Threat Detection through Big-Data Analytics: A Proactive Approach to Securing SoC IP Design Data

The Problem

As semiconductor engineering teams grow in size and become increasingly more distributed across multiple sites around the world, the proprietary and confidential System-on-a-Chip (SoC) and semiconductor intellectual property (IP) design data sets have grown to 10's or 100's of Gigabytes. Securing this data has become a huge challenge. It is no longer sufficient to secure data within the walls of a single company site. Data now must be secured within the collaborative teams that share that data across international boundaries.

Adding to the challenge is the fact that most of the current generation of security tools are perimeter-based solutions, focused on preventing outsiders from gaining access to internal company networks, file systems and databases containing sensitive, proprietary data. However, defending organizations from unwitting employee security breaches, compromised accounts, and insider attacks is becoming a growing concern.

Solving the complete IP security problem calls for technologies that protect source data from internal security weaknesses and provide early-warning alerts for risky and anomalous internal behavior. Security solutions must take into account the multi-site collaborative nature of today's design teams.

This paper outlines the characteristics and advantages of such a solution.

The Solution

To successfully protect IP design data from within, companies must look to technologies that support the concepts of both IP and file-level security and big data-centric threat detection. This requires two foundational elements:

1) IP-Level and File-Level Security through an IP management platform such as **Methodics ProjectIC**™ that provides IP-level permission assignments and tracking of design data according to IP parent/child relationships, IP branches, levels of hierarchy, and the tracking of who is using which data, where in the design, and in which geographic locations. Once the appropriate permissions are set, the IP management platform will pass this information to the underlying data management system such as **Perforce Helix,** which then assures the data is secured at repository, branch, directory and even the at an individual file level.

2) Data-Centric Threat Detection provided by **Perforce Helix Threat Detection** that offers big data behavioral analytics and identification of threats and risky behavior performed on Helix SCM repositories.

Security within the Design Hierarchy

Most modern SoCs are more than just a collection of IPs and custom development. There is typically a large number of IP blocks from a wide variety of sources, assembled into a multi-level hierarchy.

Additionally, the team of designers tasked with creating an SoC may be composed of many contributors who span multiple organizational units and geographical regions. This means that 10's, 100's, or even 1,000's of engineers might have access to proprietary design information at any given time. Restricting access to design data on a need-to-know basis is essential over concerns that range from trade secret protection to international regulations.

An IP management system is used to provide the IP-level security by selectively granting or denying who has access to which IP blocks. This is pre-determined by managers or admins through "permissions," which are set for either individual blocks or at various levels of the hierarchy as appropriate within the IP management software.

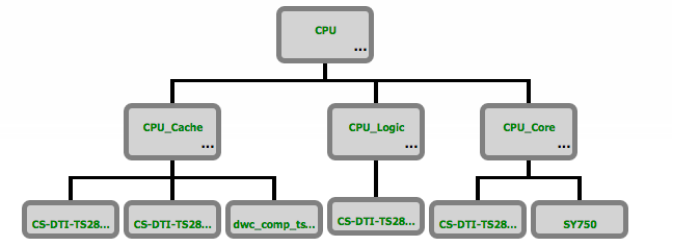

Fig. 1: Example of the hierarchy in a "CPU" subsystem

For example, in a typical hierarchical design, it may be desirable to allow certain designers to edit an IP (read/write permission), and other designers to only use that IP (read permission). As another example, one may want to grant edit access to all IPs in a library for an entire group of designers while allowing only a few of them to add or delete IPs. In a third example, entire groups of engineers may be denied access to certain IP and, in fact, may not even know it exists.

Below the IP-level, the content of IP blocks consists of a variety of data files. These files range from design data such as HDL code and mask layout to more abstract descriptive information like specifications and data sheets.

In all cases, access to the individual source files within each block is then granted or denied by the underlying design management (DM) tool - such as Perforce Helix - and the operating system of a designer's computer. This provides the file-level of security.

Security across Geographies

A second important facet of securing design data is taking into account where the data is stored and the way it is moved or copied between sites.

Fig. 2: Example of distributed multi-site data repositories

After being granted the appropriate permissions to access data, engineers must be able to access that data quickly as well as securely to meet tight project schedule demands. To achieve high-performance within today's design environment, IP and data management solutions must retain hierarchical IP-level and file-level security measures while also providing for:

Multi-Site Data Replication

In order to reduce the time taken to deliver files to user workspaces, ProjectIC supports the Perforce Helix Edge and Replica servers. These maintain a complete copy of the central repository at the local sites so that meta-data queries, commits and syncs can happen locally without the need for WAN activity. These replica servers are kept up to date in real time with no user intervention required and offer the same IP and file-level security already described.

De-centralized Data Management

Another method for reducing multi-site data latencies is to maintain the master Helix repository at a remote site, if that is where the majority of the development activity is taking place. ProjectIC allows the repository location to be defined on an IP basis so that workspace creation will query the local server and reduce WAN delays. Again, IP and file-level security is maintained.

Multi-site IP Caches

An important part of offering a high-performance multi-site solution is through the use of IP caches to maintain local read-only versions of popular IPs for consumption at remote sites. These are updated and propagated automatically as part of the Methodics IP release process. Users that need only read access to a particular IP can set that IP to "refer" in their workspace configuration and ProjectIC will manage the reference automatically. In many cases users only need write-access to a subset of the IPs in their workspace so an IP cache can represent significant disk space and IO bandwidth savings. These IP caches are managed according to the permission levels that have been set for each user.

Proactive Threat Detection

After securing permissions and access at both the IP and the file-level, the second critical element in a robust IP security solution is the ability to detect anomalous behavior and threats. This is where **Perforce Helix Threat Detection** comes in, offering a new approach to threat detection.

Helix applies advanced big data behavioral analytics to user activity to detect potential attack events, alert security teams, and quickly generate actionable reports that detail anomalous, high-risk behavior.

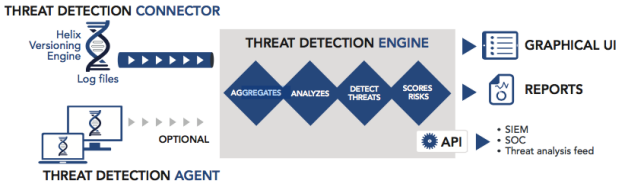

Fig. 3: Perforce Helix Threat Detection

User activity log files are ingested by the Helix Threat Detection Engine that 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. It applies analysis models (e.g., activity, statistical and clustering mathematics) to log data.

Once a threat is identified, a non-intrusive threat detection agent (endpoint sensor) can also be deployed to a laptop or desktop to capture all activity on the endpoint: file copies, cut and pastes, screen captures, printing, obfuscations and exfiltration.

A behavioral analytics engine applies machine learning to track relationships between users and projects, and defines normal baseline behavior patterns. Through a series of algorithms, the analytics engine 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, surfacing 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 cyber attacks

Once potential threats are identified, Helix Threat Detection makes it easy to quickly investigate the activity:

- Simple and clear presentation of “risk stories” over time
- Easily drill down and identify high-risk projects, people and dates and underlying data.
- Automatically generate executive summary reports

Fig. 4: An example of high-risk behavior identified by Helix Threat Detection

When a high-risk threat is identified, security teams quickly drill down to discover the machine and/or user involved, what actions created the risk and what data has been put at risk. This enhances the ability of security teams to detect and stop actual data theft.

Conclusion

Thru a combination of **Methodics ProjectIC** and **Perforce Helix Threat Detection**, companies have multiple ways of securing their proprietary IP data, at both the hierarchical block level and the file-level and, at the same time, identify and head off potential security breaches and theft of data.

For more information on advanced Threat Detection within the ProjectIC environment, or on Methodics in general, please email contact@methodics.com.

For more information on Perforce Helix Threat Detection, please email info@perforce.com.