WHITE PAPER

exabeam

UNDERSTANDING THE EXABEAM SESSION DATA FRAMEWORK

INTRODUCTION

The threat landscape has changed significantly in recent years. Attacks such as SQL injections, viruses, etc. were usually isolated transactions; they were executed, the results were typically visible, and the threat ended. While the effects could be significant (for example, customer data loss due to SQL injection), the attacks were contained. In contrast, attacks in recent years are very different. Attacks via compromised credentials are complex, spanning multiple targets, multiple accounts, multiple devices, and multiple IP addresses. Threats from malicious insiders are similar, as employees or contractors use their access rights to copy and exfiltrate data over time.

These modern attacks may last for months; they are not isolated transactions. As a result, detecting and containing them is much more difficult.

DATA STRUCTURES NEED TO CHANGE

When attacks were isolated transactions, the log event was a useful data structure. It provided enough of the *who*, *what*, and *when* of the attack to support detection and remediation. Log events on their own, however, aren't useful for detecting and understanding complex, long-running attacks. In addition to who, what, and when, security analysts need to know, for every activity, *where did this user come from*, *where did she go next*, and most importantly, *is this behavior normal*? There is no way to derive this from a log event; a new data structure is required.

The Exabeam platform creates this new data structure: ***the session***. The session object stitches together all events for each user, from session initiation to termination, and ties these to a user even if she changes accounts, changes devices, or changes IPs. Without this linking, analysts would have significant blind spots. Consider this example: An employee, Gary Hardin, is planning to resign and to sell confidential source code to a competitor. He logs into the network via his workstation using his Windows domain ID, *ghardin*. Later, he remotely accesses a Unix database server with a shared admin account *sa*. He then uses a different laptop to access GitHub with the *eng* account and pull down source code, saving it to a USB thumb drive. This session includes three different machines and services, different IP addresses, and multiple and semantically-unrelated account credentials. It is difficult to connect these simply by looking at the related event logs – the attack won't be visible.

In contrast, the session object contains each of these events, and others as well. They are all tied to Gary's identity, and arranged by time. The session contains metadata such as "user switched accounts," baseline data such as "user saved an unusual number of files to his USB drive," and external context data such as "Symantec endpoint DLP says these files are source code." (*figure 1*) As you can see, the session object is quite powerful.

It contains not only the information needed for an incident investigation, but also queryable metadata. For example, an Exabeam analyst can search for "all sessions where someone accessed a server for the first time and also came in from the VPN from a country for the first time and also generated alerts from our DLP system." This doesn't require knowledge of the underlying SIEM search language, nor understanding of the event layouts. As a result, session data objects enable questions and analytics at a different level.

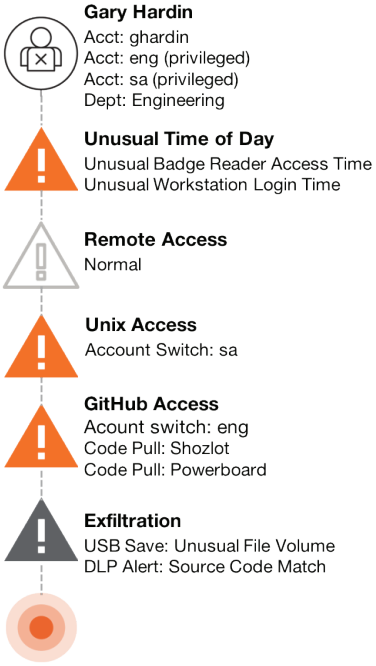

Figure 1. ***Data in the Session Object***

HOW EXABEAM BUILDS SESSIONS

The Exabeam Session Engine constructs session objects by processing existing events in a company's SIEM. It looks for events that initiate a session, such as Kerberos or NTLM logons, VPN events, application logons, physical badge-ins, etc. It then connects the events that occur for that user, following session initiation. A key point is that in the process on linking these events, Exabeam tracks state changes. In the example above, Exabeam would track Gary Hardin's remote logon to the unix server and the GitHub service, linking source IP, target IP, and both the Windows domain account and GitHub *eng* account. All would be tied to Gary Hardin, and connected within the same session. A session is built until logoff, timeout, or other signals indicate session-end. We call this process Stateful User Tracking™. The process is repeated, for every historical session for every user, until Exabeam has created a repository of session objects for each user.

SESSIONS ENABLE BASELINES AND RISK SCORING

Once Exabeam has created the set of historical sessions for each user, its Behavioral Engine creates baselines of normal behavior for each user. These baselines are used to assess new activity for each user, i.e. to determine whether activities are anomalous (unusual) for each user and for his or her peers. For example, the baseline for an HR administrator might indicate that this user never accesses the source code repository, nor does anyone else in HR. If the HR admin logged into GitHub and began pulling source code, Exabeam would detect immediately that this is anomalous behavior, for the person and also relative to his peers.

The anomalous activity is then scored by Exabeam's risk engine, which is fed by a security research team. A session's risk score is the aggregate of the risk scores assigned to each activity within the session. Sessions over a defined threshold appear at the top of the Exabeam detection dashboard, and an analyst can drill into the session to see the timeline of linked activities, i.e. a visual display of the information and metadata within the session data object.

PEOPLE AND PROCESSES

It's important to note that the session, baseline, and risk-scoring engines work for both users ("people") and machines ("processes"). Exabeam can create sessions/baselines and evaluate behavior of machine processes as effectively as users. This is important for several reasons. First, many machines use system processes that run on a regular, scheduled basis and have passwords hard-coded into the process. Second, increased growth in both cloud services and Internet of Things means that system processes will have greater access rights and be larger targets for hackers.

EXAMPLE: INCIDENT INVESTIGATION

"Here's what happens when we had an incident prior to Exabeam. We'd get a high alert, and the analyst needs to understand what the rest of this user's activity looks like, as well as whether this alert behavior is normal for this person. So, the analyst goes to our Splunk system and enters a long complex query that runs against the last 100 days of logs. That runs for at least 5 hours, then it turns out there's a mistake, so the analyst needs to fix and re-run. Another 5 hours. Then we need to change parameters and pivot and re-run; another 5 hours. Best case, a week later we have an okay view of what else this user did and how unusual it might be.

With Exabeam, we get that in 5 seconds and the picture is complete. We know exactly how risky this behavior is."

- *VP Global Security, Financial Services Provider*

USING SESSION OBJECTS IN EXABEAM

The session data layer is fundamental to the Exabeam platform. It is accessed and shared by any applications that run on the platform, such as Exabeam User and Entity Behavior Analytics (UEBA), Threat Hunter, and Analytics for Ransomware (figure 2).

- ***UEBA*** – Exabeam UEBA uses sessions and baselines to detect risky activity and to notify analysts so that they may investigate.
- ***Threat Hunting*** – Exabeam Threat Hunter uses session metadata to enable proactive search for sessions that match any arbitrary combination of activities and attributes. For example, "show me all user sessions where the user came in via VPN from a country for the first time, accessed a system for the first time, then had at least one failed login. Now show me all the users who did that and work out of a U.S. office."
- ***Ransomware*** – Exabeam Analytics for Ransomware creates session objects for system processes as well as users, to build baselines of normal behavior and to check for the early signs of a ransomware attack.
- ***Custom*** – Customer-internal applications, such as homegrown reporting or provisioning systems, can use Exabeam to support audits, take workflow actions, etc.
- ***SIEM*** – Exabeam has two-way integration with most of the leading SIEM products, and can send risk scores and session data back into existing SIEM workflows.
- ***Third Party*** – Non-Exabeam products such as adaptive authentication or remediation solutions can query Exabeam session objects, via REST API, to understand if a particular user requires additional policy controls, such as additional authentication steps. In addition, Exabeam session objects can provide greater context to alerts generated in third-party security control systems, as discussed in the next section.

ENHANCING OTHER SECURITY PRODUCTS

Third-party products, for example Symantec's Endpoint Protection or Blue Coat's Web Proxy products, often generate alerts related to a particular IP address or machine name. The IP or machine has violated some policy and the product generates an alert. However, the security analyst often does not have the context necessary to handle the alert effectively. For example, if Symantec EPP generates alert ID "1420341" an analyst may not be able to tie this to a specific user, and will almost certainly not understand how that alert relates to that user's overall activity and risk. With Exabeam in place, the analyst simply types the Symantec alert ID into the Exabeam search box, and Exabeam will respond instantly with the user, Keesha Hart, linked to that alert and the entire session that contains it (figure 3).

The alert for Keesha is automatically placed in the context of a user and a session.

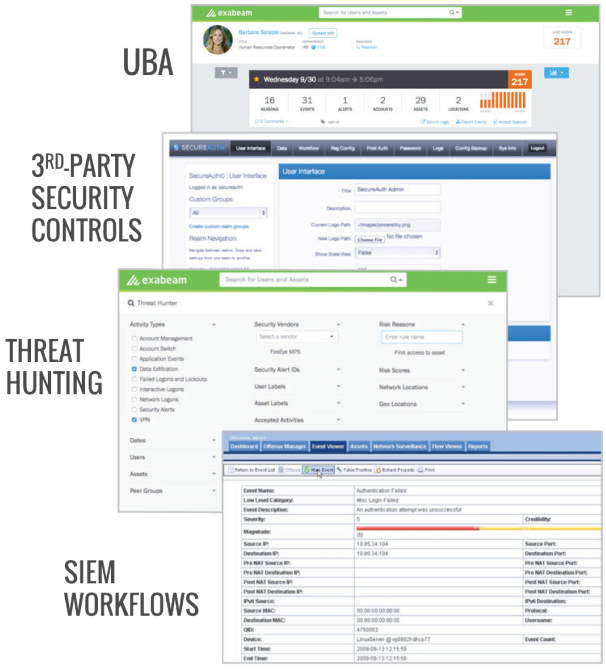

Figure 2. *Using Session Objects*

In a second example, we have a Blue Coat product generating an alert related to a machine on the network, perhaps trying to communicate with an external domain. Again, the analyst might not have the context to understand how this alert relates to a user and what that user has done. By typing the machine name (from Blue Coat) into the Exabeam search bar, the analyst instantly gets machine metadata, as well as a list of users who commonly access it (figure 4).

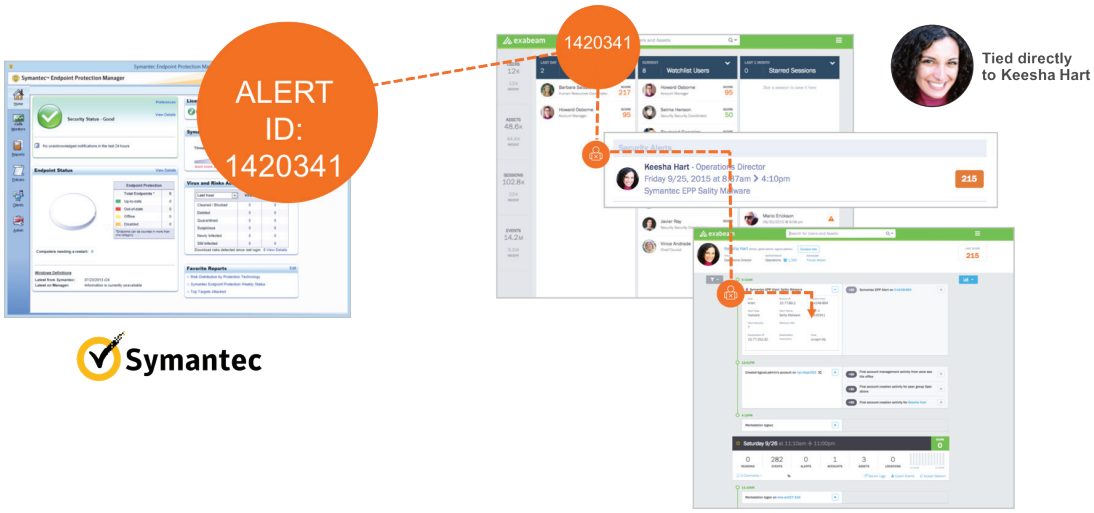

Figure 3. ***Adding Context to Symantec Alerts***

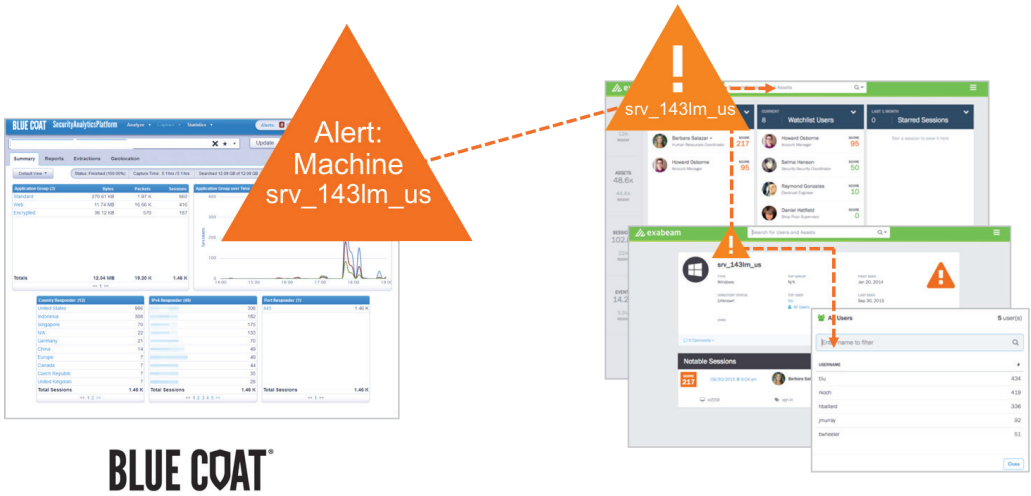

Figure 4. ***Adding Context to Blue Coat Alerts***

CONCLUSION

While the concept of a new data structure is somewhat esoteric, it is as transformative to the area of security intelligence as the data cube and star schema were to business intelligence. The session object enables analysts to ask new types of security questions, ones whose answers require data to be linked in ways that are impossible with other security products. The session data object in Exabeam enables a level of Stateful User Tracking™ that prevents malicious users from disappearing by switching accounts, machines, or IP addresses. The session is fundamental to the Exabeam platform, and feeds Exabeam UEBA, Threat Hunter, and external applications such as third-party solutions or custom reports. The Exabeam Timeline, including state changes, is not possible without this data object. Competitive products have no such data framework, and the lack of it limits these products to simple anomaly alerting.

Only Exabeam can link activity across identities and accounts, and therefore provide an end-to-end investigation automation solutions.

For more information, please visit www.exabeam.com, or send an email to **info@exabeam.com.**

ABOUT EXABEAM

Exabeam's security intelligence solution leverages existing log, endpoint, and other data to quickly detect advanced attacks, prioritize incidents and guide effective response. The company's Stateful User Tracking™ automates the work of security analysts by resolving individual security events and behavioral anomalies into a complete attack chain. This dramatically reduces response times and uncovers attack impacts that would otherwise go unseen. Built by seasoned security experts and enterprise IT veterans from Imperva, ArcSight and Sumo Logic, Exabeam is headquartered in San Mateo, California.