cyberhaven

12 Must-Have Requirements for Modern Insider Risk Management

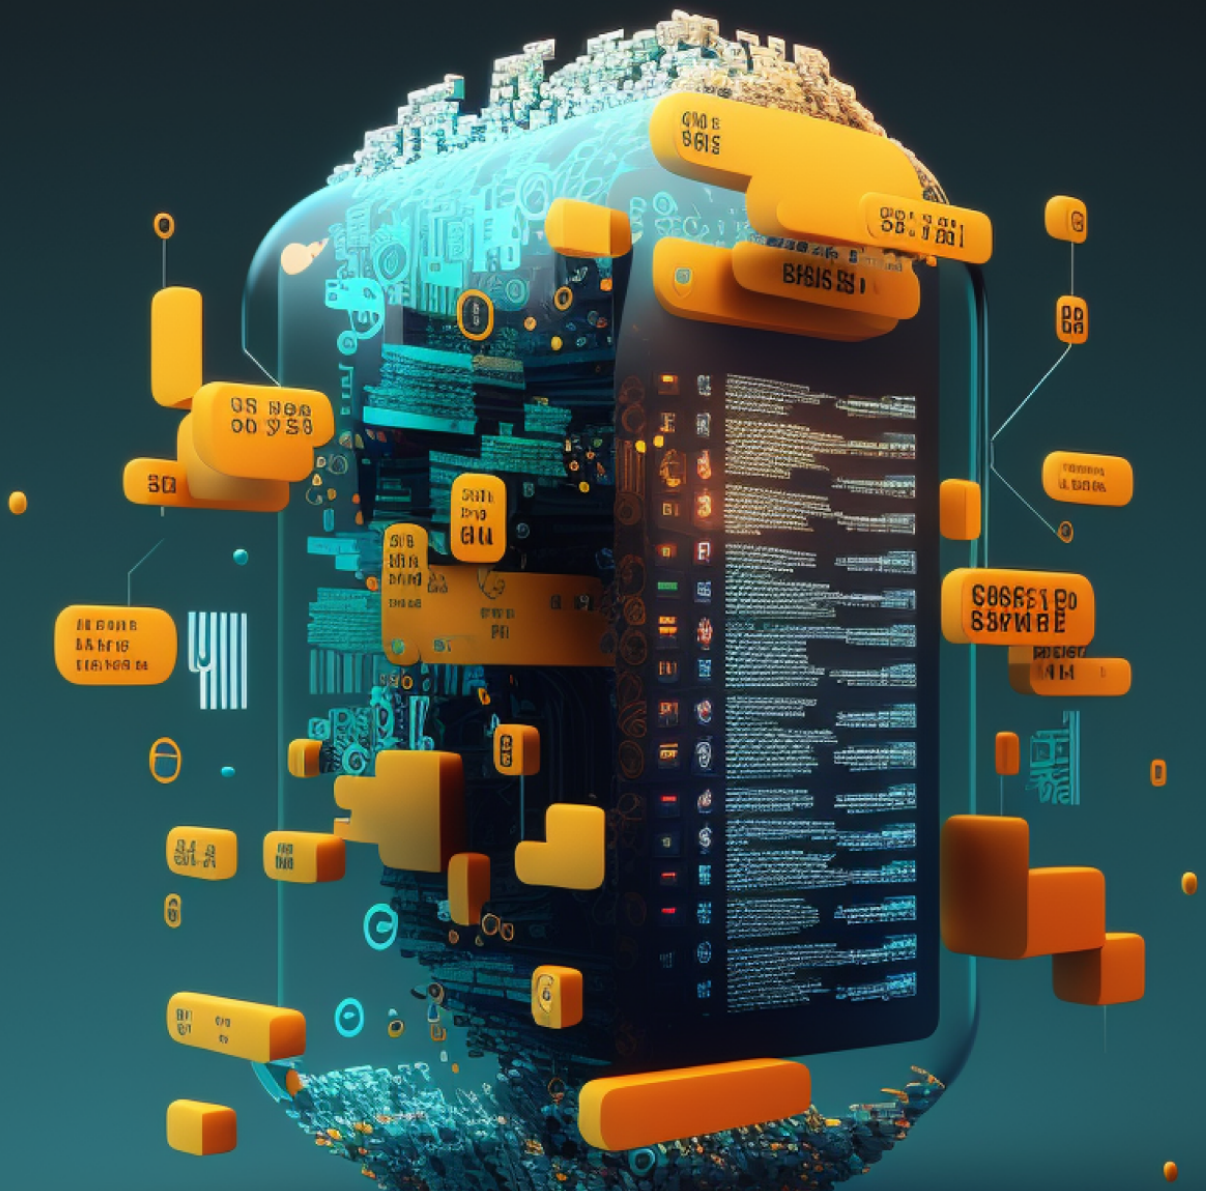

Today, the greatest risks to an enterprise's data and intellectual property typically come from inside the organization. As workers and applications become more distributed, an organization's most important secrets are constantly being put at risk by the everyday mistakes of busy employees as well as a variety of malicious insiders.

Yet protecting confidential or commercially valuable information from insider risks has proven to be a profound challenge. Organizations have traditionally been limited to two high-level approaches for managing insider risk, both with significant limitations. Data Loss Prevention (DLP) tools can analyze data but are limited to content that fits a standard pattern such as credit cards and Social Security Numbers, and struggle to classify intellectual property. More recently a wide range of behavioral analytics tools have been employed to identify unusual or suspicious behavior. However, such tools typically have no idea what the data actually is, don't block, and often require extensive manual investigation.

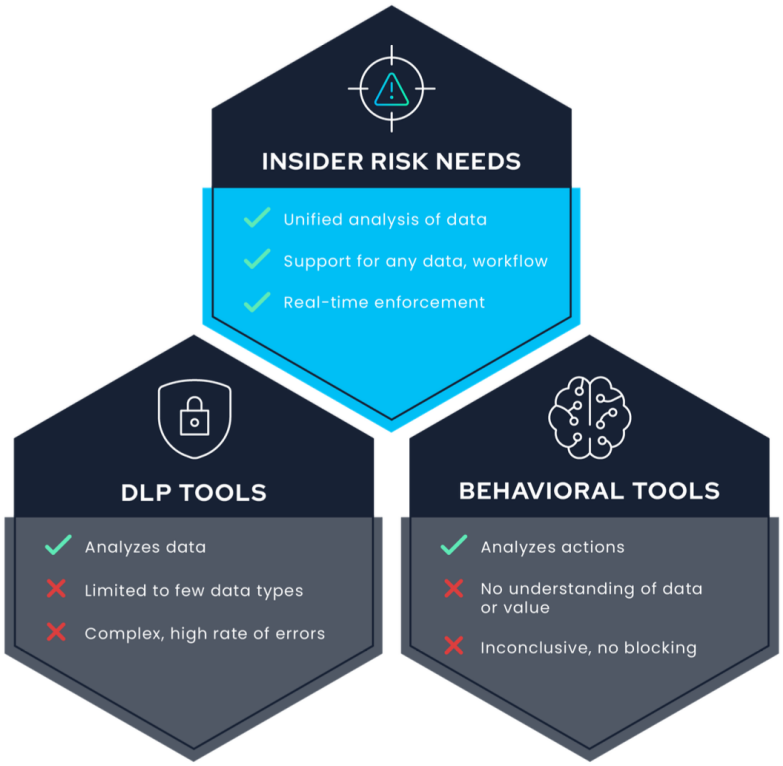

As enterprises and their risks continue to evolve, so must their requirements for mitigating insider risk. Data Detection and Response (DDR) platforms introduce a

new approach that lets security teams see and control virtually any type of data based on its risk to the business. These technologies enable security teams to see a live genealogy of virtually any enterprise data. They can see the full history of virtually any content including the user and application that created it through the many times it was shared, copied, or modified. Likewise, a DDR platform can automatically walk this data “family tree” to reveal all the “siblings” or “descendants” of the data. This means staff can quickly take a piece of data and see where the many copies and derivatives of that data are. Context is continuously maintained, and policies can proactively control how data is shared and with whom.

Let’s take a closer look at how these capabilities address some of the most important requirements in managing modern insider risks.

1 - Must Understand Both Data and Behavior Context

Most insider risk tools can either understand the data itself, or the behaviors and actions around the data, but not both. For example, a User and Entity Behavioral Analytics (UEBA) solution may see user behavioral anomalies, or a Network Detection and Response (NDR) solution may see anomalous network traffic between machines, yet neither has any idea if these events involve data that is important to the company. Did the user just download 500 MB of design files or did they download videos from the company picnic?

Conversely, DLP tools focus almost exclusively on the patterns within the content but miss all other contexts. This approach is highly prone to false positives and false negatives, and only understands a small portion of what can make data valuable or risky. Is the user moving sensitive PII that shouldn't be accessed or simply uploading personal information to sign up for healthcare benefits?

How DDR helps

In the same way a sentence needs a noun and a verb to make sense, security teams need to understand the data and the actions around it in order to truly understand risk. DDR platforms close the gap by providing an integrated view into both the data and the many actions related to it. In fact, this combined data-action context is a core trait of DDR. All enterprise data is automatically traced to its origin, and every subsequent action becomes a part of that data's lineage. This context of data (and its risk) is retained through every share, copy, or modification across users, devices, and applications. Armed with this information, security teams can take highly-informed actions based on what a piece of data is and what it means to the business.

2 - Must Protect Any Type of Data

Several trends have forced organizations to greatly expand the types of data that must be protected from insider risks. The shift to remote work has coincided with a significant increase in intellectual property (IP) theft. The rise of double-extortion techniques and leak sites both by malicious insiders and ransomware groups has put virtually any private enterprise data at risk. Virtually any data that is sensitive or even potentially embarrassing is fair game whether lost due to careless users, malicious insiders, or the increasing collaboration between malicious insiders and external threats. IP and trade secrets come in countless forms including company financials, M&A plans, legal proceedings, source code, design files, product plans, and countless others.

Traditional DLP tools simply don't detect or protect such diverse types of data. These tools are most effective for predictable forms of content that follow an alphanumeric pattern, known as a regular expression (RegEx), such as credit card numbers. In order to try to detect intellectual property, organizations must craft complex policies that attempt to classify data based on the presence of keywords. However, when directed at IP, these keyword and pattern-based policies produce numerous false positives. And they are not effective at classifying data that does not contain text–such as CAD product design files, images, videos, and other rich media.

How DDR helps

Recent advances in data tracing and lineage means that organizations can now include virtually any type of data in their insider risk management program. While not all data may fit signatures or RegEx patterns or even be text-based, all data does have a history. A DDR platform can automatically trace and analyze that history to see where data is from and who has interacted with it. This means policies can be applied to virtually any type of content. Protection is based on what data is important to the company, not what data is easy to detect.

3 - Must Protect Data Anywhere

The dispersion and spread of enterprise data has made organizations more productive, but has also increased the risk to those same assets. The rise of remote work combined with the shift to cloud applications and services means that important data is no longer tucked away in centralized databases in the enterprise. That data is also a constantly moving target as business applications foster collaboration and sharing between users and integrations foster sharing between applications. An insider risk program must be able to enable data to be used across all these locations and workflows without the organization losing visibility and control.

How DDR helps

Instead of tying individual tools to specific vectors or use cases (e.g. email, cloud, etc), organizations should focus on the ability to see and control complete workflows across the enterprise. This can include the flow of data between users, applications, fileshares, and other services. Technologies such as tagging or labeling have been used in this regard, but are often limited to certain types of files or documents and rely on the asset being properly tagged before it can be protected. Newer approaches based on data lineage can provide a far more automated approach that can be applied to virtually any type of data or workflow. In this model, security policies can automatically follow the asset itself even as it moves across the extended enterprise. Context and control are then defined at the enterprise level and are maintained across any number of steps whether from user-to- user, user-to-application, and application-to-application, whether locally or in the cloud.

4 - Must Be Able to Block/Enforce in Real Time

Many insider risk platforms lack the ability to block at all and only function as investigative tools. Protections triggered via third-party integrations are often too slow to prevent data from being lost. More importantly, most insider risk tools that rely on analytics are simply not conclusive enough for an enterprise to take action. This ultimately limits organizations to the tasks of documenting and cleaning up after insider loss events as opposed to preventing them.

How DDR helps

In order to actively control risk, organizations must be able to take action before damage is done. Ideally, such real-time enforcement should incorporate all available risk contexts for that data. Blocking should also be performed in-line and without dependence on other tools or integrations to ensure that protections are applied in real time to prevent a loss. On the other hand, it is important for organizations to have the option to employ more nuanced responses such as alerting the user to an impending policy violation providing the option to override. This can allow organizations to mitigate risk while striking an appropriate balance between prevention and productivity based on the situation.

5 - Must Prevent Loss Due to User Mistakes

According to Gartner, “The majority of insider risks are attributed to errors and carelessness; however, data theft and malicious activities are still observed.” Avoiding these user mistakes is extremely challenging both for users and security teams. Virtually every corporate and personal-use application is designed to make data sharing as easy as possible, creating countless opportunities for users to make mistakes. And since these accidental behaviors look like (and usually are) normal user behavior, they will not appear anomalous or unusual to risk monitoring or analytics tools.

How DDR helps

There are countless opportunities for users to make mistakes, and in order to truly mitigate these risks organizations must be able to establish and maintain fine-grained insight into precisely what data a user is working with and what they are doing with it. For example, to prevent data loss, you may need to recognize that a user is handling confidential financial data and is about to share with an email address or application outside of the company. Given the overlap between corporate and personal use applications (e.g. personal Gmail vs corporate) it is important that controls are extended down to the level of the user account. For example, a policy should take action if a user accidentally attempts to upload a file to a personal Google Drive account instead of the corporate account.

6 - Must Protect Encrypted or Obscured Data

Malicious insiders and attackers will naturally take steps to avoid detection. This could include moving data in a “low-and-slow” manner to avoid triggering volume-based rules, encrypting data or compressing it to a ZIP file to avoid inspection, removing tags, or altering the file extension of a piece of content to avoid content-specific rules. Once the data is obfuscated, tools today can’t “remember” what content was present before the obfuscation. Such techniques are highly effective against tools that either can only analyze content in its current form or only analyze aggregate user or network-level statistics.

Additionally, data is increasingly encrypted for completely valid reasons. While some security tools attempt to decrypt this encrypted content, the approach is increasingly unreliable, particularly when applications transmit data over a network. Applications increasingly implement certificate pinning as a security measure. This provides a more secure connection for the user but also will prevent security tools such as a network-based DLP from decrypting content for an application that doesn’t use enterprise-controlled keys.

How DDR helps

Organizations should consider controls that can operate without relying on the direct inspection of the content itself. Technologies such as data tracing and DDR can follow data and retain the context of what a piece of data is (as well as its risk) even when the content is encrypted or obscured. So if

an insider tries to hide a pilfered company roadmap by encrypting it and dropping it in a ZIP file, the security policy will still know what is in the payload and be able to control it appropriately.

7 - Must Protect All Copies and Derivatives of Data

Sensitive data often has a long life and can be modified, copied, and shared many times, and data only needs to be lost once in order to cause damage to the enterprise. While an original source file may be tagged and tightly controlled, those protections can be lost when a user copies sensitive data and pastes it into another file or application. Over time, there can be dozens of copies and derivatives of data that are completely invisible to traditional controls.

How DDR helps

Controlling all copies of sensitive data requires new types of security visibility. Instead of just enforcing how data is accessed, insider risk tools must be able to see and control the many ways that information spreads after it is accessed. Ideally, organizations are able to track information throughout the enterprise, through any number of copies and transformations independently from whatever format or file it's found in. In order to find risk tied to a derivative of a sensitive file, a solution may need to detect when a user copies/pastes sensitive data from one file or application to another.

For example, a user may copy data from a customer account in Salesforce. A solution should be able to see which customer account in Salesforce was affected and exactly where that data was pasted whether into a local file or an application such as a chat app. When this level of monitoring is performed automatically, organizations are better able to see the full scope of their data that is at risk and enforce controls to protect it and prevent unnecessary sprawl.

8 - Must Find Real Risk Not Just Anomalies

Many insider risk tools work by looking for deviations from standard behavior at the user or network level that could indicate a risk, building on metrics such as the volume of an activity or volume of data. This approach has a number of drawbacks that lead to both false positives and false negatives. First, users perform what appear to be anomalous behaviors with surprising frequency–what's known as a "long tail" in probability, creating many erroneous alerts. And behavioral tools typically don't know the context of the data and will only correlate behaviors over a set period of time.

For instance, they may trigger alerts whenever a user exports a large amount of data from Salesforce for valid reasons. An actual risk might involve the user downloading data from Salesforce and then a few weeks later, uploading it to a personal Google Drive. Without the ability to recognize that the upload contains the same data that was downloaded from Salesforce, the behavioral tool could miss the threat. This typically means organizations are forced to choose between setting thresholds high and missing risk or setting thresholds low and drowning their staff in investigative work.

How DDR helps

Reliable insider risk management requires multiple perspectives to see the full picture including the origin of the data, the content, and the full history of what users do with a piece of data as it moves through the organization. These complementary views allow security to see risk that would otherwise be missed while also providing corroboration from different perspectives. Such an ability to enrich detections across multiple contexts can also vastly reduce the overall volume of alerts by filtering out "informational" data and instead focusing on the combined risk to the enterprise. Likewise, these combined views increase the accuracy of detections, allowing security to enforce policy without worrying about inadvertently disrupting valued work. These contexts can also vastly accelerate incident response efforts by providing staff with the full lineage of the data in question, removing

the need for manual investigations. Ideally, solutions could also optionally capture screenshots and other data during a risky event to confirm the event and provide proof.

9 - Must Audit Users to See What Data They Actually Possess

Enterprise security policies often focus on controlling initial access to data. However, an organization's risk is defined by who actually has sensitive data, not how they accessed it. In today's open and collaborative workplaces and the cloud services that facilitate them, it is incredibly easy for an employee with access to a sensitive piece of information to share it with others who otherwise wouldn't be able to access it at the source.

It is critically important for organizations to know exactly what sensitive data a user has, particularly when the employee is about to leave the company. While many organizations will monitor departing employees for suspicious behavior such as downloading large amounts of data, these efforts would miss risks where the user was already in possession of sensitive data.

How DDR helps

To find any unseen risks, security teams should be able to audit the data that any user has on their devices or accounts. Teams can then compare the user's real-world possession of data to what should be expected based on the company's access policies. For example, an engineer who isn't authorized to have access to the quarterly forecast report yet has a local copy stored on a laptop.

10 - Must Control Data Coming into the Organization

While many organizations focus on the potential loss of sensitive data due to departing employees, they often overlook the risks of incoming employees bringing sensitive data from their previous employers. Introducing that data can expose their new company to legal liability, especially when the employee is bringing data from a competitor. It's important to look at data ingress not only data egress to understand this risk.

How DDR helps

By understanding the lineage of all data in the enterprise, staff can easily identify data that did not originate from within the company. Especially in an employee's first weeks at the company, data without a clear internal source or that originated externally such as downloaded from a personal Google Drive account, should be investigated. This can not only help organizations control the influx of unauthorized data into the enterprise, it can also help teams distinguish between internally created data such as source code written by internal engineers versus code imported from open-source projects.

11 - Must Train and Educate Users

To control insider risk, organizations must be able to proactively educate and instill secure work habits for their employees. Unfortunately, quarterly or annual security training lacks the regular reinforcement that is key to developing good habits, and users are always prone to making mistakes when they are busy or in a hurry.

How DDR helps

Instead of relying solely on periodic training and testing, organizations should incorporate real-time security coaching and reinforcement into users' natural workflows. For example, a user attempting to share sensitive data to a risky destination or via a risky app, can be warned of the violation and redirected to a company-approved option. Real-time coaching has been shown to reduce high-risk user behavior by more than 80% compared to periodic security training classes.

12 - Must Protect Data Without Violating Privacy

The invasive nature of many insider risk tools can put an organization in difficult positions when it comes to data privacy. Many insider risk tools require deep inspection of content on a user's machine, and that data may be uploaded to other locations for analysis. This can make it easy for organizations to inadvertently capture and store a user's personal data and violate privacy laws or the organization's internal privacy policies.

How DDR helps

The traditional trade-off between corporate security and user privacy can be achieved by understanding the data without inspecting the content itself. By tracking the full lineage and flow of data, security tools can retain a complete understanding of data and its risk without needing to access the data itself and potentially expose private information. This lets organizations actively manage and control their insider risk, without incurring new privacy risks.

Conclusions

Data lineage and DDR platforms are poised to transform how organizations view and control the risks to their data. For the first time, an enterprise can truly enable the flow of information between users and applications without losing control over the data itself. These capabilities enable security teams to tackle insider risk in a far more effective and practical way than was ever previously possible. Some of the thorniest insider risk challenges simply become irrelevant and protections can be extended to virtually any scenario.

To learn more about the Cyberhaven DDR platform and how it can help manage your insider risk, please contact us on our website at www.cyberhaven.com.