What Is Cyber Threat Intelligence? Defining the Field, Disciplines, and Intelligence Cycle

FIRST Cyber Threat Intelligence Special Interest Group (CTI SIG)

Version 1.0

Abstract

Cyber Threat Intelligence (CTI) is the systematic collection, processing, analysis, and dissemination of information about threats to an organization's operations in cyberspace, produced to inform decision-makers across all levels. It sits at the intersection of cybersecurity and intelligence studies, borrows its process and much of its vocabulary from government and military intelligence, and yet operates almost entirely in the commercial space, with consumers who may have no relationship to the producer at all. This white paper defines CTI and the terms it depends on, summarizes the government and military intelligence model that CTI inherits, and then describes what is specific to CTI: its consumers and levels, the transformation of raw data into intelligence, the origin of its sources, its technical and non-technical disciplines, the functional roles that apply them, and the CTI cycle from planning through evaluation. The distinction between disciplines (the skills) and functional roles (the way the skills are applied, legally and organizationally) is a central contribution, because in cyberspace the same technique can be lawful or unlawful depending on authorization. The paper closes by noting that most commercial CTI effort today goes into processing and exploitation rather than analysis, and that consistent training and terminology are prerequisites for moving the field forward.

Keywords: Cyber Threat Intelligence, Intelligence Cycle, structured analytic techniques, intelligence disciplines

1. Introduction

CTI is an emerging field that fuses several traditional intelligence disciplines with information security practice, and it is commercial in nature. Its purpose is to enhance an organization's understanding of the motives, targets, and attack behaviors observed on

the Internet so that defenders at the tactical level can put in context what they see on their own networks, and so that executive management can maintain a strategic outlook and allocate resources accordingly (FIRST CTI SIG, 2024; Tsantes & Ransome, 2023, Tzvetanov & Dietz, 2025) Cyber intelligence consists of actionable information used to protect an organization's cyber presence, attack a competitor's or enemy's cyber presence, or both.

The need for such a function stems from the environment in which military and commercial organizations operate. Adversaries are sophisticated and well-resourced, the technology changes quickly, and the digital ecosystem in which a company operates is interconnected far beyond its own perimeter. At the same time, the volume of entities, their transnational nature, and in many cases legal constraints mean that the defense and law enforcement institutions of any single government cannot protect their interests at scale. Consequently, while governments struggle to implement and share cyber intelligence, companies have had to build their own intelligence capabilities.

That capability has developed unevenly. CTI grew out of incident response rather than the intelligence profession; vendor marketing popularized it before practitioners defined it, and it has since suffered from inconsistent terminology, training, and reporting standards. This paper defines the field and maps its parts. It proceeds as follows: foundational terms; a summary of the government and military intelligence model and its collection disciplines; the definition and scope of CTI; the disciplines and functional roles of CTI; and the CTI cycle. Throughout, the government and military model serves as the reference against which commercial practice is compared; it is important to note where the mapping holds and where it breaks down, because several recurring CTI problems arise precisely where the two diverge.

2. Foundational Terms

Before defining CTI, we must define the terms that make it up. Each of the following has an authoritative source, and the definitions below are taken from those sources rather than from vendor usage.

2.1 Cyber and cyberspace

The word "cybernetics" traces back to Norbert Wiener's 1948 book of that name. Still, the prefix did not enter common use until the late 1980s and early 1990s, when William Gibson popularized "cyberspace" in the Sprawl trilogy. In current usage, "cyber" refers to information and communications systems and networks, and "cyberspace" is the environment consisting of electronic devices and communication networks together with the interaction of people and services within them. It is a complex environment without

a clear locality, although it is represented by physical components in specific localities (Hogan & Newton, 2015).

2.2 Cybersecurity

Cybersecurity is a multidisciplinary field focused on protecting computer systems, networks, devices, and data from unauthorized access and modification, and on ensuring authenticity, availability, and nonrepudiation. The National Institute of Standards and Technology (NIST) defines it as preventing damage to, protecting, and restoring computers, electronic communications systems and services, and the information contained therein to ensure availability, integrity, authentication, confidentiality, and nonrepudiation (Hogan & Newton, 2015). The scope matters: if only data were in scope, the term would be "information security," and if only equipment were in scope, it would be "computer security." The International Organization for Standardization (ISO) goes further, including people, society, and nations, producing an overarching definition that subsumes network, Internet, and web security (International Organization for Standardization, 2022a, 2022b).

2.3 Threat and cyber threat

A threat is any situation or event that could harm an individual or an organization, including its mission, functions, reputation, image, or assets. In the cybersecurity context, a cyber threat can compromise, destroy, disclose, or modify information, or deny service; the term also covers the likelihood that a threat actor successfully exploits a specific vulnerability in an information system (Joint Task Force Transformation Initiative, 2012).

2.4 Intelligence

Intelligence is the process of planning, collection, processing, integration, evaluation, analysis, and interpretation of available information concerning foreign nations and hostile or potentially hostile forces and elements, undertaken to support decision-making, policy formulation, and operations; the term also denotes the product of that process and the organizations that carry it out (Joint Chiefs of Staff, 2010; ODNI, 2024). Intelligence therefore focuses on understanding an adversary's capabilities, intentions, and activities. Its primary consumer is the executive branch, including the armed forces and law enforcement, and it operates at the tactical, operational, and strategic levels. The process itself is the Intelligence Cycle, summarized in Section 3.

2.5 Cyber Threat Intelligence

CTI is the process of collecting, analyzing, and interpreting information related to cybersecurity threats and vulnerabilities to give decision-makers at different levels of an

organization actionable insight into the threats it faces: malware, vulnerabilities in software and systems, emerging trends in cybercrime and nation-state activity, and the tactics, techniques, and procedures (TTPs) of adversaries. That insight allows the organization to identify and mitigate risk proactively, prioritize security investment, and improve incident response.

The most general and succinct definition comes from the Cyber Threat Intelligence Special Interest Group (CTI SIG) of the Forum of Incident Response and Security Teams (FIRST), and it is the definition adopted here:

> *Cyber Threat Intelligence is systematic collection, analysis and dissemination of information pertaining to a company's operation in cyberspace and to an extent physical space. It is designed to inform all levels of decision makers. The analysis is designed to help keep situational awareness about current and arising threats. (FIRST CTI SIG, 2024)*

Three features of this definition deserve emphasis. First, CTI is a process, not a feed or a product. Second, its scope covers the company's operations in cyberspace, including physical space to the extent it bears on those operations. Third, its consumers are decision makers at all levels, which is why CTI, like intelligence generally, is organized into strategic, operational, and tactical levels.

3. The Reference Model: Government and Military Intelligence

CTI inherits its process, its levels, and most of its vocabulary from government and military intelligence. A summary of that model is therefore necessary; a fuller treatment belongs to the intelligence literature and is not repeated here.

3.1 Definition, consumers, and levels

Joint Publication 1-02 defines intelligence as a process and as its product (Joint Chiefs of Staff, 2010); the Office of the Director of National Intelligence (ODNI) definition is outcome-oriented, emphasizing understanding of an adversary's capabilities, intentions, and activities, and extends the scope to non-state activity such as terrorism and cyber threats (ODNI, 2024). The two are complementary. In both, intelligence informs decision makers, primarily in the executive branch: the President, the National Security Council, the armed forces, and federal and local agencies, notably law enforcement, with Congress as a secondary consumer.

Because consumers differ in the decisions they make, intelligence is organized in three levels (FIRST CTI SIG, 2024; Joint Chiefs of Staff, 2010):

Strategic intelligence analyzes long-term, high-level trends with substantial implications for strategy, policy, and plans; it operates at a slow cadence, considers geopolitical, economic, social, and technological factors, and serves those who set the organization's strategy.

Operational intelligence supports ongoing operations in real or near-real time, warns of imminent events, and maintains situational awareness within a specific operational environment.

Tactical intelligence supports specific actions against specific threats or targets in a localized or time-sensitive context, for frontline personnel, field commanders, and operational units.

3.2 The Intelligence Cycle

The Intelligence Cycle is a five-step cyclical model of the intelligence process, with evaluation and feedback running continuously alongside it. Shorter tactical cycles, notably Find, Fix, Finish, Exploit, Analyze, and Disseminate (F3EAD), may be intertwined with it.

Step	What happens	Representative outputs
Planning and Direction	Objectives and priorities are set from the commander's information needs, decision support templates, and critical information requirements; resources, timelines, and roles are assigned.	Priority Intelligence Requirements (PIRs), Requests for Information (RFIs)
Collection	Data is gathered from the intelligence disciplines (Section 4) according to a collection plan; sources may be evaluated and ranked.	Raw data
Processing and Exploitation	Processing makes raw data usable (developing film, converting formats, translating, decrypting); exploitation extracts information from it (an order of battle from imagery).	Information

Step	What happens	Representative outputs
Analysis and Production	Information is analyzed for patterns, trends, and threats and becomes intelligence, using structured analytic techniques and machine analysis; output may be forensic, a current assessment, or a forecast.	Intelligence products
Dissemination and Integration	Products are packaged, tailored to the consumer, distributed with a level of precedence, and integrated into planning and operations. Uncertainty is expressed with Words of Estimative Probability (WEP) and Levels of Confidence in Assessment (LCA). This step informs the planning of the next cycle.	Reports, briefings, decisions
Evaluation and Feedback (continuous)	Personnel at every level evaluate their performance and its impact, in collaboration with consumers.	Adjustments to every step

3.3 Raw data, information, and intelligence

The distinction among these three terms is the cycle's conceptual core. Raw data is the unaltered output of a sensor: undeveloped film from an optical satellite, a field-strength measurement, or the Internet Protocol (IP) address of a node on a network. On its own, it has no meaning. Processing makes it usable (the film is developed; the IP address is enriched with its autonomous system and passive Domain Name System history), exploitation extracts information from it, and analysis turns information into intelligence: a geospatial analyst produces the disposition of adversary forces from the imagery, and a network analyst cross-references the IP address with the context in which it was observed. Data that yields no information is discarded as noise.

3.4 Information classification

To protect sources and methods, the intelligence community labels information by its sensitivity to national defense or foreign relations, which governs how it may be stored, distributed, and used: in the United States, Confidential, Secret, and Top Secret (18 CFR §

3a.11, 1982; Executive Order 12356, 1982), plus Controlled Unclassified Information (CUI). CTI has an analogous need, addressed by the Traffic Light Protocol (Section 7.5).

4. Military Intelligence Disciplines

Transforming raw data into intelligence depends on the technical particulars of the subject matter. The specialized fields that produce particular types of intelligence are the intelligence disciplines, each with methods and tools tailored to its sources (JP 2-0, 2013). The point that matters for CTI is structural: every discipline except counterintelligence derives intelligence from a comparatively narrow source type using specific and usually highly technical methods.

Discipline	Source and method
Geospatial Intelligence (GEOINT)	Imagery and geospatial data on physical features and activity on the Earth's surface. Imagery Intelligence (IMINT) interprets electro-optical, radar, infrared, multispectral, and LiDAR imagery; geospatial information adds location and boundary data managed in Geographic Information Systems (GIS).
Human Intelligence (HUMINT)	Information from human sources: interrogation of detained personnel, source operations with overt and clandestine sources (developed, unwitting, or walk-in), and debriefing of cooperative sources such as friendly forces and refugees.
Signals Intelligence (SIGINT)	Exploitation of foreign emissions: Communications Intelligence (COMINT); Electronic Intelligence (ELINT), operational and technical; and Foreign Instrumentation Signals Intelligence (FISINT) covering telemetry and command-state changes. Cues, and is cued by, GEOINT.
Measurement and Signature Intelligence (MASINT)	Analysis of the innate physical attributes of targets and events across electro-optical, radar, radio-frequency, geophysical, materials, and nuclear-radiation measurements, for example a missile's flight profile.

Discipline	Source and method
Open-Source Intelligence (OSINT) and Publicly Available Information (PAI)	Intelligence from publicly accessible information, free or purchased. Fills gaps but is vulnerable to manipulation and deception, is a primary vehicle for influence operations, and requires vetting and operational security (OPSEC).
Technical Intelligence (TECHINT)	Exploitation of foreign equipment and scientific data to assess capabilities and vulnerabilities.
Counterintelligence (CI)	Denial of the adversary's intelligence capability by deceiving, exploiting, disrupting, or safeguarding against espionage and sabotage; the only discipline with both offensive and defensive tactics.

5. Cyber Threat Intelligence

5.1 Origins

Although CTI overlaps heavily with traditional intelligence, it originated in information security. Incident response companies began grouping patterns of computer network intrusions and tracking the adversaries behind them. The first known public mention of the term was by Ian Cook at a FIRST conference in the late 2000s. The milestone that defined the field was Mandiant's APT1 report (2013), which disclosed a decade of work tracking an adversary group linked to the Chinese People's Liberation Army (PLA). Mandiant's founder, Kevin Mandia, came from the US Air Force Office of Special Investigations and had investigated this type of threat in government service; the report showed what a commercial organization could do with the same methods.

The number of companies operating in the space grew rapidly after the report, and so did their marketing teams. Terms were initially introduced haphazardly and used inconsistently between vendors; analyst training and analytical and reporting standards suffered from the same problem. This is one of the reasons the FIRST CTI SIG was founded in 2018, to create a common baseline, and eventually a standard, for analyst training and for the terminology used in the field.

5.2 Marketing: a serious challenge or an unrealized opportunity?

CTI had a rough start. As the information security market saturated, many traditional security companies and Managed Security Service Providers (MSSPs) started talking

about CTI, and their marketing teams pushed blog posts about investigations against nation-states, producing shallow research and premature publication.

Publicly exposing cyber threat data can severely damage ongoing law enforcement (LE) investigations. Because their findings must be admissible in court, LE officials must follow strict, time-consuming legal procedures—such as securing a wiretap order for a command and control (C&C) server. If a researcher casually blogs about a C&C server found in a malware sample, the criminals will likely move it. This single blog post can instantly destroy months of meticulous legal work.

However, marketing can also enhance CTI. If treated as part of the **Dissemination stage** of the cycle, an information release can be deliberately timed to serve the intelligence effort in one of two ways:

- To force attacker movement: This increases their activity and yields additional collection opportunities.
- To impose asymmetric costs: For attackers operating from jurisdictions that shield them from law enforcement, rotating infrastructure is expensive. Repeatedly exposing them drives up their operational costs, heavily increasing the likelihood of an exploitable error.

Ultimately, the difference between a harmful disclosure and a useful one is whether the release is a planned step in the intelligence cycle or a side effect of a content calendar.

5.3 Purpose and focus

Companies invest in a CTI function to manage risk by being better informed about, and better prepared for, malicious cyber activity targeting their information systems. Most of that activity occurs in cyberspace, but CTI may also cover physical components that affect the company's operations. Over the last decade, CTI has matured to the point where it tracks real-world actors on a par with several traditional disciplines, exploiting information that is predominantly in digital form.

5.4 Consumers of CTI

CTI consumers fall into two groups. Technical consumers use CTI to improve technology and specific implementations; an engineering manager who learns that a particular technology is widely exploited can focus engineering effort on securing the components that use it. A historical example is the rapid transition to chip-based credit cards once the falling cost and easy availability of card skimmers made a rise in card fraud predictable. Technical consumers usually operate at the tactical and operational levels. Non-technical consumers own business processes and overall company strategy and operate at the

strategic and operational levels; they look at long-term trends and may take actions that alter not only the company but its ecosystem, as when large corporations began funding university cyber programs to secure the supply of qualified workers.

Two differences from government intelligence follow. First, in industry the consumer may be loosely connected to the producer or not connected at all: a blog post by company A may benefit company B although the two have no relationship, or even an adversarial one, and the post may have been funded by company A's marketing department rather than its security or intelligence group. Second, consumers should not be confused with stakeholders. Some authors use the terms interchangeably; that is incorrect. A stakeholder has a stronger partnership with, and usually a financial investment in, the process.

In contrast, in the example above even the internal technical consumers of company A may have contributed nothing to the report. The conflation is understandable in a small government where the consumer writes the PIRs (Priority Intelligence Requirements) and the output is classified and never shared. Still, it does not hold in a multi-agency government such as the United States, where a PIR may originate in one agency and be satisfied by another, and it does not hold in industry.

5.5 Levels of CTI

Like intelligence generally, CTI works at the strategic, operational, and tactical levels and serves the corresponding consumers.

Strategic CTI takes a broad view of the threat landscape and delivers long-term trend analysis and prognosis to inform the organization's cybersecurity strategy and objectives. It serves senior leadership, policymakers, and strategic planners; covers evolving TTPs, trends in software development, and software supply chain dependencies; and allows executive management to plan and allocate resources, including personnel and training, for emerging threat areas.

Operational CTI is short-term and supports ongoing security activity and incident response with actionable intelligence that improves detection, mitigation, and response. It informs monitoring and threat detection, may involve ingesting feeds of technical indicators (suspicious IP addresses, domains, file hashes), and provides the TTPs of specific adversaries, which in turn inform the design of monitoring and defense. It may also include a vulnerability management program that monitors vulnerability sources to identify which of the organization's software is susceptible and potentially targeted so that it can be upgraded or protected with compensating controls. The Ivanti vulnerabilities and the SolarWinds supply chain compromise are recent examples of the

two cases: a serious bug exploited by hostile actors, and a large supplier breached so that a backdoor was distributed to many industry and government customers.

Tactical CTI serves the immediate and specific needs of threat hunting and incident response, and it deals mostly with atomic indicators: IP addresses, domain names and host records, malware signatures, and other digital artifacts.

The APT1 report illustrates how a single product can span levels: it provided operational CTI in the form of the group's TTPs and organization, and tactical CTI in the form of the specific IP addresses and domain names the group used.

5.6 Raw data, information, and intelligence in CTI

The metamorphosis of raw data through information to intelligence described in Section 3.3 applies unchanged to CTI. Some groups, including the FIRST CTI SIG (2024), have also invoked the older information-science concept of the Data, Information, Knowledge, Wisdom (DIKW) pyramid (Ackoff, 1999; Awad & Ghaziri, 2004). The parallel holds for the first three layers, but it breaks down after that, and the reasons matter for how analysts are trained.

Knowledge is not intelligence, because it has no predictive component. Wisdom translates even less well. For an individual, wisdom may be said to confer an ability to foresee events, but this is very different from producing a probabilistic estimate for a specific event from specific data; on the contrary, wisdom in that sense is a cognitive bias formed by previously observed repeated patterns, and it is precisely what a capable adversary exploits through deception. Furthermore, an individual quality is not a repeatable process. For an institution, "wisdom" resolves the process problem but starts to look like institutional bias and groupthink, which is again unsuitable for analysis. The DIKW pyramid may still serve as a teaching aid for students at an early stage, provided that the differences among the categories are made explicit.

5.7 Source types classified by data origin

Compared with governments and militaries, commercial organizations have limited resources to collect raw data and develop intelligence, and most of them outsource intelligence production to some degree. They are also at different levels of maturity, which limits what they can consume. The result is a rich partnership ecosystem that provides data, information, and analysis, with sources that fall into three categories by origin.

Open source (OSINT) originates in public channels: social media platforms, web sites, network telemetry such as Border Gateway Protocol (BGP) advertisements, messaging

applications such as Telegram, other open digital forms, and in some cases public records such as the Public Access to Court Electronic Records (PACER) system. Reliability varies widely by source: a video posted on X or Telegram can easily be a deepfake or a recording of a different event, whereas a GIS map downloaded from the State of Indiana is likely accurate.

Closed (proprietary) source information is not readily available on the Internet; it usually requires a contract that governs its use and, in many cases, a fee. It is usually reliable, and its errors result from technical limitations rather than malicious intent.

Internally generated information comes from the organization's own environment: system and security logs, incident response and threat hunting work, and vulnerability assessment and management programs.

6. CTI Disciplines and Functional Roles

All of the traditional intelligence disciplines derive intelligence from a narrow field, from specific source types, using specific and usually very technical methods; even HUMINT can be viewed that way, with the human being the source that is exploited. CI is the exception, since it focuses on denial. CTI disciplines are more nuanced, particularly in how they are applied. Exploiting a buffer overflow to gain access to a computer system is legal within the scope of an authorized vulnerability assessment and illegal under the Computer Fraud and Abuse Act (CFAA) without that authorization; the technique is the same, and only the framing differs. For that reason, CTI separates the disciplines, which are bodies of skill and method, from the functional roles that employ them.

6.1 Technical disciplines

The technical disciplines are specialized areas of expertise and analysis focused on technology. They usually deal with atomic technical indicators such as IP addresses, host and domain names, malware hashes, and file paths. Artifacts can be composite: a picture has content (what one sees), the signal characteristics of that content, and metadata. Non-technical means exploit the content under document and media exploitation; the latter two are the subject of the technical disciplines.

Discipline	Scope
Network analysis	Collection, processing, and exploitation of network data at any level, from full packet capture through NetFlow and sFlow metadata to BGP route advertisements, asset ownership, and WHOIS records; analysis ranges from link analysis to handing content to other disciplines.
System analysis	Collection of artifacts from memory and storage: socket state revealing a C&C address, an in-memory ransomware key, code that never touched disk, and traces of tactics and procedures. Overlaps with digital forensics but focuses on live systems rather than reconstruction.
Malware analysis	Exploitation of malware samples by reverse engineering, static analysis (code, string, file, metadata), dynamic analysis (behavioral, API monitoring, memory, sandbox), and hybrid approaches.
Threat data processing	Automated ingestion, enrichment, analysis, and in some cases distribution of feed data to devices; some organizations act automatically above a confidence threshold, for example by blocking C&C infrastructure at the edge.
Technical exploitation of social media	The metadata and technical indicators of social media rather than its content: usernames, posting times, link analysis, and Exchangeable Image File Format (EXIF) data in images.
Digital-forensics-enabled disciplines	Digital image and sound forensics (signal properties and metadata); storage media and file system forensics, including damaged devices; database forensics; mobile device forensics (embedded, proprietary, hardened devices); and Internet of Things (IoT) forensics. Each has a technical side covered here and a content side covered under document and media exploitation.
Operational Security (OPSEC)	A supporting discipline: every action in a network leaves a trace, and OPSEC minimizes the footprint a given type of collection leaves.

6.2 Non-technical disciplines

The non-technical disciplines do not primarily involve technical skills or tools, yet they are essential to CTI; generally, they provide context, and in some cases they provide the core intelligence.

Human Intelligence. Even in cyber, one of the highest-fidelity sources is a human, because human reporting nearly always arrives with context and some degree of vetting. Unlike the government and military world, commercial HUMINT does not require anyone to break the law, although insider threats may do so.

Synthetic identities. The commercial equivalent of undercover work. The usual targets are criminal and underground forums, often on the dark net, but the technique is also used on social media and other online communication, largely for OPSEC.

Industry groups. Formal organizations created for mutual help among corporations. FIRST is an international, cross-vertical body facilitating training and collaboration. In the United States, the Information Sharing and Analysis Centers (ISACs) serve specific verticals (aviation, automotive, communications, financial services, and many more) under multilateral non-disclosure agreements and other legally enforceable obligations. Exchange happens formally through meetings and mailing lists and informally through personal relationships; several ISACs operate infrastructure for machine-readable exchange in the form of threat feeds. Because most ISACs represent critical infrastructure, their interests align with the federal government, which explains the outreach programs of the Cybersecurity and Infrastructure Security Agency (CISA) and, long before it, the Federal Bureau of Investigation's InfraGard program.

Informal communities. Groups with no formal legal structure, governed by peer pressure and reputation, with differing levels of trust among members. One of the earliest to coin the term "trust groups" after its original name, Operations Security Trust (later OpsTrust). At the lower end of the trust scale, researchers organize around mailing lists and messaging workspaces.

Document and media exploitation. The content side of image, video, and sound forensics: analyzing a satellite image for the presence and disposition of forces, or identifying the location of a missing child from the background of a photograph. Content exploitation is a discipline of its own and requires a skill set very different from the technical analysis of the same file.

6.3 Functional roles

Functional roles group disciplines by how they are applied, and they fall into two groups.

Operational roles handle day-to-day operations and are subdivided by purpose:

Defensive operations (Blue Team) defend organizational assets and provide the incident response function, implemented in the Computer Security Incident Response Team (CSIRT).

Offensive operations (Red Team) test organizational security by simulating adversary behavior. The idea dates back to the Red Cell, created by the US Navy to practice assembling teams in hostile jurisdictions with local resources; it later moved into information security, helped create an entire industry, and many large corporations now maintain their own adversarial component. Depending on the level and intensity of engagement, the same function is called a vulnerability assessment team, a penetration testing team, or a Red Team.

Purple Team. In some organizations, a single combined team performs the blue and red functions.

Network and system monitoring usually sits outside the security and CTI organization, but its outputs feed the CTI collection systems.

Engineering roles are a major and often overlooked consumer, and in companies that develop software, they play an important part in the cycle. The first engagement is through the vulnerability management program, which aggressively collects vulnerabilities in the software and components the organization uses, brings them to engineering, and sees the software rebuilt and redistributed after correction; mature organizations implement this in a Product Security Incident Response Team (PSIRT). The second is the Security Development Lifecycle (SDL), where intelligence about new exploitation techniques and other forms of software abuse informs both the development process and technology choices.

7. The Cyber Threat Intelligence Cycle

The CTI cycle is very similar to the Intelligence Cycle. It differs mainly in its lack of rigor, due mostly to the nature of commercial organizations and partly to a lack of consistent training. The steps are the same; what differs is what each step contains.

7.1 Planning

Planning is informed by the organization's need to protect the assets identified during asset enumeration and threat assessment. The planning process considers the value of each asset, lists vulnerabilities with higher risk (those more damaging or more likely to be exploited), and ranks the resulting information needed by risk. The planning team then directs the collection effort by writing specific PIRs and incorporates feedback from

previous or similar cycles. Notably, the term External Attack Surface Management (EASM) has emerged to describe the first part of this process.

7.2 Collection

Collection is the systematic gathering of data and information about potential threats, vulnerabilities, and actors. From the PIRs, the collection team builds a collection plan to acquire the most comprehensive set of indicators, with the best coverage, as efficiently and expediently as possible. This includes selecting and prioritizing sources and requires assessing and ranking them for trustworthiness and reliability.

CTI's collaborative nature gives it sources that traditional intelligence lacks. The main source of machine-readable data is the CTI feed, produced by a wide range of entities external to the consuming organization. More than twenty years ago, non-profit organizations such as Team Cymru, Shadowserver, and abuse.ch began offering free lists of malicious IP addresses separated into categories, consumed by download from the provider; feed formats now range from simple comma-separated values to structured standards such as the Incident Object Description Exchange Format (IODEF) and Structured Threat Information Expression with Trusted Automated Exchange of Intelligence Information (STIX/TAXII). A second source is openly published material on forums, social media, and channels such as Telegram or the dark web; this material is unstructured and requires human interpretation, but it tends to carry a higher level of refinement and confidence. A third is the company's vulnerability management program, which may have access to proprietary sources and can inform both specific vulnerabilities and, from large-scale research, which classes of vulnerability are being discovered and where attackers are focusing. That, in turn, informs compensating controls and upgrade timing.

Note that the use of the words information and intelligence in the context of collection is not an error. Unlike a sensor, CTI collection can acquire and incorporate finished products.

7.3 Processing and exploitation

Processing in CTI resembles processing elsewhere, with a focus on cyber indicators. Data arrives in a variety of encoding, encryption, and packaging schemes and must first be converted into a consumable form, the CTI equivalent of developing satellite film. Conversion includes decoding (reversing an encoding applied for transmission or storage, such as base64, not to be confused with decryption), decryption (recovering cleartext with the corresponding key material), format conversion (among CSV, XML, JSON, and STIX/TAXII, and where necessary binary image formats), and language

translation. Translation is rare in CTI and is properly the task of a qualified linguist familiar with the local culture; in commercial practice, budget and availability usually force machine translation or an analyst with partial knowledge of the language, and analysts must account for the noise and bias introduced during analysis.

The converted data is then deduplicated and normalized, and individual indicators are extracted, sometimes based on context. Deduplication removes redundant entries, such as the same C&C address repeated across feeds or a single address already covered by a CIDR block in a blocklist; care is required where entries are measured to be treated statistically, since removing a repeated value alters the baseline of the series. Normalization rescales values, transforms between domains, or reconciles signals from sources with different ranges and properties. These transformations may be layered in varying order.

CTI is still a young discipline, and as a result, most work happens in this phase, with far less in analysis. This is the most important observation about the current state of commercial practice, and it is why analyst training is the priority it is.

7.4 Analysis and production

Analysis draws on the information distilled in the previous phases and produces a clear, actionable product. It is divided into human analysis, based on structured analytic techniques (SATs), and machine analysis.

Human analysis. The SATs are described in detail by Heuer (2007) and by Pherson and Heuer (2021); the table below summarizes the families.

Family	Purpose	Techniques
Basic	Insight through information transformation	Sorting, chronologies, matrices, weighted ranking, link analysis, event trees, event mapping
Advanced: contrarian	Challenge current thinking	Analysis of Competing Hypotheses (ACH), Devil's Advocacy, Team A/Team B, High-Impact/Low-Probability Analysis, "What If?" Analysis
Advanced: imaginative	Develop new insight	Brainstorming, Outside-In Thinking, Red Team Analysis, Functional Analysis using Critical Factors Analysis

Family	Purpose	Techniques
Diagnostic	Make the analytic process itself transparent	Key Assumptions Check, Indicators of Change, Quality of Information Check

Machine analysis. Machine learning (ML) enhances analysis in several modalities (supervised, unsupervised, and deep learning); artificial intelligence and Large Language Models (LLMs) are not treated separately here because, for the most part, they are forms of ML. The principal applications are: indicator extraction and text summarization, where LLMs reduce workload by summarizing large bodies of text and extracting indicators for machine processing; threat detection and classification, where models cluster large volumes of traffic, logs, and samples and, with an analyst assigning meaning to the clusters, classify new threats into known harmful categories; malware detection and analysis, including variant detection, family classification, and feature extraction; network anomaly detection against a learned baseline; correlation across diverse sources, from technical indicators to vulnerability feeds and dark web forums; and predictive analytics. Two caveats apply. ML-driven correlation suffers from a lack of explainability: the model will not provide a human-understandable reason why certain indicators cluster together. Predictive analytics rests on detected patterns, and a pattern applied to future events is the machine equivalent of a cognitive bias, which a capable adversary can exploit through carefully crafted deception.

Whether the analysis is structured or machine-assisted, by the time the analyst has correlated and reviewed the data, they have a good sense of the information's quality and the reliability of its source. Source and information reliability should therefore be re-scored in this step and fed back to collection.

7.5 Dissemination and integration

The CTI process produces briefings, human- or machine-readable reports, and CTI-specific formats, notably the threat feed: a stream of indicators regularly updated and downloaded by consumers. Format and content are tailored to decision makers at the tactical, operational, and strategic levels. This is the step in which the product sees practical application and can be validated, and, as in the Intelligence Cycle, it informs the planning of the next cycle. Two conventions govern how you communicate uncertainty and handle the product.

Words of Estimative Probability and Levels of Confidence. Several standards exist; the five-level NATO standard is recommended because it is simple, compatible with the

seven-level ODNI scale (Irwin & Mandel, 2019; Mandel, 2019; ODNI, 2022), and nearly identical in its percentages to the MITRE model for threat occurrence (The MITRE Corporation, 1999). In the NATO model, WEP is expressed as remote chance (below 10%), unlikely (10% to 40%), even chance (40% to 60%), likely (60% to 90%), and very likely (above 90%). LCA is expressed on three levels: high, where the assessment rests on good-quality information from multiple and preferably diverse sources that allows a clear judgment; moderate, where the evidence is open to several interpretations, or is plausible and credible but uncorroborated; and low, where the assessment rests on fragmentary information or on sources of dubious reliability.

Traffic Light Protocol. Like government intelligence, CTI must label products and intermediate outputs to protect them from disclosure to adversaries and to third parties without a need to know. The Traffic Light Protocol (TLP) was devised in the early 2000s by the UK National Infrastructure Security Co-ordination Centre (from 2007 the Centre for the Protection of National Infrastructure) to facilitate information exchange; it is widely adopted in industry and forms part of ISO 27010 for inter-sector and inter-organizational communication (Luiijf & Kernkamp, 2015; International Organization for Standardization, 2015). Its key principle is originator control: whoever contributes information decides how widely it may circulate. There are four levels. RED is for the named recipients only, or for those present in a meeting. AMBER permits sharing within the recipient's organization on a need-to-know basis, with any further limits set by the originator. GREEN is for community-wide use. WHITE, which FIRST now calls CLEAR, carries no distribution limit beyond other contractual or legal constraints. All information must be marked at the time of disclosure.

7.6 Evaluation and feedback

This step is nearly identical to its counterpart in the Intelligence Cycle, because it does not depend on technical specifics. It provides evaluation and feedback throughout the cycle, applies to components at any level, and requires participants at all levels to regularly evaluate their performance and its impact in continuous collaboration with intelligence consumers.

8. Conclusion

CTI is best understood as the Intelligence Cycle transplanted into a commercial organization and pointed at cyberspace. The definitions, the three levels, the transformation of raw data into intelligence, the structured analytic techniques, and the conventions for expressing uncertainty all transfer intact, and a practitioner trained in the government model will recognize every step. What does not transfer is the context, and that is where the field's persistent difficulties originate. Consumers may be unrelated to

producers, or adversarial to them; a marketing department may occupy the Dissemination step; the same technique may be lawful or unlawful depending on who authorizes it, which forces the separation of disciplines from functional roles; and the source ecosystem consists largely of feeds and partners whose reliability is rarely evaluated with the rigor the cycle assumes.

Three conclusions follow for anyone building or teaching a CTI function. First, adopt the reference model's vocabulary and use it consistently, because the alternative is the vendor-defined vocabulary that has already cost the field a decade of confusion. Second, recognize that most present-day effort sits in processing and exploitation; automation belongs there, and the analyst's time belongs in analysis, where structured techniques and disciplined expression of uncertainty separate an intelligence product from a news item. Third, treat the distinction between disciplines and functional roles as fundamental to training, because it is the point at which the legal and organizational boundaries of the commercial world enter the cycle. Each step of the cycle has clear goals and methods; some are highly technical, and some are not; some are easy to automate, and some require human judgment. A curriculum that respects those distinctions is the prerequisite for the consistent training the field still lacks.

Authors and Contributors

Principal author: Krassimir Tzvetanov, FIRST CTI SIG (krassi@krassi.biz). The principal author researched and wrote this paper and is responsible for its content and conclusions.

Contributors: Ken Dunham and Laurie Tyzenhous reviewed the draft and contributed comments that are reflected in the final text.

Suggested citation: Tzvetanov, K. (2026). *What Is Cyber Threat Intelligence? Defining the Field, Disciplines, and Intelligence Cycle* (Version 1.0). FIRST Cyber Threat Intelligence SIG, Forum of Incident Response and Security Teams.

References

18 CFR § 3a.11 - Classification of Official Information. (1982). https://www.law.cornell.edu/cfr/text/18/3a.11

Ackoff, R. (1999). From data to wisdom. In Ackoff's best: His classic writings on management (pp. 170-172). Wiley.

APT1: Exposing one of China's cyber espionage units. (2013). Mandiant. https://www.mandiant.com/sites/default/files/2021-09/mandiant-apt1-report.pdf

Awad, E. M., & Ghaziri, H. (2004). Knowledge management. Prentice Hall.

Executive Order 12356. (1982). National Archives. https://www.archives.gov/federal-register/codification/executive-order/12356.html

FIRST CTI SIG. (2024). Introduction to CTI as a general topic. FIRST CTI SIG Curriculum. https://www.first.org/global/sigs/cti/curriculum/cti-introduction

Heuer, R. J., Jr. (2007). Psychology of intelligence analysis (2nd ed.). Pherson Associates.

Hogan, M., & Newton, E. (2015). Supplemental information for the interagency report on strategic U.S. Government engagement in international standardization to achieve U.S. objectives for cybersecurity (NIST IR 8074v2). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8074v2

International Organization for Standardization. (2015). Information security management for inter-sector and inter-organizational communications (ISO/IEC 27010:2015).

International Organization for Standardization. (2022a). Cybersecurity: Guidelines for Internet security (ISO/IEC 27032:2023).

International Organization for Standardization. (2022b). Information security management systems: Requirements (ISO/IEC 27001:2022).

Irwin, D., & Mandel, D. (2019). How intelligence organizations communicate confidence (unclearly). SSRN. https://doi.org/10.2139/ssrn.3441302

Joint Chiefs of Staff. (2010). JP 1-02, Department of Defense dictionary of military and associated terms.

Joint Chiefs of Staff. (2013). JP 2-0, Joint intelligence.

Joint Task Force Transformation Initiative. (2012). Guide for conducting risk assessments (NIST SP 800-30 Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-30r1

Luiijf, E., & Kernkamp, A. (2015). Sharing cyber security information. TNO. https://publications.tno.nl/publication/34616508/oLyfG9/luiijf-2015-sharing.pdf

Mandel, D. (2019). Assessing and communicating uncertainty effectively in a rapidly changing world. SSRN. https://doi.org/10.2139/ssrn.3441329

ODNI. (2024). What is intelligence? https://www.dni.gov/index.php/what-we-do/what-is-intelligence

Office of the Director of National Intelligence. (2022). Intelligence Community Directive 203: Analytic standards.

Pherson, R. H., & Heuer, R. J. (2021). Structured analytic techniques for intelligence analysis (3rd ed.). CQ Press.

The MITRE Corporation. (1999). Risk matrix user's guide (Version 2.2).

Tsantes, G. K., & Ransome, J. F. (2023). Cybertax: Managing the risks and results. CRC Press. https://doi.org/10.1201/9781003330103

Tzvetanov, K., & Dietz, J. E. (2025, April 21–22). What is Cyber Threat Intelligence? [Paper presentation]. Purdue Military Research Institute (PMRI) Defense & Security Research Symposium, West Lafayette, IN, United States.

Document History

Version	Date	Author	Changes
1.0	September 2026	Krassimir Tzvetanov	Initial release