DARKTRACE

CRIMSON ECHO:

Understanding Chinese-Nexus Cyber Tradecraft

Through Behavioral Analysis

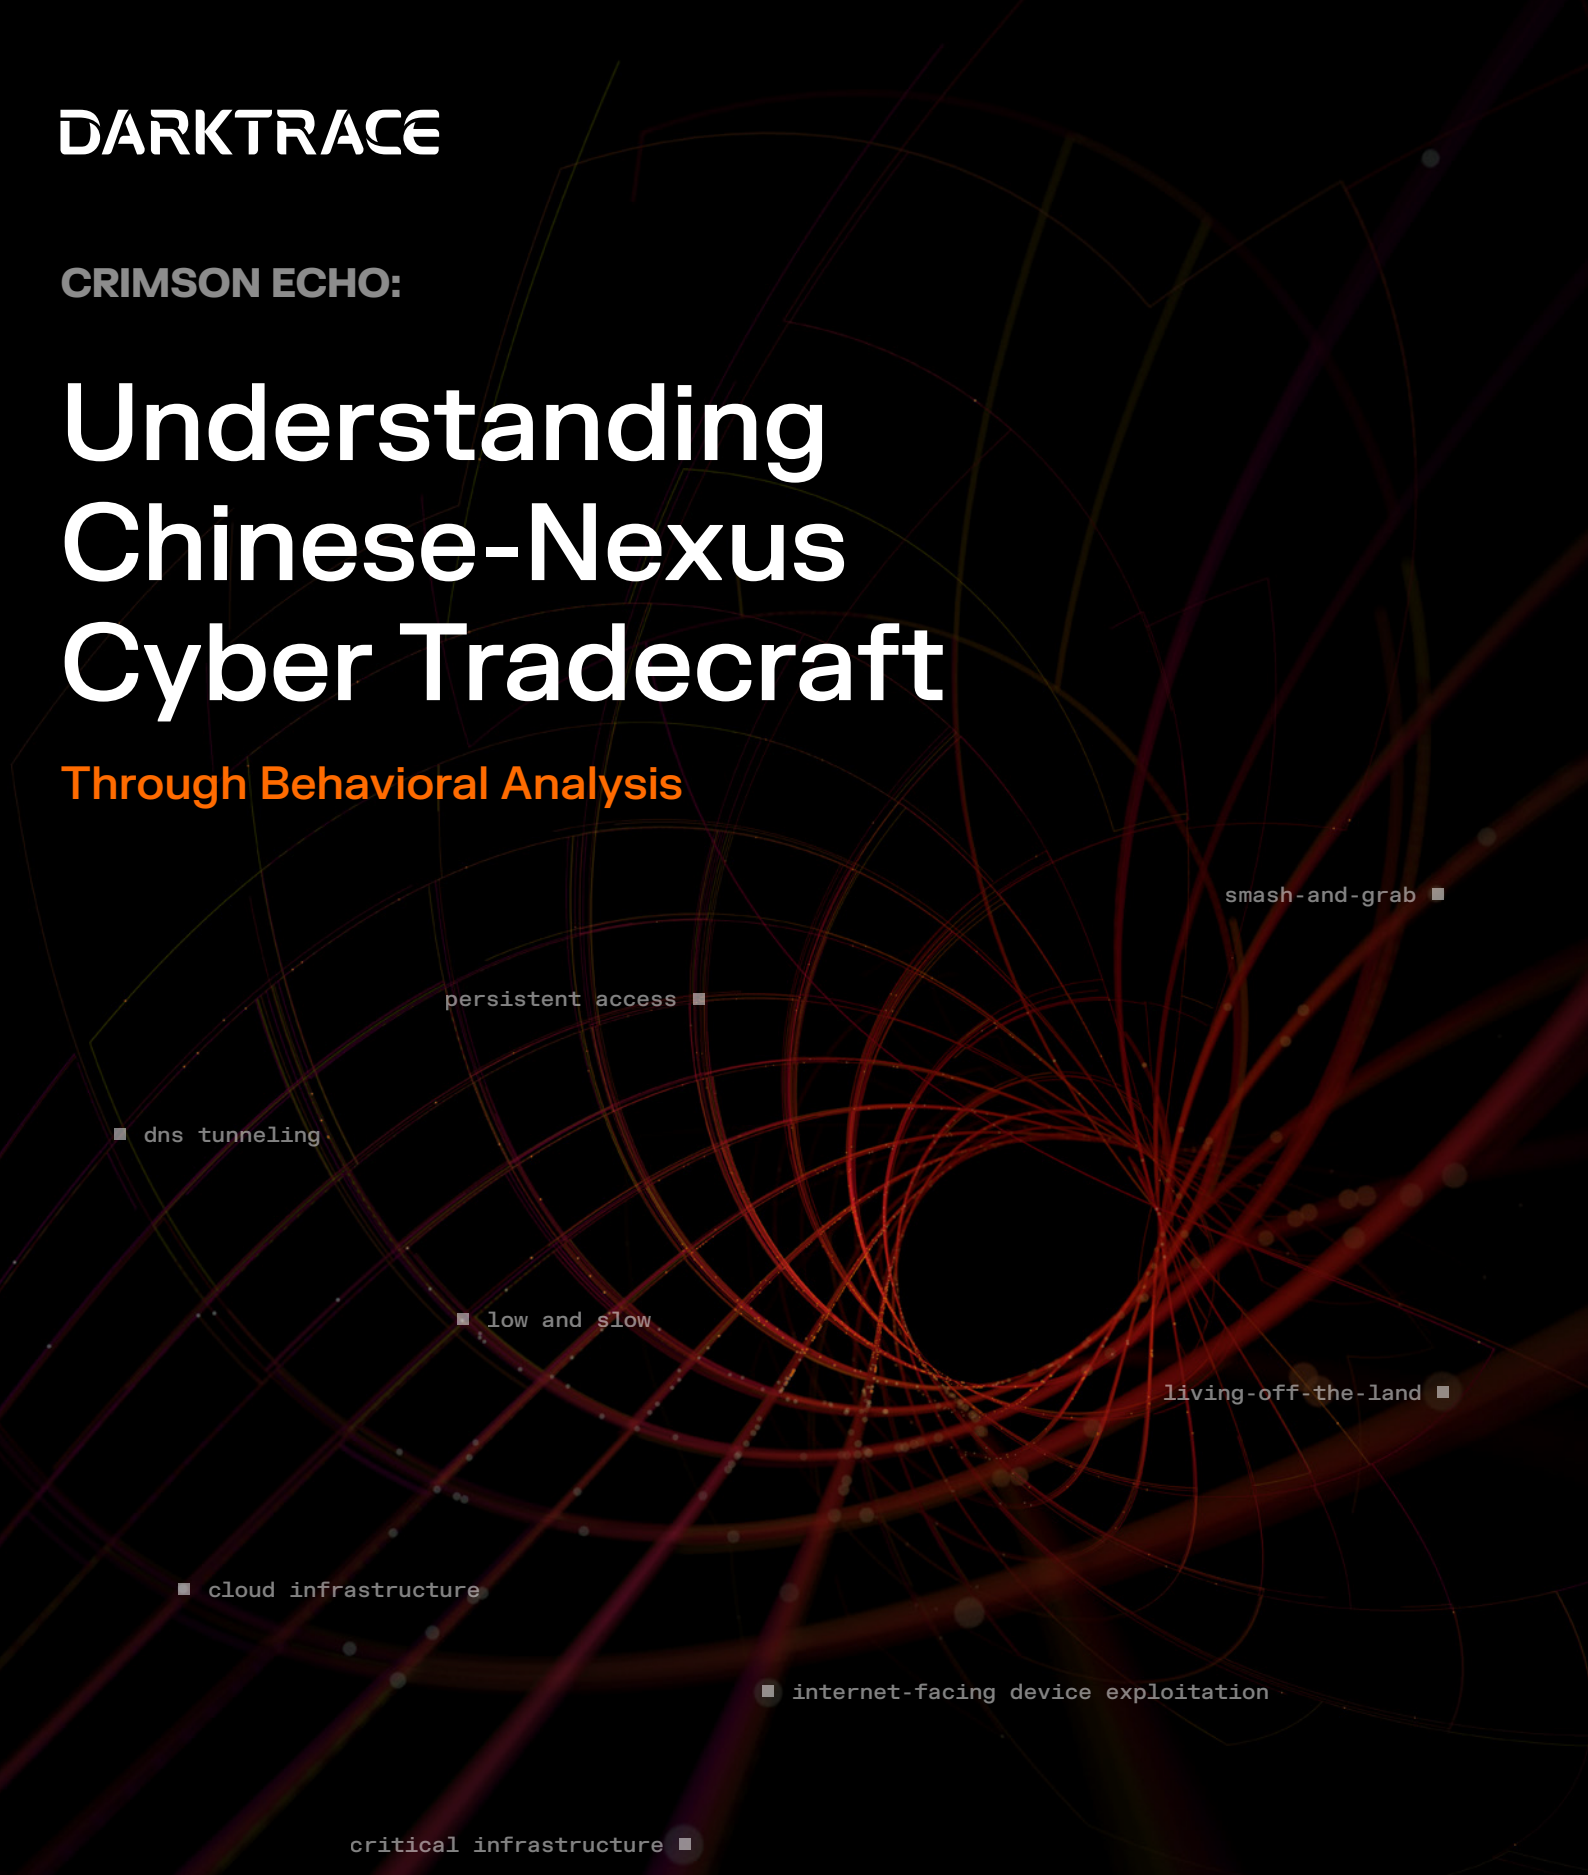

© 2026 Darktrace Holdings Limited. All rights reserved.

Contents

Introduction

In recent years, China's role in global digital and economic systems has expanded, accompanied by a steady increase in cyber activity linked to Chinese-nexus operators.

These actors routinely target public and private organizations whose networks may offer insights or access relevant to China's national objectives.

As this activity continues, organizations across multiple sectors should expect a sustained level of risk and ensure they are equipped to identify and manage intrusions associated with Chinese-nexus groups.

To better understand the risks from Chinese-nexus operators and provide insights to the community, Darktrace conducted a long-term review and threat hunt for evidence of cybersecurity incidents involving Chinese-nexus operators across its customer base over the last three years.

This review examined anomalous activity detected by Darktrace, irrespective of country and sector, from mid-July 2022 to September 2025, and featured multiple structured techniques for incident identification including hypothesis, signature, behavioral, and tactics, techniques, and procedures (TTP)-based threat hunts.

Limitations and Challenges

to Identifying Chinese-Nexus Operations

Cybersecurity professionals traditionally rely on a combination of tools, information streams, and analytic resources to combat threats posed by Chinese-nexus threats.

While cyber threat intelligence streams and signature-based detections are useful, overreliance on such methods can fall short when dealing with entities well-versed in evasive techniques. Security analysts therefore rely heavily on adversary profiles for particularly advanced and capable advanced persistent threats (APTs) such as those sponsored by national militaries or espionage agencies.

However, this reliance on APT profiling has resulted in an overproliferation of such resources. Organizations that provide these services frequently use different naming conventions resulting in redundant, yet slightly variable, understandings of the same groups.

Personnel and technical resource sharing between these threat groups renders additional challenges. The delineation of distinct organization boundaries between Chinese-nexus groups may not be entirely possible with a high degree of confidence because of the sharing of human and technical resources amongst Chinese security and military apparatuses. Cyber defenders now face the added challenge of analyzing potentially overlapping profiles of ever-increasing group delineations with similar TTPs to extract actionable insights for their Security Operations Centre (SOC).

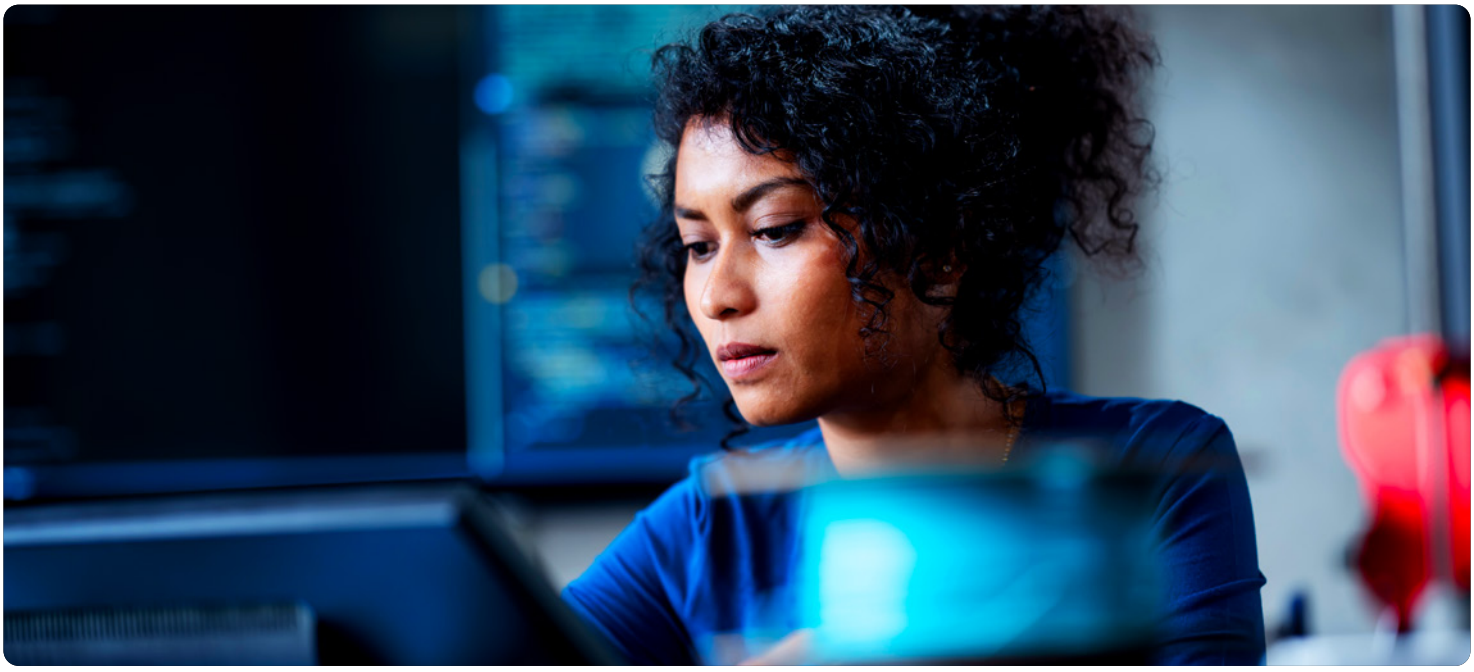

Darktrace's Approach

for Chinese-Nexus Threat Hunting

Rather than continuing to detail the characteristics, behaviors, tooling, and tactics of each Chinese-nexus actor, Darktrace opted to focus on identifying a general framework for Chinese-nexus threat hunting cyber defense, based on the Darktrace Cybersecurity Attribution Framework (Appendix 1).

This approach is intuitive for Darktrace not only given the agnostic nature of anomaly-based detection but also given the previously noted organizational structure Chinese-nexus operations detailed in the "Strategic Statecraft" appendix of this report. While the sharing of personnel and resources makes the demarcation of specific APT groupings more difficult, this reality alternatively allows analysts to aggregate activity at a higher level.

To accomplish this goal, the Darktrace Threat Research team conducted a two-phase research project. The first phase included a form of literature review: performing a long-term retrospective threat hunt over the past three years to identify as many cases of suspected Chinese-nexus activity within the customer base as possible. For the second phase, the Threat Research team then assessed the confidence of each identified case and analyzed the results to tease out general patterns in kill chain activity and TTPs.

This effort resulted in the identification of medium-to-high confidence cases of cyber intrusions involving Chinese-nexus threat actors. This resulting database of compromises yielded distinct patterns with relevance for cyber practitioners at all levels of seniority within an organization.

The report found that Chinese-nexus cyber operators during this time show broad TTP preferences, dwell time patterns, and targeting patterns based on strategic goals. The Darktrace Threat Research team used these insights to form, and continue to refine, detection mechanisms to identify future instances of potential Chinese-nexus activity within customer networks.

This report aims to contextualize the challenge of Chinese-nexus cyber activity in a broader framework given the overlapping and often ambiguous dividing lines between subsets of Chinese-nexus intrusions. In doing so, Darktrace hopes that SOC analysts and other cybersecurity professionals can utilize these insights to more easily and readily identify such instances of suspected Chinese-nexus activities within their environment and increase their cyber resilience against threats.

Project Director: Nathaniel Jones

Project Analyst Lead: Adam Potter

Project Participants: Emma Foulger, Paul Jennings, Nahisha Nobregas, Nicole Wong

Contributors: Nicole Carignan, Margaret Cunningham, Eugene Chua, Owen Finn, Samantha Gonzalez, Keanna Grelicha, Daniel Levy, Angel Lopez, Nathan Ly, Qing Hong Kwa, Will Palmer, Shawn Puckett, Tyler Rhea, Steven Sosa, Priya Thapa, Hyeongyung Yeom

Report Editors: Sarah Murphy, Ryan Traill

KEY INSIGHTS:

Chinese-nexus cyber operations are best understood as **continuous strategic planning**, not episodic campaigns.

Detection of short dwell time intrusions should not be interpreted as tradecraft failure but **deliberate operational choices.**

Western security models remain overly incident centric and **systemically undervalue persistent identity risk.**

China's cyber activity is **not just IP theft-based** but increasingly aligns with Belt and Road Initiative (BRI) dependencies and critical infrastructure leverage globally, with particular emphasis on the United States.

THE FOLLOWING PAGES INCLUDE:

A detailed **methodology** for the research

A detailed **report** on the medium and high confidence attribution cases, with clear regional and sector-level breakdowns.

An **assessment** of how observed activity aligns with and supports objectives tied to China's BRI, its flagship foreign policy initiative launched in 2013.

Appendix 1: Darktrace Cybersecurity Attribution Framework

Appendix 2: List of Indicators of Compromise

Appendix 3: Meta-Model Creation

Appendix 4: Anomaly Detection Examples Prior to CVE Disclosure

Appendix 5: Classifying Organizations as Critical Infrastructure

Methodology

To conduct a comprehensive long-term review of Chinese-nexus cyber operations across the customer base, the Darktrace Threat Research team set a clear outline of metrics, definitions, and planning to ensure relevancy. That also required an attribution framework that was focused on empirical evidence

This methodology applies a structured, multi-pillar approach to attribution, combining Darktrace's AI-driven behavioral detections with external intelligence, sector context, and comparative analysis of known threat actor TTPs. By evaluating evidence across infrastructure, tooling, TTPs, victimology, artefacts, and corroboration, analysts can build calibrated confidence levels.

This enables consistent, defensible attribution assessments even when threat actor boundaries are overlapping or ambiguous. Findings reflect alignment with known strategic priorities, not confirmation of direction or command, and are assessed by evaluating observed operational pattern and defensive risk exposure rather than tasking authority or implying any future state action.

Examples of low, medium, and high confidence assessments are presented in the Appendix to avoid false precision, preserve analytical integrity, and to invite the community to discuss Crimson Echo with Darktrace analysts. Confidence scores are comparative across cases rather than absolute measures of certainty and reflect corroboration across multiple evidence classes as described above.

The Darktrace Threat Research team then outlined a four-phased approach to plan, threat hunt, quality control, and analyze the cases of Chinese-nexus threat activity within the customer base. This section will review elements of the project plan and methodology of relevance for any third-party seeking to better contextualize the results, parameters of threat hunting, and techniques to assess confidence of the resulting cases.

Findings identify structural risk patterns rather than imminent activity, and not all observed activity reinforces actor aligned assessments.

While the findings of Crimson Echo are robust, they are constrained by the scale and structure of the underlying dataset. In total, low-medium-high cases over the last 3 years reflects upwards of 80 at the time of writing.

The meta model and subsequent analysis were restricted to the more than 50 cases deemed medium to high confidence, which provides its own challenge of both mistaken observations and analysis omittance.

This limits statistical power and means the results should be treated as descriptive rather than definitive. Summary statistics, bootstrapped confidence intervals, KDE and QQplot comparisons, and ShapiroWilk tests all supported the conclusion that dwell time distributions are right skewed with dense early activity and long tails, but the low sample size constrains how far these methods can be interpreted. Overall, the analysis reliably captures behavioral trends but cannot yet support higher resolution modeling without expanded data sets.

Timeframe of analysis

The Darktrace team set a timeframe for threat hunting of approximately three years, between July 2022 through the end of September 2025.

While a fully comprehensive review of Chinese-nexus groups could span decades, the timeframe of the investigation was limited for several reasons. Like individuals and organizations, APTs evolve over time and can refine their craft.

Preferred targets are also heavily influenced by the goals of their sponsoring governmental bodies.

While some overarching objectives at the strategic level may remain static, operational and tactical initiatives inevitably will shift. Indicators of compromise (IoCs) also have limited relevance for analysts, and malware that was once heavily relied upon may wane in usage over time. In combination, these factors render older instances of Chinese-nexus exploitation potentially less relevant for analysis.

Given these constraints, the selected timeframe ensures a sufficient sample size while minimizing the impact of outdated indicators and legacy activity.

Log Types

Like all threat hunting efforts, the nature of the threat hunting and results are impacted by the composition of the logs reviewed.

Threat hunting efforts for the project largely utilized data on behavioral deviations identified by Darktrace. Various threat hunting approaches, including IoC, hypothesis, and TTP-based investigations were used to identify relevant behavioral deviations. Therefore, activity identified during threat hunting efforts has already been flagged in the same capacity as anomalous for the noted alerting device.

Threat Hunting Phases

The Darktrace team divided threat hunting into two phases to ensure efforts were not duplicated and collectively reviewed as much as possible. These phases included the following focuses, which will be discussed further below:

Phase 1: Campaign/Malware Investigations

Phase 2: TTP and Sequence-based Behavioral Hypothesis Investigations

Each phase consisted of a **standard series of threat hunts** based around an identified campaign, malware, or core TTP associated with Chinese-nexus intrusions. Investigations contained a regimented sequence-based threat hunting framework already referenced (hypothesis-based, etc.).

Phase 1: Campaign/Malware Investigations

The first phase focused on first reviewing prior Chinese-nexus cyber activity as detailed by external security researchers. Analysts created an aggregation of open-source intelligence (OSINT) sources and articles related to Chinese-nexus groups (as assessed by third parties).

Each identified campaign or series of exploitations detailed by third parties prompted a corresponding investigation within the customer base for prior activity relating to the operations. This phase largely focused on the "literature review" element of the project. In addition, Darktrace analysts also performed specific investigations for evidence of malware known to have a high degree of affiliation with Chinese-nexus groups.

These malware variants included, but were not limited to, the following:

ShadowPad
PlugX
SnappyBee

However, the Darktrace teams recognize that malware usage alone is not a strong indicator for Chinese-nexus activity, as other nation-state organizations will use the same malware to achieve their national security goals while misleading security teams under a "false flag" approach.

As such, malware investigations were supplemented by prior Darktrace proprietary information, OSINT sources, and other intelligence based on established industry relationships.

Phase 2: TTPs-based Investigations

The second phase of threat hunting efforts focused on using identified TTPs to detect additional cases of Chinese-nexus activity not affiliated with identified campaigns or designated operations.

Analysts derived these TTPs from the insights of the previous Campaign/Malware Investigations phase. Key themes in general procedures and operational approaches began to emerge in the process of reviewing each prior campaign. These themes were then reinforced by actual identified cases of exploitation discerned within the customer base.

This combination of themes derived from the "literature review" and operations identified within model alert data served as core components of TTPs used to build threat hunting queries. A sample of the TTPs used include:

Internet-facing infrastructure
LOLBins
Dynamic-link library (DLL) sideloading and search order hijacking
DNS for tunneling and command-and-control (C2)

Although each tactic or technique has variations in how it was carried out (which made it difficult to build precise threat hunting queries), researchers relied on common core elements of the activity to look for evidence of exploitation using sequence-based threat hunting methods.

Quality Control and Confidence Ascertainment

All identified cases in both phases of threat hunting underwent at least two rounds of quality control. Analysts first reviewed activities to determine if the activity could reasonably be attributed to, or associated with, any other threat actor, malware, or operation not likely to feature Chinese-nexus intrusions. All network-based IoCs were reviewed via multiple OSINT methods and sources to determine potential associations with unrelated malware.

Analysts also reviewed cases where full kill chains were present to remove any cases of non-malicious activity (i.e., pen testing/attack simulation activity) or unrelated malicious operations (i.e., ransomware by criminal groups). Activities whose malicious nature or affiliation remained uncertain underwent additional rounds of quality control.

Once all unrelated events were removed, cases received an overall confidence level for likelihood of affiliation with Chinese-nexus activity. Confidence levels corresponded with discrete score values, which were derived using a combination of inputs that serve as the basis for attribution frameworks. The Darktrace Cyber Attribution Framework can be found in Appendix 1.

Results and Key Findings

Summary Statistics

from Dataset of Compromises

The resulting database of compromise events yielded key insights into operational trends of Chinese-nexus threat actors since July 2022. These figures will help contextualize more specific insights derived from dwell time, model alert sequence, and TTPs data.

All figures and insights are derived from cases with at least moderate confidence as detailed by the DCAF scoring system (Appendix 1). While low confidence cases can still provide useful information for SOC analysts and security researchers, these cases were excluded due to the increased likelihood of potentially erroneous attribution.

Data summarized in the following section will be impacted by the nature of the current and previous composition of the Darktrace customer base. Despite this limitation, the resulting dataset of compromise cases serves as a solid informational base to derive insights into China's strategic objectives and operational kill chain patterns.

Regional and Sector Victimology and Targeting Trends

US-based organizations made up the single largest portion of observed cases, representing nearly a quarter (22.5%) of all global events. US customers also comprised over 60% of events across the broader Americas (AMS) region. While the composition of Darktrace's customer base may impact the results, the high prevalence of US customers in the dataset likely reflects the US position as China's main geopolitical rival. The US is also home to the largest economy in the world by Gross Domestic Product (GDP) , underscoring the degree to which industrial espionage has become a focus of Chinese-nexus intrusions.

This logic is reflected in the data when considering the sector and critical national infrastructure (CNI) designation of the cases within the US. This project relied on the Cybersecurity and Infrastructure Security Agency (CISA) framework for identification of CNI sectors, and the overwhelming majority of compromise events involved Darktrace customers classified as CNI (88%). Even using the more stringent European Union (EU) NIS2 framework, nearly three quarters (72%) of the dataset cases involved either "critical" or "important" sectors as classified by the European standard.

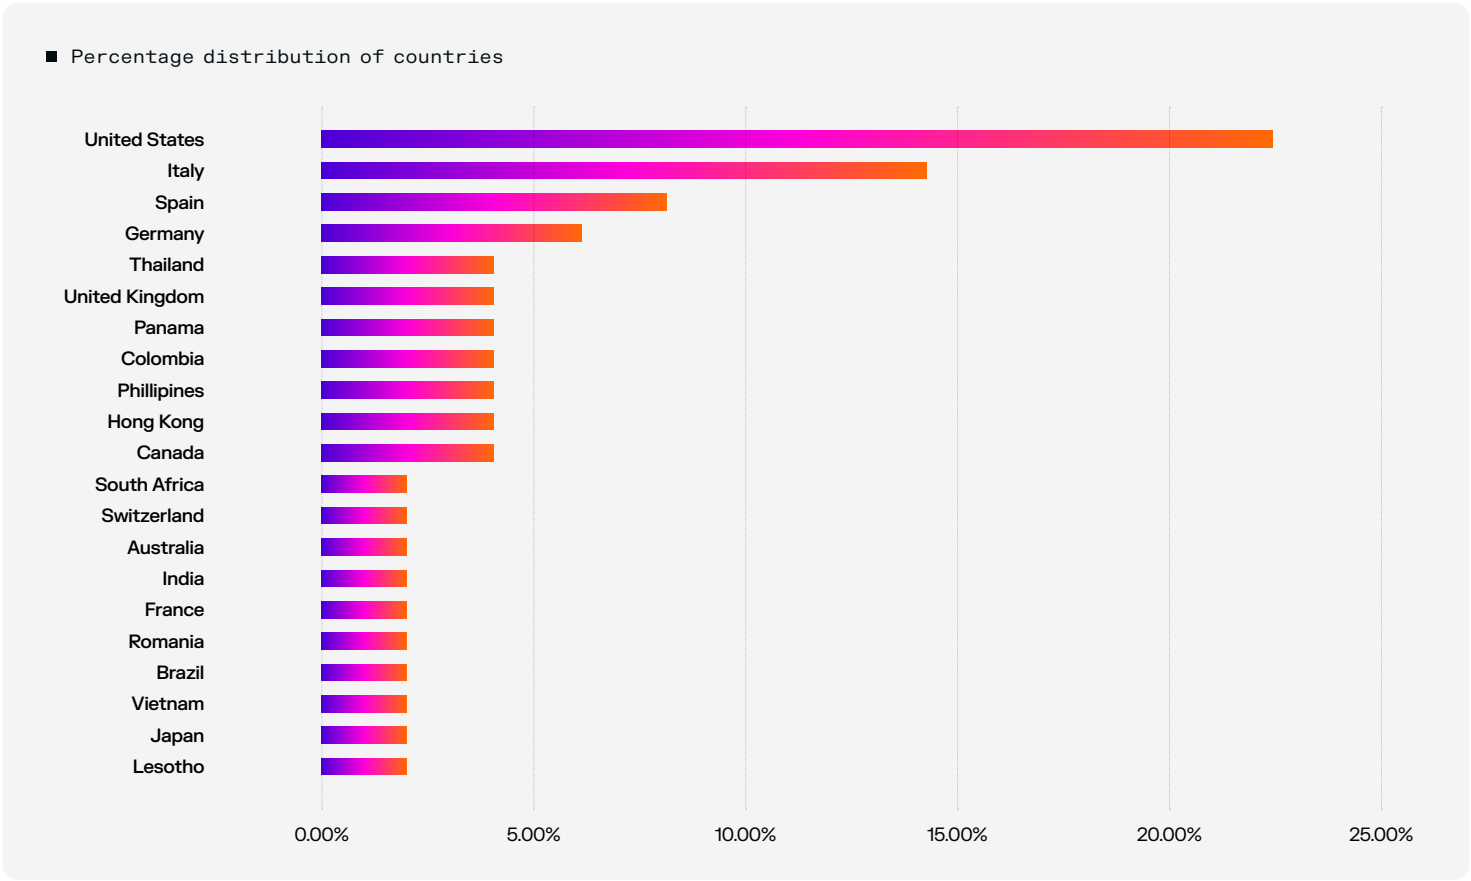

Regional Trends: US

The US accounted for approximately one-fifth (~20%) of all cases involving a CNI customer and involved the following sectors:

- Transportation Systems Sector
- Healthcare and Public Health Sector
- Government Services and Facilities Sector
- Information Technology Sector

Regional Trends: EMEA

Beyond the US, three of the top five most prevalent countries within the dataset are all EU members: Italy, Spain, and Germany.

These countries include some of the largest economies in the Eurozone and many of the organizations identified during the threat hunting process operate in core sectors of economic interest for the Chinese government including digital infrastructure, advanced technology, manufacturing, synthetic materials, and agriculture technology.

Within the Europe, the Middle East and Africa (EMEA) region more broadly, over three quarters (>75%) of the CNI cases involved the following CISA sectors:

- Transportation Systems Sector
- Communications Sector
- Critical Manufacturing Sector
- Information Technology Sector
- Food and Agriculture Sector

Together with the US, these countries account for nearly half of all compromise cases involving CNI sectors.

China has long exhibited an extensive history of cyber operations within the digital and economic backbone of countries that trade heavily with and/or receive economic funding for infrastructure development through the BRI. The general focus of targets in such CNI sectors in both the US and EMEA regions provides further evidence of these objectives. In the EMEA region, around half of the observed cases were concentrated in three sectors: **manufacturing, telecommunications/digital infrastructure, and shipping/logistics, irrespective of whether the affected organizations were classified as CNI.**

The targeting of telecommunications and transportation infrastructure in both the US and EMEA also reveals the focus of Chinese-nexus actors to preposition themselves in the event of direct confrontation with Western countries.

80% of cases involving telecommunications infrastructure occurred in the EMEA region, and all cases specifically involving transportation infrastructure (railways, airports, ports, etc.) occurred within the US, UK, and Canada. The US led this category, comprising 60% of such events.

Regional Trends: APJ

The least prevalent region within the dataset, even when including low confidence events, was Asia-Pacific and Japan (APJ).

Of the identified compromise events, cases appear to have been fairly evenly distributed amongst customer countries within the region. However, one should note that all recorded countries within the APJ region are members of either the Association of Southeast Asian Nations (ASEAN) or Quadrilateral Security Dialogue (QUAD). Targeted entities within this region underscore the importance of cyber operations to buttress China's security posture in the region, specifically within the South China Sea and with respect to Taiwan.

However, the targeting of IT services within Hong Kong more likely reflects internal security goals.

Cases in APJ also skew more heavily towards the public sector and media sectors. Whereas half of the cases in the APJ region consist of government or communications related entities, the EMEA region includes a broader range of critical economic targets including manufacturing, health systems, shipping and logistics, and defense. Again, analysts should keep in mind the potential skew within the Darktrace customer base.

This pattern seemingly underscores how targeting is in part determined by a strategic objective. While cyber activity in the APJ region appears to focus on more traditional espionage for China's strategic advantage in security issues in the region (South China Sea, Taiwan, etc.), cyber intrusions in EMEA reflect potentially the broader focus of supporting BRI initiatives and economic espionage.

The large percentage of cases in the US, Canada, and EMEA specifically involving the transportation systems sector (CNI) may underscore how China is prepositioning itself to cause disruptions in the case of a direct military confrontation over Taiwan.

Goal/Objective Variation

Short and Long Duration Cases

The approach of this project is grounded in the idea that a generalized framework for threat hunting Chinese-nexus groups is possible.

However, Darktrace analysts do not assume that every compromise will have identical kill chain sequences or attack elements. There still exists a high degree of variability in the resulting compromise database. Attack chains vary in duration, techniques, tactic sequence, and activity observed.

While a single factor accounting for all variability was not identified, a general discrepancy in the data set centered on total dwell time. Specifically, cases of a shorter total duration tended to feature a more characteristic grouping of model sequences and tactics/techniques when compared to compromises of a longer duration. The cut-off point to determine short- and long-term durations was set at one month, given the higher variability of cases beyond this dwell time. Although deviations still exist within each grouping, generally shorter-term cases tend to feature initial access via internet-facing device exploitation with quick tooling ingress, the establishment of C2 communications and/or egress activity.

In comparison, cases that have longer dwell times tend to highlight wider network penetration beyond the initial cluster of devices.

When considering the nature of Chinese-nexus cyber operations and the victimology patterns, Darktrace analysts believe this discrepancy results from a split in objective sets for Chinese-nexus actors. Operations aimed at industrial and traditional espionage may naturally result in longer observed compromise durations and persistence in the network for continued data egress. This contrasts "smash-and-grab" style operations that are more opportunistic, or cases of "lay-and-wait" prepositioning. The nature of this bifurcation will be explored in the following sections.

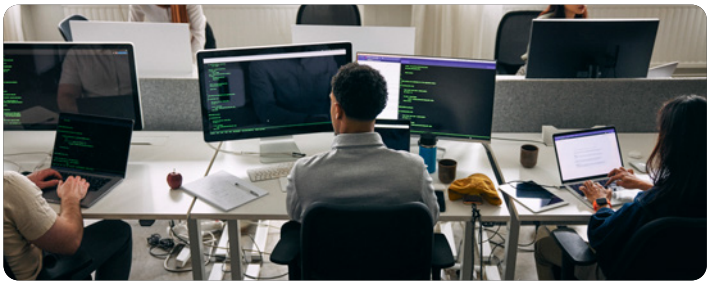

Technique Distribution and Co-Occurrences

Cases identified with at least medium confidence to be associated with Chinese-nexus activity were also evaluated based on a series of tactics and techniques of particular importance. These specific implementations of broader tactics were identified during both a review of existing open search research and pattern identification noticed during threat hunting efforts. While there are some limitations for some subsets of activity, particularly host level data, broader trends can still be revealed.

Technique Frequency and Distribution by Region

A majority of cases (63%) involved exploitation of internet-facing infrastructure, typically for initial access into the environment.

While it is possible that other cases featured interaction via public facing infrastructure, events associated with this technique triggered specific model alert activity in response to observed or implied Common Vulnerabilities and Exposures (CVE) exploitation. Direct cases and model detection evidence associated with pre-CVE exploitation are included in the Appendix. Other prevalent techniques observed in the overall data set include explicitly observed command line utility, usage of cloud providers for C2 activity and/or payload hosting, and employment of the DNS protocol for tunneling, each individually occurring in 40% or more of all cases.

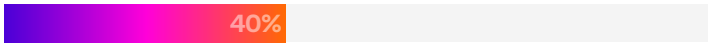

The regional distribution of tactics/techniques also evidences key trends. By percentage, cases of Chinese-nexus compromise initiated via the targeting of perimeter devices were more common within the EMEA and AMS regions. Over 75% of EMEA cases and 60% of AMS cases include this technique. When separating the US from the grouping, the pattern also remains consistent. This reliance of perimeter device exploitation contrasts with the 30% of APJ cases commenced via this initial access method. While initial access operations were not always directly visible, cross-referencing potential activity with concurrent campaigns reported in OSINT sources suggests that these cases may have originated from phishing attempts or broader search engine optimization (SEO) hijacking operations.

Compromise events within the APJ and EMEA regions also share unique parallels not seen within the AMS region. Irrespective of compromise duration, events involving customers within EMEA and APJ were more likely to involve application or transport layer reconnaissance operations, potentially suggesting a potential bias towards quicker zero/n-day compromise versus "low and slow" goals for US and AMS customers when compared with operations in other regions. When examining targets within the US specifically, discrepancies appear further magnified. The US had the lowest percentage of cases involving DLL sideloading or search order hijacking, with fewer than 10% of such cases featuring explicit evidence for the technique. This figure contrasts the frequency within EMEA and APJ, with around 40% and 60%, respectively.

DLL sideloading and search order hijacking were primarily observed during two phases within the kill chain: initial tooling download and extensive lateral movement. Consequently, duration may still have an effect on the prevalence of the technique, albeit with less of an impact compared with internal reconnaissance-related techniques. Again, though, this would potentially suggest a preference for compromise and persistence activity, with less widespread network penetration within the US target base. This parallels trends noted in the sector distribution of cases noted earlier.

Technique Co-Occurrences

Analysts also evaluated the compromise dataset for co-occurrence of tactics and techniques.

Co-occurrence metrics were first assessed across all cases and then reevaluated when splitting the data into short-term and long-term groupings. When reviewing the entire compromise dataset, analysts identified two key figures: an elevated co-occurrence of internet-facing device exploitation with tooling ingress as well as a moderate co-occurrence of tooling ingress with explicit command line utility. A majority of cases within the overall dataset featured some form of perimeter device exploitation. The resulting heightened co-occurrence of this technique with tooling ingress therefore is not unexpected as many such compromises would likely begin with such a sequence of events. Moreover, given that server exploitation frequently would involve remote command execution for file retrieval, the observed co-occurrence of tooling ingress activity with explicit command-line user agents also comports with expectations.

Like the undifferentiated dataset, the short-term matrix notes a strong co-occurrence of internet-facing device exploitation with tooling ingress, as well as observable tooling ingress with command line utility.

However, shorter term cases also specifically exhibit a heightened co-occurrence of internet-facing device exploitation with both DNS tunneling and HTTP/SSL beaconing. The higher frequency of DNS tunneling in cases involving internet-facing device exploitation in part stems from the usage of DNS tunneling and bin services as a means of exploit validation. Compromises involving a shorter dwell time also interestingly have higher co-occurrence of cloud provider usage with techniques such as tooling ingress and DNS tunneling.

There were no clear trends regarding data egress, Active Directory (AD) reconnaissance, and Registry/Service operations and moderate, albeit relatively weaker co-occurrence of network scanning with LOLBin usage. However, the lower co-occurrences for lateral movement and reconnaissance related procedures may stem from the nature of shorter duration compromises. Such cases may not have progressed to later stages of the kill chain and/or do not include broader network compromise as a goal.

Co-occurrence of internet-facing device exploitation with tooling ingress, and tooling with command line utility is higher in this dataset when compared to the baseline.

Shorter-term cases, potentially those involving “low and slow”, or “smash-and-grab”-style operations, therefore show a potential bias towards quick exploitation for remote code execution and trojan download. Shorter dwell times also feature an elevated co-occurrence of cloud provider usage alongside tooling ingress not seen in the long-term data.

This could suggest a reliance of cloud providers for C2 and payload hosting in such cases of quick exploitation and persistence establishment in shorter duration events. Short-term cases also exhibit elevated DLL sideloading paired with anonymizing infrastructure, suggesting these services are being used to host payloads specifically intended for DLL-based execution.

The highest co-occurrence within any derivation of the dataset appeared within the long-term cases, specifically involving the appearance of network reconnaissance (a combination of application layer reconnaissance and transport layer scanning), with nearly 80% of network scanning instances also featuring DLL sideloading or search order hijacking in compromise.

As noted, longer compromise duration may result in a higher co-occurrence for reconnaissance and lateral movement activity, simply due to the nature of longer kill chains.

Yet this high co-occurrence also indicates that DLL sideloading, particularly when it features as part of the lateral movement phase, is seen more commonly with widespread network penetration rather than smaller scale operations for access and persistence. DLL sideloading also appears more commonly with DNS tunneling within the dataset.

Egress related activity appeared more frequently alongside cloud provider usage, suggesting that in longer duration intrusions, Chinese-nexus groups may increasingly leverage cloud services for exfiltration and C2 rather than solely for payload hosting.

Remote management tooling also co-occurred with multiple techniques—such as DLL sideloading, internet-facing device exploitation, and broader tooling activity—indicating greater reliance on Remote Monitoring and Management (RMM) capabilities during extended operations and potentially reflecting both actor adaptability and the characteristics of targeted organizations.

Longer-term cases further showed elevated co-occurrence between lateral movement, reconnaissance, and other TTPs.

While some of this is attributable to extended dwell time, the breadth of overlap suggests that these groups deliberately expand and vary their technique set during prolonged compromises, using the additional operational window to pivot, deepen access, and widen their footprint across the environment.

Dwell Time Metrics

The Darktrace Threat Research team analyzed all compromise cases with respect to timing and duration, collectively referred to as "dwell times."

Threat hunters and cyber defenders frequently rely on dwell times to set timeframes for threat hunts and scope the breadth of investigations. Metrics related to the duration of identified compromises also inform timing delays in future models and detection heuristics to allow for appropriate detection. Therefore, to support both use cases, analysts focused on three subsets of "dwell time": overall, egress, and lateral movement durations.

All Darktrace model alerts retain a standardized timestamp corresponding to the moment of logging within a customer's environment.

Each dwell time category was simply calculated by taking the difference between the identified endpoint timestamp, and the first model alert identified as the initiation of the kill chain. While many compromise cases did start with model alerts traditionally related to initial access, a subset of cases featured two or more model alerts one could reasonably argue constituted the start of the kill chain. In instances of ambiguity, the more conservative estimation of start time was utilized as the starting point, in order to avoid overinflation of dwell time metrics.

All starting model alerts were also manually assessed and quality controlled to ensure consistency. However, this approach assumes that the earliest observable model alert accurately reflects the true initiation of malicious cyber activity, which may not hold in cases where earlier actions occurred in aspects of the customer's environment not covered by Darktrace.

Overall Dwell Time

In the context of this report, overall dwell time is simply the duration between first and last model alert within the kill chain for each compromise case. The distribution of overall dwell time is right-skewed, with a median value of 237.63 hours (approximately 10 days), but with a smaller subset of higher values extending as far as over 600 days. Bootstrapping the median and the 95th percentile yields a standard confidence interval of 3.5 to 26 days and 246 to 584 days, respectively.

■ OVERALL DWELL TIME

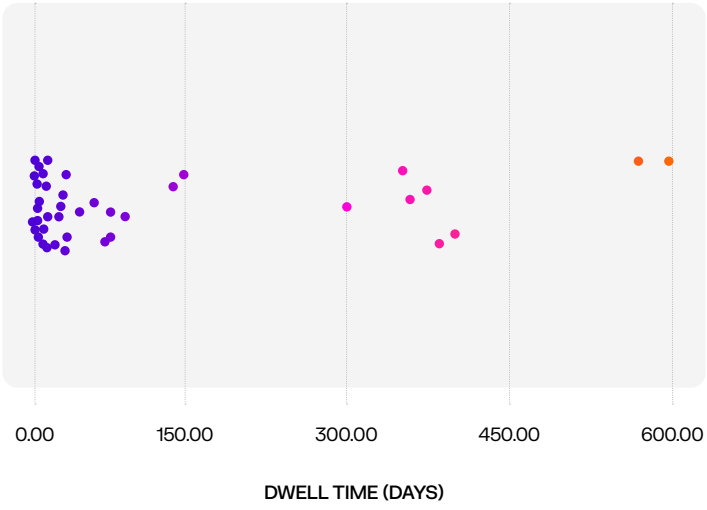

■ OVERALL DWELL TIME KDE

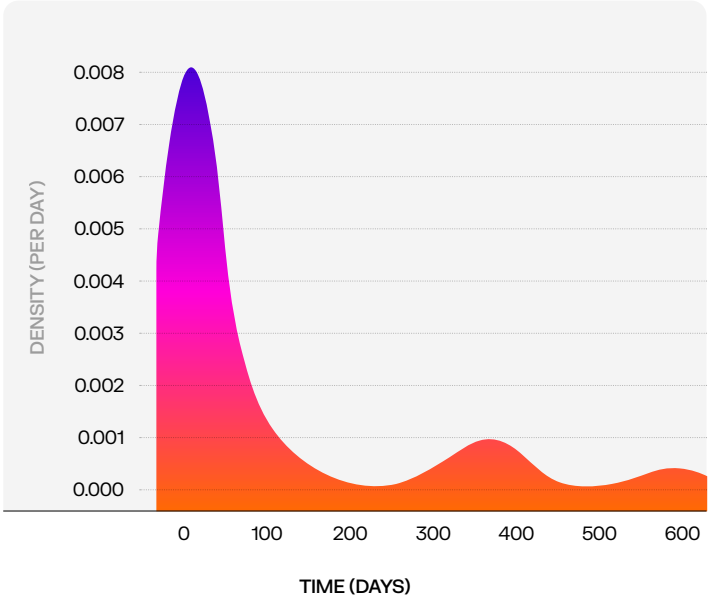

The contrast between these intervals indicates a distribution with most observations concentrated in a relatively narrow central region, with a long tail. Around half the cases occur within 0.5 hours to 10 days. Theoretical models were proposed to describe the overall data distribution, including log normal, multi-modal etc., yet hypothesis testing proved inconclusive. Regardless, the rightward skew of the data generally reflects the anticipated pattern, whereby the data highlights a plethora of more opportunistic, short-term compromise cases that may involve quicker, more opportunistic n-day exploitation, and a smaller subset of cases featuring long term persistence for strategic targets.

Egress Dwell Time

Egress dwell time again utilized the same starting timestamp but used the timestamp of the first egress related model alert as the endpoint. While some models related to C2 or beaconing could reasonably identify activity involving some forms of data egress, only models specifically focused on detecting exfiltration activities were utilized for this figure. Like the overall dwell time distribution, the distribution of egress dwell time is right skewed, with a median value of 48.54 hours but an arithmetic mean of 842.81 hours (35 days).

The 95% confidence interval for the median and the 95th percentile are 0.5 to 7.1 days, and 16.9 to 341.1 days, respectively. Again, one can see a distribution which has a high density of points at low values with a very light left tail and a few isolated extreme high values dominating the right tail. Unlike the distribution of overall dwell times, the length of time between first alert and first egress model reasonably follows log-normal distribution; running the Shapiro-Wilks test on the log-transformed data yields a p-value of 0.36 meaning that we cannot reject the assumption of lognormality.

These insights again support the working hypothesis/model for generalized Chinese-nexus activity: where egress and data exfiltration are the key goals, a majority of compromises feature quicker "smash-and-grab"-style operations. However, for industries and targets of higher value and strategic importance, long-term persistence is established first, and then later features egress over a longer timeframe.

■ Egress Dwell Time KDE

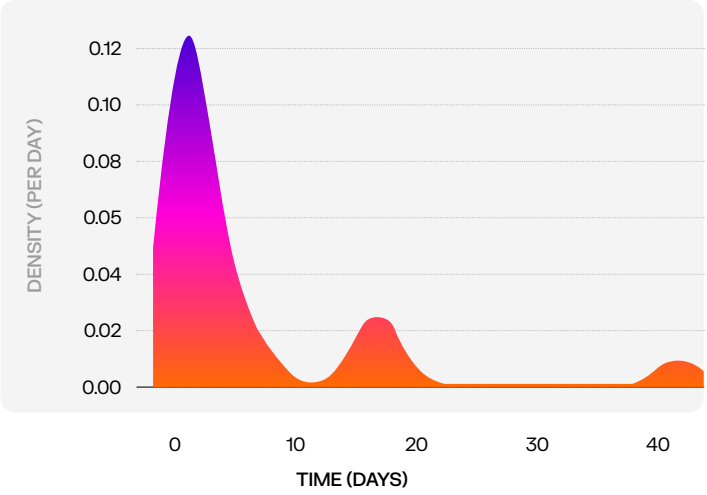

■ Lateral Movement Dwell Time KDE

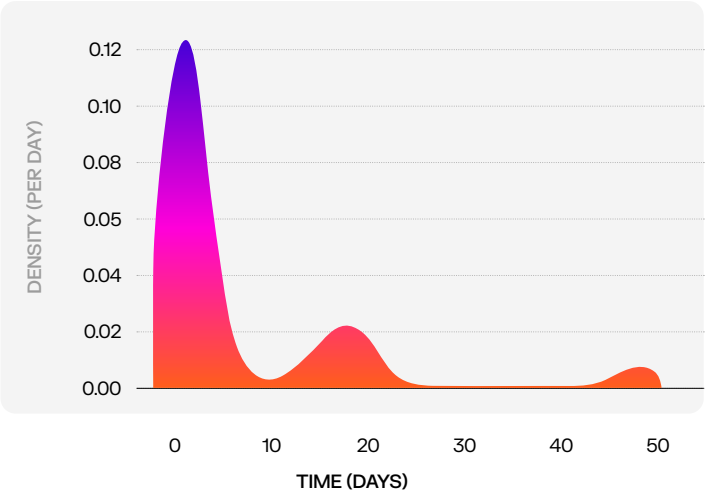

■ Histogram of Lateral Movement Clusters

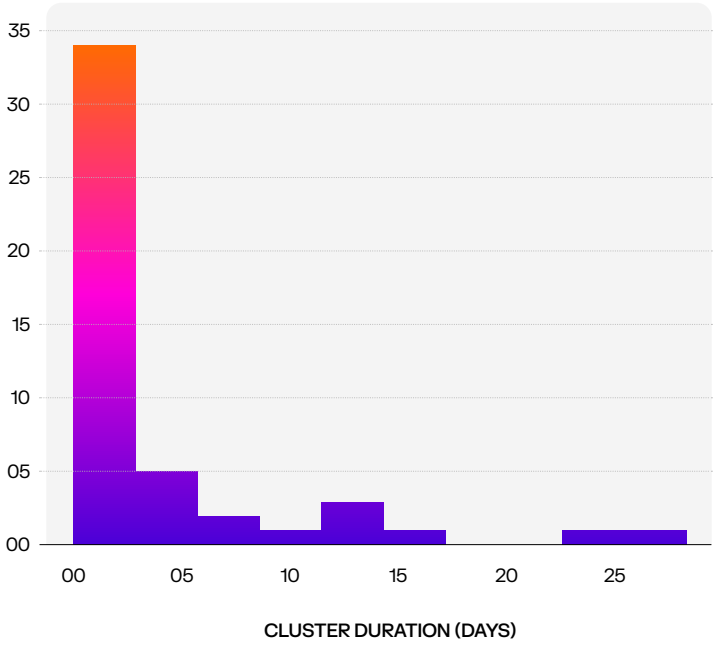

Lateral Movement Dwell Time

Rather than marking the first lateral movement alert as the start of activity, analysts identified dense "bursts" of lateral movement models to avoid misclassifying early probing or reconnaissance as full network penetration. Models were clustered using HDBSCAN, and the start of the first cluster served as the dwell time anchor. While initial lateral movement alerts can still be useful for threat hunters, analysts judged the interval between initial access and the first sustained lateral movement cluster to be a more meaningful metric.

The resulting distribution is right skewed, with a median of 27.7 hours, a mean of 147.6 hours (6.15 days), and a 95th percentile of 478 hours (19.9 days). The data is broadly consistent with a lognormal pattern, though the fit is weaker than for egress.

Compared to egress, the lateral movement distribution is narrower, with fewer extreme values: its median is roughly half that of egress, and its average is nearly six times smaller. Put simply, attackers usually move laterally more quickly and with less variability than they conduct egress.

The rightward skew was not necessarily expected, as one might assume sophisticated actors delay lateral movement to avoid detection. However, cases involving clear lateral movement clustering—indicative of deeper network compromise—show that egress is often delayed until after broad internal spread. In 75% of intrusions featuring both lateral movement clusters and egress, widespread lateral movement occurred first, suggesting that in higher value operations, actors emphasize full network compromise and long-term positioning before extraction or action on objectives.

As noted, the egress dwell times shows a larger variance, and wider gaps for tail-end data. However, when removing these outlier cases, interesting patterns emerge when comparing the two datasets.

Looking at the density graphs for egress and lateral movement dwell time, one can see the noted general spike of initial activity typical for a lognormal distribution, however, in both graphs, a peculiar spike in activity occurs in the 10 to 20-day time range. Irrespective of whether the activity involves egress or lateral movement, should a longer duration compromise be under way, a spike in activity appears to be somewhat expected potentially in the 10 to 20-day timeframe. This pattern suggests that once a compromise persists beyond the initial stabilization period, threat actors often escalate activity around the 10–20 day mark—likely when they have mapped the environment enough to proceed with expansion or data staging.

Finally, analysis was conducted on the lateral movement clusters specifically. The median cluster duration was around one day, and where clusters were detected, the average number of clusters was four. The median cluster duration was 0.96 days with a 25th and 75th percentile mark of 0.1075 and 2.965 days respectively, while the standard deviation was 6.13 days. This pattern suggests that although lateral movement activity is typically executed in tight bursts, certain operations require extended internal maneuvering—often in larger or more complex networks—resulting in a much longer tail.

Model Alert Activity Clusters

The final stage of analysis applied unstructured clustering to identify "bursts" of alert activity across entire kill chains.

This approach helps threat hunters distinguish high density subcomponents within compromises, even those with long dwell times, and allows SOC defenders to better interpret when and where operational tempo increases. Darktrace analysts used the same clustering technique applied to lateral movement analysis but broadened it to all model alerts regardless of kill chain phase.

The clustering algorithm did not identify bursts in every compromise, as some intrusions simply did not exhibit dense enough activity to meet the criteria. This variability underscores both the diversity of intrusion styles and the adaptability of threat actors. Where clusters were present, they were typically dominated by a single tactic—most frequently C2, internal reconnaissance, lateral movement, or tooling—and the next alert in a sequence was most often the same tactic repeated. Most clusters contained between five and fifteen alerts and lasted fewer than fifteen hours.

Analysts also examined tactic-to-tactic transitions by condensing consecutive chains. Across all clusters, tooling, privilege escalation, and egress were most commonly followed by C2, while internal reconnaissance and lateral movement frequently transitioned into each other. These patterns align with both the design of the Darktrace model deck and common intrusion workflows. Notably, no relationship was found between cluster length and cluster entropy, indicating that long sequences of alerts did not necessarily involve more diverse TTPs.

When entropy was nonzero, no single tactic overwhelmingly dominated the cluster beyond its co-occurring counterparts.

Analytical Takeaway 1

Clusters tend to reflect short, tactic-focused bursts of attacker activity, suggesting that operators often execute rapid, concentrated phases of work rather than broad, multi-technique surges.

Analytical Takeaway 2

The absence of strong correlations between cluster size and tactic diversity indicates that escalation in activity does not necessarily equate to more complex or varied behavior. In many cases, attackers simply intensify the same technique during high priority phases.

These findings show that alert density clustering provides a meaningful lens into attacker workflow, highlighting when adversaries shift into higher intensity operational phases and enabling defenders to detect and respond to these surges more effectively.

CASE STUDIES

Chinese-Nexus Attacks in Focus

The following section details the specific kill chains and attack sequences element of three cases within the overall dataset. A review of these events serves to add context for the patterns and assertions made in earlier sections and will provide insights into how such analysis was derived.

Moreover, these cases highlight key trends in relation to co-occurrence, dwell time, tactics sequencing, as well as the sector/regional breakdown mentioned previously. Beyond methodological purposes, this section should also assist cyber defenders by showcasing potential implementations of a potential Chinese-nexus operation in varying environments.

■ **In practice**

Rapid Exploitation of SAP NetWeaver and Cloud-Hosted Tooling Delivery

In April 2025, Darktrace identified a server device located on the network of a European transportation organization exhibiting early signs of exploitation.

The system began performing unusual DNS querying and tunneling to bin services, including requests for Out-of-Band Security Testing (OAST) domains. While security professionals can use OAST domains for testing activity, these endpoints also frequently appear alongside compromise events as a means of exploit validation. The subsequent behavior from the server would later confirm this usage of the platform. Analysts later concluded that the likely initial access occurred via exploitation of CVE-2025-31324, a vulnerability that allows unauthenticated remote code execution via SAP NetWeaver Visual Composer.

Additional tooling activities commenced within a few days of exploit validation and initial access. The server proceeded to make several outbound HTTP GET requests for binaries hosted on Amazon S3 buckets. The curl command line tool appeared as the designated user agent for these HTTP requests for data that appears to have potentially been masquerading as a different file type. Additional binary/octet-stream download requests from cloud servers continued for several additional days, and included the additional usage of PowerShell, as noted within the user agent HTTP header field.

The total duration of the compromise event appears to have lasted around four to five days and did not feature explicit egress, reconnaissance or lateral movement activity.

Some of the additional HTTP requests may have resulted from an inability for previous file downloads to install on the affected system.

Regardless, the multiple binaries received likely represent tooling ingress for persistence and additional tooling to maintain access and achieve further goals within the environment. Given the customer's status as a key transportation sector organization, this would appear consistent for such cases for longer term, lay-and-wait goals. The compromise also exhibits core technique co-occurrences for shorter term compromises, specifically internet-facing device exploitation resulting in command line utility execution for tooling download.

This contrasts with other shorter term cases featuring tooling download, cloud service providers and egress operations. One of the shorter duration cases within the APJ region featured tooling download and cloud provider usage. This compromise involved explicit data egress to a cloud storage service and did not include multiple rounds of binary tooling featuring command line utility. Notably, this event occurred within the network of a public sector customer and likely featured phishing as an initial access vector.

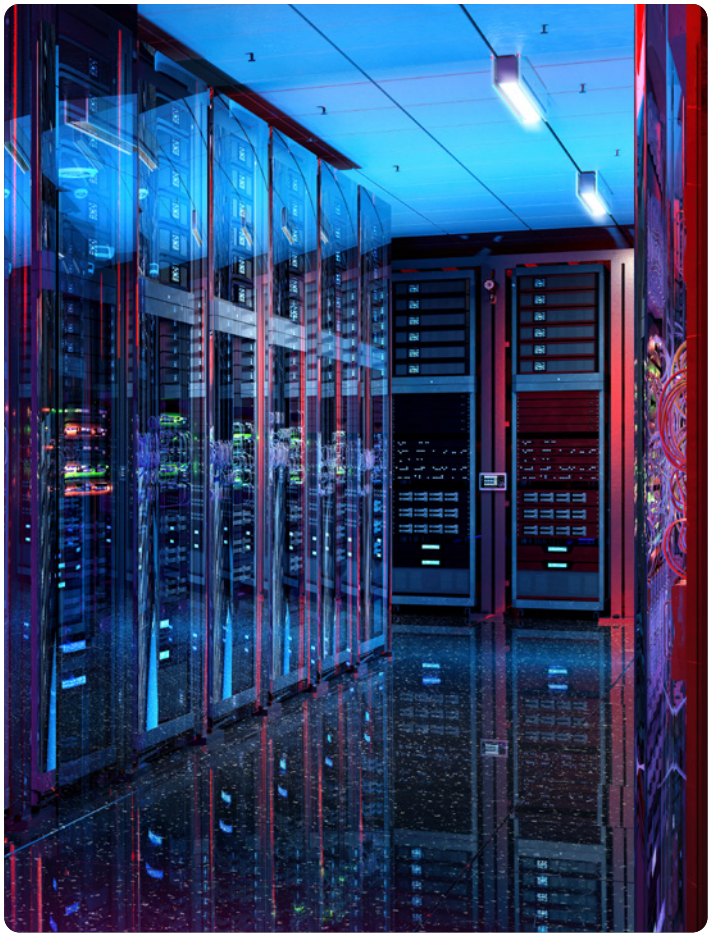

■ In practice

Gradual Reconnaissance and Credential Abuse in a US CNI Environment

Darktrace identified model alerts associated with exploitation of CVE-2025-0994 Trimble Cityworks 18 days prior to public disclosure, with additional alerts flagging malicious activity even earlier.

This highlights the value of behavioral visibility in detecting adversary activity against environments that underpin national-level functions. The affected environment qualifies as CNI due to its role in maintaining continuous aviation operations that support national mobility, emergency response, and defense logistics. Because it is integrated into wider multimodal transport and logistics systems, any disruption could create cascading economic, operational, and security impacts—an inherent feature of high-value CNI targets.

The combination of the early detection window and the criticality of the operational environment underscores both the strategic importance of the customer and the advantages of behavioral-based detection ahead of formal vulnerability disclosure.

As early as mid-December 2024, Darktrace observed a device within a US public sector customer's network deviating from its expected behavior, indicative of potential device exploitation. The device displayed multiple indicators later associated with exploitation of Trimble Cityworks CVE-2025-0994. As Trimble Cityworks can be installed on a range of devices with varying degrees of internet exposure, it is notable that Darktrace detected prior internet exposure weeks before deviations from established patterns of life began.

The affected host initiated uncommon WMI requests in the early-to-mid December timeframe. While this alone does not confirm compromise, the use of LOLBins such as WMI instrumentation appears frequently across analyzed cases due to the difficulty of distinguishing it from legitimate traffic.

One week later, the host issued DRSGetNCChanges requests to a domain controller, and by the end of the month leveraged an administrative credential uncommon for that device.

Analysts assessed that the DRSGetNCChanges activity likely enabled the credential's subsequent use. With this credential, the host performed an SMB drive write of an .ini file into the destination server's TEMP folder. Although .ini files can legitimately appear in this directory during software installation, analysts deemed the behavior unusual due to the newly used credentials and the atypical filename.

One week later, the threat actor began internal reconnaissance, which unfolded in two one-week phases.

The initial phase consisted of SMB scanning focused on identifying devices with open port 445. The following week, the host-initiated RDP-related internal connections, likely informed by results from the prior scan. In the final stage of the kill chain, the host downloaded additional tooling by issuing outbound HTTP GET requests for an alphanumeric executable hosted at a rare external IP address. The presence of a PowerShell user-agent in the alerts indicates command-line activity, while the use of an uncommon destination port (3219) for HTTP suggests an attempt at defense evasion by circumventing perimeter-based detections.

This case reflects several attributes of longer-term compromises identified across co-occurrence, tactic-sequence, and dwell-time metrics. The customer operates within the US transportation sector and is a key CNI entity, a factor likely contributing to the absence of distinct egress-related model alerts—unlike long-term intrusions in sectors more commonly targeted for economic espionage.

The case also exemplifies the heightened prominence of domain-enumeration and reconnaissance behaviors, particularly DRSGetNCChanges requests, in longer-duration operations. Although two potential lateral-movement alerts were observed, analysts did not identify evidence of widespread lateral movement. As such, the absence of DLL sideloading activity via SMB writes is unsurprising, as is the lack of broad network penetration.

Interestingly, command-line utility does appear in the context of tooling downloads; however, the HTTP requests occur toward the end of the observable compromise rather than at its outset. This case highlights that sophisticated intrusions rarely look dramatic at first—they build quietly over time—so the ability to detect small behavioral shifts is critical for protecting high-value systems.

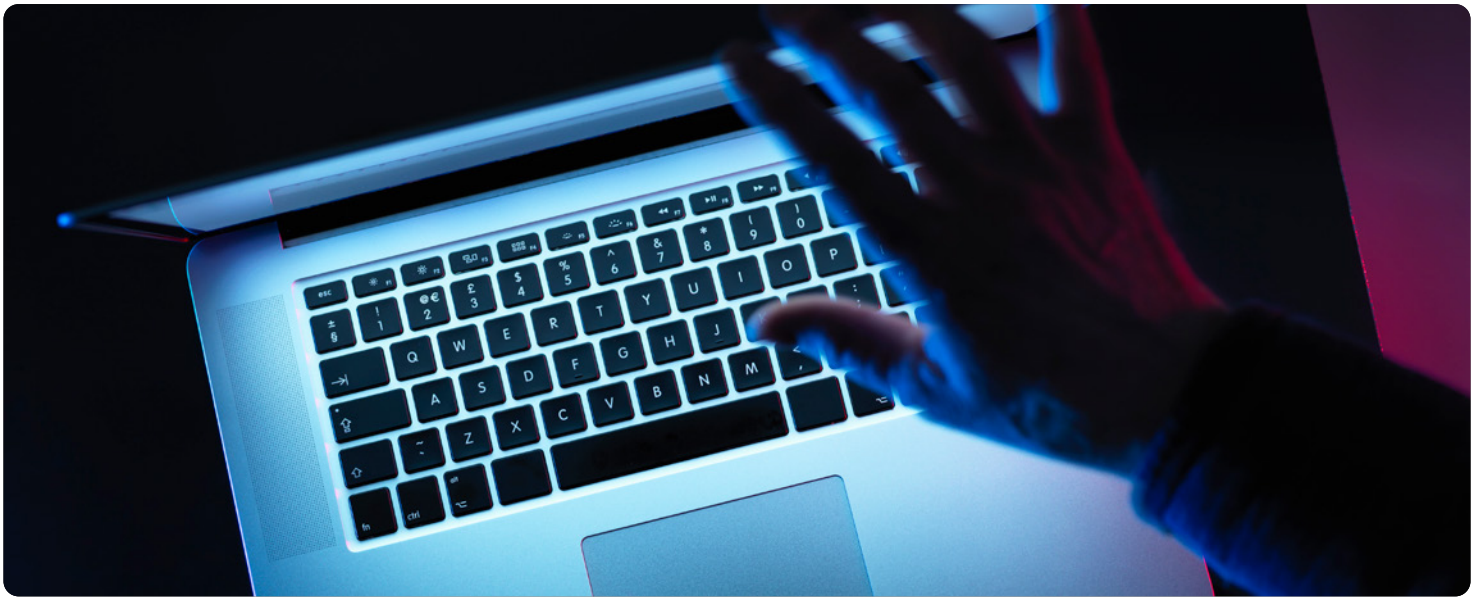

Deep Network Penetration in a Critical Manufacturing Environment

Several devices within the corporate network of a critical manufacturing customer in the APJ region showed initial signs of infection in October and November 2022.

Although direct exploitation and initial access events were not directly seen, the affected hosts-initiated application layer reconnaissance operations by attempting anonymous NTLM authentications and widespread internal connectivity over port 445. This initial scanning and network enumeration may have represented initial assessments for wider network compromise.

By December, one of the desktops utilized a legitimate administrative account, likely stolen from the local host, to establish SMB sessions and write files to the DISK shares of several network servers. Specifically, the system transferred suspicious DLL and a likely XML configuration file to the "vss" folder on the DISK share of destination servers. Analysis of the network indicators during this time suggests both DLL sideloading/search order hijacking techniques as well as likely use of the PCExter malware.

Darktrace analysts detected additional binaries and files written via SMB during this time including potential additional tooling and/or configuration/encrypted payload content.

Some such executables written include binaries that likely enabled data exfiltration, including those named "Onedrive.exe" as well as various archive files likely containing chunked data for egress.

This activity continued throughout December 2022 and included usage of LOLBins native to the customer's environment including the PSEXESVC administrative tool. The threat actor appears to have initiated egress activities in mid-December.

Affected systems, many of which wrote and received the noted binaries sent large volumes of TLS encrypted data likely aggregated in November to a legitimate cloud service provider. The threat actor appeared to go dormant following initial egress operations.

A pattern then emerged whereby devices seen operating during the first round of operations between autumn and winter of 2022 engaged in occasional bursts of reconnaissance and or lateral movement activity during 2023 and into at least July 2024. Analysts identified at least two rounds of DLL and RAR file writes of similar files over SMB, LOLBin usage (WMI), and domain enumeration via DRSGetNCChanges requests. Analysts also identified subsequent rounds of data exfiltration using variable cloud storage platforms within this cycle.

Again, the duration, techniques and tactics progression identified within the kill chain reflect the strategic relevance of the target.

The operators behind the intrusion-maintained access within the network for several months. Compared with "smash-and-grab" or "lay-and-wait" style attacks involving other critical infrastructure sectors, the specific strategic relevance for the Chinese government likely reflects more so the customer's situation in an industry of relevance for economic espionage and/or BRI goals; in this case, the ICT/ semiconductor industry.

This strategic outlook likely reflects the longer dwell time in conjunction with specific techniques/tactics associated with deeper network penetration and access. Widespread scanning, multiple rounds of DLL sideloading featured during the lateral movement phase, and multiple tooling transfers all appear during this case over several months.

Moreover, as noted in cases with longer dwell times, egress events are somewhat delayed. The first evidence of data exfiltration occurs weeks after initial signs of access. Moreover, subsequent rounds of egress model alerts involving cloud providers also occur well beyond preliminary lateral movement activities. The delay in egress operations may reflect how threat actors prioritize broader network compromise in the initial phase of operations for targets slated for long term persistence and iterative egress objectives.

Conclusion & Community Discussion

This report highlights the high-level patterns of Chinese-nexus cyber-activity over the past three years. Rather than focusing on individual threat groups or campaign names, the research contributes to a generalized operational framework that can help defenders recognize Chinese-nexus activity earlier and with greater confidence.

While the findings are relevant for security researchers, the implications extend across organizations—from executive leadership and CISOs to SOC analysts responsible for day-to-day monitoring.

Understanding the strategic logic behind these operations can support organizations as they interpret anomalous activity and align defensive priorities accordingly. This report is complimented by short outlooks providing that perspective through follow-on discussion for CEOs, CISOs, analysts, and policymakers.

To recap, the key insights the Darktrace team derived from this research include:

Chinese-nexus cyber operations are best understood as **continuous strategic planning**, not episodic campaigns.

Detection of short dwell time intrusions should not be interpreted as tradecraft failure but **deliberate operational choices.**

Western security models remain **overly incident centric** and systemically undervalue persistent identity risk.

China's cyber activity is **not just IP theft-based** but increasingly aligns with BRI dependencies and critical infrastructure leverage globally, with particular emphasis on the US.

Strategic Implications

The dataset reinforces that Chinese-nexus operations show a distinct preference for organizations within critical national infrastructure categories and strategically important sectors such as transportation, telecommunications, manufacturing, healthcare, and digital infrastructure.

These targeting patterns broadly appear to reflect goals set by the Chinese state to support both traditional espionage for strategic advantage, and BRI/industrial espionage goals.

This contextual data emphasizes how important it is that C-suite executives understand how and where in the CNI framework their organization resides. Investment into cyber defense resourcing, focus on proactive threat hunting, and the cadence of IT security control reviews may need to be reassessed given the intent to act on strategic objectives by Chinese nexus actors in such critical sectors.

Implicit in this data is the fact that organizational risk will be directly impacted and better mitigated by an understanding of medium- and long-term strategic planning by the Chinese state apparatus. CISOs and decision makers can benefit from a more nuanced understanding of how their organization may be viewed as a potential target for such objectives.

Operational insights from this report—including dwell-time trends, technique co-occurrences, and kill-chain sequencing—can help SOC teams refine threat-hunting parameters including time frames and prioritize defensive monitoring.

Rather than relying solely on static indicators or actor-specific profiles, defenders can focus on recurring patterns of behavior such as anomalous credential use, reconnaissance activity, unusual cloud connections, and bursts of lateral movement.

The findings also underscore the growing importance of anomaly-based detection. Chinese-nexus actors frequently employ living-off-the-land (LOTL) techniques, cloud infrastructure, and legitimate administrative tools that evade traditional signature-based controls.

Detecting these operations therefore depends increasingly on identifying deviations from established patterns of network behavior.

Community Discussion Questions

The findings of Crimson Echo raise several strategic questions for discussion across the cybersecurity, business, and policy-maker community:

How can organizations shift from incident-centric security models to frameworks focused on persistent access risk?

What level of identity governance, telemetry retention, and monitoring visibility is required to detect long-horizon intrusions?

How should governments and industry address the systemic exposure created by digital supply chains and cloud dependencies?

What defensive strategies are most effective when adversaries prioritize access preservation over immediate disruption or theft?

How can organizations better align executive-level risk understanding with the operational realities faced by security teams?

Final Observations

Chinese-nexus cyber activity increasingly reflects long-horizon strategic planning rather than isolated campaigns. These operations demonstrate a consistent behavioral cadence: attackers establish access, evaluate its strategic value, and maintain it patiently until conditions justify escalation.

For defenders, the challenge is not simply preventing individual intrusions but managing continuous exposure across complex digital environments. By shifting from actor-centric attribution toward behavioral pattern recognition, organizations can better detect emerging threats, prioritize defensive investments, and strengthen resilience against the evolving dynamics of strategic cyber competition.

The goal and insights from this research hopefully turn a complex threat ecosystem into actionable defensive guidance — for policy makers, CEOs, CISOs, analysts, and businesses — helping organizations anticipate how Chinese nexus actors operate and strengthen against the next phase of strategic cyber competition.

1 International Monetary Fund: Datamapper [referenced Jan. 2026] https://www.imf.org/external/datamapper/NGDPD@WEO/OEMDC/ADVEC/WEOWORLD

APPENDIX 1

Darktrace Cybersecurity Attribution Framework

Navigating the Complex Terrain of Cyber Attribution

Cyber attribution—the process of identifying the origin and perpetrators of malicious cyber activity—is one of the most complex and consequential tasks in modern threat intelligence. It is not merely a technical exercise but a strategic and often political decision, especially when state actors are involved.

Attribution underpins accountability, informs defensive and offensive cyber strategies, and can trigger diplomatic, legal, or even military responses.

National Frameworks and Strategic Imperatives

Nation-states have developed robust attribution capabilities within their national security apparatus. For example, the US National Security Agency (NSA) and the UK National Cyber Security Centre (NCSC) possess extensive human, technical, and financial resources dedicated to attribution, often working in tandem with private-sector threat intelligence vendors. Nation-state actors have increasingly moved to public attribution of threat actors, particularly other nation-state groups, as a strategic tool in the international arena.

This geopolitical undercurrent has provided ample conversation within the industry around public attribution, and various frameworks have been publicized with many private companies sharing their own. Similarly, the IISS Cyber Power Matrix outlines a six-tier framework—from non-state actors to state-integrated entities—used to assess actor–state relationships and guide attribution decisions. Darktrace is publishing the attribution framework to promote transparency, support analyst reproducibility, and align with public-sector expectations for defensible methodology.

Challenges in Attribution

Despite the existence of national and private-sector attribution frameworks, attribution remains fraught with challenges

- **Technical complexity:** Threat actors use obfuscation, false flags, and multi-stage operations to mask their identities.
- **Legal ambiguity:** International law lacks clear standards for cyber attribution, especially for operations below the threshold of armed conflict.
- **Political risk:** Misattribution can escalate tensions or undermine credibility, making governments cautious about public declarations.
- **Resource intensity:** Attribution demands deep forensic analysis, geopolitical context, and corroboration across intelligence domains.

These challenges often result in responsibility gaps, where malicious activity is detected but cannot be confidently linked to a perpetrator or state.

The Role of the Darktrace Cyber Attribution Framework

Attributing cyber threats requires a multi-dimensional approach that blends behavioural analytics, external intelligence, and contextual overlays. We are not assessing intent. We are describing repeatable behavioral patterns that persist across years, sectors, and tooling changes.

The Darktrace Cyber Attribution Framework supports this by aligning by aligning Darktrace's proprietary AI-driven detection models with known threat actor TTPs enabling analysts to build attribution confidence across six core pillars:

01 INFRASTRUCTURE

This pillar focuses on the network infrastructure used by threat actors—such as IP addresses, domains, and VPNs. Analysts assess whether the infrastructure:

- Has been reused across campaigns
- Is linked to known threat actor clusters
- Matches patterns from previous incidents

Example Actor:

APT41 (aka Barium / Winnti) Reference: CISA (2021) – Chinese exploitation of Pulse Secure VPN

Example Darktrace Models:

- Compromise::Beaconing to Rare Endpoint
- Anomalous Connection / Unusual External Connection

02 MALWARE / TOOLING

This pillar examines the malicious software and tools deployed during an attack. Analysts look for:

- Known malware families (e.g., PlugX, gGh0st RAT)
- Custom tooling or side-loaded DLLs
- IoCs

Example Actor:

APT41 (Winnti family clusters) Reference: Mandiant (2020) – Dual espionage and cybercrime operations

Example Darktrace Models:

- Compromise::Unusual Process Execution
- Anomalous File Transfer
- Anomalous File Download

3. TTPS

TTPs reflect how an attacker operates. This includes activity such:

- Credential theft
- Scheduled task creation
- Data exfiltration

Example Actor:

APT29 (Nobelium) Reference: Palo Alto Unit 42 (2021) – SolarStorm supply chain attack

Example Darktrace Models:

- Unusual Admin Credential Use
- Compromise::Suspicious Scheduled Task
- Unusual Data Exfiltration

04 VICTIMOLOGY

This pillar analyzes who is being targeted—by sector, geography, or strategic value. Attribution is strengthened when:

- Victim profiles align with known actor interests
- Targeting patterns match geopolitical motives

Example Actor:

- APT10 (Cloud Hopper) Reference: UK NCSC Advisory (2020) – Targeting of telecoms

Analyst Context / Enrichment

- Analyst overlays context manually (sector, geography)

05 LANGUAGE / ARTEFACTS

This pillar focuses on linguistic and technical artefacts such as:

- Lure documents in specific languages
- File naming conventions
- Compile times and metadata

Example Actor:

- Mustang Panda - Reference: CrowdStrike (2021) – Chinese-language lures and malicious RAR archives
- Analyst enrichment (e.g., lure docs, filenames, compile times)

06 EXTERNAL CORROBORATION

This final pillar involves cross-validating findings with:

- Vendor threat reports
- Government advisories
- Intelligence assessments

Example Reference:

US DNI (2021) – People's Republic of China (PRC)persistent targeting of healthcare

Analyst Overlay:

Analyst compares findings with external sources (vendor/government reports)

Confidence Levels in Attribution

Attributing a cyber incident to a specific threat actor or nation-state is rarely a binary decision. Instead, it involves assigning a confidence level based on the quality, quantity, and convergence of evidence across multiple investigative pillars. Intelligence agencies like the NSA (US) and NCSC (UK) often use tiered confidence frameworks—such as low, moderate, and high confidence—to communicate how strongly the available data supports a given attribution.

These levels are informed by technical indicators (e.g., malware signatures, infrastructure reuse), behavioral patterns (TTPs), geopolitical context, and corroboration from external sources.

A low confidence attribution might rely on circumstantial evidence or weak correlations, while a high confidence attribution typically involves multiple independent sources, consistent actor behavior, and strategic alignment with known motives. Importantly, confidence levels help analysts and decision-makers calibrate their responses, ensuring that defensive actions, public statements, or diplomatic consequences are proportionate to the certainty of the attribution.

By embedding confidence scoring into the investigative workflow, analysts can maintain transparency, reduce bias, and support repeatable, defensible conclusions—especially when using structured frameworks, which guide attribution through checkpoints aligned with these principles.

- **High Confidence →** Multiple independent, corroborating indicators align (e.g., infrastructure reuse + malware family + victimology).
- **Moderate Confidence →** Several indicators align, but evidence gaps remain.
- **Low Confidence →** Indicators weak, circumstantial, or based primarily on third-party reporting.

EXAMPLE DECISION TREE ATTRIBUTION WORKFLOW

Step 1: Infrastructure

If “Compromise::Beaconing to Rare Endpoint” or “Anomalous Connection” → pivot to WHOIS/OSINT to check infra reuse.

Step 2: Malware/Tooling

If “Unusual Process Execution” or “Anomalous File Transfer” → check file hashes & code similarities against public reporting.

Step 3: TTPs

If “Unusual Admin Credential Use” or “Suspicious Scheduled Task” → map to MITRE ATT&CK, compare w/ known APT tradecraft.

Step 4: Victimology

If targeting aligns with historical APT priorities (e.g., telecommunications, defense, non-governmental organizations (NGOs)) → strengthens attribution.

Step 5: Artefacts

Analyst checks lure documents, filenames, compile times → supports attribution narrative.

Step 6: Corroboration

Compare findings with external threat reporting (CISA, NCSC, Mandiant, CrowdStrike, Palo Alto).

Confidence Assignment: Simple Calculation

- 3+ strong pillars align → High Confidence
- 2 pillars align → Moderate Confidence
- 1 weak pillar → Low Confidence

Pillar	Spoofability	Adjusted Weight	Tier
Infrastructure	Low	0.80	Tactical
Malware / Tooling	Medium	0.80	Tactical
TTPs	Medium	0.85	Operational
Victimology	Low	0.70	Operational
Language / Artefacts	High	0.40	Tactical
External Corroboration	Low	0.95	Strategic

Analyst Confidence Levels

Confidence assessment:
Introduce weights for each step based on:

- Ease of spoofing (e.g., language artefacts are easier to fake than infrastructure reuse).
- Historical reliability (e.g., malware/tooling tends to be more actor-specific but not necessarily).
- Analyst confidence (based on past investigations).

4. EXAMPLE (HIGH CONFIDENCE) WORKFLOW: DARKTRACE → ATTRIBUTION

Scenario: Darktrace flags beaconing inside a telecom.

Step 1: Infrastructure → Compromise::Beaconing to Rare Endpoint. Infrastructure overlaps with CISA Alert A21-110A.

Step 2: Malware → Compromise::Unusual Process Execution. DLL sideloading matches Mandiant APT41 report.

Step 3: TTPs → Suspicious Scheduled Task. Persistence technique noted by Palo Alto Unit 42 in SolarStorm.

Step 4: Victimology → Telecom sector, aligned with NCSC Cloud Hopper advisory.

Step 5: Artefacts → Chinese-language decoy file, similar to CrowdStrike Mustang Panda campaign.

Step 6: Corroboration → ODNI 2021 ATA confirms PRC interest in telecommuncations.

Example Formula

Score = (Adjusted Weight × Evidence Strength) + Boost

For Infrastructure:

- Adjusted Weight = 0.80
- Evidence Strength = 1.00 (strong match with CISA alert)
- Boost = 0.10 (due to corroboration with external campaign data)

Each pillar was supported by strong evidence:

- **Infrastructure:** Overlap with CISA alert → +0.10 boost
- **Malware / Tooling:** DLL sideloading matches Mandiant → +0.15 boost
- **TTPs:** Matches Palo Alto SolarStorm tradecraft
- **Victimology:** Telecommunications sector targeted, aligned with NCSC
- **Language / Artefacts:** Chinese-language lure → +0.10 boost
- **External Corroboration:** ODNI confirms PRC interest
- **0.90+0.95+0.85+0.70+0.50+0.95=4.85**
  4.85 = High Confidence

- `Score = (Adjusted Weight x Evidence Strength) + Boost`

Pillar	Adjusted Weight	Evidence Strength	Boost
Infrastructure	0.80	1.00	0.10
Malware / Tooling	0.80	1.00	0.15
TTPs	0.85	1.00	0.00
Victimology	0.70	1.00	0.00
Language / Artefacts	0.40	1.00	0.10
External Corroboration	0.95	1.00	0.00

Example Final Attribution Statement (High Confidence):

"We assess with high confidence this intrusion is linked to PRC-affiliated espionage (APT41 cluster), based on infrastructure overlap (CISA), malware/tooling (Mandiant), persistence tradecraft (Palo Alto Unit 42), victimology (NCSC), and artefacts (CrowdStrike)."

Example (Moderate Confidence) Workflow: Darktrace → Attribution

- Infra overlaps with known actor infra (CISA alert)
- Malware partially matches known tooling (Mandiant)
- No TTPs observed
- Victimology aligns with sector targeting
- No artefact match
- Strong external corroboration
- **Raw Score: 3.16**
  Max Score 4.85
  Normalzed score of .65 > Moderate

Pillar	Weight	Strength	Boost
Infrastructure	0.80	1.00	0.10
Malware / Tooling	0.80	0.80	0.15
TTPs	0.85	0.00	0.00
Victimology	0.70	0.60	0.00
Language / Artefacts	0.40	0.00	0.10
External Corroboration	0.95	1.00	0.00

Example (Low Confidence) Workflow: Darktrace → Attribution

- No infrastructure or malware match
- Weak victimology and artefact indicators
- No TTPs or external corroboration
- **Raw score .90**
  Max Score 4.85
  Normalised score of .19 = Low confidence

Pillar	Weight	Strength	Boost
Infrastructure	0.80	1.00	0.10
Malware / Tooling	0.80	0.80	0.15
TTPs	0.85	0.00	0.00
Victimology	0.70	0.60	0.00
Language / Artefacts	0.40	0.00	0.10
External Corroboration	0.95	1.00	0.00

Why This Framework Works

- **Darktrace-first:** Rooted in real model detections analysts see daily.
- **Reference-backed:** Linked to authoritative sources (CISA, NCSC, Mandiant, Palo Alto, CrowdStrike, ODNI).
- **Confidence-calibrated:** Uses structured IC-style language for consistent attribution statements.

Darktrace built this framework now because the threat landscape has fundamentally shifted. What we're seeing across our customer base is no longer a collection of disconnected incidents — it's sustained, state-aligned activity that blends espionage, supply-chain compromise, and long-term operational positioning. In that environment, attribution isn't a political exercise; it's a risk-governance requirement.

Organizations need to understand whether they're dealing with opportunistic intrusion, commercially-motivated compromise, or a campaign aligned to a nation-state's strategic objectives. That distinction shapes everything from how you prioritize controls to how you think about resilience, disclosure, and long-term exposure.

The strength of this framework is that it's grounded in real behavior, not static indicators.

Threat actors can spoof infrastructure, reuse malware, or plant misleading artefacts — but they struggle to mask the deeper patterns of how they operate over time. By anchoring attribution in those durable behavioral signals and cross-checking them against public-sector and vendor intelligence, we can build confidence based on convergence, not guesswork.

The result is a process that's structured, repeatable, and transparent: analysts move through a defined set of pillars, understand why each matters, and arrive at confidence levels that stand up to scrutiny.

We're also explicit about the limitations. Attribution in cyberspace will always involve ambiguity: shared tooling, actor overlap, false flags, and deliberate misdirection are all part of the terrain. That's why we validated this framework through retrospective case reviews — rerunning historic intrusions through the six-pillar method to confirm it produces consistent, defensible outcomes.

The goal isn't absolute certainty; it's disciplined, calibrated judgment. This gives defenders a reliable way to understand the strategic context behind intrusions and respond with the appropriate level of urgency, investment, and executive attention.

APPENDIX 2

List of Indicators of Compromise

Note: TLP Clear Indicators are below. Some indicators we have kept internally for operational security. If you feel your organization would benefit from TLP Amber indicators and provide reasonable justification, please email crimsonecho@darktrace.com.

IoC	Type	Confidence
shell.cdn-sina[.]tw	Hostname	High
xxl17z.dnslog[.]cn	Hostname	High
asljkdqhkhasdq.softether[.]net	Hostname	High
aar.gandhibludtric[.]com	Hostname	High
plugins.jetbrians[.]net	Hostname	High
dscriy.chtq[.]net	Hostname	High
cybaq.chtq[.]net	Hostname	High
ns1.akacur[.]tk	Hostname	High
ns2.akacur[.]tk	Hostname	High
trkbucket.s3.amazonaws[.]com	Hostname	High
tnegadge.s3.amazonaws[.]com	Hostname	High
fconnect.s3.amazonaws[.]com	Hostname	High
times.windowstimes[.]online	Hostname	High
applr-malbbal.s3.ap-northeast-2.amazonaws[.]com	Hostname	High
beansdeals-static.s3.amazonaws[.]com	Hostname	High
bringthenoiseappnew.s3.amazonaws[.]com	Hostname	High
brandnav-cms-storage.s3.amazonaws[.]com	Hostname	High
abode-dashboard-media.s3.ap-south-1.amazonaws[.]com	Hostname	High
maxdesigns[.]top	Hostname	High

IoC	Type	Confidence
asdasw21[.]icu	Hostname	High
micheeasodh[.]top	Hostname	High
a.micheeasodh[.]top	Hostname	High
lsls.casacam[.]net	Hostname	High
meetls.kozow[.]com	Hostname	High
vals.bumbleshrimp[.]com	Hostname	High
4.232.170[.]137	IP	High
185.238.251[.]244	IP	High
45.251.240[.]55	IP	High
137.175.30[.]36	IP	High
192.74.254[.]229	IP	High
107.148.219[.]227	IP	High
107.148.149[.]156	IP	High
107.148.219[.]54	IP	High
107.148.219[.]55	IP	High
103.27.108[.]62	IP	High
103.27.110[.]83	IP	High
103.13.28[.]40	IP	High
89.31.121[.]101	IP	High
16.162.188[.]93	IP	High
5.181.132[.]95	IP	High
63.245.1[.]34	IP	High
158.247.213[.]167	IP	High
158.247.199[.]185	IP	High
94.131.110[.]28	IP	High
193.56.255[.]214	IP	High
15.188.246[.]198	IP	High
149.104.23[.]171	IP	High
192.210.239[.]172	IP	High

IoC	Type	Confidence
107.155.56[.]87	IP	High
156.244.28[.]153	IP	High
154.90.63[.]250	IP	High
45.76.191[.]59	IP	High
45.77.170[.]188	IP	High
38.54.29[.]25	IP	High
154.205.139[.]12	IP	High
45.76.209[.]205	IP	High
64.176.59[.]232	IP	High
149.28.28[.]9	IP	High
17d65a9d8d40375b5b939b60f21eb06eb17054fc	SHA1	High
da23dab4851df3ef7f6e5952a2fc9a6a57ab6983	SHA1	High
fa645f33c0e3a98436a0161b19342f78683dbd9d	SHA1	High
4a8b8a164e20748e23fbded8b048bacb9c3d715c	SHA1	High
b5367820cd32640a2d5e4c3a3c1ceedbbb715be2	SHA1	High
8da0489e4d6307b461cb4090dd661d0fbeef9928	SHA1	High

APPENDIX 3

Meta-Model Creation

This project aims to provide a preliminary framework for identifying Chinese-nexus activity, irrespective of specific subgroup delineation. In pursuit of this objective, analysts used the output of threat hunting efforts to create Darktrace / NETWORK™ models for generalized Chinese-nexus activity in a network environment. The model is a "meta-model" signifying that specific combinations of other underlying models within a delineated period will result in an alert.

The models that compromise the subcomponents were identified from noted model alerts and kill chain progressions of medium to high confidence cases and specific techniques and procedures that featured commonly in OSINT sources as part of the "literature review" process.

The pairings of models within each subcomponent, indicating what combination of activities can contribute to the model alerting, were informed by technique co-occurrence metrics. The duration of time Darktrace allows for subcomponents to fire also reflects dwell time data. More limited indicators such as IP ranges/ASN blocks, cloud providers, file naming patterns, and more also feature in some component parameters to prevent erroneous alerting.

The two models detect core kill chain elements seen frequently across higher confidence cases. The first model focuses on initial exploitation, execution, and foothold establishment operations. While this model aims at encapsulating "shorter term" activity, this does not limit detection should the customer reside in a sector more likely targeted for long-term persistence. Different combinations of underlying models can result in an alert of the model, and each combination identifies different short-term breach initiation sequences identified previously.

For example, alerts likely to feature in perimeter device exploitation, such as those related to internet-facing system file downloads and new user agents feature within a core component of the model. An alert within this category could trigger the overall model if additional model alerting involving tooling ingress suggestive of future DLL sideloading, certutil/BITS utility for tooling ingress, and DNS tunneling activities. Other components models identify activities like tooling ingress suggestive of DLL manipulation techniques, DNS tunneling, and outbound remote protocols such as SSH, TLS related beaconing, and new user agents. Again, more signatures-based indicators are used to prevent erroneous and unrelated alerting.

The longer-term model uses the same approach, whereby underlying model combinations within a certain timeframe will trigger the overall meta-model. While the short-term model focuses solely on preliminary kill chain elements like initial access and foothold establishment, the subsequent model can only alert if both a preliminary element and at least one secondary kill chain element are observed. These subsequent components focus on egress and reconnaissance/lateral movement activity.

It is assumed that activity triggering the short-term model, even if featured as part of an attempted longer duration compromise, will be detected and addressed. Therefore, the preliminary element for the longer duration model has broader parameters.

This detection variation also allows for repeated HTTP/S beaconing related activity, in conjunction with other variables to satisfy the initial element. These beaconing/C2-related alerts rely on Darktrace's behavioural analytics and communication-pattern monitoring. This provides resilience in scenarios where additional signals that typically strengthen a China-nexus assessment are absent from the available telemetry.

Due to the nature of Darktrace's model detection mechanism, models trigger when relevant behaviors appear, not after prolonged durations, so ultra-long persistence sequences aren't treated as a single detection. In these cases, this use case model is intentionally designed to provide early warning rather than cover every phase of a full attack lifecycle. Again, the goal in this instance is not a model that detects the entire sequence of an attack involving these groups. Rather, it is meant to provide a confident, early detection mechanism for such operations.

Therefore, the duration (i.e., model delay to allow for either egress or recon/lateral movement components to alert) reflects the dwell time metrics for initial activity to the first instance of egress or lateral movement. Egress activity identified within compromise cases had a narrower band of corresponding model alerts. These raw detections will satisfy the second component for egress activity. Lateral movement model alerts largely focus on LOLBin usage, Active Directory reconnaissance operations such as uncommon LDAP activity and DRSGetNCChanges requests, and registry/service/task operations for persistence, amongst other factors.

It is worth reemphasizing that the underlying components themselves represent anomaly-based detection mechanisms. Therefore, not only is Darktrace able to detect particular combinations of kill chain activity, but also these components avoid the pitfalls of signatures- or heuristic-based detection. In practice, this combination of tactic/technique and sequential alerting takes into account normal patterns of life for the device. For example, the use of a LOLBin alone is not sufficient to satisfy a lateral movement component. Rather, the usage of the LOLBin itself must also be unusual/unexpected for the device. Similarly, reliance on specific cloud providers for tooling download, by itself, will not trigger a parameter. Instead, the cloud provider endpoint must be rare for the network, and the specific download/ingress activity must differ from the device's established pattern of life. In doing so, this meta-model encompasses a broader set of potential deviations in activity, while also prevents unrelated activity from alerting.

APPENDIX 4

Anomaly Detection Examples Prior to CVE Disclosure

More than 48,000 Common Vulnerabilities and Exposures (CVEs) were reported for 2025 (a 20.6% increase year-on-year). The gap between exploitation of a zero-day and vulnerability disclosure can often be considerable, making retroactive identification of successful exploitation within networks a persistent challenge.

Abnormal behaviors within networks or systems, such as unusual login patterns or unexpected data transfers, can indicate attempted cyberattacks, insider threats, or compromised assets. As Darktrace does not rely on predefined rules or known signatures, it is able to detect malicious activity that deviates from established behavioral norms, even when full context about a specific device or asset is unavailable.

By continuously analyzing behavioral patterns, Darktrace enables organizations to identify and contain potential exploits at an early stage. Leveraging these anomaly detection parameters, Darktrace analysts conduct retrospective analysis to better understand detections across the broader threat landscape and to enrich findings with additional context. This behavioral approach also supports pre-CVE detection, allowing Darktrace to identify emerging or previously unknown exploitation techniques based on attacker behavior, before vulnerabilities are formally disclosed or assigned a CVE.

Darktrace demonstrated its ability to identify malicious activity prior to public disclosure in several notable cases relevant to this research, as outlined in the table below.

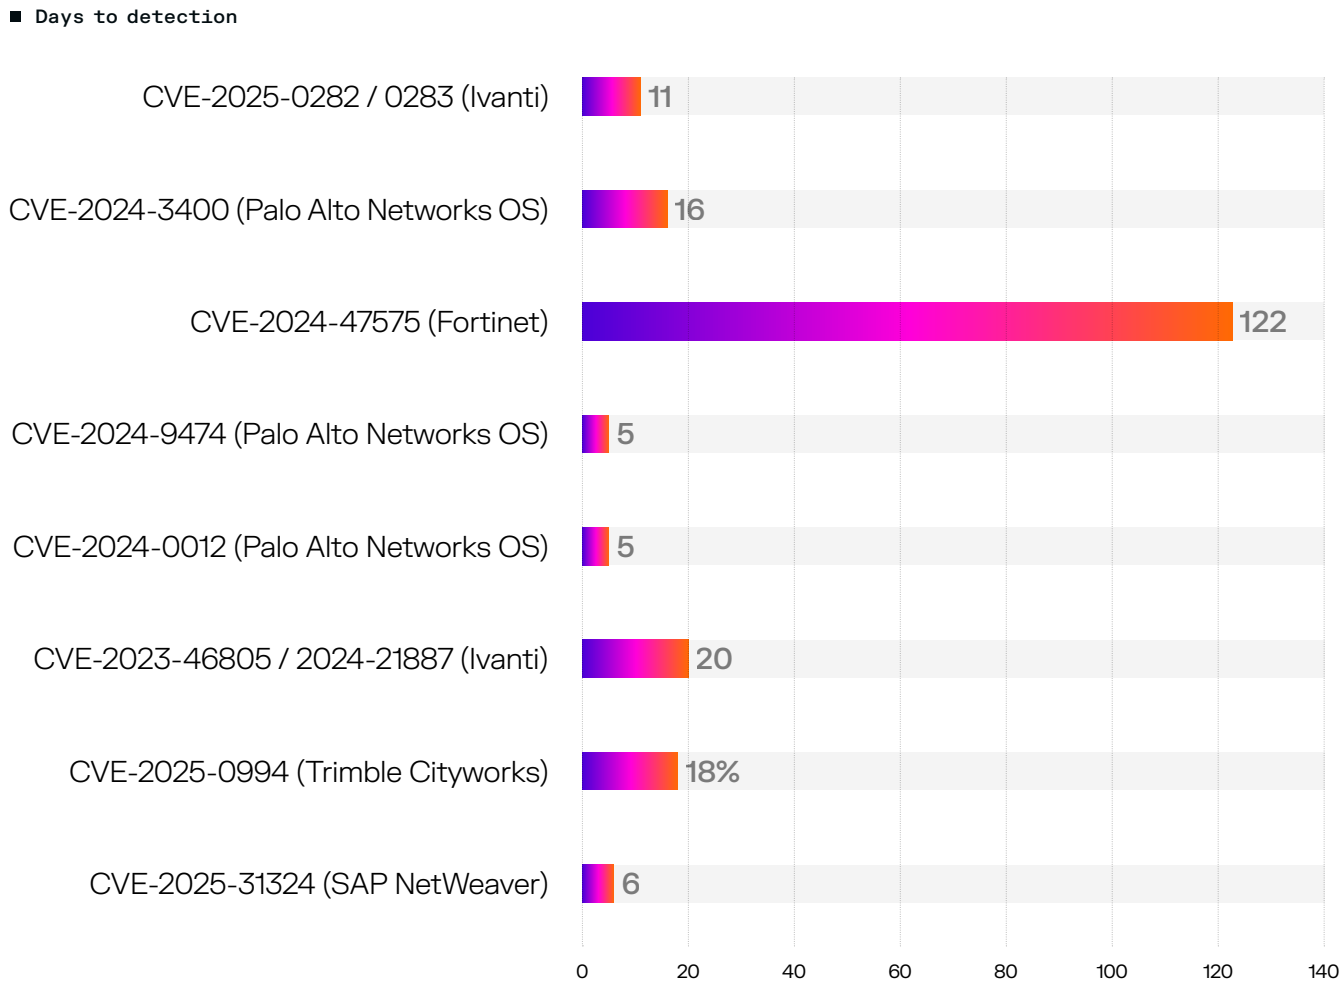

APPENDIX 5

Classifying Organizations as Critical Infrastructure

To classify organizations as critical infrastructure, Darktrace used the victim country of origin's national or regional policy. For example, for cases in the United States, the 16 critical sectors in national security policy and how it is applied from the CISA

In EU cases, Darktrace utilized the European Union Agency for Cybersecurity's taxonomy.

Sector Designations:

Information Technology

- IT products and services including software/hardware products, non-critical Ssoftware-as-a-service (SaaS)/platform-as-a-service (PaaS)/infrastructure-as-a-service (IaaS) providers, managed service providers

Manufacturing

- Critical and non-critical manufacturing
  - Critical manufacturing including steel plants and metal goods, industrial machinery, cement/construction/synthetic material producers, high tech hardware producers including processors, microchips etc.
  - Non-critical manufacturing including textiles, consumer goods, and secondary/packaged food and beverage producers

Arts, entertainment and recreation

- Movie/television/music studios, cinema related services including audio engineering, visual effects, production companies and transcription services, entertainment and sporting venues, sporting leagues and governing bodies, gaming/casinos, museums and cultural sites/institutes

Education

- Public and private primary, secondary, and tertiary educational institutions

Agriculture, Forestry and Fishing

- Larger scale, industrial farming, primary food producers, parks, wildlife management, agricultural/livestock machinery

Professional, scientific and technical services

- Consulting firms, technical advisory organizations, non-health related research institutes

Electricity, gas, steam and air conditioning supply

- Any energy, water or cooling entity including power grids, energy markets, oil and gas producers/refineries, water treatment and power plants, solar panel operators and producers, windfarms, nuclear power plants

Media and Broadcasting

- Broadcasting and print media agencies and organizations, publications, newspapers, television and radio stations

Financial and insurance activities

- Banks, credit unions, lender institutions, credit agencies, stock markets, insurance agencies, and other financial adjacent organizations

Transportation Infrastructure

- Airports, airlines, ports, railways/operators, bridge/light-house/port operators, traffic systems, and airlines, bussing organizations

Public Administration and Government Services

- Federal, provincial, state, local, tribal, and territorial government agencies, publicly funded projects and organizations

Legal Services

- Law firms, advisory organizations, and other legal services

Health Products and Services

- Providers including hospitals anc clinics, healthcare services and administration, pharmaceutical manufactures and biomedical research organizations

Telecommunications / Digital Infrastructure

- Telecommunications products and services including phone and internet providers, critical cloud operators, public key infrastructure entities (certificate authorities, intermediate authorities, key management providers etc.)

Shipping, Storage, and Logistics

- Trucking, railway, and nautical shipping/transportation organizations for goods and products

Mining and Quarrying

- Raw mineral and ore mining and extraction
  - Includes both rare/critical and non-critical elements and ores

Other service activities

- Services not directly related to health, professional/scientific, legal, financial, or business needs and concerns

Defense Sector

- Weapons and munitions manufacturers, security contractors, public sector/defense consulting advisory services

Construction

- Construction and design firms

CISA CNI Categories:

Sectors and organizations grouped according to CISA CNI sector designations:

- Information Technology Sector
- Critical Manufacturing Sector
- Commercial Facilities Sector
- Food and Agriculture Sector
- Financial Services Sector
- Energy Sector
- Communications Sector
- Transportation Systems Sector
- Government Facilities Sector
  - Government Facilities Sector - Education Facilities Subsector
- Healthcare and Public Health Sector
- Defense Industrial Base Sector

Reference: https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors

About Darktrace

Darktrace is a global leader in AI cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,300 employees who protect nearly 10,000 customers globally. To learn more, visit www.darktrace.com.

North America: +1 (415) 229 9100 | Europe: +44 (0) 1223 394 100 | Asia-Pacific: +65 6804 5010 | Latin America: +55 11 4949 7696

darktrace.com | info@darktrace.com

© 2026 Darktrace Holdings Limited. All rights reserved.