has interacted with the Assistant; they do not include interactions from the "Try it out" panel. Conversation history can offer insight on how to improve Assistant understanding and respond better to user requests. Basic usage metrics can be gleaned easily from the Overview panel, including total conversations, average messages per conversation, maximum messages per conversation, and an estimated level of cognitive understanding. Interactions between users and the SOC chatbot were observed within the Conversation tab, enabling identification of messages with unrecognized intents and misclassified or unknown entities, as shown in Figure 6.

Author Name, email@address

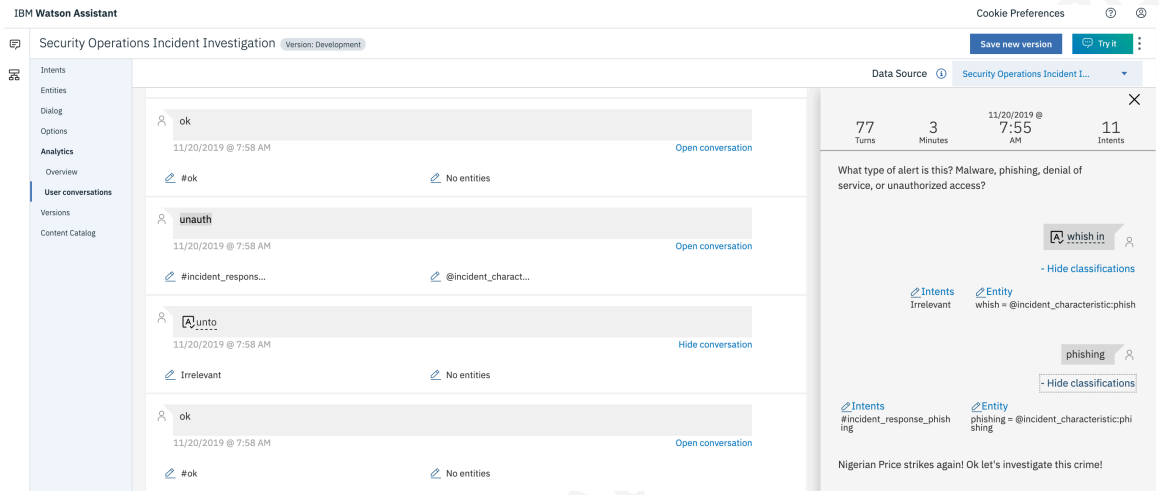

Figure 6. Watson Assistant Dialogue Analytics, User Conversations: categorized intents from the unknown entity (phishing).

Language variable classification needs were identified in the initial trial, and adjustments included additional synonyms for context variables and uncategorized response handling. This initial conversation assessment also identified the need to build out entity diversity beyond incident categorization and response, to include explanation nodes for terminology used (e.g., credential stuffing, password spray).

On day two, each user was asked to participate in the same 3-4 conversations again. Total interactions between users and the SOC chatbot were observed within the Conversation tab. Despite node refinement inspired by initial trial data and additional user interaction, the cognitive understanding was “weak” for more than 30 messages, suggesting that intent classifier nodes need continued improvement to increase conversational proficiency. According to IBM, individual messages with weak understanding are not classified by intent and do not contain known entities. These weak messages should steer dialog improvement efforts to remediate potential dialogue problems (2019). Most frequently occurring intents and entities were examined to prioritize classification adjustments (see Figure 8).

Author Name, email@address

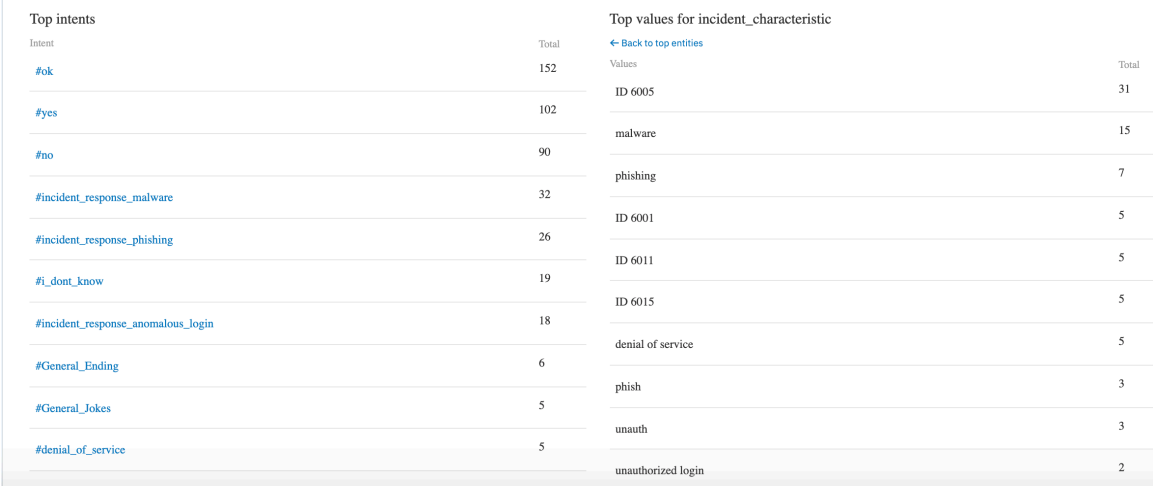

Top intents

Intent	Total
#ok	152
#yes	102
#no	90
#incident_response_malware	32
#incident_response_phishing	26
#i_dont_know	19
#incident_response_anomalous_login	18
#General_Ending	6
#General_Jokes	5
#denial_of_service	5

Top values for incident_characteristic

← Back to top entities

Values	Total
ID 6005	31
malware	15
phishing	7
ID 6001	5
ID 6011	5
ID 6015	5
denial of service	5
phish	3
unauth	3
unauthorized login	2

Figure 8. Watson Assistant Dialogue Analytics, User Conversations: Top intents, Top entities.

2.1.2. Chatbot Enhancement

Interactions between the chatbot and users reveal several dialogue improvement opportunities. Findings from the conversation logs indicate that some users are quite knowledgeable and specific when describing cybersecurity events. This vocabulary should be considered during intent creation and synonymization. Additionally, disambiguation, digressions, and autocorrection identified in multiple dialogue nodes should be pragmatically addressed. As a feature enhancement, the SOC chatbot could ask the user to copy and paste data results (from malware scanners, or cyber threat intelligence sources) to store as variables to offer an "investigative summary" at the end of a conversation. This would timestamp the investigation, aggregate investigative data, and prompt the user if data fields are required before closing or escalating the incident.

Future SOC chatbot skill integrations include Watson Services and Watson Assistant Search to enhance bot understanding of user input and to improve the relevancy of provided responses. Dialogue enhancements would leverage Natural Language Classifiers (NLC) and Natural Language Understanding (NLU). The NLC Watson Service may shorten the chat bot's training phase (less training data is needed for use in ensemble machine learning) and would enable multilingual functionality (Watson, 2019). Watson NLU could offer industry vertical specific entities and relations (Watson, 2019),

which would address current performance limitations of the SOC chatbot as cybersecurity entities are being hand-curated. Additionally, the ability to perform NLU upon unstructured data shows valuable potential for enriching cyber threat data with contextual information to assess actual risk (e.g., attributing malware usage to a threat actor, or associating a tactic with a critical vulnerability). Assistant Search skills could also be incorporated into the SOC chatbot to provide relevant information from an external data source, extracting specified data and returning it to the security analyst to include in the incident ticket (e.g., IP reputation score, the maliciousness of a hash, or general threat intelligence enrichment).

The SOC chatbot could be deployed via Slack, given its multitude of compliance certifications proving its data encryption and protection posture (Slack, 2019), and further trained by SOC analysts to curate additional intents and cast relevant entities. Conversation log data collected from the beta deployment would be further curated (synonymization and disambiguation performed) to ensure a robust dialogue aptitude for successful live implementation. Additionally, SOC chatbot integration with other security investigation platforms (e.g., ticketing system, SIEM) may result in valuable trial data.

3. Conclusion

The SOC chatbot demonstrated adequate navigation through four basic security response workflows (malware, phishing, denial of service, and unauthorized access) in accordance with threat enumeration logic derived from SANs Network Security Essentials (2017), SANs Incident Handling (2018), and the Incident Response Consortium (2019). According to a trial user survey (Annex B), dialogue efficacy was gauged to be [adequate/useful yet needs improvement/dysfunctional, not useful] in soliciting information needed to satisfy basic questions about a security incident, but requires significant error handling, digression, and disambiguation improvements. The SOC chatbot implementation was not without challenges. Deployment presented an initial challenge; when attempting public link generation, IBM warned of cost accrual for hosting a stand-alone integration in the Cloud. Required retraining time for dialogue node modifications resulted in slow versioning times, and given the element of human

interaction, significantly more Bot Control dialogue nodes should be implemented to navigate within a conversation effectively (e.g., clarification, confirmation, response satisfaction).

Building and implementing a chatbot for cybersecurity investigations was a first step in leveraging cognitive computing; the following serve as potential recommendations for an optimized SOC chatbot:

- Script execution capability (execute scripted commands to fetch external data, process and format results, and populate an incident service ticket)
- Non-human language processing ability (cybersecurity data can contain encoded or unstructured data, so the ability to discern and handle these data types would offer pre-analysis support)
- API call integration (ability to make calls to and retrieve data from API services to enrich cyber-related data points)

The SOC chatbot is not intended to replace human security analysts, only to assist and guide them. There are significant automation efforts occurring in the cybersecurity industry (Security Orchestration and Response (SOAR)) and analyst playbooks (e.g., Phantom, Demisto), which aim to strengthen security analyst impact and increase response capacity. The SOC chatbot, and other AI and machine learning initiatives for cyber defense, will remain tied to human responsibility and serve as assistive tools to meet the cybersecurity threat demand.

Author Name, email@address

References

Cisco. (2018). Cisco 2018 Annual Cybersecurity Report.
Retrieved from https://www.cisco.com/c/dam/m/hu_hu/campaigns/security-hub/pdf/acr-2018.pdf

Fidelis Security. (2019). The State of the SOC. Retrieved from
https://www.fidelissecurity.com/resource/report/the-state-of-the-soc/

Gartner. (2019). Gartner Top 7 Security and Risk Trends for 2019.
Retrieved from https://www.gartner.com/smarterwithgartner/gartner-top-7-security-and-risk-trends-for-2019/

IBM Cloud. (2019). IBM Cloud Docs: Watson Assistant. Metrics Overview.
Retrieved from https://cloud.ibm.com/docs/services/assistant?topic=assistant-logs-overview

IBM. (2019). IBM Products: Natural Language Classifier.
Retrieved from https://www.ibm.com/watson/services/natural-language-classifier/

IBM. (2019). IBM Products: Natural Language Understanding.
Retrieved from https://www.ibm.com/watson/services/natural-language-understanding/

IR Consortium. (2019). Playbook Policy Engine.
Retrieved from https://www.incidentresponse.com/playbooks/

MITRE. (2014). Ten Strategies of a World-class Security Operations Center.
Retrieved from https://www.mitre.org/sites/default/files/publications/pr-13-1028-mitre-10-strategies-cyber-ops-center.pdf

SANs. (2017). Automating Enterprise Security Response Webinar.
Retrieved from https://www.sans.org/webcasts/live-demonstration-automating-enterprise-security-response-105185

Slack. (2019). Security at Slack.
Retrieved from https://slack.com/security

Author Name, email@address

Annex A

Draw.io SOC Chatbot dialogue conceptual mapping:

https://drive.google.com/file/d/1j_AkYbDD2TywJ2L8UbjhpSL0bka8KuV-/view?usp=sharing

Annex B

Trial user survey:

1. Did the SOC chatbot seem to replicate a valid security investigation workflow?
2. Was the SOC bot dialogue adequate, useful yet needs improvement, or dysfunctional/not useful?
3. Did the SOC chatbot handle unexpected inputs gracefully?
4. Would you utilize cognitive computing to assist in security investigations?

Annex C

SOC Chatbot:

https://assistant-chat-us-south.watsonplatform.net/web/public/468015c2-c0eb-48a7-9487-20a0be429d60

Author Name, email@address