CLAROTY

TEAM82

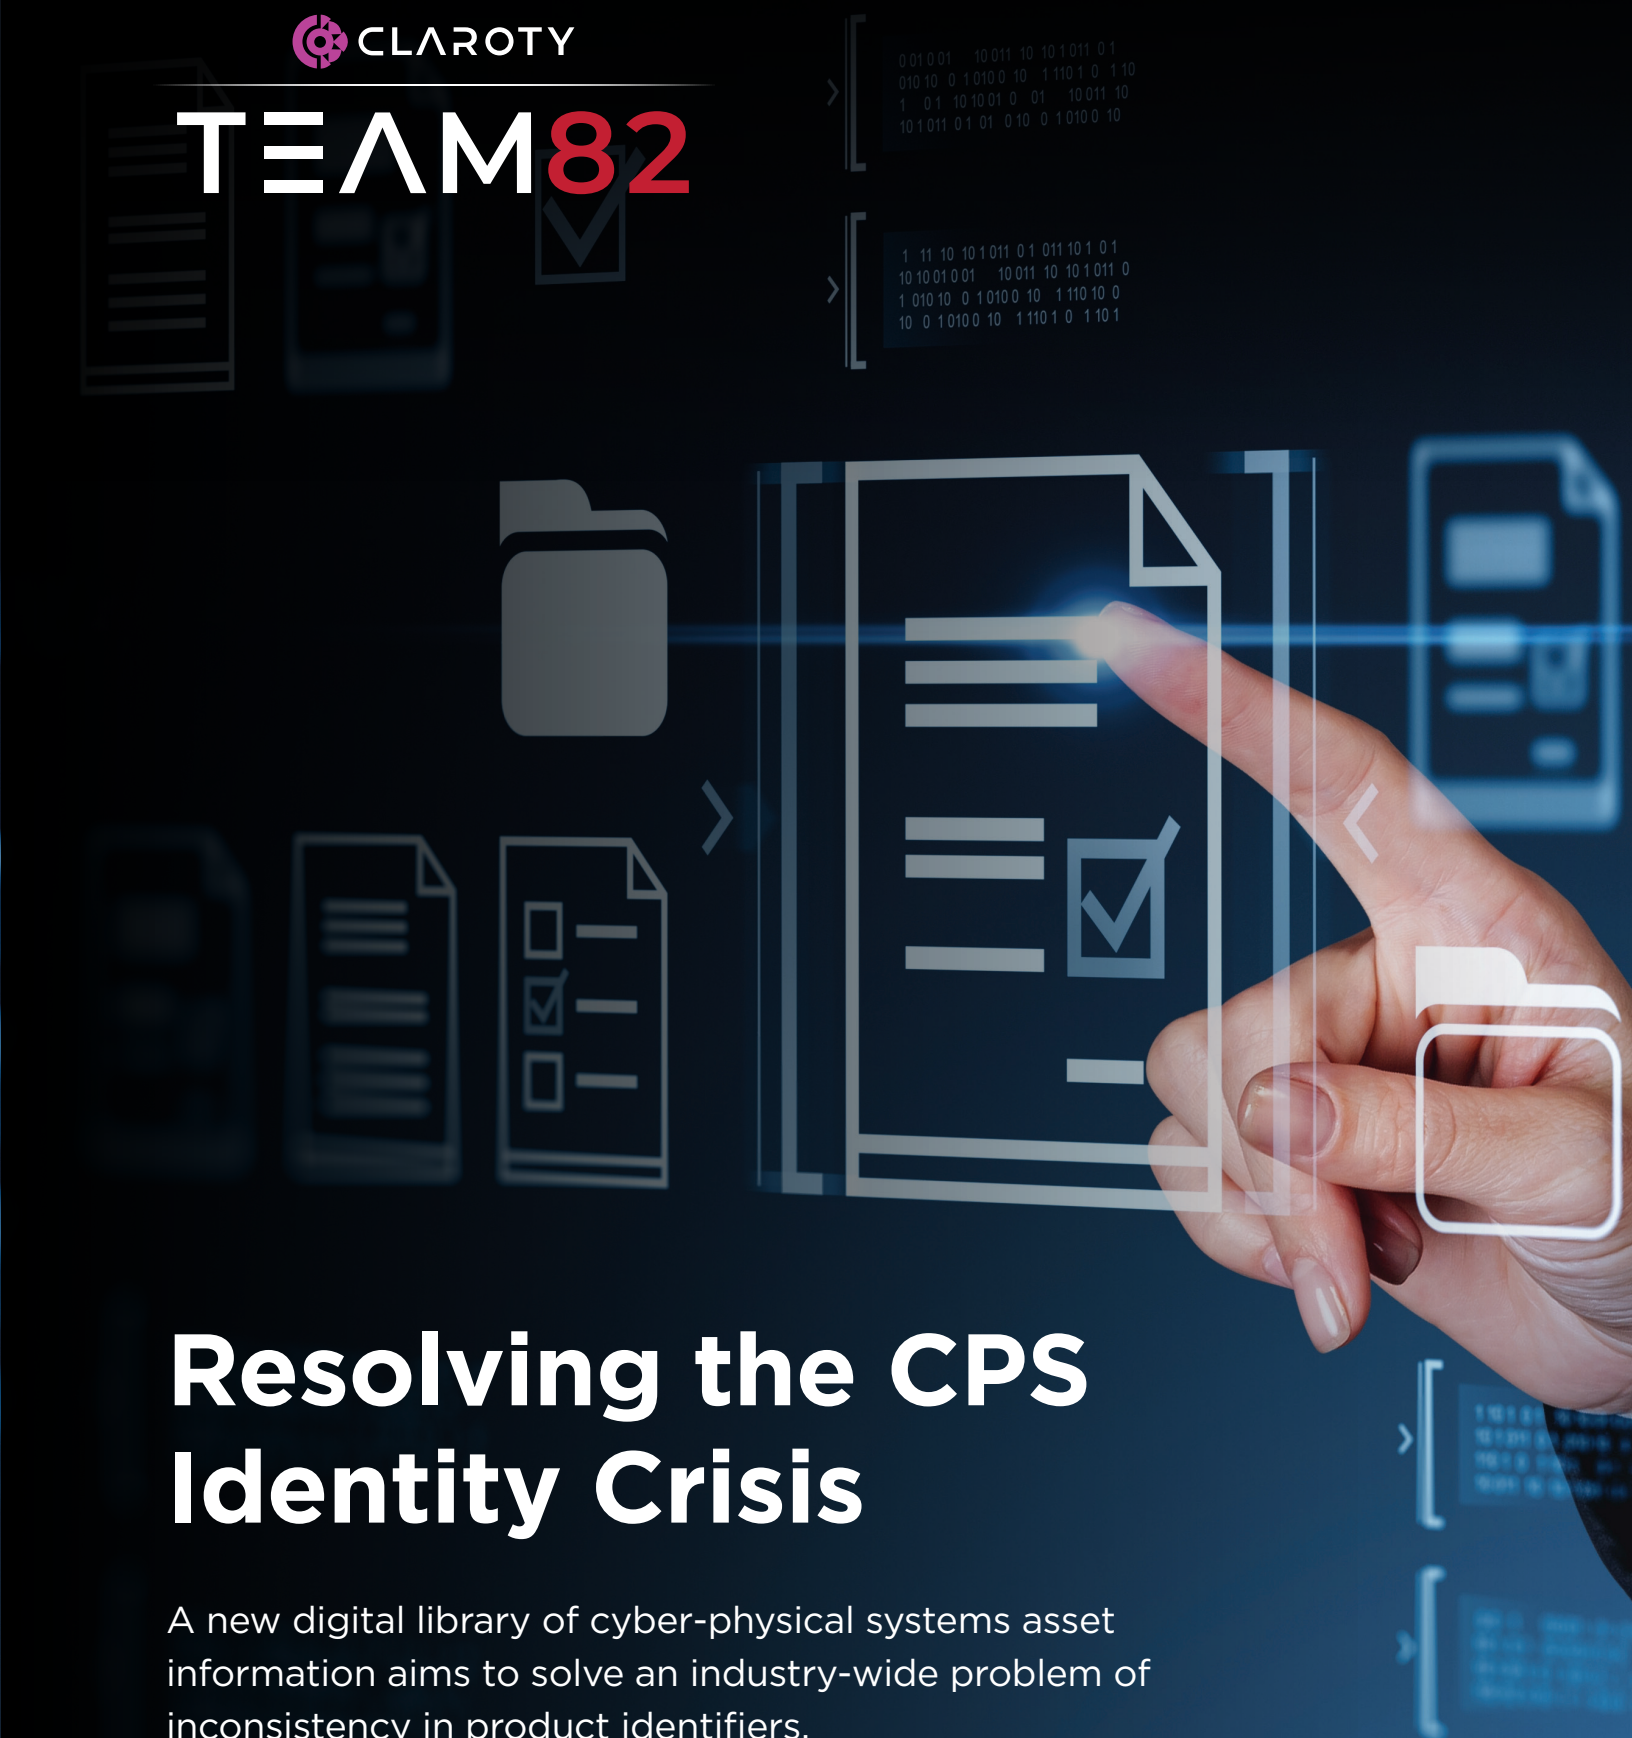

Resolving the CPS Identity Crisis

A new digital library of cyber-physical systems asset information aims to solve an industry-wide problem of inconsistency in product identifiers.

Introduction

Cyber-physical systems (CPS) assets have a myriad of variations and configurations due to the broad variety of customer and regional requirements they serve. While each medical device manufacturer (MDM) or original equipment manufacturer (OEM) has a well-structured way of precisely identifying each variant in the physical world, it is clear that their naming conventions in the cyber world were an afterthought.

With inconsistent aliases reported from the same asset based on how the information was collected, security teams and network administrators tasked with protecting CPS are at risk for partial matches when it comes to common vulnerability and exposure (CVE) attribution, impairing their ability to accurately assess the scope and remediation of cyber risk in their environments.

Our data-driven examination of the problem exposes the scale of this issue by demonstrating the variances in available product information. It also presents the need for a centralized repository that makes OEM and MDM default configurations available, as well as current vendor-approved patch levels, whether devices are shipped with default or known credentials, and helps identify other risks that impede last-mile remediation.

To illustrate: Programmable logic controllers (PLCs) from the same product family, for example, may be configured with different network interface cards (NICs), modules, or central processing units (CPUs) that may introduce their own sets of software and firmware vulnerabilities or insecure configurations. Active and passive traffic collection methods may also obtain different identifiers depending on the information source and protocol in use, often making it challenging to correlate this information.

This disparity can also then trickle down to CVE authorities, which rely on information provided by affected vendors to craft advisories; some internal device mapping may not be enriched by the MDM/OEM over time, and we've discovered that many advisories often lack the relevant context for a security team to fully understand which products, versions and configurations may be impacted by a newly reported vulnerability.

Data inconsistencies such as multiple or missing operating system names and versions, or diverse model number prefixes, will often require security and network teams to manually attempt to puzzle this information into a complete picture from resources provided by their respective vendors.

©Copyright Claroty Ltd. All rights reserved

Within our data set, for example, we can see some OEM product codes and model numbers that are never published within CVE advisories. The vendor's product catalog may list the product name along with numerous model numbers it supports, but again, much of this information is incomplete within CVE advisories, leaving security teams to infer from available information whether a device is vulnerable to a particular CVE. As a result, multiple vulnerabilities are possible for the same asset, depending on how devices are set up and whether vulnerabilities affect products only in certain configurations.

Key Findings

From a dataset of 17 million CPS assets, our research concluded that:

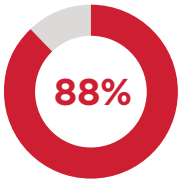

of CPS assets do not transmit an exact product code from data collection

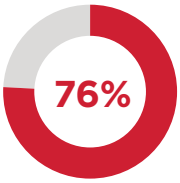

of devices transmit product codes that differ from the vendor's official record

To tackle the fragmentation of device data and work toward overall CPS risk reduction, we used our experience in protocol research combined with an automated AI-driven analysis of our dataset to fill in some of these gaps, remedy inconsistencies, and considerably improve asset visibility.

- For one large industrial automation vendor, our process for mapping accuracy showed a significant improvement, increasing product code identification from **4%** to **83%**
- With a newly matched product code, **56%** of devices received new or updated security recommendations for outdated firmware
- As a result, we were able to improve the accuracy of identifying vulnerabilities by **25%**

Users can quickly get stuck when it comes to already-complex processes associated with last-mile remediation of vulnerabilities. In the end, how can asset owners be sure the practical work of applying software patches, firmware updates, or compensating controls is complete, and risk is reduced?

This report examines these nuances of asset visibility and the impact on overall CPS protection, breaks down the core challenges to CPS asset identification, and then introduces an AI-driven approach to device information mapping that dramatically improves this aspect of asset visibility.

Let's first look at the core challenges facing security teams prompting this research.

©Copyright Claroty Ltd. All rights reserved

Inconsistent Naming Conventions

Product codes are an important standard for device identification, yet they are rarely available digitally over operational technology (OT) and Internet of Medical Things (IoMT) protocols. Within our data set, we found that **88%** of CPS assets do not transmit product codes; inconsistencies among other naming conventions such as model names and product numbers that don't match vendor catalogs also complicate the complete identification of assets in the field.

For example, our analysis uncovered that many CPS devices do not regularly transmit operating system (OS) names or versions, which may create a significant exposure management gap for organizations wishing to correlate this information with known vulnerabilities. Most security teams, therefore, cannot completely connect vulnerabilities to individual assets, creating blind spots, prolonged exposure to attacks, and incomplete remediation.

41% of CPS devices do not have an OS version available

24% of CPS devices do not have an OS name available

One Asset, Multiple Aliases

Industrial control systems (ICS) and legacy medical devices that we analyzed illustrate another challenge to accurate asset identification: a vulnerable device such as a controller, for example, could have multiple aliases depending on the protocol in use or collection method.

These aliases can take the form of the product family name (more than three-quarters of devices in our data set have multiple name variants), model number, or version, among other identifiers. Furthermore, the listing of the same affected controller may differ from what's listed in a National Vulnerability Database advisory, the vendor's own security advisory, or the OEM/MDM's product catalog.

Security teams are often left to infer the accuracy of a device's product family, model, or version from numerous sources simply because product codes are hardly ever transmitted and made available to security tools for correlation.

33% of models have multiple name variants depending on the protocol or integration in use

76% of devices have model names that differ from vendor names

©Copyright Claroty Ltd. All rights reserved

Introducing the CPS Library: Last-Mile Remediation Relief

Last-mile remediation is a significant hurdle for security teams because there is no central library or repository of certified OEM and MDM patch levels available to them to confirm whether their assets are indeed affected by a CVE. Right now, this requires a manual effort on the part of the security teams to access vendor resources and match their potentially affected products with the vendor's catalog. In an enterprise with potentially dozens of vendors of CPS assets, this is a time-consuming and costly process that must regularly be scheduled.

Without adequate visibility and device identification, organizations are fundamentally unable to manage exposures completely. Without a standardized way for vendors to identify products at the model and component level, the issue cascades down to inconsistencies in vulnerability advisories: partial asset-to-vulnerability matches exist, and CVE attribution is inconsistent.

Patching challenges, meanwhile, are severe for CPS, including OT, ICS, healthcare systems, and medical devices, and leave the window of exposure open for excessive periods of time for many organizations. Medical devices in the United States, for example, cannot be updated without the U.S. Food & Drug Administration's (FDA) review and approval of any updates that impact cybersecurity. In addition, MDMs require devices run only at certain patch levels, introducing another layer of complexity. This dynamic is similar for some of the complex industrial automation systems of assets running within CPS environments. Some OEMs, for example, require that control systems and other connected assets run only at certain patch levels.

The availability of a centralized library of necessary identifiers to properly match assets to vulnerabilities allows organizations to hurdle these last-mile remediation challenges.

©Copyright Claroty Ltd. All rights reserved

We analyzed a major OEM's device catalog and deployed our AI-driven techniques combined with our OT protocol knowledge to fill in these gaps and reconcile missing data to dramatically improve our mapping and identification accuracy.

The work produced a dramatic improvement in mapping accuracy of **83%** of its assets.

Improving Mapping Accuracy for One OEM

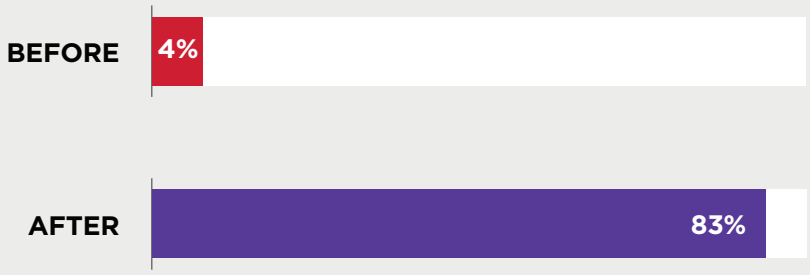

Only 4% of one OEM's devices were identified

Using our advanced techniques, 83% of the OEM's devices were identified

From there, we were able to provide a new or updated recommendation for outdated firmware in the environment in 56% of the devices we analyzed.

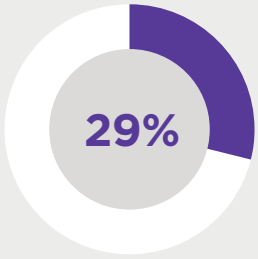

of devices received a new remediation recommendation for previously unknown firmware

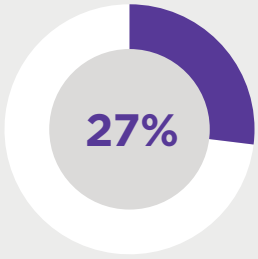

of devices that received an updated remediation recommendation after newer firmware versions for particular models were identified

Ultimately, this analysis and research resulted in a better vulnerability assessment of the vendor's devices.

15% False Negative Reduction

Previously unmatched vulnerabilities identified with improved mapping accuracy

10% False Positive Reduction

Improvements in device-to-vulnerability matching

©Copyright Claroty Ltd. All rights reserved

Advanced AI for CPS Product Code Mapping

Claroty's CPS Library tackled one of the most complex and real-world problems in CPS security: reliably mapping discovered assets to their authoritative product codes from a leading vendor. This is foundational for accurate asset, vulnerability, and exposure identification, which are fundamental to risk management strategies. We delivered a solution that goes beyond simple lookup: blending multi-source data collection, evidence graph modeling, robust statistical ensemble learning, and domain expert calibration—showcasing our deep data science expertise in real operational environments.

Building the Evidence Graph: Ground Truth Calibration

Our methodology starts by systematically collecting asset data from network traffic, vendor catalogs, configuration management records, and security advisories. These diverse inputs are interconnected in a comprehensive "evidence graph," which models the universe of possible relationships, including aliases, series, attributes, replacements, and legacy codes. Critically, we flag specific "ground truth" connections curated and validated by CPS experts. These connections become the calibration points that our AI agents learn from, retrain against, and use to adjust their mapping strategies in real time, guaranteeing that every inference is benchmarked against actual field data rather than cookbook generalities.

©Copyright Claroty Ltd. All rights reserved

Multi-Agent AI: Specialized CPS Intelligence

Within the CPS Library, Claroty deploys multiple CPS-AI agents. Imagine CPS experts with years of experience, working on the data, each focused on a unique slice of device identification:

NLP Models: Parse mixed-format, protocol-derived naming strings, software markers, and code fragments.

Statistical Reasoners: Apply confidence scoring and sophisticated statistical tests to weigh evidence and discriminate signal from noise, leveraging empirical relationships learned from ground truth corrections.

Domain-Guided CPS Rules: Implement field-proven logic, recognizing nuances such as hardware generations, firmware compatibility, and replacement cycles, ensuring no "lookalike" asset is misclassified.

All agents continually recalibrate as new ground truth data enters the evidence graph, a closed-loop system for real-world reliability and ongoing improvement.

Ensemble Methods: Data Science Expertise in Action

No single model suffices, especially in the noisy, often contradictory reality of vendor product code mapping. Here's how our ensemble system sets a new standard:

Weighted Voting & Robust Noise Calibration

To combine all the results, we employed a data science technique known as ensemble learning. Each AI agent or model in our architecture brings specialized strengths but also its own blind spots, whether that's struggling with ambiguous catalog entries, limited protocol context, or data sparsity for certain device types.

Recognizing this, our ensemble methodology goes beyond simple aggregation: it is engineered to actively compensate for the unique gaps in each agent's knowledge. By having multiple, diverse models weigh in, each using different perspectives, data sources, and inferential methods, we meaningfully reduce risk that would occur if any single model dominated the decision. This is particularly powerful in areas of "white noise," such as regions of incomplete or conflicting data, where ensemble consensus helps suppress random or directionless errors. The overall result is a more robust output: the ensemble fills in the gaps left by individual blind spots, while also minimizing error rates in ambiguous or noisy data environments, demonstrating the depth of our data science and practical expertise in managing real-world CPS complexity.

©Copyright Claroty Ltd. All rights reserved

Grounded, Expert-Validated Consensus

Every decision from the ensemble is ultimately anchored in our evidence graph's ground truth connections, making the process both statistically robust and operationally accountable. This "calibrated consensus" means our recommendations align with what real operators encounter in the field, not just theoretical optimality.

Verification and Continuous Improvement

Ambiguous or low-confidence mappings are flagged for human-in-the-loop expert review. Manual benchmarking, feedback cycles, and data enrichment further refine the AI models, ensuring continuous accuracy gains and defensibility in new environments. Statistical analysis of error distributions helps target the most impactful curation efforts for future releases.

©Copyright Claroty Ltd. All rights reserved

Wrapping Up

Security teams are significantly challenged with regard to CVE attribution and having a full understanding of the risk within their CPS environments because of inconsistencies in the way product information is delivered over various communication protocols. This creates gaps in reconciling missing or myriad naming conventions to vulnerabilities and firmware levels.

Our AI- and data-driven analysis of the problem exposes the scale of this issue by demonstrating the variances in available product information. We demonstrate how infrequently vendor product codes—an important standard data point for device identification—are made available digitally. Other asset information such as OS, version, patch levels, default configurations, and more must also be gathered, often manually from various OEM and MDM resources.

Security teams are left to wade through disparate model names and product configurations as they attempt to match CVE information to assets in their environments. The modularity of CPS assets also means that some vulnerabilities are only present in certain configurations, compounding the complexity of this exercise. These multiple aliases often force users to infer the accuracy of product information from numerous sources. CVE information, meanwhile, is also often incomplete because it works from the same inconsistent product information.

Our new approach to tackling the accuracy of asset mapping resulted in dramatic improvements. Our AI-driven techniques combined with our OT protocol experience improved the accuracy of asset identification for one major OEM product catalog we analyzed. We were able to provide a new recommendation for outdated firmware in the environment in 56% of the devices we analyzed, as well as reduce false negatives and positives in asset-to-vulnerability matching.

About Claroty

Claroty has redefined cyber-physical systems (CPS) protection with an unrivaled industry-centric platform built to secure mission-critical infrastructure. The Claroty Platform provides the deepest asset visibility and the broadest, built-for-CPS solution set in the market comprising exposure management, network protection, secure access, and threat detection – whether in the cloud with Claroty xDome or on-premise with Claroty Continuous Threat Detection (CTD). Backed by award-winning threat research and a breadth of technology alliances, The Claroty Platform enables organizations to effectively reduce CPS risk, with the fastest time-to-value and lower total cost of ownership. Claroty is deployed by hundreds of organizations at thousands of sites globally. The company is headquartered in New York City and has a presence in Europe, Asia-Pacific, and Latin America. To learn more, visit claroty.com.

About Team82

Team82, the research arm of cyber-physical systems (CPS) protection company Claroty, is an award-winning group of researchers known for threat research, OT and medical protocol analysis, and discovery and disclosure of industrial, healthcare, and commercial vulnerabilities. Committed to strengthening CPS cybersecurity and equipped with the industry's most extensive testing lab, the team works closely with leading vendors to evaluate the security of their products. As of October 2025, Team82 has discovered and disclosed more than 700 vulnerabilities. Learn more at claroty.com/team82.

©Copyright Claroty Ltd. All rights reserved