WHITE PAPER

A Unified Approach to Securing and Protecting IP

Securing intellectual property (IP) and confidential product data is a challenge for many corporations.

Major data breaches have occurred at large corporations in financial services, high tech, gaming, healthcare, telecom, media, and retail, as well as in the government and military. Perimeter-based security solutions and the current generation of security tools have not succeeded in mitigating these major attacks.

To successfully protect critical IP, companies must look to new technologies that support the creation of data-centric threat detection so they can uncover anomalous risky behavior and stop attacks before data breaches happen.

This report concludes with a brief overview of Perforce Helix Threat Detection.

PERFORCE

Contents

Cyber attacks and data breaches of large corporations continue to become more common. No longer confined to the retail segment, the world's biggest data breaches (those involving >30,000 records) now touch all major industries, including financial services, technology, gaming, healthcare, telecommunications and media, in addition to the military and government.[1]

1 http://www.informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks/

Although most security efforts, such as network defense and signature-based products, are focused on preventing outsiders from gaining access to internal file systems and databases containing sensitive information, defending organizations from insider attacks is becoming a growing concern. In a recent **Insider Threat Report from Vormetric**, 89 percent of global senior business managers and IT professionals polled felt that their organizations were now more at risk from an insider attack; 34 percent felt very or extremely vulnerable. Surprisingly, this report found that the most dangerous insiders are those who have privileged access, followed by contractors and business partners.

For example, Edward Snowden obtained administrative passwords from co-workers at the NSA, while misused or stolen user credentials were responsible for major data breaches at Target, Home Depot, Sony and Anthem. In the Anthem case, hackers were able to steal the credentials of five different technical employees during their attacks and accessed records of up to 80M U.S. account holders.

Solving the IP security problem calls for looking at it from a different perspective. It requires securing the data source in addition to the perimeter.

To better understand the issues pertaining to IP security, let's examine the key scenarios that put your data at risk.

Challenges to Securing IP

Today's product teams share and collaborate on a variety of assets and intellectual property. Their contributions might include source code, media (graphics, music and images), business docs (presentations, business plans and contracts), hardware design specs, environment artifacts and more. Unfortunately, most of these contributions are stored in unsecured repositories and network shares, while access is distributed across local and remote employees, some of whom might be temporary. Modern organizational practices frequently put your IP at risk.

Disparate Authentication and Access Tracking Methods

Software developers commonly use source code repositories for code collaboration, whereas marketers may use shared web portals for business docs and graphic designers prefer cloud services to share media assets. Unfortunately, each system may have its own authentication method for log-in and access control. And even if there is federated identity enabling authentication using a single system (such as LDAP) that enables access to multiple internal systems, access control and asset access tracking vary widely in terms of capabilities. Not all collaboration systems are flexible enough to support advanced security features such as two-factor authentication or access logs.

Multiple disparate systems make it difficult for security teams to determine and track who has access to specific files and content, and detect such unusual activity as sudden and large downloads from inactive projects that may point to potential data theft.

In addition, corporate IP is often spread across many different systems, presenting multiple targets for cyber attackers to use phishing or other social engineering methods that obtain employee credentials. Many of these collaboration systems don't provide detailed audit logs, making it nearly impossible for security teams to monitor access and protect mission-critical IP.

To successfully protect critical IP, organizations must have a deeper understanding of what's happening with their important data, so that they can see and understand what's truly at risk.

Geographically Dispersed and External Teams

Many, if not most, product teams have members or collaborators who work in different global locations. These members include not only internal company employees but also contractors and external service providers and business partners.

Project collaboration with diverse teams requires fine-grained access control. It is not sufficient for groups of users to get access to entire projects or repositories. Sometimes access to a specific

highly confidential asset needs to be restricted on a per-file basis to only a few select individuals and rigorously tracked with full audit logs. A collaboration tool that leverages IP address-specific access control rules to enforce access from only authorized locations, or to selectively grant access to users in different regions, is useful for partner companies and external service providers. For example, such a tool would allow limiting access to external collaborators to only a specific section of a repository or set of files.

Ensuring data security calls for encrypting data at rest as well as data in transit. Increasing levels of security may also be desirable, such as use of time-limited authentication, strong password enforcement policies, restrictions on password storage locations and authentication for all users of replicated remote repositories or depots.

Careless and Compromised Employees

Careless employees violate corporate policies by moving sensitive data to unprotected locations such as laptops or public cloud storage. Employees who move data to unsecure locations in order to ease their work processes create risk by unwittingly exposing this data to external hackers or data thieves who work within their companies, at supply chain partner companies or among contractors.

Even the best authentication and access control policies can't protect corporate IP from compromised accounts with privileged access rights. Compromised employees are those who unwittingly steal data for an external source; they are a very common source of data loss. This data exfiltration is typically a long-term process moving small amounts of data over an extended time. A collaboration tool that provides detailed audit logs is necessary to track access. However, a more pressing challenge is finding a security solution that can leverage log data to quickly detect small amounts of data exfiltration, without requiring time-consuming manual forensic work.

Leaving Employees

"Leaving" employees who take sensitive data with them are also a major problem. Studies consistently find that almost 60 percent of former employees have taken sensitive company data when they departed an organization—regardless of the reason they left. One Symantec study found that 56 percent of workers believe it is okay to take data with them and use it at a competitor. This includes not only customer lists but also the IP and trade secrets related to the programs with which these employees were involved.[2]

2 http://www.symantec.com/about/news/release/article.jsp?prid=20130206_01

Limitations of Current Collaboration Tools

Although some source code management (SCM) and content collaboration solutions have strong version control and collaboration features, not all have the integrated security features the enterprise requires.

Authentication and Access Control Restrictions

Some open-source version control solutions lack standard enterprise authentication support, such as Active Directory or LDAP. Some don't have built-in security in a repository and no concept of access control, users, or groups. To authenticate a user and to determine whether the user has access to a directory, some of these tools also require installing and configuring Apache, which adds extra complexity and requires additional resources. Apache access control is also typically limited to a whole repository, unless special path-based rules are defined, which may impact server performance.

Limited Access Tracking/Audit Logs

For audits, the following types of information must be provided: file change history, access history, new files added, and changes to files. Some open-source collaboration tools rely on different third-party statistics tools that provide some of this information with different levels of report completeness. Or additional steps are necessary that involve parsing Apache server logs to see actual access to an SCM system.

Also, some enterprise version control and collaboration tools provide limited audit logs, which identify only download and upload of source code data between clients and servers. To better identify anomalous behavior, it is necessary to use more detailed logs listing more granular actions, such as file syncs, diffs and examining content, which can be used to better identify anomalous behavior.

Perimeter-Based Security Solutions Lose Effectiveness

With the prevalence of BYOD, third-party apps, wireless transfer and cloud storage, it's impossible to inspect all the traffic coming and going within an organization. In addition, perimeter-based security solutions can't prevent insider attacks or attacks that result from compromised user credentials. In an RSA 2014

presentation[3], Jason Clark (Accuvant CSO) stated that 80 percent of security spend is going toward firewalls, IDS and anti-virus solutions, despite the fact that these are only effective 30 percent of the time. Clark notes that little security spend has been put toward user activity and data protection and that most organizations are immature in understanding user and data behavior. He also notes that the ultimate goal is to ensure security with a risk-based and data-centric approach that enables organizations to identify threats to their critical data determined by risk-based scoring and threat modeling. According to Clark, organizations need to identify the data/assets they most want to protect, assess threats in a consistent way and develop meaningful risk metrics.

Extreme Visibility is Key to Stopping IP Data Exfiltration

The biggest challenge with these ever-changing advanced new security threats is proactively identifying them when they occur. What typically happens is that attacks are discovered months after significant data breaches have already occurred. To successfully protect critical IP in this new threat environment, organizations need to adopt a data-centric approach to identifying suspicious activities and potential threats. The key to detecting risk is to quickly identify the users and projects associated with this anomalous behavior and then proactively address theft very early in the process.

Effective security tools must also be able to overcome security alert noise (sometimes thousands of false alerts per day) and accurately identify real threats in a prioritized manner.

Protecting IP with Perforce Helix

Perforce Helix is an SCM and content collaboration platform that enables users to store any type of assets for any number of users. It puts no limits on file size. The platform includes industry-leading version control capabilities, built-in code collaboration facilities and a complete Git management solution for the enterprise. With its support for a wide variety of workflows, it enables all members of a product development team to be productive and efficient.

To secure and protect your content, Helix provides comprehensive IP protection and security capabilities:

- Flexible authentication
- Strong password security features
- Fine-grained authorization/access control
- Detailed logs for high-resolution visibility
- Ability to work with encrypted file systems
- Network security support
- Secure replication
- Behavioral analytics-based threat detection

Flexible Authentication

Helix supports native authentication using Active Directory or LDAP, as well as external triggers, which enable customers to implement their own custom authentication methods. Other authentication methods, such as two-factor authentication, can be easily implemented in this manner.

Strong Password Security Features

Users can also be authenticated using their Perforce group accounts, which provide password strengthening options including minimum length, maximum login attempts, password reset upon log-in and password expiration time frames. Helix offers several security levels that control password policies The highest security (level 4) requires the use of time-limited authentication tickets, enforces strong password policies, and prevents the use of passwords stored in configuration files, the environment, or the Windows registry. Service users (not standard users or operator users) are intended for inter-server communication and can issue a limited set of commands and allow replicated and multi-server environments to authenticate themselves.

Fine-Grained Authorization/Access Control

With Helix, it is possible to restrict access by IP address, user and group. Actions can also be restricted by successive access levels: list metadata, file read, file open and edit, write (edit, delete or add plus additional commands), branch, review, admin and super. These access levels can be specified in a protections table.

Access controls can be applied to a repository, branch, directory or individual file. IP-address specific or subnet-specific (CIDR) access control rules can be used to enforce access to only users

3 http://www.rsaconference.com/writable/presentations/file_upload/dsp-w03-castles-in-the-air-data-protection-in-the-consumer-age.pdf

in different geographical regions or authorized locations. This capability can be used to restrict external collaborators based on their network IP address to only a specific section of a repository or set of files.

Detailed Logs for High-Resolution Visibility

The Helix Versioning Engine's auditing feature enables the creation of a server log that tracks individual user access to all digital assets stored in the Helix Versioning Engine with very fine granularity. Detailed audit logs are useful in determining who accessed which assets and when. In addition, each working file has an md5 digest file that is stored in the server, and the md5 checksum can be compared to reveal any tampering of files. Very detailed logs are necessary for effective threat detection, as described in a later section of this paper.

Ability to Work with Encrypted File Systems

Repository data stored within the Helix Versioning Engine can be secured by encrypting the file system volumes that the repository data is stored on. If the server store is encrypted, then user protection is used when the client workspace is on an encrypted file system as well. Helix works with encrypted file systems that are supported by industry-standard operating systems.

Network Security Support

To encrypt data that is transported over a network or over the Internet, companies can use ssh or VPN software. Helix supports enabling secure SSL connections among different Helix Versioning Engine components, including the Helix Broker, Proxy, and command-line client.

Secure Replication

For remote offices, external contractors and third parties, replica servers can be implemented to facilitate replication. For both performance and security reasons, it can be advantageous to filter what is sent to replica servers. For example, it is easy to configure a replica so that only some revisions of a file are transferred to the replica. Data transferred between the master and replica servers is also secured via SSL.

Behavioral Analytics-Based Threat Detection

Helix Threat Detection, a component of Helix, applies behavioral analytics models to user interactions with source code, product designs and related assets managed in the Helix Versioning Engine. It detects attack events in real time, alerts security teams and quickly generates actionable reports that detail anomalous, high-risk behavior.

A threat detection connector automatically passes user activity logs to the Threat Detection Engine for analysis. It applies machine learning to track relationships between users and projects and define normal baseline behavior patterns. Through a series of algorithms, the analytics engine then surfaces anomalous activities and applies a risk score based on the anomalous nature of the event, the importance of the project and the riskiness of the user. The result is an accurate and prioritized list of threats.

More Detailed Logs Improve Threat Detection

Unlike other collaboration tools providers whose log data identifies only download and upload of source code data between clients and servers ("receive-pack" and "upload-pack"), the Helix Versioning Engine logs record all actions that transfer content between servers and clients. The Threat Detection Engine correlates and analyzes: login and logout, project and file access (folder, file, path, etc.), amount of data moved or synchronized (get, commit, sync, etc.), timestamp and user data. With this greater amount and better fidelity of log data, threats can be detected more quickly and accurately than otherwise possible. Other repository-related actions within the Helix Server logs (annotation, diff, integration, etc.) are also visible to the Threat Detection Engine, and this enables faster convergence and better-quality threat detection models.

Detects Anomalies Other Security Solutions Can't

Helix Threat Detection finds anomalies by:

- Comparing access patterns, data usage patterns and data movement patterns against historic behavior
- Determining similar user patterns across the Helix environment and comparing behavioral patterns between users and groups of users
- Detecting dissimilar patterns among members of the same project group or job role
- Comparing individuals against the entire user group

These anomalies are leading indicators of threat activity; they also require no foreknowledge or configuration to detect. Using a weighted anomaly approach in combination with machine learning effectively minimizes and, over time, reduces the noise and false positives that plague security teams.

Types of Threats That Can Be Detected

Helix Threat Detection can surface a wide range of threat scenarios, including:

- Compromised, careless and departing employees who download large amounts of data from sensitive projects
- Insiders who slowly take small amounts of data over a long period of time
- Machines compromised by stealth malware that are siphoning data
- Outside or advanced persistent threat (APT)

Conclusion

Today's cyber breach headlines underscore the sophistication of attacks and the failure of major organizations to protect their sensitive IP. Traditional perimeter-based security and the current generation of security tools have not succeeded in preventing insider attacks or attacks that use compromised credentials. Security teams are simply too overwhelmed by the volume of security alert noise and the dearth of effective security solutions that can detect slow and small data leaks or large data breaches quickly enough to stop data theft.

A Unified Approach to Securing and Protecting IP

Helix offers integrated IP protection capabilities that include flexible authentication, fine-grained authorization, network security and secure replication. By leveraging the most detailed and complete audit logs and advanced behavioral analytics models available in the industry, Helix Threat Detection can quickly surface hard-to-detect insider attacks, compromised accounts, botnets that are siphoning data, as well as advanced persistent cyber attacks.

Unique Advantages of Helix Threat Detection

- **Comprehensive.** The Threat Detection platform collects extensive log data for analysis, offering unparalleled visibility into all types of potential threats.
- **Advanced Analytics.** Unlike other solutions, Helix Threat Detection leverages advanced math to provide probabilistic risk scores (no manual setting of thresholds).
- **Accuracy.** Analytics tuned specifically for the Helix platform detect anomalous behavior and automatically prioritize risk and threat, minimizing noise and false positives.
- **Simplicity.** Intuitive graphical interface and plain-language reporting, linked to analytics and machine learning, enable IT teams to successfully protect valuable IP assets.
- **Fast and Current.** Real-time data collection, analytics, and reporting, coupled with automated batch processing, ensure ongoing visibility of threats as they develop.
- **Agentless.** All threat-detection results can be delivered without agents. End-point agents are required for only post-detection forensics.
- **Built for the Enterprise.** Helix Threat Detection is a fast, flexible and robust security solution that scales in large enterprise environments. It processes tens of millions of events collected from tens of thousands of users to generate a prioritized list of "risk stories" (threats) with fine-grained context (e.g., specific weeks or days and scope of anomalous behavior).

North America
Perforce Software Inc.
2320 Blanding Ave
Alameda, CA 94501
USA
Phone: +1 510.864.7400
info@perforce.com

Europe
Perforce Software UK Ltd.
West Forest Gate
Wellington Road
Wokingham
Berkshire RG40 2AT
UK
Phone: +44 (0) 845 345 0116
uk@perforce.com

Australia
Perforce Software Pty. Ltd.
Suite 3, Level 10
221 Miller Street
North Sydney
NSW 2060
AUSTRALIA
Phone: +61 (0)2 8912-4600
au@perforce.com

PERFORCE

perforce.com