Cloudbreak: Answering the Challenges of Cyber Command and Control

Diane Staheli, Vincent F. Mancuso, Matthew J. Leahy, and Martine M. Kalke

Lincoln Laboratory's flexible, user-centered framework for the development of command-and-control systems allows the rapid prototyping of new system capabilities. This methodology, Cloudbreak, effectively supports the insertion of new capabilities into existing systems and fosters user acceptance of new tools.

» **As the number and size of networks** maintained by the Department of Defense (DoD) continue to grow, concerns about the complexity of providing cyber security for these networks have mounted. In 2012, then Secretary of Defense Leon Panetta established the Joint Cyber Centers (JCC) at U.S. geographic combatant commands (COCOMs) to coordinate cyber activities within each command's area of responsibility (AoR) and to apprise combatant commanders of the impacts of the cyber landscape to their missions [1]. The JCCs were instituted to resolve the lack of coordinated cyber security within and across all the COCOMs.

The JCCs charted their own paths for defining the structure of their organizations, determining their work processes, and procuring the tools and capabilities necessary to accomplish their missions. To help address the COCOMs' capability needs and improve upon their model for technology delivery, leadership at the JCCs turned to Lincoln Laboratory's Cloudbreak[1] initiative, which had been sponsored by the Assistant Secretary of Defense for Research and Engineering. During its four-year tenure, the Cloudbreak program successfully filled critical gaps in COCOMs' cyber situational awareness by utilizing an iterative user-centered design process to rapidly deploy cyber capabilities to the warfighter.

The Cloudbreak process is designed to address near-term capability gaps once for all COCOMs rather than once for each COCOM. The overall goal of Cloudbreak

[1] The name Cloudbreak was selected as the next in a series of program names inspired by weather terms; it does not imply a connection to cloud computing.

is to rapidly deliver technologies to confront emerging and unanticipated threats. By allowing operators to drive technology development rather than giving them predefined solutions, the Cloudbreak approach aims to provide agile, interoperable, and reusable applications. This article describes Cloudbreak's genesis and its successful technology development and insertion process. Case studies demonstrate how the Cloudbreak process was applied to the implementation of two cyber security tools: the Cyber Analytical Station and Cyber Dashboard.

Cyber Challenges for Combatant Commands

The COCOMs are responsible for maintaining command and control of U.S. forces in their AoR during military operations, in times of conflict and peace, and during crisis interventions, such as humanitarian relief or disaster response activities. Two critical ingredients to any successful military operation are timely, reliable situational awareness and efficient, secure communication of that information to all participants in the operation. The cyber challenges to the realization of those ingredients fall into two main categories: mitigating difficulties caused by the inability of multiple users to share information over disparate computing systems and addressing problems caused by either a lack or overabundance of data relayed to COCOMs during operations.

As an illustration of these challenges, consider the difficulties faced by the U.S. disaster relief operation launched in response to the 11 March 2011 Great East Japan Earthquake, which led to a tsunami with waves higher than 40 m that traveled up to 10 km inland and that caused a major nuclear meltdown at the Fukushima Daiichi Nuclear Power Plant [2]. Operation Tomodachi, under the control of the U.S. Pacific Command (USPACOM), spanned nearly two months and involved multiple organizations responsible for the 24,000 U.S. service members, 189 aircraft, and 24 naval ships deployed in the mission [3]. During Operation Tomodachi, USPACOM found that existing military network resources were inadequate to keep pace with evolving situations and activities. Because the existing software tools and computing procedures were stove-piped (designed for specific organizations' needs) and not interoperable, they did not enable USPACOM to efficiently gain sufficient situational awareness of the mission and the environment, and did not support on-the-fly acquisition or development of software tools better suited to the tasks at hand. Situational awareness also suffered because the information sent to command varied in quantity ("drought or deluge") and tools varied in their ability to process data.

The solution to the problem of stove-piped, incompatible tools is not simply providing access to more tools, and the ready availability of data is not necessarily an advantage. With the advancement of sensor systems for gathering data and the expansion of computing resources for processing, storing, and distributing data, operators have more access to more information than ever before. With this deluge of information comes the risk of information overload. The vast amounts of diverse information (e.g., text, video, imagery) that are disseminated daily throughout DoD commands and organizations strain the ability of analysts to develop a comprehensive picture of evolving situations. When the current tool set does not support the goals of the command or the individual operators, these drawbacks may become greater than the benefits of the expanded toolsets and datasets.

Challenge of the COCOM Acquisition Process

Currently, COCOM acquisitions are conducted through Integrated Priority Lists (IPL). These lists represent an individual COCOM's most important capability needs prioritized across military service and function lines, risk areas, and long-term strategic planning issues [4]. These IPLs are then used to inform the programming and budgeting processes about COCOM needs. Each IPL represents the needs of an individual COCOM (e.g., USPACOM, U.S. Southern Command [USSOUTHCOM]) and is developed to satisfy the particular requirements and procedures of each COCOM's branches. This compartmentalization can lead to a lack of awareness of the overall capability needs across COCOMs, tools that are not generalizable across COCOMs, and redundant functionalities. Current tools are often stove-piped for individual threats and organizations, and updates are infrequent and difficult. Optimally, COCOM development and acquisition should provide agile, user-centered decision support tools that are (1) composable capabilities that can be built and modified on the fly and (2) interoperable, reusable applications that are generalizable across commands and threats.

Cloudbreak and User-Centered Design

The Cloudbreak process has origins in both user-centered design and agile software development. Human-centered design, defined in the International Organization for Standardization's standard ISO 9241-210 [5], is an approach to interactive system development. Typically, this process uses the characteristics of relevant stakeholders (e.g., users) and their environment to define a set of requirements for design solutions; the tools developed to meet those requirements undergo user evaluations that then inform subsequent iterations of the tools. User-centered design requires significant upfront research and analysis of user needs, resulting in a longer time to deliver a working product. Agile methods, on the other hand, focus on rapidly delivering small sets of features onsite to customers, iteratively updating using a feedback loop between the developers and the users.

Traditionally, user-centered design has been seen as incompatible with the agile development process [6]. However, if the two are aligned, user-centered and agile methods can be used to maintain a close connection to users while rapidly iterating on system design and requirements [7, 8]. This hybrid strategy is flexible and holistic, taking into account the entirety of the problem space and allowing for incremental development that can make system modifications based on evolving circumstances.

While many developers in industry and academia have been reluctant to combine the two approaches to system design, researchers at Lincoln Laboratory have championed taking an agile, user-centered approach to aid in building effective, practical tools and visualizations that satisfy the requirements of their users [9]. In their review of user-centered design in cyber visualizations, Staheli et al. found that in the majority of visualization developments described in the published research, users were not even consulted during the design process [10]. Additionally, in the efforts discussed in that research, post-design evaluation of the visualizations was mainly limited to high-level qualitative analyses, such as surveys. During the development of the Extreme Malicious Behavior Viewer, Yu et al. interviewed users to understand how they interact with cyber data [11]. While the geographic locations of malicious cyber events may not seem to yield adequate information for cyber defense (most attacks will likely be clustered in populous locations), the team found that geolocation was a simple, intuitive option for conveying relevant information to users with limited cyber knowledge. The team's interviews revealed that a map displaying network activity in relation to geopolitical entities was helpful to users' decision making in identifying threats that target specific regions, employ language or culture-specific social engineering, or exploit localization or pirated software. When developing Macroscope, a network-based intrusion-detection system, Cunningham et al. based their design of the system display, RapIDisplay, on interviews with intrusion-detection analysts [12]. These interviews led to the incorporation of display features that are not common in many intrusion-detection systems: a presentation that allows rapid access to documentation and report generation, and a visualization of the confidence of an attack happening.

The Cloudbreak program applies methods used in previous successful system implementations to the development of rapid, composable designs and software. Cloudbreak's efficient, flexible iterative process is better suited to quickly and effectively providing tools to meet the complex and emergent needs of COCOMS than are the current models for technology delivery (Figure 1).

Cloudbreak Process

The Cloudbreak process leverages current capabilities to quickly deploy to operators newly composed software solutions for responding to emerging threats. Cloudbreak's approach is a cycle of problem definition, identification of relevant existing capabilities and solutions, and deliveries of new tools in spirals, each of which is informed by ongoing observations of the tools' productivity (Figure 2). The outcome of the mission for which the new capabilities were created and the lessons learned from their implementation are used to determine if the new capabilities can be applied, perhaps with modifications, at other COCOMs. This focus on post-deployment assessments enables a successful reuse of newly designed software across multiple COCOMs and mission areas. For example, a capability deployed for unclassified information sharing to support nation building in one COCOM's AoR could be repurposed to coordinate a response during a humanitarian assistance crisis or a disaster recovery operation.

The Cloudbreak process allows for activities to be completed concurrently and in various orders to address changing needs. The first step, defining the problem and software gaps, starts with assessing each COCOM and

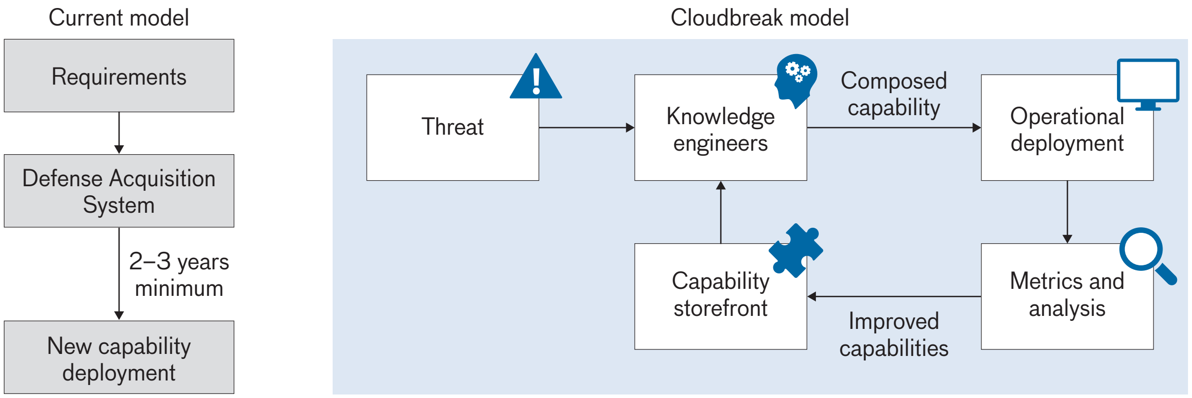

FIGURE 1. This comparison of the Cloudbreak model to the current model for technology delivery shows the iterative nature of Cloudbreak: knowledge engineers assess the nature of a threat, compose from available technologies a solution to mitigate the threat, provide the solution capability to operational users, analyze the capability's effectiveness, make improvements to the capability on the basis of the analysis, and add the new capability to the storefront for future use. The current Department of Defense acquisition cycle, which can take a minimum of two to three years, is a complete design and build of a new system for supplying the requisite functions.

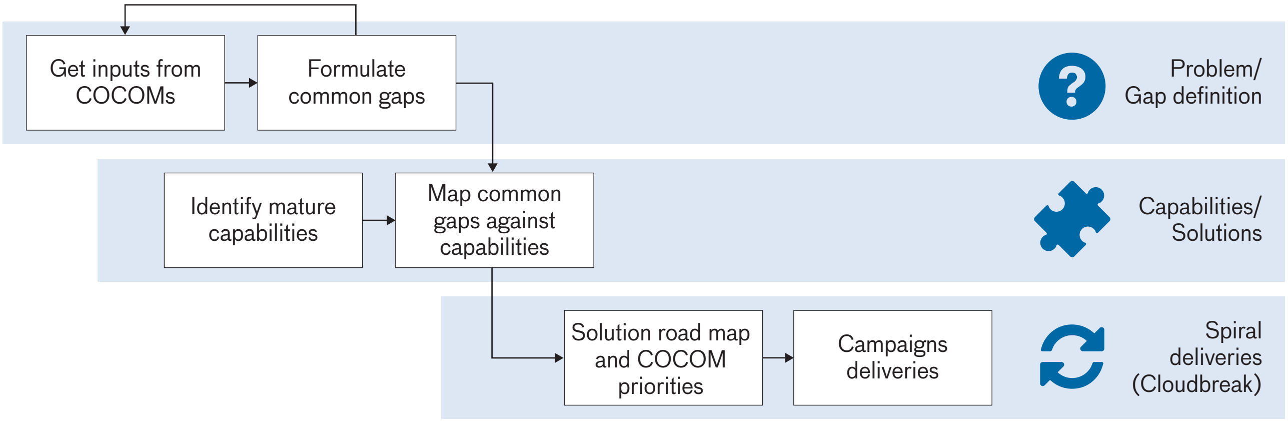
Get inputs from COCOMs
Formulate common gaps
Problem/ Gap definition
Identify mature capabilities
Map common gaps against capabilities
Capabilities/ Solutions
Solution road map and COCOM priorities
Campaigns deliveries
Spiral deliveries (Cloudbreak)

FIGURE 2. The three-tiered Cloudbreak process begins with defining the needs of combatant commands (COCOMs). Existing software and mature capabilities that may provide solutions to those needs are evaluated against the gaps in the command's current software tools. Developers propose a plan for creating new tools and deploy usable solutions that may require multiple development spirals to fully answer all requirements.

using the results to formulate a set of common gaps, which are then presented back to each COCOM to ensure that the descriptions of the gaps adequately summarize the COCOM's specific problems.

Next, Cloudbreak practitioners identify a set of mature capabilities that map against the common gaps. These capabilities are essentially building blocks that can be acquired from a de facto "storefront" and can be composed into a new unified capability. Available capabilities include systems from other domains; previously matured technologies, including government and commercial off-the-shelf (GOTS and COTS, respectively) systems; and additional data sources.

Once technology capabilities are mapped against needs, COCOMS are provided with a solution road map that links needs to available services and capabilities in the storefront. This plan provides information on how many and which COCOMs have a specific capability gap and what the implementation of the plan will cost and involve. Knowledge engineers, i.e., multidisciplinary engineers armed with an understanding of the needs and operational environment of the COCOMs and experienced with diverse technology solutions available through the storefront, use common architectures to synthesize and tailor a capability that supports the operational need.

During the composition of a new capability, it is critical that the knowledge engineers involve the operators in an iterative development strategy to promote strong acceptance of the new tool. Formal and informal assessments of the new software can generate feedback that allows the engineers to remain responsive to operator needs. When the knowledge engineers insert the new technology into the operational environment, they can concurrently evaluate the technology's utility. Finally, the engineers transfer the newly composed capability, with improvements developed on the basis of lessons learned, back into the storefront to be used to address other gaps and emerging incidents across multiple commands.

Elements of Successful Technology Insertion

The Cloudbreak process is a general framework for inserting new technology into previously composed systems. Through our work with the COCOMs, we have identified several elements critical to a successful technology insertion.

- Collaborating across organizations. Exchanging new, useful services and applications is a critical aspect of the Cloudbreak model. However, many command centers that have similar needs and missions do not currently take advantage of each other's capabilities. Cloudbreak provides a platform on which COCOMs can build an awareness of the capabilities and current gaps of other commands. Once COCOMs are cognizant of each other's systems and technology gaps, knowledge engineers can work across the commands to ensure that work is not repeated or further resources are not spent on solutions already available. COCOMs can leverage the previous work from other commands and work together to develop unified capabilities.
- Leveraging existing capabilities. COCOMs do not have access to unlimited resources; therefore, it is important that they maximize their resources. Currently, COCOMs are required to invest in GOTS and COTS systems, but these systems may not fully meet their requirements. Cloudbreak offers a process and platform for COCOMs to pull previously developed, known capabilities from one command to another; thus, individual COCOMs can insert high-quality solutions into their systems while expending significantly less time and resources.
- Customizable tools and composable architectures. Cloudbreak allows tools to be adapted to individual problems and commands. From the storefront, knowledge engineers can obtain capabilities to reach a 90% solution and then tailor those capabilities to attain a COCOM's goals. Each COCOM can utilize its own data feeds and develop other, low-level customizations to achieve a 100% solution. These customizations can then be integrated back into the storefront for use by other commands that may have similar requirements.

Operational Deployment Case Studies

As a part of the Cloudbreak program, researchers from Lincoln Laboratory visited USPACOM, USSOUTHCOM, and the Defense Information Systems Agency to interact with users and understand the current technology needs and deficiencies across their organizations. During these visits, which have occurred regularly since 2012 and typically last at least a week, researchers interview numerous analysts to better understand their command-level technology gaps and analyst-level needs.

Once the problems were identified, the Cloudbreak team focused on cataloging capabilities and potential solutions available from Lincoln Laboratory and COTS and GOTS providers. As a part of the initial assessment of the COCOMs' cyber programs, an exhaustive list of available capabilities was compiled and organized according to the mission area utilizing the capabilities, COCOM deploying the capabilities, and utility accruing from the capabilities. Tools identified as important for the cyber mission were cataloged on the basis of their applicability in improving situational awareness and the analytical process. In addition to identifying the tools and their primary usage, all tools were cataloged according to their current usage across COCOMs, estimated costs, designations as enterprise-level software, composability, availability in the Cloudbreak storefront, and maturity.

By identifying available technologies and aligning them with the current needs, the Cloudbreak methodology can support the combination of the latest technologies with existing tools, datasets, and capabilities. The key to effectively compiling capabilities is to understand the operational relevance of a technology.

The following case studies from the Cloudbreak program illustrate the practical implementation of the process.

Case Study: Cyber Analytical Station

Operators at JCCs had employed commercial cyber defense tools to log network activity, monitor the network for anomalies, and generate alerts upon detection of potentially malicious activity. However, the information assurance policies that dictated monitoring on a wide range of activities resulted in millions of alerts each day. This abundance of data far exceeded the capacity of the JCC teams to effectively monitor network traffic. The JCCs needed a way to prioritize the incoming alerts so that operators could direct resources to processing the most relevant data.

The Cloudbreak initiative identified seven requirements for an improved system for the JCCs:

1. Enable users to analyze cyber event data via a prioritized dashboard of critical and/or alarming events
2. Support the analysis of tens of millions of cyber events daily
3. Quickly identify the most important vulnerabilities
4. Facilitate timely creation of remediation plans to prevent escalation of cyber threat activity
5. Provide an interface with event data organized for easy exploration
6. Be easily and quickly learned
7. Enable forensic analysis of cyber event trends

These requirements were mapped against the current capabilities for operator situational awareness so knowledge engineers could identify technologies to integrate into a new tool—the Cyber Analytical Station. Typically, only mature capabilities are considered for integration, but in the case of the Cyber Analytical Station, no suitable capabilities existed, necessitating a custom development effort.

From the gap analysis, three predominant requirements for the system were apparent: the technology would have to (1) perform automated triage, (2) enrich data and apply context, and (3) support the investigation and analysis process. Each of these requirements contains several subrequisites as delineated in Figure 3. These requirements were then mapped to the available Lincoln Laboratory and COTS and GOTS technologies.

The Cyber Analytical Station aims to provide operations centers at military commands with the cyber situational awareness necessary for monitoring and managing the performance, security, and integrity of computer networks. To compose the final operational prototype, knowledge engineers integrated several mature research-grade capabilities:

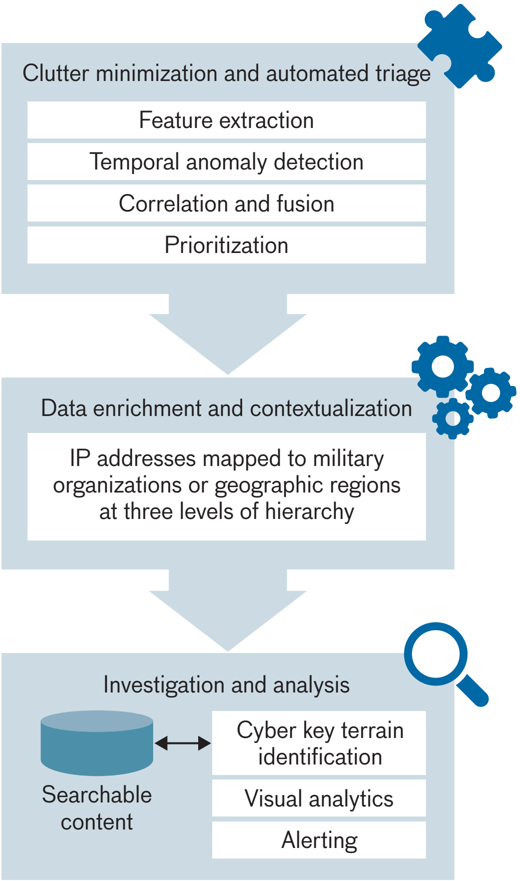
Clutter minimization and automated triage
Feature extraction
Temporal anomaly detection
Correlation and fusion
Prioritization
Data enrichment and contextualization
IP addresses mapped to military organizations or geographic regions at three levels of hierarchy
Investigation and analysis
Searchable content
Cyber key terrain identification
Visual analytics
Alerting

FIGURE 3. The above data analytical features required by a system to achieve satisfactory situational awareness were identified by the Cloudbreak researchers. During clutter minimization and triage, the system should be able to extract key features, perform anomaly detection, correlate and fuse data from multiple sources, and prioritize alerts. Next, the system should map Internet protocol (IP) addresses to known entities and/or geographic regions to contextualize the incoming data. Finally, the data should be stored in a way that allows the cyber analyst to interact with the data either in a raw form or through a visual analytic.

- Ingest and enrichment. The final composed application was able to ingest and enrich necessary cyber data with geolocation and organizational data from known sources, execute automated analyses, and deliver the outcomes through visualizations on a user-friendly web interface. The ingest capability is responsible for loading, enriching, and storing cyber data, which are principally acquired from network intrusion-detection systems. Because the Cyber Analytical Station can perform an automated enrichment while loading and storing data, it can help cyber operators by providing the context they need to more quickly categorize and interpret the data.

- Event detection and prioritization. This automated data analysis capability provides users with rapid detection and prioritization of anomalies in the data. The station proactively identifies significant events rather than just offering alerts that are based on predetermined, rigid heuristics, also known as "trip-wire" conditions. This capability was easily addressed via a prototype application previously developed at Lincoln Laboratory for anomaly detection [13]. Finally, outcomes of the ingestion, enrichment, and analysis are presented to users through an interactive interface that enables analysts to review prioritized anomalies and drill down to the underlying data (Figure 4).

Even though several of the capabilities within the Cyber Analytical Station were already mature, through the Cloudbreak technology insertion process, we were able to deliver a solution that was customized for the needs of the COCOM. Early on, while Lincoln Laboratory staff were working to solicit specific requirements for transitioning the Cyber Analytical Station to COCOM operational use, JCC operators had a great many firm requirements and new questions. These questions and requirements were then spiraled back into the prototype application in an iterative strategy, supporting the development of a system that met the evolving needs of the JCCs and attained strong user acceptance.

During the initial stages of Cloudbreak, primary efforts centered on developing the detailed measures of a system's effectiveness and performance that would be used in conducting formal assessments. For the Cyber Analytical Station project, formal evaluations helped provide overall impressions on the station, but the most useful feedback came from simply observing operators and recording their behaviors and subjective commentaries as they used the system. For example, during one informal observation period, we watched an operator manually perform a copy-and-paste operation. To capture the details of an anomaly for inclusion in a cyber-incident report, the operator used the mouse to highlight 10s of table rows and paste them into a spreadsheet. The development team captured this behavior as a new requirement, executed a feature "quick turn" to implement a new functionality for converting tables to spreadsheets, tested this software update, and delivered it to multiple COCOMs within three days.[2]

RESULTS OF THE CYBER ANALYTICAL STATION TECHNOLOGY INSERTION

As demonstrated in the Cyber Analytical Station case, the Cloudbreak process can be used to mitigate risk within the resource-constrained JCCs. A lean, efficient interdisciplinary team was able to execute at a low cost, and the providers of existing capabilities benefit from interacting with a broader and more diverse user base. Involving cyber operators early and reacting to their requirements with an iterative strategy were critical to gaining strong user acceptance of the new capabilities. Once the initial new technology was delivered, assessments of it, both formal and informal, were invaluable in generating useful feedback. For example, the need for an automated copy-and-paste function may have not been identified if we had not informally observed how operators actually interact with the software. Finally, by remaining responsive to the requests of users and to the continuing evolution of the JCCs, the Cloudbreak team was able to deliver quick-turn features and to modify the tool set in a matter of days.

During its initial deployment, the Cyber Analytical Station quickly demonstrated utility during a brute-force attack (i.e., an exhaustive trial-and-error method to breach password or cryptographic protections) against a public-facing file-transfer server. The Cyber Analytical Station provided enhanced cyber situational awareness by enriching the data with context and enhancing data exploration. The Cyber Analytical Station provided interactive access to data not previously available and allowed operators to focus on high-priority tasks, thus improving operators' efficiency and accuracy.

Case Study: Cyber Dashboard

The Cyber Dashboard was conceived as a means to integrate and visualize disparate data sources (e.g., cyber, operational, and intelligence data) to support information exchange and commanders' cyber situational awareness.

Achieving operational cyber situational awareness requires the integration of data from six classes of information: the current and near-term threat environment, anomalous network activity, vulnerabilities, key cyber terrain, current operational readiness, and ongoing operations [14]. In the COCOMs, the providers of this information were stove-piped within each organization, requiring the Cloudbreak team to work across joint services to locate and obtain access to the authoritative sources. In many

[2] The update did not include security-relevant changes and did not trigger an information assurance reaccreditation.

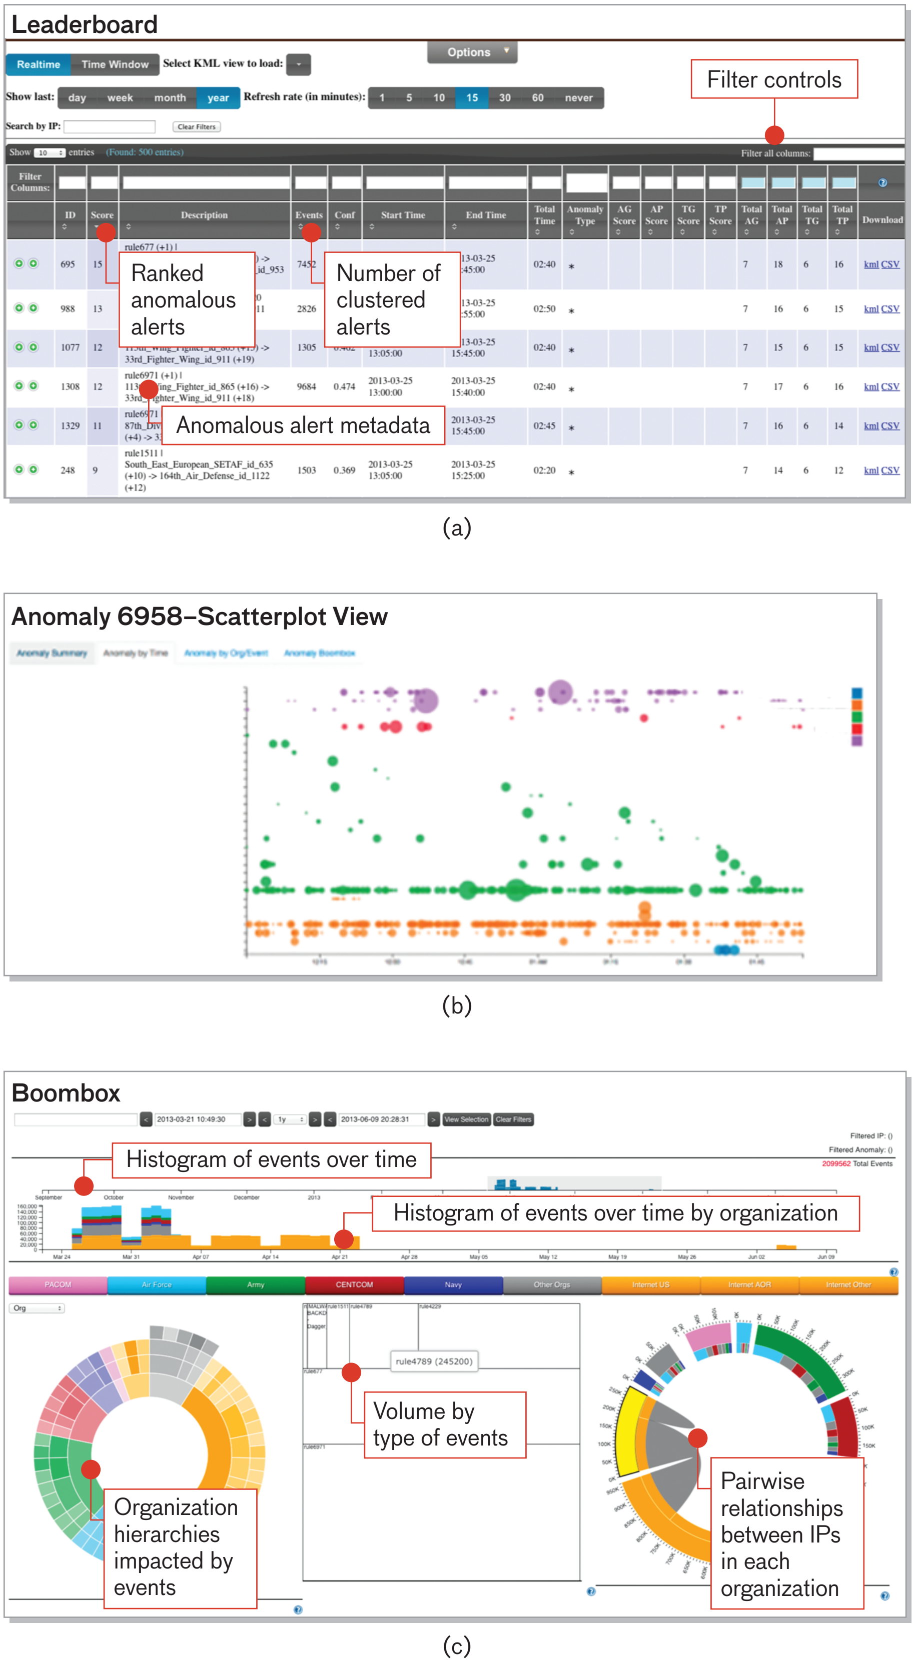

Leaderboard

Realtime | Time Window | Select KML view to load: | Options

Show last: day week month year | Refresh rate (in minutes): 1 5 10 15 30 60 never

Search by IP: | Clear Filters

Show 10 entries (Found: 500 entries) | Filter all columns:

ID	Score	Description	Events	Conf	Start Time	End Time	Total Time	Anomaly Type	AG Score	AP Score	TG Score	TP Score	Total AG	Total AP	Total TG	Total TP	Download
695	15	rule677 (+1) |) -> id_953	7452			13-03-25 :45:00	02:40	*					7	18	6	16	kml CSV
988	13	0 11	2826			13-03-25 :55:00	02:50	*					7	16	6	15	kml CSV
1077	12	) -> 33rd_Fighter_Wing_id_911 (+19)	1305		13:05:00	13-03-25 15:45:00	02:40	*					7	15	6	15	kml CSV
1308	12	rule6971 (+1) | 113 ing_Fighter_id_865 (+16) -> 33rd ghter_Wing_id_911 (+18)	9684	0.474	2013-03-25 13:00:00	2013-03-25 15:40:00	02:40	*					7	17	6	16	kml CSV
1329	11	rule6971 87th_Div (+4) -> 3				2013-03-25 15:45:00	02:45	*					7	16	6	14	kml CSV
248	9	rule1511 | South_East_European_SETAF_id_635 (+10) -> 164th_Air_Defense_id_1122 (+12)	1503	0.369	2013-03-25 13:05:00	2013-03-25 15:25:00	02:20	*					7	14	6	12	kml CSV

(a)

Anomaly 6958–Scatterplot View

(b)

Boombox

(c)

FIGURE 4. This interactive visualization tool consists of an anomaly leaderboard (a); a screen for anomaly inspection (b); and a visualization of all alerts on a display called the Boombox (c). The leaderboard display allows the watch-floor manager to focus on high-priority alerts and to group and sort those alerts on the basis of various criteria. The scatterplot view shows the appearance, prevalence, and disappearance of alerts on the network. The Boombox display is designed to allow analysts to explore the network data: it color-codes attacks in terms of an organization's hierarchy; its top third shows histograms of overall attacks experienced by organizations and of attacks encountered by each organization over time; its middle region shows the volume of attacks presented as a treemap in which the rectangles for different attack scenarios are proportional to the total number of attacks; and the right circular graphic displays the pairwise communications occurring between systems in each organization. Collectively, the three displays (a–c) help analysts to prioritize incoming events, to understand why events are considered anomalous, and to discover relevant information about events and put it into an organizational context.

cases, because data providers were not under the jurisdiction of the COCOMs, the Cloudbreak team had to make connections to external data-providing entities, such as the Joint Improvised-Threat Defeat Agency.

Analysts rely on a variety of data sources to maintain accurate cyber situational awareness. The initial focus for the dashboard was on the display of anomalous activity; subsequent development spirals incorporated additional data sources into the display. Analysts at the COCOMs also require the ability to overlay additional arbitrary data sources on the map to visually fuse information from various sources. This capability allows analysts to connect information across multiple sensors and visually inspect and analyze the relations and interactions between the data. For example, a network outage located in the same geographical area that is experiencing an unusually high number of alerts could be a cause for concern; a display coordinating those two pieces of information allows operators to identify a situation that may need further investigation. Additional data sources could number in the 100s, depending on the situation.

The Cloudbreak team identified the following additional capability needs. Analysts desired functionality for preserving the current state of the dashboard and sharing it with others. Two reasons drove their request for a shareable dashboard: (1) analysts wanted to share a link for a particular finding in the data with other analysts or managers, and (2) analysts wanted to create custom dashboards to address emerging situations. The needs and the respective dashboard approaches to meeting those are illustrated in Table 1.

The Cyber Dashboard was built as a series of Microsoft's SharePoint Web Parts to best leverage existing capability. Four types of SharePoint libraries compose the capability: the Map libraries render the main map canvas and geospatial data, the Tree/Graph libraries render hierarchical data, Timeline libraries render temporal data, and Data libraries transform, correlate, and archive the original data sources. Each of the Web Parts requires configuration of a data source from a common data format: Keyhole Markup Language (KML), Extensible Markup Language (XML), Comma Separated Values (CSV), JavaScript Object Notation (JSON), or SharePoint lists are all supported. Data are not stored or managed by the Cyber Dashboard, but remain in their original location and under their existing access control policies. Dashboards can be customized by modifying the configuration files and can be shared via unique Uniform Resource Locators (URL). The architecture is illustrated in Figure 5.

The design of the dashboard visualization utilizes a canvas-palette metaphor; a geographic map serves as the background of the browser window and as a canvas upon which geospatial data are depicted. Multiple

Table 1. Dashboard Capability Needs of COCOMs and Cloudbreak Solutions

NEEDED CAPABILITY	SOLUTION
Commanders need a flexible, customizable dashboard that presents a common operational picture of cyber situational awareness	Provide an agile display that has a "brief from tool" capability
Operators need a system that enables them to react promptly to evolving situations (e.g., Ebola outbreaks or disaster response efforts)	Supply a dashboard whose easy configuration and customization enable a new dashboard to be created and shared within minutes
System must be capable of integrating data sources from within and outside COCOMs	Eliminate the use of back-end databases; allow data to be accessed from original authoritative sources
Display must be easy to interpret for operators who may have limited experience with and knowledge of the onscreen visualization	Create a display that uses a geospatial background (a familiar reference point for users) and that supports multiple data formats (potentially 700 data sources)
Capabilities must mitigate problems of COCOMs' existing limited infrastructure, hardened systems, long acquisition process	Utilize in-house SharePoint infrastructure and expertise as much as possible

floating palettes are then layered on top of this canvas for non-geospatial data. Analysts acknowledged that a geospatial map may not be the most suitable display paradigm for all cyber data, but the map provides a good start as a common visual representation that is familiar and accessible to all audiences. The flat design style of the map also helps remove visual clutter, such as the representation of terrain features, and allows the data points to be viewed more clearly.

A permanent, large palette on the left contains the master list of data sources; other palettes can be drawn on demand and positioned as needed. Palettes can display data in a number of conventional visualizations (e.g., tree map, node-link diagram, sunburst chart, timeline). Each visualization palette has basic parameters that can be configured: data sources, transformations of the data sources (correlation, georeferencing), and graphical elements, such as sizes or colors. The configuration for the entire dashboard can be saved and shared. The final design is illustrated in Figure 6.

RESULTS OF THE CYBER DASHBOARD DEVELOPMENT

The dashboard, using operational datasets, was demonstrated for analysts in three cyber operations centers to solicit feedback on the display and to gauge its operational utility. Analysts and managers provided qualitative feedback via comments, both as a group during the demonstration and in private conversations after the presentation. Developers then worked with analysts individually to identify new requirements, deliver software updates, incorporate new data sources, and gather further feedback. As a result of the interactions with users, the team delivered 92 iterations of the software within the 2015 calendar year. These deliveries included two significant product features that were incorporated on the basis of user feedback: (1) visualization network and circuit dia-

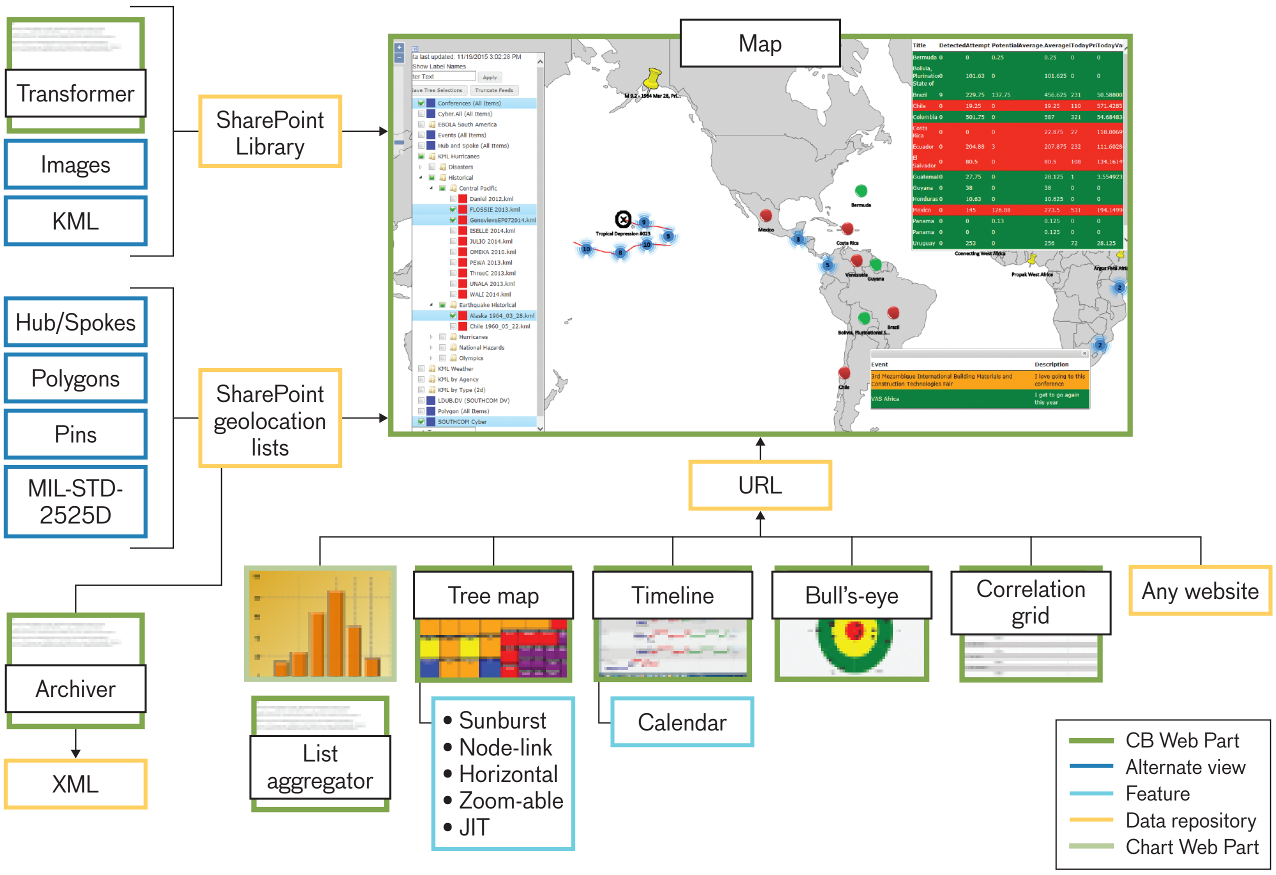

FIGURE 5. The Cyber Dashboard architecture is based on SharePoint Web Parts (green), which render authoritative data types, i.e., data that have been verified as coming from an official trusted source, (yellow) on a geospatial canvas. Many of the Web Parts offer multiple options (e.g., node-link or horizontal charts, calendar) for rendering visualizations (blue).

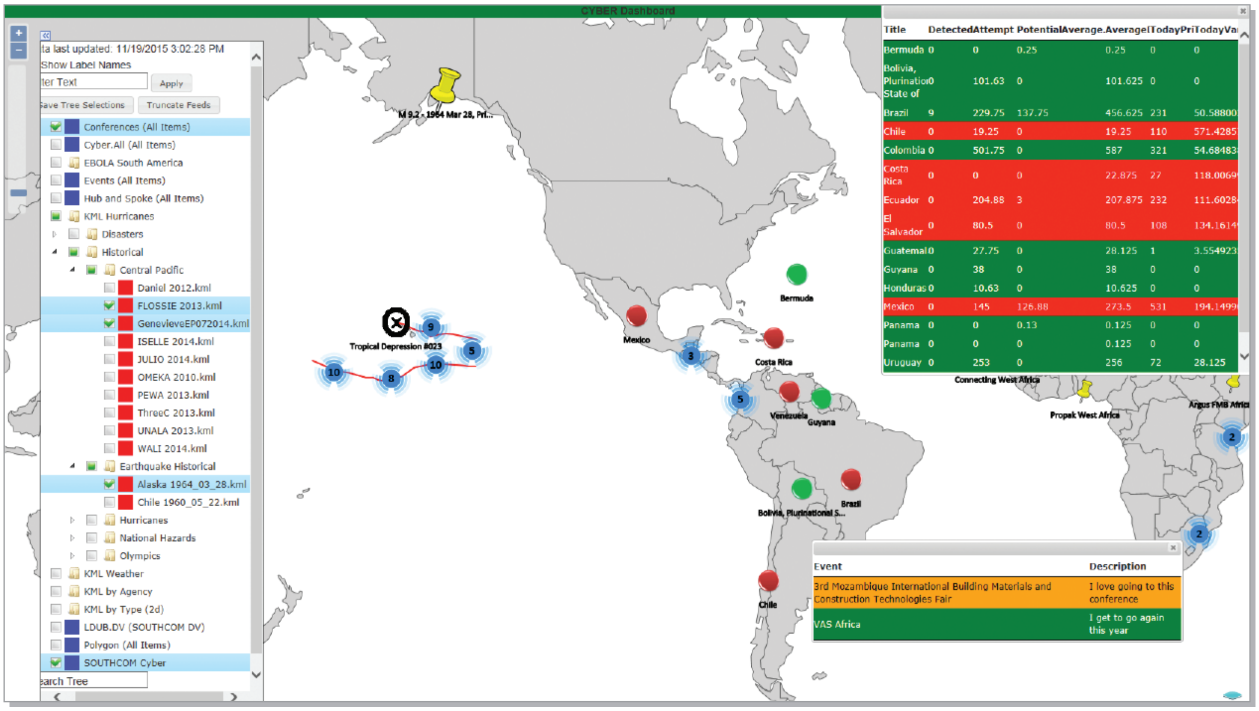

ta last updated: 11/19/2015 3:02:28 PM

Show Label Names

ter Text | Apply

ave Tree Selections | Truncate Feeds

- ☑ Conferences (All Items)
- ☐ Cyber.All (All Items)
- ☐ EBOLA South America
- ☐ Events (All Items)
- ☐ Hub and Spoke (All Items)
- ■ KML Hurricanes
  - ☐ Disasters
  - ■ Historical
    - ■ Central Pacific
      - ☐ Daniel 2012.kml
      - ☑ FLOSSIE 2013.kml
      - ☑ GenevieveEP072014.kml
      - ☐ ISELLE 2014.kml
      - ☐ JULIO 2014.kml
      - ☐ OMEKA 2010.kml
      - ☐ PEWA 2013.kml
      - ☐ ThreeC 2013.kml
      - ☐ UNALA 2013.kml
      - ☐ WALI 2014.kml
    - ■ Earthquake Historical
      - ☑ Alaska 1964_03_28.kml
      - ☐ Chile 1960_05_22.kml
  - ☐ Hurricanes
  - ☐ National Hazards
  - ☐ Olympics
- ☐ KML Weather
- ☐ KML by Agency
- ☐ KML by Type (2d)
- ☐ LDUB.DV (SOUTHCOM DV)
- ☐ Polygon (All Items)
- ☑ SOUTHCOM Cyber

earch Tree

Title	Detected	Attempt	PotentialAverage.	AverageI	TodayPri	TodayVa
Bermuda	0	0	0.25	0.25	0	0
Bolivia, Plurinatio State of	0	101.63	0	101.625	0	0
Brazil	9	229.75	137.75	456.625	231	50.58800
Chile	0	19.25	0	19.25	110	571.4285
Colombia	0	501.75	0	587	321	54.68483
Costa Rica	0	0	0	22.875	27	118.0069
Ecuador	0	204.88	3	207.875	232	111.6028
El Salvador	0	80.5	0	80.5	108	134.1614
Guatemal	0	27.75	0	28.125	1	3.554923
Guyana	0	38	0	38	0	0
Honduras	0	10.63	0	10.625	0	0
Mexico	0	145	126.88	273.5	531	194.1499
Panama	0	0	0.13	0.125	0	0
Panama	0	0	0	0.125	0	0
Uruguay	0	253	0	256	72	28.125

Event	Description
3rd Mozambique International Building Materials and Construction Technologies Fair	I love going to this conference
VAS Africa	I get to go again this year

FIGURE 6. The Cyber Dashboard design uses a background of a flat geographic map over which are laid palettes that present various types of data. In this image, the palette on the left side enumerates the data sources available. Users can interact with icons on the map to obtain details on demand; the drill-down information is displayed in the floating palettes depicted on the right side of the screen. The colors in the palettes can be defined for each data source; in this example, color relates to severity of the event, with green being low and red being high.

grams as a geospatial overlay and (2) a what-you-see-is-what-you-get (WYSIWYG) editor to assist novice users in managing and configuring their data sources.

FUTURE DESIGN CONSIDERATIONS FOR THE CYBER DASHBOARD

During the development of the final dashboard design, we identified several considerations to address in future versions. These items were considered out of scope for the original project but remain important operational considerations:

- Standardization. When users are incorporating 100s of data sources in a single geospatial display, it is difficult for a human to keep track of the provenance of the data. Color-coding, iconology, and taxonomy standards all play a role in helping a human in the loop to distinguish between data sources and perform visual search tasks. Conventions for how to standardize these elements must be included in future design guidelines.
- Data source "freshness." When using multiple data sources, analysts must understand how current the information is. Standard conventions for how to convey the timeliness of data need to be developed.
- Representation of complex relationships. For the cyber security domain in particular, complex dependencies beyond geospatial coordinates exist between data points. Understanding how these additional visual conventions can work in conjunction with map-based representations would improve situational awareness.
- Interpalette interactions. A natural next step for this project is to explore interaction paradigms by using the map-canvas metaphor to produce a generic framework that can be applied to interactions among arbitrary data sources.
- Intelligent fusion. Our current design allows analysts to do basic data management tasks but requires manual integration of cyber data across sensors and visualizations. Providing the ability to automatically tie together data sources on the basis of common fields or other dimensions would be beneficial.

Lessons Learned

Over the past three and a half years, the Cloudbreak team has learned many valuable lessons that researchers can apply to future DoD applications of the Cloudbreak approach.

Informal discussions and interviews emerged as a rich source of data to identify key behaviors and needs of the operators, enabling the team to quickly turn around new functionalities to improve operators' subjective experiences and performance while using the software. These interactions with the operators also were useful in gaining an understanding of the training, expertise, and experiences of the current operators. For example, during our discussions, we found that current Internet applications (e.g., Google, YouTube, social media) define operators' expectations of how DoD information technology (IT) systems should work. Operators expect speed and behaviors consistent with these applications; however, developing tools that accommodate both DoD system restrictions and operator expectations is a challenging balancing act. As an example, for a large text-analytic system, we updated the search interface from a powerful, expressive Boolean language to one that behaved similarly to popular Internet search engines.

Because of budget and security constraints, web browsers in use by the DoD tend to lag behind those found on the Internet. The latest web technologies, such as HTML5, are often simply not supported by DoD systems. For instance, we could not use the HTML5 canvas and scalable vector graphics elements to drive rich visualizations in all environments. We had to adjust our mindset and strike a balance between advanced technology and compatibility when we designed the web-based interfaces.

We found that operators were quick to discard or ignore capabilities and features that required them to employ many steps to accomplish a task; therefore, we worked hard to design utilities that eliminated excess steps from operator workflow and to take advantage of familiar interaction paradigms. For example, many workflows at COCOMs revolve around sharing data stored in a COTS enterprise content management system. Operators were more likely to use a capability when it was integrated into their content management system. We noticed a similar reaction to integrating an existing system with a public key infrastructure (PKI): By eliminating the need for additional username and password combinations and integrating the PKI with the operators' existing authentication system, we removed another barrier to their adoption of new tools.

From an organizational point of view, one cannot make assumptions that all operations centers share common processes. While most COCOMS have a common goal, there are variances in the battle rhythms based on the preferences of the leader and the current available skillsets within the command. Each leader has specific requirements for the way he or she prefers to consume information. The systems we provide must help operators prepare ahead of time for their commanders' needs. This task does not, however, equate to flooding leaders with information. The solution is to supply commanders with timely, mission-relevant data and to preserve other details for on-demand access.

The ultimate success of Cloudbreak is the users' adoption of the new technology. Because COCOMs are extremely busy, and the operators' time is split in many different directions, changes to systems must demonstrate immediate, recognizable benefits. Adoption results when the new capability demonstrates that it has practical value. Without this value, every capability is "just another tool" and will sit idle. Our example of eliminating the manual copy-and-paste operation is a great example of the addition of an immediate, clear benefit. Furthermore, we found that once an initial value of a tool or concept had been established, operator enthusiasm for new capabilities and the Cloudbreak process increased significantly.

Furthermore, aside from contractors, most operators at COCOMs rotate duty assignments every two to three years. Therefore, system developers cannot assume the operators' level of technical expertise or their familiarity with software and existing systems to remain constant. Continuing to communicate with operators and leadership to evolve systems as COCOM personnel, work practices, and goals change is critical to long-term adoption of systems.

Once adopted, capabilities cannot be set up and left to run indefinitely. Improvements, bug fixes, constant security monitoring, and hardware concerns all drive the need for a clear operation and maintenance plan. Cost, particularly that related to operations and maintenance, is a major consideration for DoD leadership. The COCOMs do not have the resources to either take on additional IT responsibilities or hire external IT support services. Consequently, in addition to providing capabilities, we were charged with determining how the capabilities would be sustained. Working with multiple government organizations and leveraging common interests, we were able to construct ways to distribute the costs and responsibilities for operating and maintaining systems among developers and users.

Future Work

The Cloudbreak initiative has been successful within a resource-constrained DoD. Through technology reuse and composability, we enabled cost savings. Our process for successful capability insertion has its foundation in a strong relationship with operators. The connections and trust we developed with them helped us discover the true areas where they most needed help. We allowed the operators' needs to drive the technology development, thus supporting operators by rapidly filling critical COCOM technology gaps. The Cloudbreak process also led to a new relationship with operators that could inform other collaborative projects and provide us access to operational datasets for future research, development, and experimentation.

Our experience with the Cloudbreak model leads us to endorse its continuation and replication across other areas of the DoD and within Lincoln Laboratory. The individual capabilities we have delivered, such as the Cyber Analytical Station, have reduced risk for systems currently being developed by identifying and validating requirements. The workflows and the features that we helped define have the support of JCC operators and COCOM leaders; therefore, the process of developing requirements for future systems does not need to start from scratch.

Moving forward, we will continue to apply an agile, user-centered research and development model. Currently, we are using the Cloudbreak approach in several ongoing efforts with USCYBERCOM, the U.S. Transportation Command, and the U.S. Navy. Building off lessons learned and the relationships with and access to operators developed under Cloudbreak will enable current and future Lincoln Laboratory programs to effectively align capability development with user requirements and to accomplish successful technology insertions. ■

References

1. T.J. Dosher, "NORAD, USNORTHCOM Joint Cyber Center Stands Up," U.S. Northern Command website, 2012, available at http://www.northcom.mil/Newsroom/tabid/3104/Article/563711/norad-usnorthcom-joint-cyber-center-stands-up.aspx.
2. "2011 Tōhoku Earthquake and Tsunami," Wikipedia, Jan. 2016, available at https://en.wikipedia.org/wiki/2011_T%C5%8Dhoku_earthquake_and_tsunami
3. "Operation Tomodachi," Wikipedia, Dec. 2015, available at https://en.wikipedia.org/wiki/Operation_Tomodachi
4. J.H. Pendleton, M. Morgan, N. Bleicher, M. Jones, M. Silver, J. Spence, and K. Williams, "Defense Management: Perspectives on the Involvement of the Combatant Commands in the Development of Joint Requirements." Washington, D.C.: Government Accountability Office, 2011, available from the Defense Technical Information Center at http://www.dtic.mil/dtic/tr/fulltext/u2/a546159.pdf.
5. "Ergonomics of Human-System Interaction—Part 210: Human-Centred Design for Interactive Systems," ISO 9421-210. Geneva: International Organization for Standardization, 2010.
6. P. McInerney and F. Maurer, "UCD in Agile Projects: Dream Team or Odd Couple?" *Interactions*, vol. 12, no. 6, 2005, pp. 19–23.
7. D. Sy, "Adapting Usability Investigations for Agile User-Centered Design," *Journal of Usability Studies*, vol. 2, no. 3, 2007, pp. 112–132.
8. D. Fox, J. Sillito, and F. Maurer, "Agile Methods and User-Centered Design: How These Two Methodologies Are Being Successfully Integrated in Industry," *Proceedings of the Agile 2008 Conference*, 2008, pp. 63–72.
9. S. McKenna, D. Staheli, and M. Meyer, "Unlocking User-Centered Design Methods for Building Cyber Security Visualizations," *Proceedings of the 2015 IEEE Symposium on Visualization for Cyber Security*, 2015, pp. 1–8.
10. D. Staheli, T. Yu, R.J. Crouser, S. Damodaran, K. Nam, D. O'Gwynn, S. McKenna, and L. Harrison, "Visualization Evaluation for Cyber Security: Trends and Future Directions," *Proceedings of the 11th Workshop on Visualization for Cyber Security*, 2014, pp. 49–56.
11. T. Yu, R. Lippmann, J. Riordan, and S. Boyer, "Ember: A Global Perspective on Extreme Malicious Behavior," *Proceedings of the 7th International Symposium on Visualization for Cyber Security*, 2010, pp. 1–12.
12. R.K. Cunningham, R.P. Lippmann, and S.E. Webster, "Detecting and Displaying Novel Computer Attacks with Macroscope," *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, vol. 31, no. 4, 2001, pp. 275–281, doi:10.1109/3468.935044.
13. K.M. Carter and W.W. Streilein, "Probabilistic Reasoning for Streaming Anomaly Detection," IEEE Statistical Signal Processing Workshop, Ann Arbor, Mich., 2012, doi: 10.1109/SSP.2012.6319708.
14. J. Dressler, C.L. Bowen, W. Moody, and J. Koepke, "Operational Data Classes for Establishing Situational Awareness in Cyberspace," *Proceedings of the 6th International Conference on Cyber Conflict*, 2014, pp. 175–186.

About the Authors

Diane Staheli is a member of the technical staff in the Cyber Systems and Operations Group at Lincoln Laboratory. Her current projects focus on field research with cyber security analysts and operators and the translation of user needs into visualization capabilities. Her research interests include cyber decision making, cyber human cognition, visual analytics, visualization evaluation, human-computer interaction, and emerging user interfaces. She joined the Laboratory in 2010, bringing 10 years of experience in industries ranging from a small home-networking startup to a global information security company. She serves on the board of the New England Chapter of the Human Factors and Ergonomics Society and is a current program committee member for the IEEE Visualization for Cyber Security Symposium, the VAST Challenge, and the Visualization for Data Science Workshop. She holds a bachelor's degree in communication, studio art, and film from the University of Massachusetts, Amherst, a master's degree in human factors in information design from Bentley University, and a master's degree in information technology and software engineering from Harvard University.

Vincent F. Mancuso is a member of the technical staff in the Cyber Systems and Operations Group at Lincoln Laboratory. His research interests include exploring issues of human factors in cyber operations and team cognition in distributed environments. Prior to joining the Laboratory, he was a postdoctoral researcher in the Human Performance Wing's Applied Neuroscience branch at the U.S. Air Force Research Laboratory, conducting research on cyber operator performance monitoring and optimization. He holds a bachelor's degree in information systems and human-computer interaction from Carnegie Mellon University and a doctoral degree in information sciences and technology from the Pennsylvania State University.

Matthew J. Leahy is a member of the technical staff in the Cyber Systems and Operations Group. He currently works on projects dealing with operations and cyber situational awareness. Previously at Lincoln Laboratory, he has worked on distributed test beds for live testing of ballistic missile defense systems, radar systems development, network communications, and distributed software systems. He holds a bachelor's degree in computer science from Worcester Polytechnic Institute and a master's degree in information technology with a specialization in software engineering from Carnegie Mellon University.

Martine M. Kalke is an assistant leader in the Cyber Systems and Operations Group. She currently directs projects to improve network operations and cyber situational awareness. At Lincoln Laboratory, she has previously worked on the development and evaluation of advanced algorithms for ballistic missile defense and intelligence, surveillance, and reconnaissance applications. She also was involved in designing a system architecture for the Missile Defense Agency while she was stationed in Washington, D.C. She holds a bachelor's degree in physics from Carleton College, and a master's degree in physics and a doctorate in condensed matter physics from Indiana University. Her doctoral thesis focused on computational condensed matter physics.