

SANS

Smart Enterprise Visibility with DTEX InTERCEPT

Written by **Matt Bromiley**

December 2020

Sponsored by:
DTEX Systems

There's one common weak point among every digital enterprise in the world. Whether your organization provides healthcare, processes financial transactions or moves freight from one place to another, the least common denominator remains the same: the user. An unsuspecting user may inadvertently provide an attacker with access to an environment, whereas a malicious insider intentionally seeks data to steal and damage to cause. Many information security teams take steps to limit what the user can do, while monitoring their technology with as much visibility as possible. However, attacks still occur, and we continue to see high-profile, public breaches month after month.

Does this mean the user is inherently bad or security is destined to fail? Quite the opposite. It means *users need different types of protection so that security has a fighting chance to succeed*. In this whitepaper, we review a platform that seems to have taken that statement to heart: DTEX InTERCEPT. By focusing its efforts on user behavioral patterns, DTEX has uncovered a core truth of cyber intrusions: Threat actors typically need accounts to achieve their objectives, and those accounts seldom resemble normal activity. InTERCEPT is an (extremely lightweight!) agent-based platform that, via holistic visibility, provides unique insight into user behavior with its business-critical classifications.

Specifically, in this paper, we review:

- How InTERCEPT provides correlated, intricate data in a consumable and actionable manner
- How the platform is powerful enough for strong security analysts while flexible enough to provide data for security leadership and executives who are making business-critical decisions

- InTERCEPT’s user behavior scoring engine, a truly unique feature that provides DTEX an advantage in analyzing past and future user behavior
- The capability to get to raw data and perform rapid investigations and threat hunting, found via InTERCEPT’s custom dashboards, correlations and analytics

As you work your way through this whitepaper, we encourage you to examine how DTEX utilizes, correlates and delivers on user behavior activity. We found that InTERCEPT took a unique approach toward detecting suspicious user activity, providing us with enough context to understand the risk to the organization and the next steps to take. We also encourage you to examine your own capabilities in the space of user behavioral analysis and detections—are you leaving an incredibly useful detection vector on the table?

Working with InTERCEPT

When assessing holistic security platforms, we often focus on first impressions when getting hands-on. The initial dashboard or screen is often what analysts see before they begin working on their tasks. Leadership and stakeholders who utilize the platform will also see the same data; thus, the initial screen must be able to effectively communicate to users of all levels within an organization’s security team.

Visibility into the Enterprise

Luckily for us, the initial InTERCEPT dashboard, known as the “Threat Overview” screen, holds little back. Packed full of data, it contains multiple representations of threats within the environment. As shown in Figure 1, the dashboard provides insight into platform visibility right up front. Capturing the number of alerts, users, endpoints and activity gives a quick glance at how much of the environment is currently being observed.

Takeaway

It is easy to see that DTEX understands that organizations need lots of visibility to remain vigilant about their security programs. From high-level, impactful data points to granular searching options, DTEX provides multiple ways that analysts can examine the data within their environment.

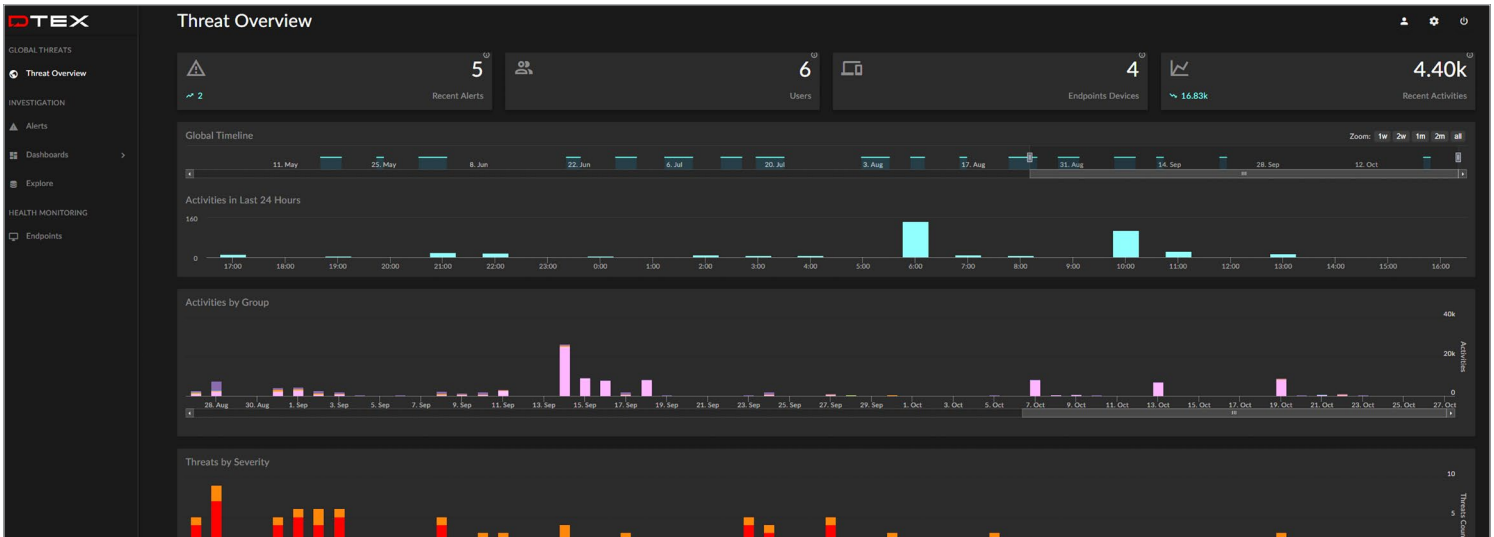


Figure 1. Threat Overview Dashboard

After assessing initial visibility, InTERCEPT provides multiple collated vantage points. It is worth noting here—as we will again—that the initial platform is entirely data-driven and interactive. Furthermore, the initial dashboard simultaneously provides historical *and* up-to-date threat details. This happens via multiple data panels that depict observed activities and threats, both of which are controlled via a single Global Timeline, as shown in Figure 2.




Figure 2. Threat Overview Dashboard: Global Timeline Scrollbar

The Global Timeline enables a user to look at patterns of activity over the *history* of InTERCEPT’s data capture. There are also quick links for weekly and month-based visibility. We would argue that this feature—the capability to quickly customize representation of threats over time—is one of the most necessary features right out of the box. From an analyst’s perspective, historical correlation and activity within the environment are often dependent on recordkeeping and note taking by the security team. If management wishes to know activity for the previous three or six months, that knowledge is often buried in complex SIEM queries or a ticketing system that hopefully has been kept up to date—if retrievable at all. It was refreshing to see that DTEX thought of this up front and allowed for an intuitive yet highly granular opening screen that can easily switch from historical to up-to-the-minute context with a simple slide of a scrollbar.

Thus far, everything we have examined has simply been the opening dashboard. It should provide actionable and impactful data. With time-based customizations, DTEX provides the data via correlated panels that show trends or summarized data points. Our favorites were the combined, correlated and interactive views of Activities by Group and Threats by Severity. (See Figure 3.)

As mentioned earlier, the dashboard updates in real time and is interactive. In Figure 3, we hovered over activity from September 14, 2020—a date that showed a clear spike in grouped activity from our timeline. Both panels work together from a hover-over perspective, and the analyst can quickly determine that there were four threats (2 high and 2 urgent) extracted from *thousands* of activity data points.

Figure 3 is also our first preview into the data types that InTERCEPT receives, analyzes and correlates on the back end. Activities such as network connections, file system and process operations, and external device and clipboard usage are just a few examples of powerful data points that InTERCEPT uses in its detection and alerting. We will look at these data points and InTERCEPT alerts in more detail in the next section.

Takeaway

It cannot be overstated how many of the “simplistic” visualizations offered in InTERCEPT are based on significant data collection and correlation behind the scenes. One of our favorite features of the platform was how complex data were represented in an easy-to-consume format.




Figure 3. Activities by Group and Threats by Severity Panels

One for All, All for One

It is obvious at this point that we found a lot of value in the Threat Overview dashboard. This is a nod toward holistic security platforms that provide useful data points for all levels of enterprise security, from analysts to C-suite executives. These users are all looking to extract different data summaries from the same environment—not an easy undertaking. We have observed many products that provide fantastic analyst insight, while leaving much to be desired for management reporting and enterprise metrics.

Figure 4 illustrates how DTEX found applicability for all within the enterprise security ecosystem.

As we briefly saw earlier, InTERCEPT is collecting, analyzing and correlating myriad activity types. This creates a ripe

opportunity for even junior analysts to gain granular visibility within their enterprise environment. In the next section, we will make use of the data and review investigation and hunting capabilities within InTERCEPT. But for now, let us shift our focus to examining how InTERCEPT simultaneously positioned itself as an accessible and easy-to-use tool for both security management and leadership.

Enterprise Security Data—Intuitive, Insightful and Accessible

The role of security managers and leadership is often wrapped up in metrics and reporting, as they work to ensure that top threats are handled effectively and the organization remains secure. We found that sitting on top of a mountain of data enables InTERCEPT to provide extraordinary insight and answers to some of the most complex questions.

Within the main navigation pane, we grew fond of the Alert Heat Map dashboard, which provides insight into and classification of user activity. At this point, we must pause to reiterate the power behind InTERCEPT’s analysis, correlative and data visualization capabilities. In many platforms or data-aggregation SIEMs, thousands of data points are seldom represented in an easily consumed format or with the granularity we see here. DTEX thought through this issue and provided useful visualizations.




Figure 4. InTERCEPT is applicable for all users, from daily analysts to C-suite members.

Takeaway

Many of the dashboards geared toward management and leadership not only provide high-level statistics and summaries, but also wrap them up in categories such as data loss, compromised and/or malicious. With advanced behavioral analytics behind the scenes, this requires little interpretation, allowing for quick decision making. We love it!

What does this mean for the security analyst? Without leaving the initial dashboard, the analyst can zoom in on concerning periods of time and immediately see the data points behind each period. However, within the specific dashboards that InTERCEPT provides, analysts can also view their environment in a new and unique way: one that focuses on user behaviors, scoring and actual threats to the environment. One of our favorites, shown in Figure 5, is the Alert Heat Map dashboard, which provides a comprehensive, data-packed viewpoint into the indicators detected within the environment.

The scores represented in Figure 5 are the result of InTERCEPT's user activity monitoring, which sets it apart from other platforms. As we previewed earlier, InTERCEPT is continuously collecting loads of metadata points from endpoints within the environment. Much of the metadata center around system activity and user behavior, the latter of which comes with some very telling signs of account compromise or malicious behavior. It is very, very difficult to move around an enterprise network without an account. Often, one of the first things attackers do is compromise an account so they can perform their next actions.

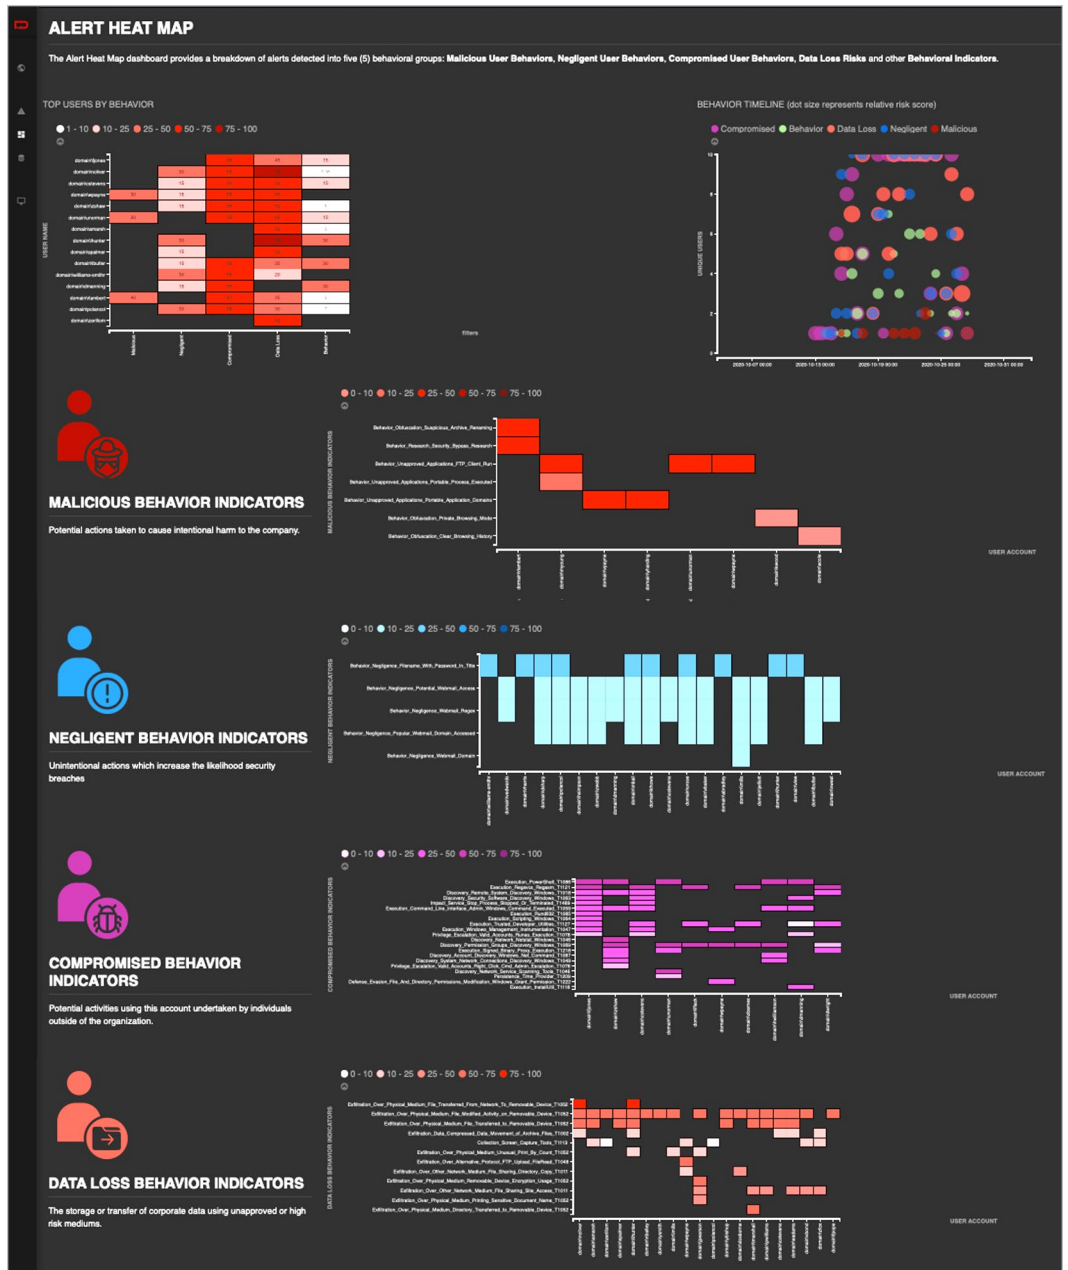


Figure 5. Alert Heat Map Showing User Activity Classification and Risk Scores

Takeaway

InTERCEPT is constantly collecting metadata from endpoints, much of which includes user behavior. The metadata are analyzed to detect suspicious user behavior, which is subsequently scored, alerted on and provided to you via the InTERCEPT dashboard. No guesswork, no confidence scoring for your security team—the platform does it all, so your team can focus on protecting the organization.

When InTERCEPT detects a pattern or sequence of activity deemed to be suspicious, it will score it and alert appropriately. For example, in the Alert Heat Map shown in Figure 5, we can zoom in on the Top Users by Behavior diagram, shown in Figure 6. Observed enterprise accounts have their behavior scored and aligned with behavioral filters, including indicators of the following types of activity:

- Malicious
- Negligent
- Compromised
- Data Loss




Figure 6. Alert Heat Map Focused on Top Users by Behavior

How does this help? For starters, InTERCEPT has taken the guesswork out of analysts needing to write their own user behavior detections. It is the very nature of this platform to analyze and score that behavior. In addition, analysts can utilize the severity of scores to know where to prioritize their efforts. Let us drill down further, for example, into users shown in Compromised Behavior Indicators (see Figure 7).

How does InTERCEPT determine “compromised behavior” activity?

By aligning it with the MITRE ATT&CK® Matrix, of course. Notice in Figure 7 that InTERCEPT has detected activity that aligns with certain malicious activities, such as execution of scripts within Windows (T1064), stopping or terminating services (T1489) and usage of data encryption tools (T1486). Aligning observed behavior with industry-standard techniques takes a huge step toward normalizing alert details for analysts.




Figure 7. Compromised Behavior Indicators from the Alert Heat Map Dashboard

InTERCEPT also observes and scores behavior that may not be indicative of malicious intent (another feature we loved), but instead focuses on future employee activities. In Figure 8, for example, we can see that InTERCEPT detected a user who was visiting recruitment websites and uploading their resume. Although changing jobs is not a bad thing, it is useful for the security team to have insight into future potential threats.



Figure 8. General Behavior Indicators from the Alert Heat Map Dashboard

With enough visibility, DTEX has realized that it can take its user behavior scoring engine a step further: to determine user *intent*. Think about the following example:

1. A user receives and clicks through a spearphishing email, inadvertently launching malware on their system.
2. An attacker gains remote access to the system and, via the context of the victim account, begins performing intrusive activities.
3. The attacker escalates privileges (taking over yet another account) and begins performing more intrusive activities, potentially moving laterally to other systems and deploying additional malware.

Many security solutions that focus on exploits, malware and malicious code will (hopefully!) detect malware deployment and perhaps attacker activities such as enterprise reconnaissance or lateral movement. DTEX, however, has realized that all the activity represented in the preceding steps *requires* user account activity. Although some solutions might detect in step 3, DTEX would detect in step 2 and perhaps even step 1—when a user clicks a malicious email. The point remains strong: User accounts are involved in intrusions almost 100% of the time. InTERCEPT answers the question: What if we focus our detection efforts on user behavior, rather than waiting for malicious code to execute?

The CISO Scorecard

Another extremely impressive feature of InTERCEPT is its appeal to higher executive and C-suite security leadership. Built into the platform, within the main navigation bar, is the aptly named CISO Scorecard. This executive overview (shown in Figure 9) provides a one-stop click for leadership to gain a viewpoint into the current state of the environment.

As shown in Figure 9, the CISO Scorecard not only provides the numbers of alerts and users, but also averages and summarizes the risks the environment has experienced. Looking for a quick way to determine if your security investment is working? DTEX has you covered: Grab metrics from month to month and watch to determine whether your security posture is improving.

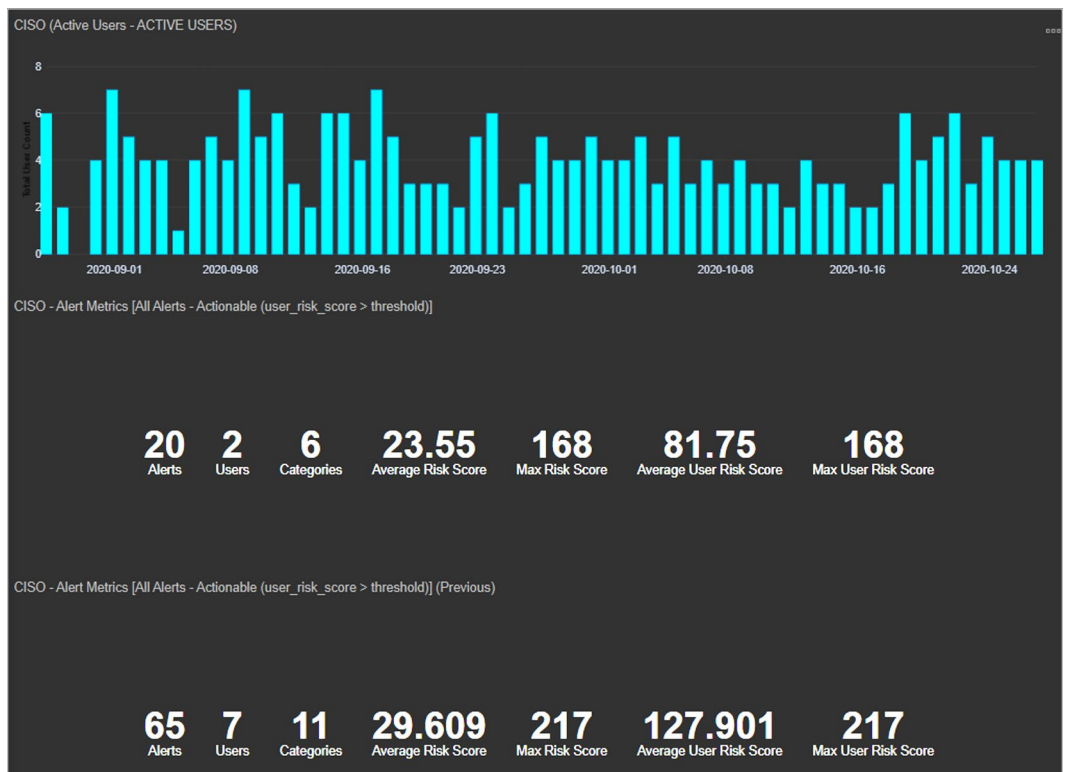


Figure 9. CISO Scorecard in the InTERCEPT Platform Dashboard

Scrolling down provides more granular details, such as the types of alerts observed, number of alerts and users over a period of time and, of course, associated risk scores. Figure 10 shows an example of more granular data points, still from the CISO Scorecard.

Perhaps one of the best features we uncovered in the platform is the CISO Recommendations portion of the scorecard. This list of recommendations, a snippet of which is provided in Figure 11, describes a list of categories, risks to the business and recommendations for handling these types of risks.

Executive insight does not stop there. Built into the platform (and applicable for many other dashboards) is the ability to export data into concise, informative single-page reports.

Figure 12 on the next page shows a copy of the CISO Executive Overview report.

The CISO Executive Overview provides the data executives need at a glance:

- The quantity of workstations and active users within the environment
- An internal risk scorecard, assessing incidents against DTEX's benchmarks of user behaviors, indicating malicious, negligent, compromised and data loss behavior
- A summary of recommendations, focused on technology, people and process

CISO - Category Alerts

Risk_Area	Alert_Count	Unique_Users	Max_Risk_Score
Malicious	9	2	75
Behavior	7	2	5
Data Loss	4	2	30

CISO - Category Alerts - All

Category_ID	Category	Risk_Area	Alert_Count	Unique_Users	Max_Risk_Score
AGG-AL-DTACP	Aggregation - Data Archive Creation - Correlation	Behavior	7	2	5
OBF-AL-ARCRN	Obfuscation - Suspicious Archive Rename	Malicious	7	2	33
EXF-AL-PRSWML-T1011	Exfiltration - Over Other Network Medium - Personal Webmail [T1011]	Data Loss	2	2	30
EXF-AL-SHRWEB-T1011	Exfiltration - Over Other Network Medium - File Sharing Site [T1011]	Data Loss	2	1	25
EXF-AL-OBFUP	Exfiltration - Obfuscated Internet File Upload	Malicious	1	1	75
EXF-AL-PSTDAT-T1011	Exfiltration - Over Other Network Medium - Posting Data To Website [T1011]	Malicious	1	1	20

Figure 10. CISO Scorecard: CISO Category Alerts and Breakdown

INTERCEPT CISO RECOMMENDATIONS

category_id	risk	recommendation
BAC-AL-ADMUSR	potential misuse of a privileged account	Review and revoke access to unapproved local, service or domain administrative accounts
BAC-AL-LIADM	access to an unauthorized admin account	Review and revoke access to unauthorized local, service or domain administrative accounts
BAT-AL-LUPOFF	large uploads to offnetwork sites	Investigate the confidentiality of the files transferred to the offnetwork sites
BFR-AL-FLTRSK	flight risk early indicators	Review the flight risk indications and escalate individuals to the Persons of Interest list
BNG-AL-INGTNT	access to inappropriate internet content (e.g. pornography) with increased risk of drive-by-	Remind users of IT Acceptable Use Policies

Figure 11. CISO Scorecard: CISO Recommendations

With these offerings, InTERCEPT has created a platform that provides access to data for all levels of enterprise security. It has also solved perhaps one of the hardest but most common questions in security monitoring: We have an alert. What do we do now? DTEX provides a list of recommendations that can be actioned upon immediately to help defend the organization from any subsequent damage. *We love it when tools bundle recommendations with observed activity!*

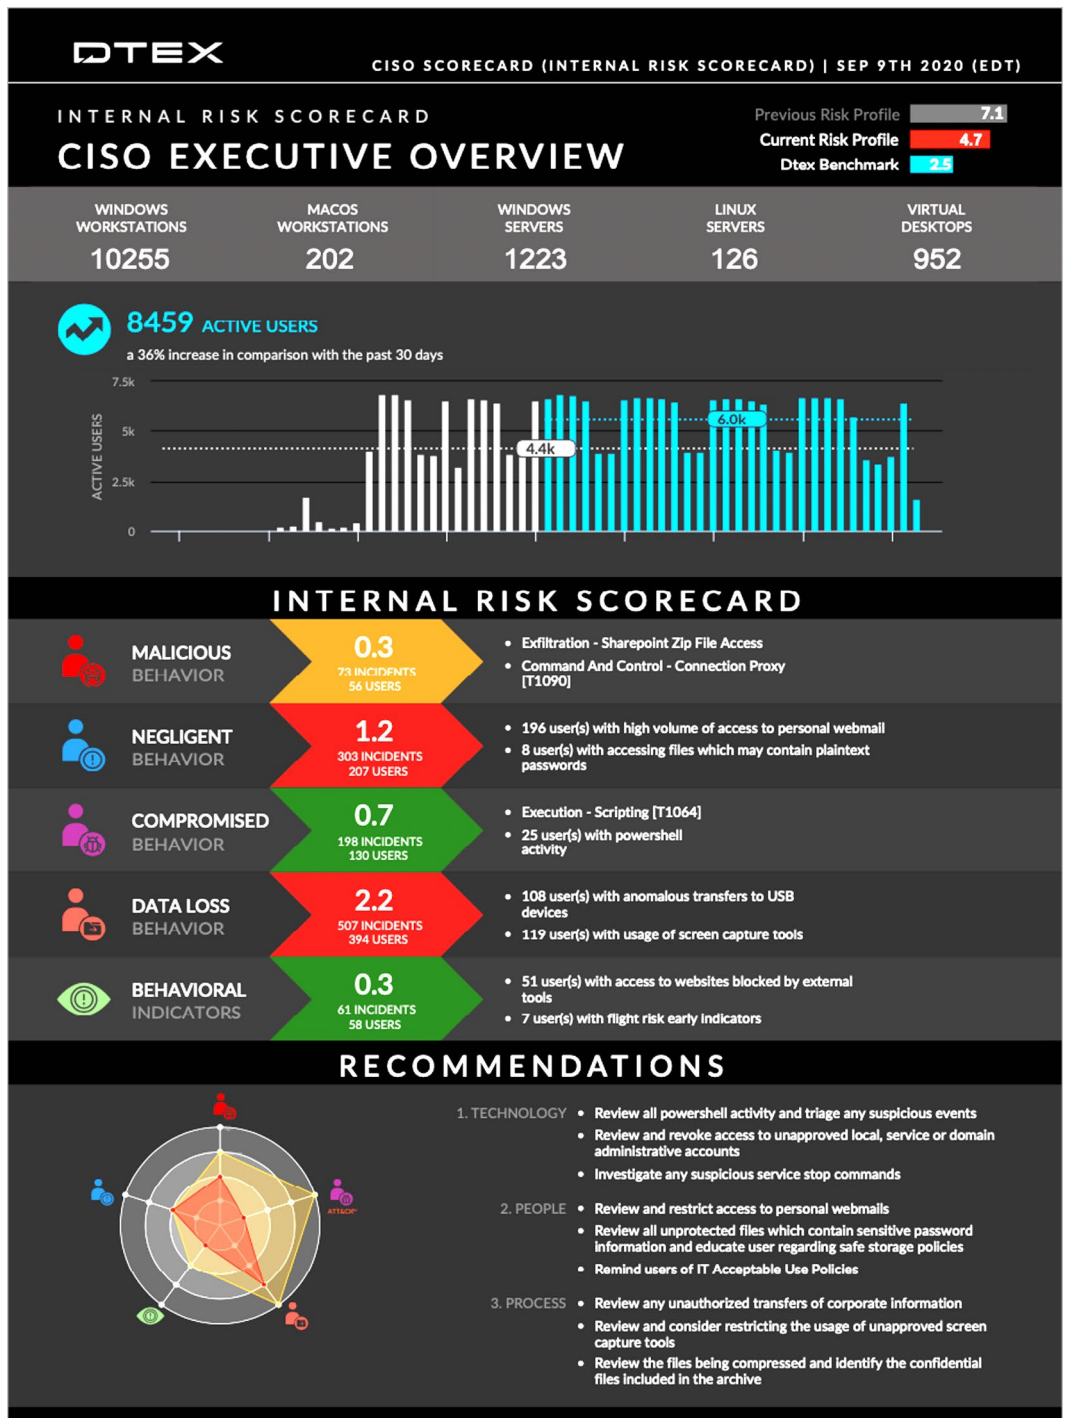


Figure 12. CISO Executive Overview Report

Investigating and Hunting with InTERCEPT

In the previous section, we spoke at length about the benefits that InTERCEPT provides to management and executive leadership. Although the reporting and metrics are useful, the real power of InTERCEPT is experienced by the day-to-day security analyst and/or incident responder.

Alerts, Analysis and Investigation

Within the initial dashboard, analysts can use the Alerts section (shown in Figure 13) to drill down into detections and analyze user activity.

Figure 13 contains multiple alerts that were correlated into three “events.” Notice that our scoring behavior is present, showing that our offending user account received a combined score of 135 on 2020-10-29, for example. Remember, the alerts and activity we will be exploring are based on *user account activity*.

The “state” of each correlation shows the status of analysis, indicating whether an analyst has picked up and is examining the activity further. From here, an analyst can take *multiple* routes. For example, exploring the category further (remember, almost every text field in the Alerts screen is a link) provides more context on the category that InTERCEPT observed. Figure 14, for example, provides insight into the **OBF-AL-ARCRN** category, also known as Potential Obfuscation Behavior via Suspicious Archive Renaming.

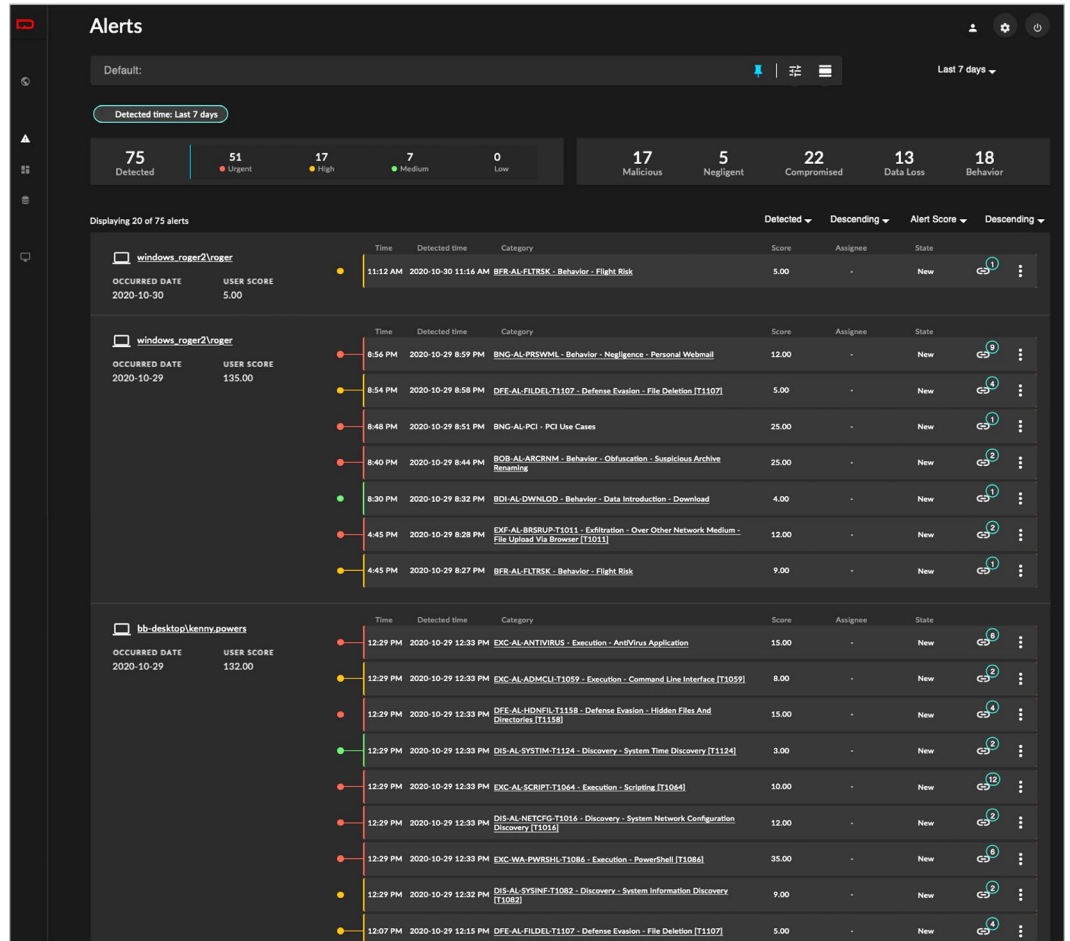


Figure 13. Alerts Dashboard from the InTERCEPT Main Screen

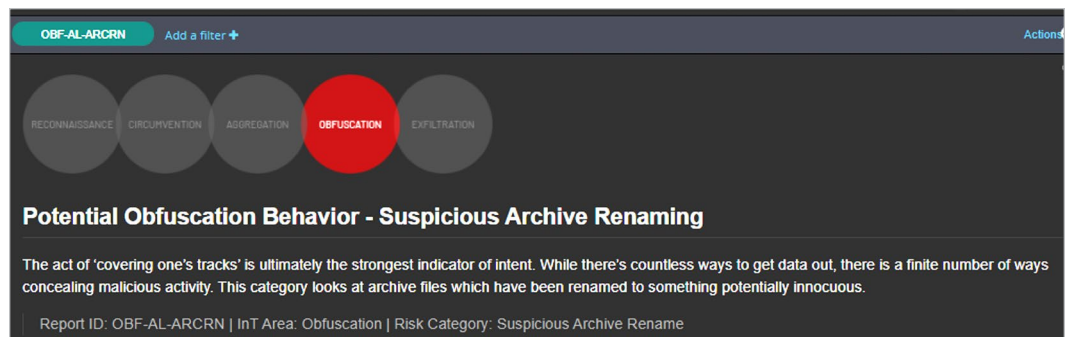


Figure 14. Explanation of **OBF-AL-ARCRN** Category (Potential Obfuscation Behavior)

Time	User_Name	Device_Name	Activity_Type	Process_Name	Activity_Details	Process_Parameters	Parent_Process_Name	Parent_Process.CommandLineParameters	tags.rule	tags.category_id
Q October 28th 2020, 14:25:30.751	windows_roger2\roger	WORKGROUP\Windows_Roger2	WindowGolfocus	VeraCrypt.exe	VeraCrypt	"C:\Program Files\VeraCrypt\Veracrypt.exe"	Explorer.EXE	-	Impact_Data_Encrypted_For_Impact_Encryption_Tools_T1486	IP1-AL-DATENG-T1486
Q October 28th 2020, 14:25:16.493	windows_roger2\roger	WORKGROUP\Windows_Roger2	WindowGolfocus	VeraCrypt.exe	VeraCrypt	"C:\Program Files\VeraCrypt\Veracrypt.exe"	Explorer.EXE	-	Impact_Data_Encrypted_For_Impact_Encryption_Tools_T1486	IP1-AL-DATENG-T1486

Figure 15. Investigation into User roger

An analyst can also expand details of the alert, getting to raw audit data that InTERCEPT ingested and analyzed. Clicking one of the numbered links takes the analyst to the dashboard shown in Figure 15.

In Figure 15, we can see that InTERCEPT identified usage of **VeraCrypt.exe**, a data compression tool. This is subsequently tied to MITRE Technique T1486, which highlights data impact via use of encryption tools. If this dashboard looks familiar, there is a chance you have seen this interface before. One pleasant surprise we kept encountering during our review is that DTEX sits on top of an Elasticsearch instance, which helps drive its analytics, visualization and data representation.

We absolutely love this architecture! Although many security solutions do a good job of summarizing data and providing context to the analyst, getting access to the raw data that generated an alert often requires

some back-end access or a secret API call. DTEX decided not to go that route, instead providing analysts with access to the raw data and back end. An analyst can stay in the InTERCEPT dashboard or move around the raw data, investigating and hunting as they go. And the raw data are accessible at any time from the main screen, provided in the Explore link in the navigation pane.

Although we love the access to raw data, we wanted to explore other ways the tool could mold and shape the data for us. We quickly learned that DTEX thought of this, too.

Within the list of available out-of-the-box dashboards is one called User Investigation. This dashboard, a custom overlay of the same Elasticsearch instance mentioned previously, gets right to the point of identifying suspicious data within the environment. Figure 16, for example, provides a sample user investigation dashboard for the user **windows_roger2\roger**.

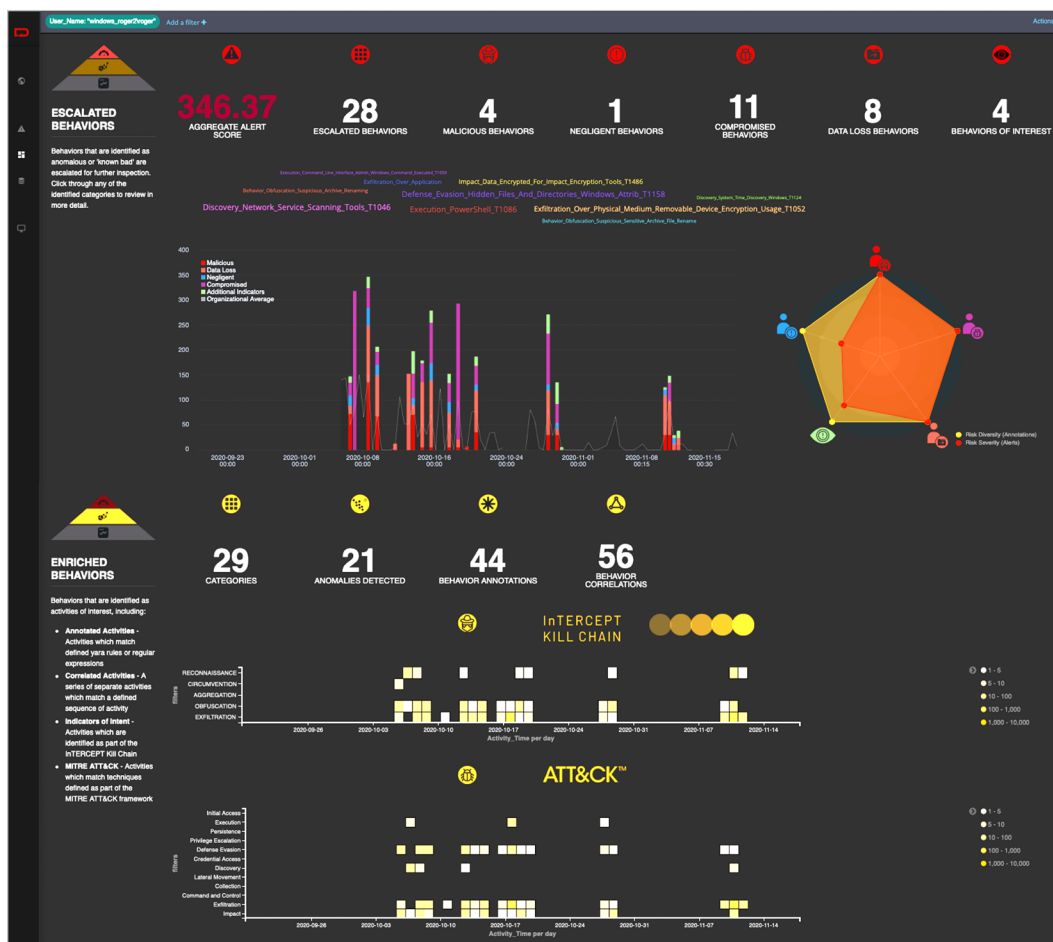


Figure 16. User Investigation Dashboard: Summarized Metadata

Security analysts can also use the User Investigation dashboard to highlight users that have generated the most activity, yet another avenue for detecting suspicious user activity within the enterprise. Figure 17 provides an example of cumulative InTERCEPT user-based data; Figure 18 provides the same data but with a focus on key devices.

User Name	Count
dtexsalesm01stleve etcher	84,049
russell-centos-3 dtex.lan/root	11,264
desktop-roger1vroger	9,343
desktop-sp43sas\demo_user_al	8,386
bill-breacherbill.breacher	5,787
tiffany-trader\tiffany.trader	4,835
desktop_demo3tricky bobby	1,858

Figure 17. User Investigation Dashboard: Aggregate User Account Counts

The User Investigation dashboard is not just a display of data; it is actually a custom Elasticsearch interface, meaning that each and every data point can be filtered in or out to perform live searching and correlation. Applying a filter for the top offending user account, for example, modifies our live content to show us the applications that the user account in question used and the websites visited. Figure 19 provides another example, highlighting the context of an account of interest.

Device Name	Operating System	OS Version	Device Type	Count
workgroup\dtexsalesm01	Windows	10 (10.0.17763)	Desktop	84,135
russell-centos-3 dtex.lan	Linux (CentOS Linux)	7 (Core)	Virtual Machine	11,489
workgroup\desktop-roger1	Windows	10 (10.0.19041)	Laptop	7,464
workgroup\desktop-roger1	Windows	10 (10.0.18363)	Laptop	2,227
workgroup\desktop-sp43sas	Windows	10 (10.0.18363)	Laptop	8,310
workgroup\bill-breacher	Windows	10 (10.0.15063)	Desktop	5,896
workgroup\tiffany-trader	Windows	10 (10.0.15063)	Desktop	4,877

Figure 18. User Investigation Dashboard: Aggregate Device Counts

Of course, InTERCEPT is not the first to utilize these available data points. However, to view them in user-based context, brought together and correlated, allows for an extremely impactful way to view and analyze user behavior without needing to seize a system or collect additional artifacts. The telemetry that the InTERCEPT agent captured was enough to analyze behavior, drill down deeper and make an appropriate decision about further investigative needs.

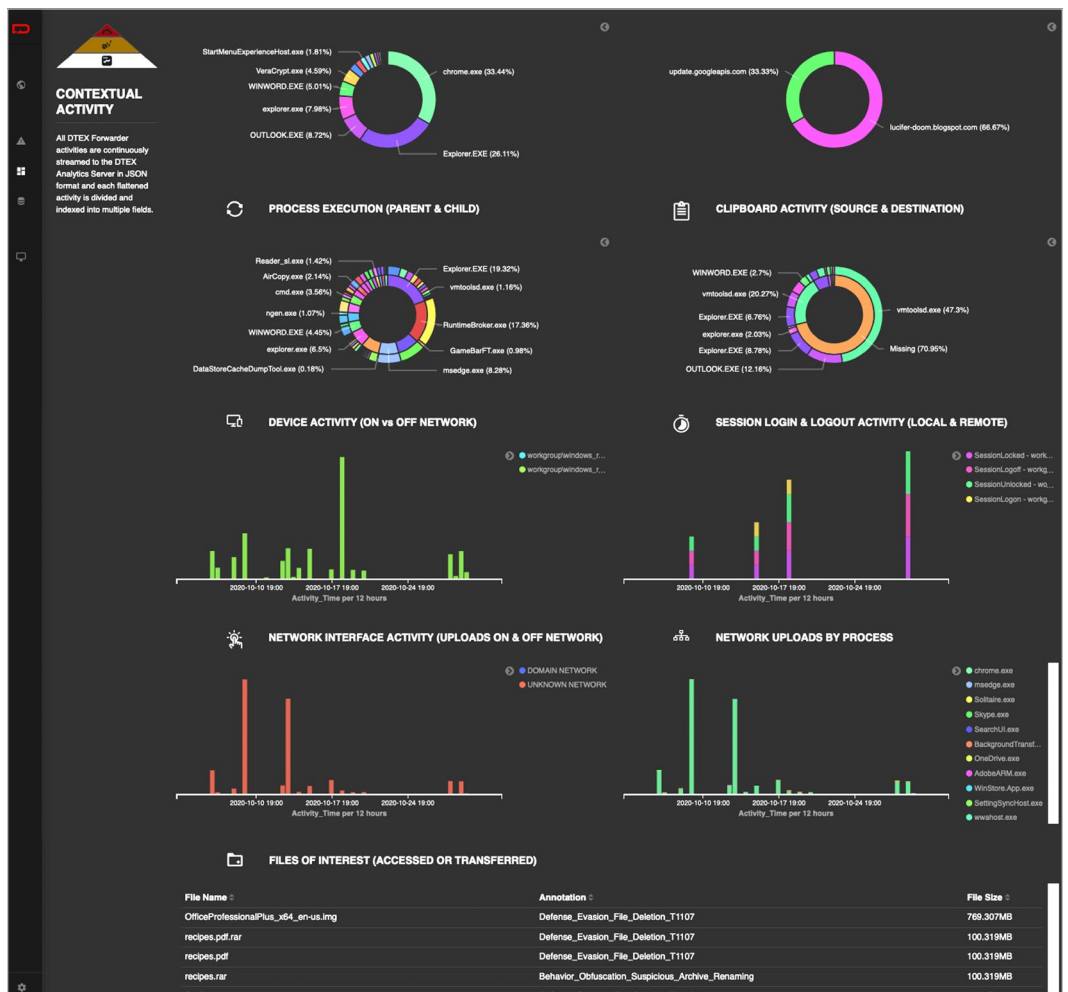


Figure 19. User Investigation Dashboard

And of course, as we have previously mentioned, InTERCEPT's reporting capabilities are seldom matched within the industry. Figure 20, for example, provides a User Investigation Report for the user **DESKTOP-ROGER1\ROGER**, a user we have seen in multiple previous examples and one that we should certainly be investigating.

In previous dashboards and screens, we were hopping among alerts, events and correlated user behavior. As shown in Figure 20, we have a comprehensive viewpoint of investigative activities, allowing for a quick snapshot of what the security team has observed. We loved this output!

Finally, sometimes detected activity is simply a user either doing their job or inadvertently tripping alarms for non-tuned alerts. Tuning security tools is perhaps one of the most crucial functions of a security analyst. DTEX makes this easy as well. From the original alert screen, an analyst can also quickly turn observed activity into a policy via the intuitive Add Policy option for each alert. Shown in Figure 21 on the next page, this feature allows for quick user-based policy exemptions.

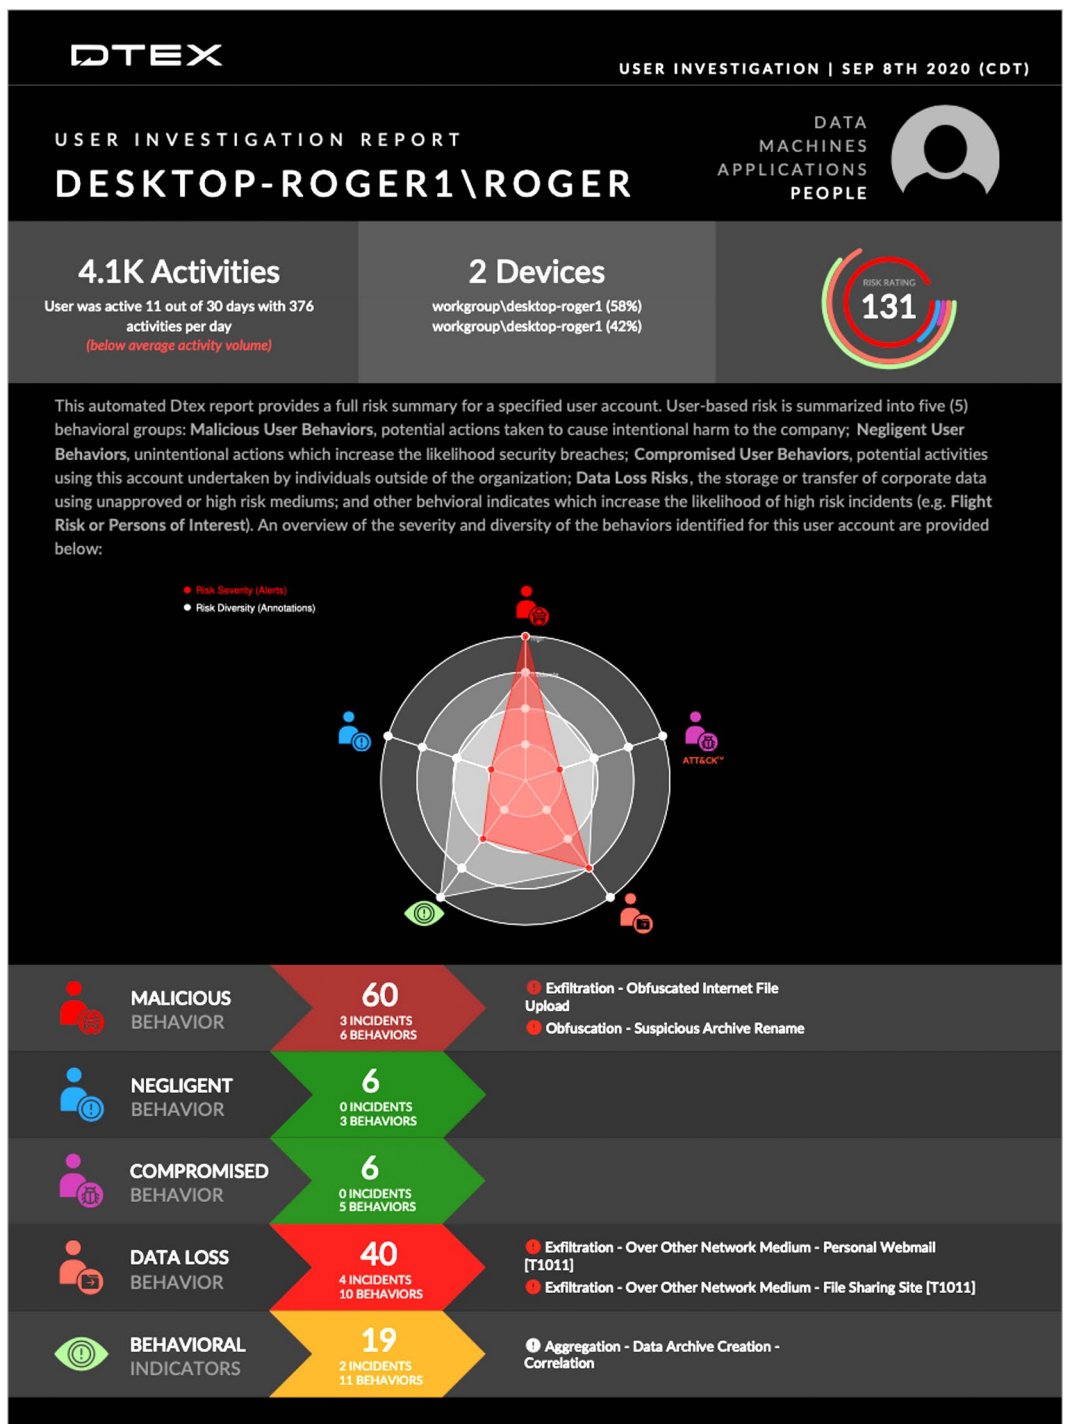


Figure 20. User Investigation Report for the User **DESKTOP-ROGER1\ROGER**

Policies can also be created with start and end dates to help minimize noise and reduce false positive detection within the environment. We always appreciate a feature that makes analysts' work easier, and while InTERCEPT does a fantastic job of rolling up and scoring suspicious user behavior, we like that analysts are able to tune InTERCEPT to the needs of their specific environment.

Looking at Data Differently

Last, but certainly not least, another feature of InTERCEPT we have raved about and loved was its capability to contextualize metadata into business risk and impact. For example, when we began to explore the different dashboards available within InTERCEPT, we came across parent categories such as "Compliance" and "Remote Worker," as well as insights such as daily active users or VPN user trends (see Figure 22).

It is likely that your enterprise looks very different now than it did a year ago. We were pleasantly surprised to see that InTERCEPT recognizes this and, via its vast telemetry and endpoint metadata, can provide useful insights into workers on and off the corporate network.

Conclusion

In this whitepaper, we spent some time with InTERCEPT, a holistic platform that offers incredible granularity and visibility into modern enterprises. Via a lightweight agent and collection of myriad metadata points, InTERCEPT can, with very strong confidence, build on and detect user account behavior and anomalies. As we identified in the opening, most (if not all) threat actors rely on accounts to maintain a foothold in a network and achieve their objectives. Unfortunately for them, their activity seldom represents legitimate user behavior.

InTERCEPT captures, analyzes and alerts on this behavior with an intuitive, Elasticsearch-based platform that provides rich insight into observed behavior. However, InTERCEPT is more than alerting. Through usage of the tool, we discovered a platform that has all security users in mind, from the starting analyst to management and executive leadership. By creating a robust platform, DTEX appeals to a wide audience and provides a novel user behavior detection engine to boot.

Finally, InTERCEPT answered a question that we have been thinking about for a while: Are current detection techniques working—and, if so, why do we continue to see breach after breach? Perhaps it is time to shift detections and instead focus on catching the attacker before they can launch an attack. With InTERCEPT, it is the attacker's own behavior that fires an alert and gives the defenders an actual advantage.

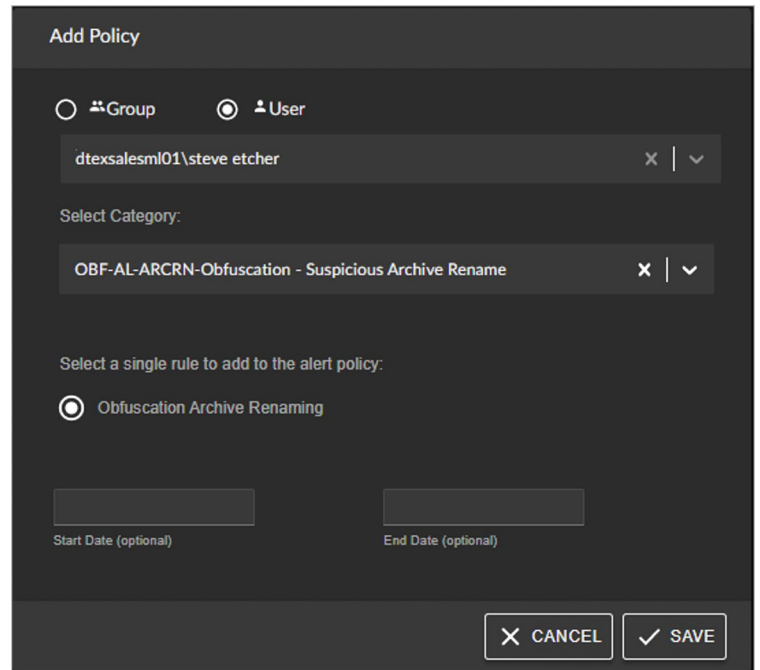


Figure 21. Alerts Panel: Add Policy Submenu



Figure 22. Sample of Dashboards in InTERCEPT, Providing Granular Insight into the Enterprise

About the Author

Matt Bromiley is a SANS digital forensics and incident response instructor, teaching [FOR508 \(Advanced Incident Response, Threat Hunting, and Digital Forensics\)](#) and [FOR572 \(Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response\)](#).

He is a principal consultant at a global incident response and forensic analysis company, combining his experience in digital forensics, log analytics, and incident response and management. His skills include disk, database, memory and network forensics; incident management; threat intelligence; and network security monitoring. Matt has worked with organizations of all shapes and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor:



THE WORKFORCE CYBER INTELLIGENCE COMPANY™