A "Psychological Warfare" to Show Off Cyber Capabilities: A Comprehensive Analysis of SentinelOne's Exposure of fast16

Antiy Security Research and Emergency Response Center

First draft completed: April 24, 2026

First published: April 25, 2026

Update time of this edition: April 25, 2026

Table of Contents

1 Background and Debunking Misinformation

1.1 Background and Focus of the SentinelOne Report

In April 2026, the U.S. cybersecurity firm SentinelOne presented a technical report titled "Fast16: The Silent Saboteur" at the Black Hat Asia Conference in Singapore [1] (hereinafter referred to as the "SentinelOne Report"), revealing a nation-state-level cyber sabotage framework known as fast16, which is alleged to date back to 2005. The report was co-authored by SentinelOne researchers Vitaly Kamluk and Juan Andres Guerrero-Saade. Its central argument is that fast16 emerged approximately five years earlier than the Stuxnet worm, which was made public in 2010, thereby "pushing back the timeline for the detection of nation-state-level cyber sabotage operations by five years." Antiy Security Research and Emergency Response Center (Antiy CERT) has been tracking and analyzing this report and the associated samples. It concludes that both the research and the report demonstrate a very high level of technical expertise. At the same time, the timing and motives behind their publication clearly align with the U.S. side's efforts to conduct "psychological warfare" in the Middle East.[1]

1.2 Debunking Misinformation: Clarifying the Timeline

A key issue with the timeline presented in the SentinelOne report is its claim that fast16 pushes the historical starting point of nation-state-level cyber sabotage five years earlier than Stuxnet, moving it from 2010 to 2005—the year the fast16 sample was compiled. This statement is factually incorrect. The earliest known attack attributed to Stuxnet was not in 2010, but in 2008. According to a June 2012 investigative report by The New York Times and an in-depth 2019 investigation by the Dutch newspaper De Volkskrant, the U.S. government, working through the Dutch General Intelligence and Security Service (AIVD), bribed maintenance engineers for Siemens industrial control systems at Iran's Natanz nuclear facility in 2007. These engineers used USB-based transfer devices to implant Stuxnet version 0.5 [2] into the target industrial control system. The year 2010 is widely recognized as the time when Stuxnet version 1.x was modified to propagate via a worm-like mechanism. After spreading and infecting systems, it was first detected in June 2010 by the Belarusian security firm VirusBlokAda. Subsequently, leading cybersecurity vendors such as Symantec, Kaspersky Lab, ESET, and Antiy Labs conducted follow-up analyses and publicly released research reports [3]. In other words, 2010 was the year Stuxnet was publicly exposed, not the year it first carried out an attack. Based on the timeline of Stuxnet, preparatory activities such as the registration of associated

attack domains had already begun as early as 2005, with some resource files being compiled as early as 2003. It is certain that Fast16 was deployed earlier than Stuxnet. However, based on the preparation period, sample timestamps, and operational timeline, the interval was not five years but significantly shorter.[2][3]

Table 1-1: Fast16.sys Sample Labels11

Malware name	Trojan/Win32.Fast16
MD5	0FF6ABE0252D4F37A196A1231FAE5F26
Original file name	07c69fc33271cf5a2ce03ac1fed7a3b16357aec093c5bf9ef61fbfa4348d0529
File size	43.54KB (44,580 bytes)
File format	BinExecute/Microsoft.SYS[:X86]
Processor architecture	Intel 386 or later, and compatibles
Timestamp	2005-07-19 15:15:41
Compiled language	Microsoft Visual C/C++
Shell type	None
Signature Status	None
Information source	Virusview.net[4][4]

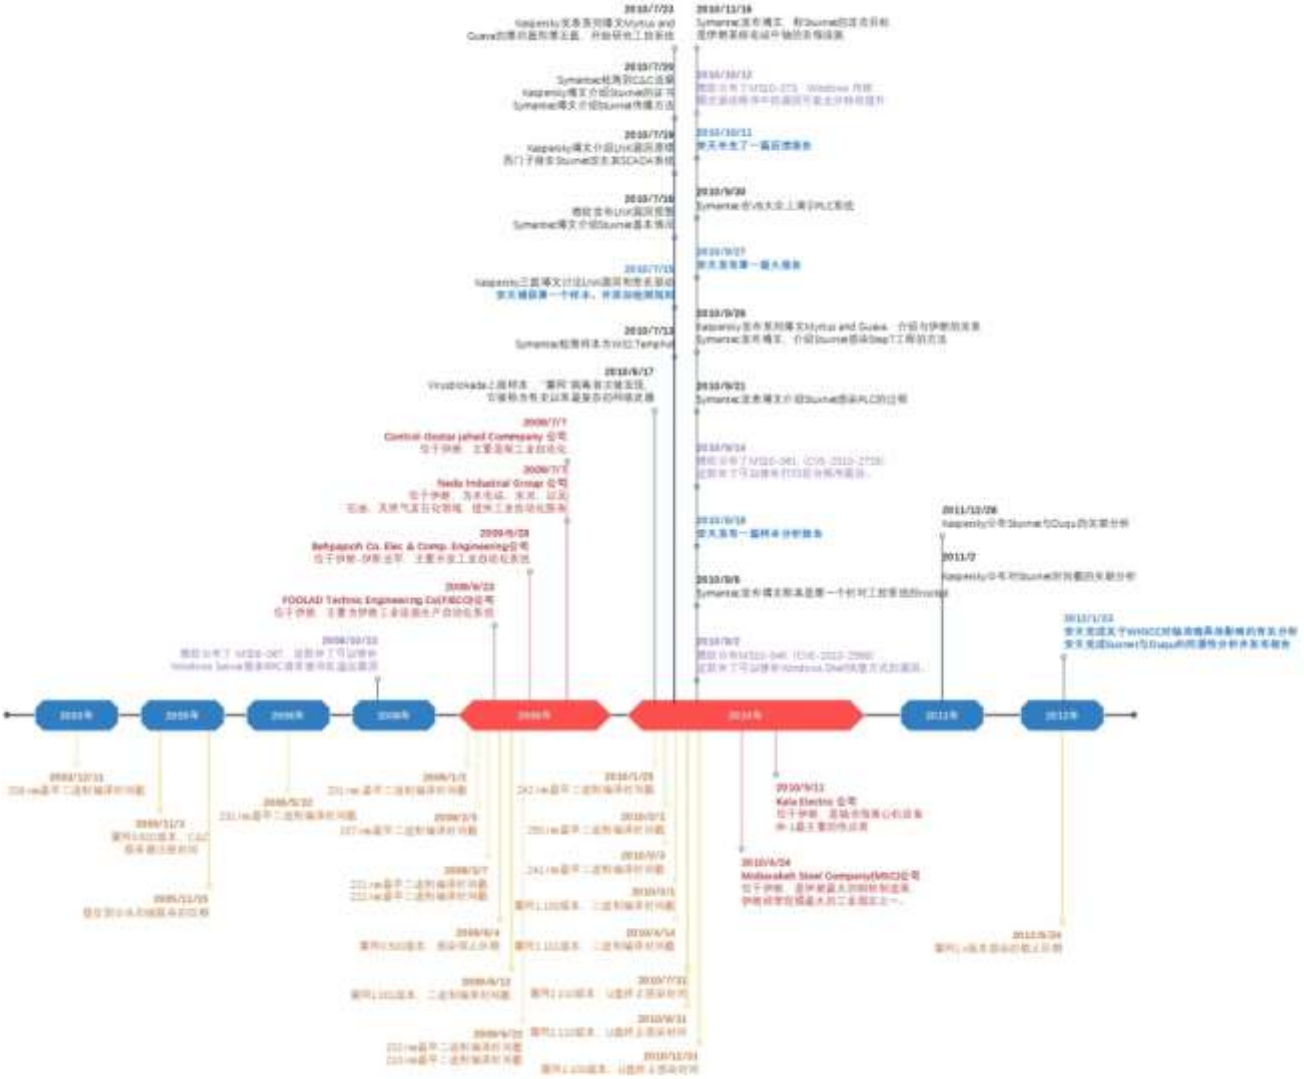

Figure 1-1: Stuxnet Incident Timeline11

1.3 Beyond the True "Early" and "Complex"

A common misinterpretation among some independent media outlets and certain industry insiders regarding FAST16 is that it marks the U.S. government's initiation of offensive cyber operations in 2005. This, however, is clearly a mistake. The U.S. National Security Agency's (NSA) dedicated cyber-attack unit, TAO (Tailored Access Operations), was established as early as 1998. By 2000, it had already compiled a clear list of operational achievements targeting multiple sectors of global information infrastructure [5]. In its 2016 report "From Equation to System of Equations" [6], Antiy Labs pointed out that the Equation Group has the capability to deploy attack payloads across multiple platforms, including Windows, Linux, Solaris, macOS, and iOS. The organization's malicious code was first observed in active operations as early as 2002, and there may even be earlier, unknown

variants. Kaspersky believes that the Equation Group (also known as TAO) began deploying the EquationLaser sample targeting Windows 9x as early as 2001.[5][6]

SentinelOne disclosed that the key value of fast16 lies not simply in its earlier emergence, nor in the complexity of its technical mechanisms, but rather in what the report describes as “elevating fast16 from the realm of generic espionage tools into that of strategic disruption.” This does not constitute a simple intrusion for data acquisition; nor is it merely an intelligence-gathering operation. Rather, it is a full-fledged military operation: by silently modifying, via a kernel-mode driver, the floating-point computation results of high-precision engineering simulation software, it introduces systematic errors over periods ranging from several months to several years. Such manipulation can lead to catastrophic engineering failures or mislead scientific research. This is also the reason the report believes it will find that nation-state cyber-attack campaigns began five years earlier than Stuxnet. From the perspective of Chinese researchers studying the intelligence activities and operational tactics of long-term cyber threat actors, we should adhere to a dialectical approach that combines strategic contempt with tactical vigilance. The findings of FAST16 should not be reduced to a competitive narrative about being “the earliest” or “the most complex”; rather, they should serve to further deepen our understanding of the principles governing cyberspace confrontation, enable us to grasp these principles, and enhance our capabilities.

2 fast16 Corresponds to a Military Operational Action

2.1 U.S. Terminology Distinguishes between Cyber Intelligence and Cyber Attack.

To accurately understand FAST-16’s operational characteristics, it is first necessary to clarify several core concepts at the theoretical framework level. Due to differences between Chinese semantic conventions and Western military terminology, China commonly refers to all activities that impair or compromise network infrastructure as “cyberattacks.” However, in the US military terminology system, the continuous information and intelligence theft activities supported by network intrusion, persistence and other capabilities are clearly called Cyber Intelligence Operations (CIO) or translated as network intelligence utilization, and its corresponding traditional abbreviation is CNE(Computer Network Exploitation, computer network utilization). Only actions that cause a substantial disruption to the operation of a target information system or network are classified as “cyberattacks,” i.e., CNA (Computer Network Attack).

In the U.S. military's latest cyberspace operations doctrine, JP 3-12 (Joint Publication 3-12, Cyberspace Operations),[7] the U.S. terminology system has undergone a significant update. This directive was first issued in February 2013 and updated to its second edition in 2018. It explicitly replaced the previous term "Computer and Network" with "Cyber," and the corresponding acronyms were revised to CE (Cyber Exploitation) and CA (Cyber Attack). Specifically, CE corresponds to the traditional CNE function of intelligence collection, while CA aligns with the conventional CNA function of destructive strikes. CD, on the other hand, refers to defensive operations aimed at protecting one's own cyberspace assets from adversary attacks. It is worth noting that JP 3-12 places particular emphasis on the "isomorphism" and "convertibility" between Cyber Exploitation (CE) and Cyber Attack (CA). The persistent access privileges and deep situational awareness of the target environment established by a given cyber operations unit during CE missions can be rapidly translated into precise strike capabilities when conducting CA operations.[7]

2.2 The Technical Essence of the Fast16 Compromise Mechanism

The key takeaway from SentinelOne's report on fast16 is not merely its alleged links to Stuxnet or the Shadow Brokers leak, nor the attribution claims themselves, but rather the precise data-exfiltration mechanisms it reveals for three highly sophisticated engineering and simulation toolsets. The three software packages are LS-DYNA 970, PKPM, and the MOHID hydrodynamic modeling platform. They respectively cover nuclear engineering physics simulation, civil engineering structural analysis, and environmental water resources modeling—three highly sensitive fields of great strategic importance. This analysis of the sabotage mechanism directly demonstrates that FAST16's design intent was not information theft, but rather to manipulate the core computational results of target software, thereby inducing erroneous engineering decisions, structural design flaws, or catastrophic failures in physical systems, ultimately achieving strategic-level destructive effects.

According to a report by SentinelOne, the overall architecture of fast16 comprises two core components: a self-replicating propagation module (a worm module) and a kernel-mode driver named fast16.sys. The worm module is responsible for lateral movement within the local area network via Windows network sharing, scanning for potential target hosts. Once it successfully infects a target system, the worm module checks whether any security software—such as antivirus programs or HIPS host-based intrusion prevention systems—is present. If the environment is deemed "safe," the module installs the fast16.sys kernel driver into the system. This "environmental reconnaissance–condition-triggered" deployment strategy reflects an exceptionally high level of operational

security awareness: the adversary is not aiming to maximize the number of infections, but rather to establish a stable, covert, and long-term destructive capability on critical targets.

The core function of the fast16.sys kernel driver is to silently tamper the floating-point arithmetic results of the target engineering simulation software. Specifically, the driver monitors all executable code loaded into memory within the system and identifies the target software by applying a sophisticated set of “pattern-matching rules.” These rules may be based on the target software’s specific code signing, memory layout characteristics, dynamic link library (DLL) loading patterns, or the call sequence of the floating-point unit. Once the target software (such as LS-DYNA, PKPM, or MOHID) is identified, fast16.sys will hook its floating-point arithmetic functions and, during the computation process, make small but systematic modifications to either the intermediate results or the final output. The magnitude of such modifications is extremely small—often beyond the fifth or sixth significant digit—making them virtually undetectable in a single computation. However, as the problem size grows and the number of iterations increases, the cumulative effect of these tiny deviations can lead to catastrophic systemic errors.

Kamluk described this very vividly in his speech at the Black Hat Asia conference: "It focuses on making slight changes to these calculations so that they lead to failure-very subtle failures that may not be immediately apparent. Systems may wear out, collapse, or break down more quickly, and scientific research may come to the wrong conclusion, potentially causing serious injury. Frankly, it's a nightmare. "This description is highly consistent with the definition of the" inaccurate "strike pattern of Antiy's cyberspace combat: not directly destroying physical facilities, but by tampering with the decision basis (calculation results), allowing the attacked to make wrong engineering decisions without knowing it, eventually leading to the destruction of the physical world. The stealth and delayed nature of this kill chain render it virtually impossible to attribute and trace the attack prior to its execution, thereby constituting a highly strategic asymmetric capability.

The three simulation software packages targeted by FAST16 are all highly specialized and strategically sensitive. Their selection was by no means accidental; rather, it was the result of a carefully crafted, intelligence-driven decision to launch a precision strike. LS-DYNA is a multi-physics explicit dynamics simulation software originally developed by scientists at the Lawrence Livermore National Laboratory (Lawrence Livermore National Laboratory) in the United States. It is widely used in collision testing, structural analysis, explosion effect simulation, and metal materials. The behavior simulation under extreme conditions and other national defense and nuclear energy fields. According to a public report by the Institute for Science and International Security (ISIS), Iranian scientists have

extensively used LS-DYNA in their nuclear-weapons-related research, including simulations of metal interactions within nuclear devices and the impact effects experienced by nuclear warheads during ballistic-missile reentry into the atmosphere—topics of extreme sensitivity. PKPM (Peking University Program for Structural Design) is a structural analysis and design software developed by the China Academy of Building Research. It enjoys an extremely high market share in China's civil engineering sector and is widely used for the design and analysis of critical infrastructure, including high-rise buildings, bridges, and dams. MOHID(Modelo Hidrodinamico) is a hydrodynamic modeling platform created by Portuguese researchers, mainly used for environmental modeling, ocean flow simulation and water resources system management. These three software packages respectively cover nuclear engineering simulation, civil structural design, and environmental water resources—three critical domains. FAST16's targeted attacks on them clearly reveal a distinct national strategic intent: by undermining the target country's nuclear engineering capabilities, infrastructure safety, and water resource management, it seeks to create long-term, systemic strategic disadvantages in the physical realm.

2.3 The Transition from Intelligence Operations to Combat Operations

A technical comparison between FAST16 and Stuxnet can more clearly elucidate the differences in their operational roles and destructive mechanisms. The core destructive mechanism of Stuxnet lies in its ability to tamper with the variable-frequency drive control commands of Siemens programmable logic controllers (PLCs), causing IR-1 centrifuges at Iran's Natanz nuclear facility to operate at rotational speeds far exceeding their design limits. This rapid over-speed condition results in immediate physical damage and failure of the centrifuge rotors within a short period. This is a quintessential "hard-kill" paradigm: the attack delivers direct, visible destruction within a narrow time window. However, precisely because of this directness of damage, the target can relatively quickly determine that it has been attacked. In contrast, Fast16 has adopted a more covert "soft-kill" approach: rather than physically destroying equipment, it manipulates simulation results to induce engineers to design products or infrastructure with systemic vulnerabilities—without their knowledge. Such defects may not become apparent for months or even years: a deviation in the load-bearing design of a bridge could lead to fatigue failure and structural collapse several years after it opens to traffic; an inaccurate calculation of the material composition in a nuclear device might result in unpredictable physical consequences over prolonged operation; and systematic errors in the stress analysis of a dam complex could trigger a catastrophic breach under extreme weather conditions. This combination of "delayed damage" and "difficult attribution" gives FAST16 a strategic deterrent value that surpasses that of Stuxnet.

From the perspective of military operations theory, FAST16 perfectly aligns with the defining characteristics of CA (Cyber Attack) as outlined in U.S. Joint Publication 3-12: “A cyber attack is an operation conducted in cyberspace that produces denial, degradation, disruption, manipulation, or destruction to achieve a desired effect.” The “manipulation” effect created by fast16—precise control over floating-point computation results—is precisely the most sophisticated and strategically concealed form of attack as defined in CA. It is neither a simple denial-of-service (DoS) attack nor direct physical sabotage; rather, it undermines the credibility of the system’s outputs and the value of its decisions from within, all while maintaining the appearance of normal operation. This operational paradigm represents a novel and highly challenging strategic threat vector for modern engineering systems that heavily rely on computer simulation and numerical analysis.

3 “Inaccuracy” is the Operational Disruption Effectiveness Achieved by FAST16, but Implantation and Persistence are its Foundational Elements.

3.1 Antiy’s definition of the impact of cyberspace military operations

In its research titled “An Analysis of the Essential Mechanisms Underlying Cyber-Conflict Effectiveness from a Whole-Domain Competition Perspective,” Antiy Labs categorizes the effects of cyber military operations into five distinct forms: “loss of connectivity, loss of control, loss of momentum, loss of precision, and loss of functionality.” In subsequent reports, the concept of “loss of assets” was further introduced. Specifically, “loss of connectivity” refers to the disruption of communication or command-and-control links, thereby depriving the target system of its ability to exchange information; “loss of control” denotes the adversary’s seizure of control over the target system, resulting in the legitimate administrator’s loss of dominion; “loss of momentum” signifies the destabilization of the system’s operational or flight dynamics, leading to corresponding physical consequences; “loss of precision” occurs when the system’s outputs or decision-making inputs are tampered with, inducing directional errors and causing judgments and decisions based on that system to systematically diverge from reality; “loss of functionality” means the system is rendered incapable of providing services, or its power and propulsion support is severed. “Loss of assets” pertains to the destruction or unauthorized transfer of the owner’s or carrier’s physical electronic resources. These six offensive postures are not isolated; rather, they serve as mutually convertible enablers within

complex cyberspace operational missions. Establish a cascading kill chain that progresses from the information domain to the physical domain. The interference effect achieved by fast16 precisely corresponds to the “inaccuracy” strike type.

Among them, “inaccuracy” stands as an independent form of cyberspace attack. Its theoretical significance lies in demonstrating that cyberattacks can not only inflict visible damage but also achieve covert strategic objectives by undermining the credibility of information. In traditional military strike theory, the effectiveness of an attack is typically assessed by the degree of physical destruction. However, in cyberspace operations, information itself constitutes both the operational environment and a key operational resource; manipulation of information can yield strategic consequences that are no less significant than those resulting from physical destruction. The manipulation of floating-point computation results by fast16 in simulation software such as LS-DYNA is, at its core, an attack on the “decision-making basis”: it does not destroy the computational code itself, but rather causes the code to produce incorrect outputs; it does not physically damage the bridge, but instead leads engineers to design and construct a bridge based on flawed stress analyses—resulting in a structure that is destined to collapse. The high level of abstraction and generalizability of this attack logic implies that “inaccurate” strikes can be applied to nearly all engineering disciplines that rely on numerical simulation and computer-aided design.

Antiy has previously identified and publicly released research findings on similar behavioral patterns through its historical analysis work. In the 2012 XCon Security Summit report, Antiy Labs demonstrated an attack vector that leverages a printer’s USB interface to paralyze a host system [8]. Printers are often an overlooked security vulnerability in office networks. Their drivers typically run with high privileges at the operating system kernel level, while printer firmware lacks robust security verification mechanisms. Antiy’s research has revealed that attackers can execute arbitrary code on host systems by spoofing print jobs, tampering with printer driver buffers, or exploiting vulnerabilities in printer firmware. Moreover, they can use the printer as a stepping stone to conduct lateral movement and compromise core assets within the internal network. The deeper significance of this research lies in the following: Once a seemingly harmless peripheral device—such as a printer—is compromised by an attacker, the resulting “inaccuracy” effects—for example, incorrectly printed engineering drawings, tampered financial reports, or forged contract documents—can directly influence decision-making and actions in the physical world, all without the attacker needing to establish any direct interaction with the targeted core system.[8]

In the XCon 2012 report published the same year, Antiy Labs also disclosed the technical feasibility of interfering with and forging long-wave time synchronization signals, such as China's BPL long-wave time signal [9]. Long-wave time signal is an important time reference source for clock synchronization in power dispatching systems, financial trading systems, communication networks, and critical infrastructure. Once the time synchronization signal is maliciously tampered with—for example, by using a rogue base station or signal-injection techniques to advance or retard the system clock by several seconds or even minutes—distributed systems that rely on precise time synchronization will face a systemic risk of time skew. In the financial sector, clock skew can lead to incorrect order sequencing in high-frequency trading systems, resulting in substantial financial losses. In the power industry, clock asynchrony may disrupt the timing of protective relays, potentially triggering widespread blackouts. In telecommunications, clock drift can cause key negotiation failures in encrypted sessions, leading to communication outages or reduced security. In its report, Antiy Labs specifically points out that the attack surface of long-wave time-and-frequency dissemination signals has long been severely underestimated by the security community. Its protective mechanisms have primarily relied on security assumptions at the physical layer—such as signal propagation delay and transmission power—yet these assumptions no longer hold in the face of modern software-defined radio (SDR) technology.[9]

In the XCon 2013 report, Antiy further analyzed the safety issues of 3D printers [10] and extended the research on the "misalignment" strike form to the field of additive manufacturing. The core workflow of a 3D printer includes 3D modeling, slicing (which breaks down the 3D model into two-dimensional layers), and G-code generation. Antiy's research reveals that attackers can carry out "misalignment" attacks on the 3D printing process through various ways: first, tampering with 3D model files (such as STL files) and modifying internal structural parameters without changing the appearance; Second, invade the slicing software and adjust the layer thickness, filling density or supporting structure algorithm to generate unexpected stress concentration when the printed product is stressed. Third, directly modify the G-code instructions, introducing microscopic defects or altering the material deposition path during the printing process. A common characteristic of these attacks is that the printed parts appear flawless during visual inspection; structural defects are only revealed when the components are subjected to their design loads in service. This can lead to severe consequences in critical applications such as aerospace components, medical devices, and automotive parts. In the report, Antiy used aviation turbine blades as an example to show how by modifying the internal honeycomb structure parameters in the slicing algorithm, the blades can undergo resonance fatigue fracture when subjected to normal aerodynamic loads, while surface inspection cannot find any abnormalities. These XCon

research findings have, from a methodological perspective, validated the universality and harmfulness of "inaccuracy" as an independent form of attack.[10]

3.2 Examining FAST16's Tactical Orchestration Through the Lens of Effectiveness Within the NSA/CSS Cyber Threat Framework

The NSA/CSS Technical Cyber Threat Framework (NTCTF), released in 2018 [11], further refined the theoretical framework for cyberspace operations. NTCTF provides a structured description of cyberspace threats across three dimensions: Actor, Action, and Effect. In the Effect domain, NTCTF clearly categorizes operational objectives into five types: "Monitor, Exfiltrate, Modify, Deny, and Destroy." Among these, monitoring and data exfiltration fall under the CE/CNE domain—i.e., achieving intelligence value through covert information gathering or data extraction—while modification, denial, and destruction clearly belong to the CA/CNA domain—i.e., delivering operational effects by tampering with data, launching denial-of-service attacks, or physically destroying assets. The "modification"-type effectiveness demonstrated by fast16—systematic tampering with the computational results of high-precision engineering simulation software—is explicitly classified within the NTCTF framework as belonging to the CA/CNA category, rather than CE/CNE. This official framework serves as an authoritative reference point for characterizing FAST16 as a military operational activity rather than an intelligence operation.[11]

Returning to the operational mechanism of FAST16: it tampers with the floating-point computation results generated by simulation software such as LS-DYNA, which, at its core, constitutes a highly precise "spoofing" attack. Unlike Stuxnet's "disabling" attack, which directly damaged the centrifuge hardware, Fast16 opted for a more covert and harder-to-trace path of destruction. It does not directly destroy physical facilities, but allows engineers and scientists to make wrong design decisions based on tampered calculations-deviations in the load-bearing design of a bridge, inaccurate calculation of the material ratio of a nuclear device, and systematic errors in the stress analysis of a dam. These errors do not become immediately apparent while the malicious code is active; instead, they may surface months or even years later in the form of engineering quality failures, thereby significantly complicating efforts to attribute and trace the root cause. What is even more difficult is that when the damage finally occurs, the investigators will first suspect the designer's professional level, material quality or construction specifications, and will not first think of some malicious code more than ten years ago. The calculation results of the simulation software have been slightly tampered. This is the most strategic deterrent in the operational design of

fast16: it extends the attribution window from "when the damage occurs" to "when the attack is carried out", possibly more than a decade apart, far beyond the conventional traceability of any judicial and intelligence investigation.

Based on a SentinelOne report, Antiy Labs has developed the FAST16 cyber threat framework for NSA/CSS - affiliated cyber operations. FAST16 encompasses six stages, 12 objectives, and 29 specific behaviors.

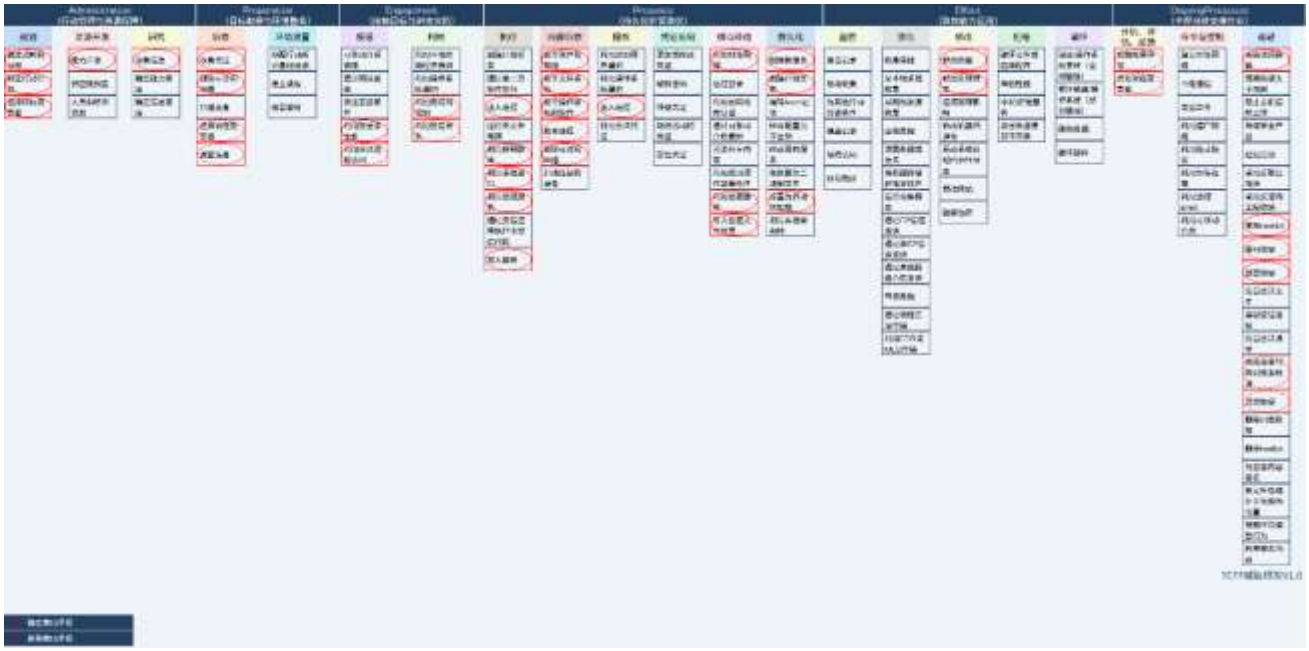

Figure 3-1: Tactical Annotations of fast16 Based on the TCTF Threat Framework (Generated from a SentinelOne Report)3

Table 3-1: TCTF Cyberspace Threat Framework Mapping for fast16 (Generated Based on the SentinelOne Report)31

Stage	Goal	Behavior
Action Management and Resource Assurance	Planning	Define strategy and objectives
		Develop an action plan
		Select the target victim
	Resource development	Capability development
	Research	Collect information
Targeted Site Investigation and Environmental Remediation	Reconnaissance	Collect vouchers
		Mapping the accessible network
Target Engagement and Offensive Penetration	Submit	Exploiting infected hosts
Persistent residency and	Execute	Inject into process

latency		Using Interpret Scripts
		Utilize the system interface
		Utilizing remote services
		Write to disk
	Internal reconnaissance	Enumerate accounts and permissions
		Enumerate the file system
		Enumerate operating systems and software
		Mapping the accessible network
	Privilege escalation	Inject into process
	Lateral movement	Utilizing peer-to-peer networks
		Utilizing remote services
		Write to remote file share
	Persistence	Create New Service
		Create a scheduled task
		Set to load at startup
Application of effectiveness	Modify	Modify data
		Modify process result
End-to-end continuous support operations	Avoidance	Direct disk access
		Using a rootkit
		Encrypted data

3.3 Isomorphism and Transformation Relations between CE and CA

However, to achieve this sustained and stable “malfunction” effect, FAST16 must first satisfy two prerequisite conditions: successfully implant itself into the target system and establish a long-term persistence mechanism. This is a quintessential manifestation of the isomorphism and transformational relationship between CE/CNE and CA/CNA, and it is also key to understanding the systemic nature of cyberspace operations. In U.S. military cyber operations doctrine, the persistent access established by CE operations, coupled with comprehensive situational awareness of the target environment and an in-depth understanding of the adversary’s defensive posture, constitute

the foundational conditions that enable the successful execution of subsequent CA operations. Without prior intelligence gathering and persistent foothold establishment, a precision data-tampering payload cannot be delivered to the intended target environment, nor can it remain undetected for an extended period to span the entire engineering analysis cycle.

From the perspective of fast16’s concrete implementation, its persistence mechanism demonstrates a very high level of engineering sophistication. The installation process of the fast16.sys kernel driver employs multi-layered camouflage and anti-detection techniques. First, the propagation component (the worm module) conducts a thorough environmental reconnaissance before entering the target system, checking for the presence of known antivirus processes, HIPS drivers, or behavior-monitoring tools. If any security software is detected, the worm module will automatically abort the kernel-driver installation and instead maintain a low-profile, dormant state, waiting for a more opportune moment. Secondly, kernel-mode driver development prior to the introduction of Driver Signature Enforcement in Windows leveraged the pre-Windows XP mechanism that allowed unsigned drivers to be loaded. Furthermore, fast16.sys employs a variety of anti-debugging and anti-forensic techniques during execution, such as memory code encryption, runtime self-modification, and detection and evasion of debugger symbols, in order to prolong its persistence within the target environment.

The fast16 sample was uploaded to VirusTotal in 2016, but the compilation timestamp of its core driver, fast16.sys, indicates a build date of 2005. During its extended, multi-year stealth phase, the framework must remain completely undetected within the target environment, avoiding triggering any alerts from security software, while also ensuring that its tampering logic remains effective across successive versions of the target application. This requires attackers to not only possess exceptionally sophisticated malware engineering skills, but also a deep understanding of the target software’s internal architecture, floating-point computation pipeline, and memory layout. It serves as an example of the continuous transformation of CE into CA capability.

4 The Contents of the SentinelOne Report Substantiate the U.S. Side's View That the Attack Constitutes an A²PT.

4.1 SentinelOne Has Previously Questioned the Use of the Term A²PT by Antiy Labs.

SentinelOne's peers had previously strongly questioned the A2PT terminology coined by Antiy Labs. SentinelOne researcher Dakota Cary posted exclusively on her X (formerly Twitter) account about A2PT, citing two 2024 Global Times articles that covered technical reports by Antiy Labs [12][13]. She specifically emphasized that Antiy uses A2PT to refer exclusively to "cyberattacks originating from the United States," thereby portraying the U.S. as "more powerful and more malevolent." Their post seeks to suggest that the term was "constructed" to align with Chinese state media's messaging, rather than being an academic concept grounded in objective technical standards. This challenge is not an isolated academic debate; rather, it is part of the long-standing propaganda campaign waged by the U.S. cybersecurity think tank known as the "Anti-China Chorus," in which SentinelOne has been actively involved.[12][13]

In response, our Russian colleague Oleg Shakirov, a specialist at the Cybersecurity Center of the Russian Academy of Foreign Affairs, commented as follows: Antiy has been using the term A²PT since at least 2016. He cited Antiy's 2016 report, "From Equation to System of Equations: An Analysis of the Cross-Platform Capabilities of Advanced Malware Used by the EQUATION Group" [6], as evidence to illustrate the evolutionary history of this technical concept. Shakirov pointed out that A2PT is not a concept hastily constructed to align with media coverage; rather, it is an academic term that Antiy has progressively developed and solidified over its long-term tracking and analysis of exceptionally capable cyber threat actors. This process was based on a systematic assessment of attack technique characteristics, the scale of their operational infrastructure, and their persistence capabilities. In its annual report released in December 2015 [14], Antiy Labs also categorized APT attack capabilities into several tiers: A2PT, APT, quasi-APT, and lightweight APT. A2PT is not about label politics; it is a capability-based classification of exceptionally sophisticated cyber threat actors.[6][14]

4.2 The True Inspirational Source of the A²PT Terminology

The terminology of Antiy A²PT was directly inspired by U.S. researchers. In 2014, Michael, a researcher at the Lockheed Martin Research Institute, published the article "Why Stuxnet Isn't APT?" [15]. This paper is a seminal

work in Western security scholarship that, for the first time, systematically challenges the practice of categorizing Stuxnet as a mere APT (Advanced Persistent Threat). Michael points out in the article: “The Stuxnet worm is extremely complex.” Its code is highly resistant to reverse engineering; it incorporates a PLC rootkit and several zero-day vulnerabilities, and can run on processors with different chipsets. In many cases, the binary files used in APT attacks are relatively simple and typically exploit a single vulnerability within client-side applications. Michael’s central argument is that Stuxnet is not a typical APT attack, because its objective was not long-term persistence and intelligence gathering (CNE/CE), but rather the direct physical destruction of infrastructure (CNA/CA). “The claim that Stuxnet was a military operation and should be classified as CNA, while APT attacks fall under CNE,” is precisely aligned with the analysis of FAST16’s operational characteristics presented earlier in this paper. It also served as the initial theoretical impetus for Antiy to distinguish between “ordinary APTs” and “ultra-capable, nation-state-level attacks.”[15]

In 2015, Antiy officially introduced the term A^{2}PT (Advanced Advanced Persistent Threat) to clearly distinguish cyberattack campaigns launched by ultra-capable state or sub-state actors—equipped with cross-platform payload delivery capabilities, substantial financial resources, and extensive infrastructure—from traditional APT attacks. This distinction is not merely a rhetorical enhancement; it is grounded in a thorough examination of the structural differences in threat configurations. Traditional APT attacks typically exhibit the following characteristics: relatively single or limited attack targets (such as specific enterprises, government agencies, or industries); moderate technical complexity of the payload (often leveraging combinations of known vulnerabilities or a single zero-day exploit); a low degree of sophistication (lacking systematic payload evolution and cross-platform capabilities); and controllable financial investment (usually in the range of tens of thousands to hundreds of thousands of U.S. dollars). In contrast, A2PT-level attacks exhibit a fundamentally different scale: their targets span multiple countries worldwide, various industries, and a wide range of infrastructure types. The attack payloads are cross-platform, supporting Windows, Linux, macOS, iOS, Android, embedded systems, industrial control systems, and more. The operational framework is highly sophisticated and modular, comprising dozens or even hundreds of independent components, each of which can be upgraded or replaced independently. Moreover, these campaigns require massive resource investments, often involving annual budgets in the hundreds of millions of U.S. dollars and research-and-development teams numbering in the hundreds.

In 2016, when Antiy Labs published its cross-platform analysis report on the Equation Group, it defined A2PT as "a threat actor with full-platform payload-based attack capabilities" and designated comprehensive cross-platform payload support as a key distinguishing characteristic of A2PT organizations. Since then, Antiy has consistently employed and refined the definition of this term in its annual threat reports and specialized analysis papers. The technical criteria for identifying A2PT actors have been progressively articulated along five key dimensions: First, the degree of systematization and modularity of the attack toolkit—A2PT groups typically maintain a toolset comprising dozens to hundreds of discrete components that are interconnected via standardized interfaces, allowing for flexible composition to suit diverse targets and operational environments. Second, the breadth of cross-platform payload delivery capabilities—A^{2}PT organizations possess full-stack offensive capabilities spanning mainstream operating systems, embedded platforms, and industrial control systems, rather than being confined to a single platform. Third, the depth and stealth of persistence mechanisms—A2PT groups' persistence modules often penetrate deep into the operating system kernel, firmware, and even hardware layers, enabling them to remain undetected within the target environment for several years or even over a decade. Fourth, the strategic intent and physical damage potential of attack operations—A^{2}PT campaigns frequently go beyond intelligence gathering, aiming instead to inflict physical destruction and achieve strategic objectives. Fifth, the iterative evolution of the operational infrastructure—A^{2}PT toolkits exhibit continuous updates and generational advancements, with each new iteration incorporating emerging technologies and circumventing evolving defenses, thereby establishing a sustained technical advantage in the cyber arms race.

4.3 The High Complexity of fast16 Confirms A^{2}PT's Assessment.

The technical characteristics of fast16 precisely provide another strong piece of evidence supporting the term A2PT. SentinelOne's report outlined several key differences between fast16 and the common attack tools of the time: "In the early 2000s, there was a proliferation of network worms." Most are written by enthusiasts, spread rapidly, and carry virtually no payload. Fast16 originated in the same period, but its emergence as a national-level tool followed a completely different pattern. "fast16 exhibits multiple typical characteristics of A2PT-level, nation-state-grade cyberattack tools.

First, a highly targeted, modular payload design. Fast16 does not spread infection indiscriminately; instead, it precisely targets specific engineering software such as LS-DYNA, PKPM, and MOHID. This level of targeting indicates that the attacker has previously conducted extensive reconnaissance and intelligence analysis—gaining a

deep understanding of the target organization's software stack, the distribution of software versions, the characteristics of its operational environment, and the vulnerabilities inherent in its computational workflows. Preparing such intelligence inherently requires substantial resources and time, far exceeding the reconnaissance costs that typical APT actors are willing to incur.

Second, deep kernel-level persistence. fast16 achieves silent manipulation of computation results by leveraging a kernel-mode driver, thereby evading detection by user-mode security software. User-mode security monitoring tools typically cannot directly inspect kernel-level memory operations or function hooking activities. Consequently, fast16's tampering operations inherently enjoy an evasion advantage over traditional security defense systems. This kernel-level persistence mechanism requires the attacker to have extensive experience in operating system kernel development, making the technical barrier significantly higher than that of typical user-mode malware.

Third, high levels of concealment. The actual deployment environment of fast16 targets host systems such as Windows XP and Server 2000 within Iran's industrial infrastructure, many of which were still in active operation and use as late as 2010. However, after the Stuxnet incident was exposed that year, it was not detected during Iran's associated forensic investigations. More critically, this sample was already submitted to the VirusTotal platform as a public file in 2016. Despite continuous static and dynamic analysis, metadata vector sorting, and threat intelligence correlation efforts by security firms worldwide, it remained undetected and unflagged. This indicates that fast16's architecture effectively circumvents conventional feature detection and common kill mechanisms. Its designers, in the course of its development, employed techniques such as kernel-driver obfuscation, anti-static-analysis measures, and modular engineering to keep it under the radar for an extended period.

Fourth, precision engineering development. SentinelOne's code analysis revealed that the fast16 binary contains traces of a version control system, indicating that it underwent multiple rounds of iterative development and version management. This engineering practice is highly aligned with enterprise-level software development workflows, in stark contrast to APT tools typically developed by individuals or small teams. The use of a version control system signifies that the development of fast16 is an organized, process-driven engineering effort, rather than an ad-hoc creation by a lone hacker or a small team.

Fifth, the intent of physical destruction and the logic of strategic-level disruption. fast16 is not a typical spyware designed to steal data or gain system control; rather, it is a "logic bomb"-style attack framework aimed at inflicting indirect physical damage in the real world. Its core mechanism involves using a kernel-level driver to dynamically

patch floating-point computation code in memory, thereby injecting small but systematic errors into the computational results of specific engineering simulation software. This type of tampering does not cause the software to crash or trigger standard security alerts; instead, it allows the computation to appear to be functioning normally while the output data has been deliberately manipulated.

In summary, the five technical features of fast16-load modularity, kernel persistence, high degree of concealment, and engineering development of physical damage intent-are highly consistent with many of the characteristics of A2PT defined by Antiy. This alignment is no mere coincidence; rather, it reflects a common pattern in the development of offensive cyber capabilities by nation-state-level, highly capable cyber threat actors: they all adhere to a developmental logic characterized by massive resource allocation, systematic infrastructure building, protracted adversarial engagement, and strategic effectiveness. As an international cybersecurity vendor that has long tracked nation-state-level threats, SentinelOne’s report highlights the fast16 indicators of compromise, which, from an external perspective, validate the objectivity and broad applicability of Antiy Labs’ A^2PT framework. Dakota Curry’s critique of the A^2PT terminology, in essence, reveals that it is a product of her own ideological stance rather than an inherent flaw in the A^2PT concept itself.

5 Geopolitical Security Context Analysis of SentinelOne’s Report

5.1 The Professional Standard and Dual Facets of Technical Reports

First, it must be affirmed and acknowledged: from a purely technical perspective, SentinelOne’s report is an analysis of exceptionally high professional caliber. Its ability to pinpoint the target sample from a vast pool of data, conduct reverse engineering of the fast16 kernel driver’s tampering mechanism, map out the target software’s identification logic, and reconstruct the relationship between the attack payload and propagation components all demonstrate a solid technical foundation. In private discussions, some industry insiders believe that these findings can be seen as a response by threat researchers to the notion that "AI can completely replace security researchers." “Human experts’ analytical capabilities and insights still hold immense, irreplaceable value in the AI era,” a Silicon Valley peer told us. “Although models like Mythos can discover a vast number of vulnerabilities with exceptional efficiency, it’s difficult for them to independently produce an equivalent volume of analytical results without human experts.”

As a security research team that also conducts threat actor attribution and analysis, Antiy Labs respects this type of technically sophisticated work. However, we maintain our own political stance. We maintain that endorsing the technical rigor of the report and analyzing and critiquing the motivations behind its publication are two distinct issues. The former should not be undermined by the latter, nor should the latter be obscured by the former.

As a company characterized by a “revolving door” dynamic, SentinelOne’s reporting practices cannot be analyzed in isolation from its inherent organizational structure and vested interests. SentinelOne was founded in 2013 and is headquartered in Mountain View, California. Its product, ActiveEDR, is an endpoint detection and response platform based on a behavioral AI model. The company exhibits a very typical “revolving door” dynamic characteristic of U.S. intelligence agencies. The so-called “revolving door” refers to a two-way mobility mechanism between U.S. government officials and intelligence agency personnel on the one hand, and senior executives in the private sector on the other: after leaving government service, officials join private companies as top managers, leveraging their government connections and access to intelligence resources to secure contracts for those firms; meanwhile, private-sector executives influence government procurement decisions and the direction of industry regulation by providing policy advice and technical services to the government. This mechanism is particularly prevalent in the U.S. cybersecurity industry and has become a key conduit for intelligence agencies to extend their influence into the private sector.

The "revolving door" feature of the SentinelOne is reflected in multiple levels. Its Chief Trust Officer, Alex Stamos, has served as a member of the Aspen Institute’s Cybersecurity Working Group, a member of the Cybersecurity and Infrastructure Security Agency’s (CISA) Advisory Council, and a member of the Advisory Board of Estonia’s NATO Cooperative Cyber Defence Centre of Excellence (CCDCOE). Stamos’s career trajectory is a textbook example of the “revolving door” phenomenon: after serving as Facebook’s Chief Security Officer and handling the Russian interference in the U.S. presidential election, he moved into academia (at Stanford University’s Freeman Spogli Institute for International Studies), then joined SentinelOne as an executive, while continuing to serve as a consultant to government and quasi-governmental organizations such as CISA and NATO. In January 2023, SentinelOne officially became a member of CISA's Joint Cyber Defense Defense Collaborative (JCDC); in November of the same year, the company acquired Krebs Stamos Group (Krebs Stamos Group) and changed its name to PinnacleOne Strategy Consulting Company, completing an important financial transaction process in the revolving door.

In the lengthy Observer Network report titled "The 'Anti-China Chorus' of U.S. Cybersecurity Think Tanks" [16], the U.S. side's mechanism for manipulating anti-China public opinion—primarily staged at hearings held by the U.S.-China Economic and Security Review Commission (USCC) [17]—is systematically documented. The key points are summarized below. At the USCC hearing on "China's Cyber Capabilities" held on February 17, 2022, Dakota Cary served as the "lead vocalist," joining forces with John Chen, a fellow at the American Enterprise Institute, Winnona DeSombre, a former cybersecurity prosecutor at the U.S. Department of Justice, and others to form a "chorus of anti-China rhetoric." Cary's testimony articulated a comprehensive narrative aimed at demonizing Chinese cybersecurity firms and advancing efforts to exclude China from international technological cooperation frameworks. SentinelOne is precisely the technical enforcement layer of this orchestration mechanism in cyberspace. The technical reports it has released serve the value narrative of U.S. monopoly capital and intelligence agencies, echoing the political posturing at the USCC hearings and together forming distinct links in the chain of cyber containment against China. Based on corporate operational data, SentinelOne's share of U.S.-based official orders surged from 12% in 2021 to 37% in 2025. This explosive growth has been highly correlated with the company's anti-China actions. In the second quarter of 2023, following SentinelOne's release of a report that specifically named Antiy, 360, and Qi Anxin as Chinese cybersecurity firms, government contracts increased by 15%. In the fourth quarter of 2024, after the company first proposed excluding China from Microsoft's MAPP (Microsoft Active Protections Program), government orders rose by another 22%. These orders include a $57 million contract from the U.S. Department of Defense for a vulnerability monitoring platform, as well as a $21 million-per-year subscription service for cyber threat intelligence provided to the Federal Bureau of Investigation (FBI). "Publishing technical reports that smear foreign countries, hiring executives with intelligence agency backgrounds to serve as revolving-door hires, and securing defense and government contracts" has become a well-established growth trajectory for U.S. cybersecurity firms. The essence of this trajectory is to clothe geopolitical narratives in the professional veneer of technical reports, thereby amplifying one's standing in government budgets by framing external threats. This ultimately gives rise to a positive feedback loop of "threat amplification – security spending – corporate profit." [16][17]

5.2 The Strategic Timing of the Report's Release and its "Psychological Warfare" Intent

Against this backdrop, the timing of the FAST-16 report's release and the strategic choices underlying it warrant close examination, as their underlying motives are open to scrutiny. The report was released in April 2026,

at a critical juncture marked by heightened tensions in the Middle East, the near-collapse of related negotiations, and the real possibility of a renewed U.S.-Israeli war against Iran. During this highly sensitive strategic window, SentinelOne's decision to unveil fast16 at the Black Hat Asia conference was no accident.

SentinelOne's report release strategy, in essence, relies on the aftereffect of the earthquake network incident and makes a "military parade" demonstration of the US cyber attack capability. As the first publicly confirmed national-level cyber physical attack in human history, Stuxnet's international influence and psychological deterrent effect on countries around the world, especially countries in the Middle East, have not disappeared. By comparing fast16 to Stuxnet and emphasizing that fast16 is "earlier" and "more covert" than Stuxnet, SentinelOne are actually sending a clear strategic signal that the United States has a longer, more advanced and more covert cyber sabotage capability than the public knows, and that any country that is an enemy of the United States should maintain full "awe" of it ". This operation, which disintegrates the opponent's psychological defense line and resistance will by displaying overwhelming technological superiority, is called "spiritual warfare" (Psychological Operations) in military psychology, that is, to weaken the enemy's fighting will, undermine the enemy's decision-making confidence and reduce the enemy's resistance ability through information manipulation and psychological influence.

The specific narrative strategies employed in the report further substantiate its character as a form of psychological warfare. The technical complexities of Fast16 throughout the report-"nightmarish," "imperceptible," and "subtle failure"-are actually sending a psychological hint to potential adversaries (especially Iran) that even if your defenses can withstand known threats, the United States still has more advanced capabilities than you can imagine and defend against. This narrative strategy mirrors the reporting pattern of Western media when Stuxnet was exposed in 2010: at the time, Western outlets similarly emphasized Stuxnet’s “unprecedented” nature and its supposed “undefendability” to reinforce the U.S. narrative of technological hegemony in cyberspace.

More concerning is that the report clearly smuggles political biases into its organizational structure. On the one hand, the report sidesteps the fundamental fact that FAST16, like Stuxnet, is a cyberspace operation initiated and led by the U.S. government against the critical infrastructure of a sovereign state, and it offers no ethical or international-law-based assessment of this fact. The implicit message of this linguistic strategy is that FAST-16 constitutes a neutral demonstration of technical capability, rather than a challenge to international law and the principle of sovereignty. On the other hand, the report particularly emphasizes Iran’s “alleged violation of Article T of the JCPOA,” directly linking the exposure of FAST-16 to the issue of Iran’s compliance with the nuclear agreement. This is a

textbook example of weaponizing technical reports to advance a geopolitical narrative—building professional credibility through an accumulation of technical details, then using political implications to steer the reader's value judgments.

5.3 "Neutral" Packaging and Commercial Considerations

Meanwhile, in its report, SentinelOne also adopts a deliberate "neutral" technical stance, seeking to position itself as an independent, third-party cybersecurity vendor free from U.S. government influence, in order to gain the trust of a broader international customer base. The report employs a "suggestive but non-conclusive" rhetorical strategy in its attribution analysis—described by Guerrero-Saade as "hardly far-fetched"—which guides readers to recognize the technical links between FAST16 and both the NSA and the Stuxnet operation while refraining from reaching a definitive, official attribution. This style of presentation is common in commercial security reports: on the one hand, it demonstrates the depth of the analysis by correlating technical indicators; on the other hand, it maintains flexibility in wording to accommodate varying levels of political sensitivity among clients in different regions while safeguarding the firm's commercial interests. For a cybersecurity vendor operating in the global market, this prudent attribution strategy is both an adaptation to a complex geopolitical landscape and a necessary approach to safeguarding cross-regional business relationships.

6 Concluding Remarks

Strategic composure must not be undermined by "psychological warfare."

In its 2025 report, "A Reconsideration 15 Years After the Disclosure of the Stuxnet Attack" [18], Antiy Labs reviewed the full scope of this landmark cyberspace operation and conducted a dialectical analysis of the intrinsic characteristics of ultra-capable cyber threat actors. In the technical section, Antiy Labs argues that, with regard to A2PT campaigns: "From both a tactical and a technical perspective, their capabilities and operational patterns warrant close scrutiny and analysis; however, there is no need to exaggerate them into a new myth." "This assessment is equally relevant to the current public discourse surrounding FAST16: By selectively presenting facts and employing emotionally charged narratives, SentinelOne has succeeded in framing U.S.-led cyberattacks as an 'unfendable nightmare.' The true objective of this psychological warfare tactic is to undermine the defenders' confidence and resolve." Antiy Labs believes that, in the face of such a public-opinion offensive, China's

cybersecurity industry must maintain clear strategic awareness: it should both acknowledge the true extent of its adversaries' capabilities and avoid falling into the psychological traps they have set.[18]

Cyber threat actors with exceptionally high capabilities are "real tigers, yet ultimately paper tigers." The so-called "real tigers" refer to adversaries that genuinely possess unprecedented advantages in technical capabilities, resource allocation, engineering infrastructure, and organizational structure. Judging from Fast16's more than decade-long dormant period, its precise tampering with floating-point computation kernels, and its targeted exploitation across multiple engineering simulation platforms, the operation clearly reflects the U.S. side's sustained modus operandi. The resource base underpinning a superpower's military intelligence apparatus comprises top-tier talent recruitment, cutting-edge technology R&D, a global intelligence network, and diplomatic cover. When confronting an adversary of this caliber, both the illusory confidence embodied in "physical isolation to keep the enemy outside the gates" and the nihilistic attitude of giving up resistance on the grounds that "it's impossible to defend anyway" are extremely dangerous and misguided. Attackers are not mythical ghosts; rather, they are real-world actors with well-defined organizational boundaries, engineering cycles, and cost constraints. Their work also has underlying assumptions that it must rely on, and they too can make mistakes.

In our defensive analytics work in the cybersecurity domain, we have repeatedly engaged in sudden, ad-hoc engagements. However, the broader historical trajectory of this effort is undeniably a protracted campaign that demands endurance, steadfastness, and coordinated collaboration. On this battlefield, attackers are continuously developing and deploying new tools and techniques, while we remain committed to iterative improvements in our capabilities for detection, analysis, attribution, and defense—driven by both feature engineering and knowledge engineering. Meanwhile, in an era where artificial intelligence has simultaneously become a comprehensive accelerator of cyber threats, a source of risk, and a prime target, it is also rapidly reshaping the landscape of cybersecurity. We will strive to realize a new vision of the seamless integration of productivity, safety, and combat effectiveness.

Appendix 1: Antiy AVL Code’s Automated Analysis Results for Fast16 Samples

The automated analysis results for the Fast16 sample were pushed to the Antiy Vertical Response Platform today, April 26, 2026. Please follow the “Antiy Vertical Response Platform” official account to read.

Appendix II: References

[1] SentinelOne. fast16 | A mysterious Shadow Brokers reference reveals high-precision software sabotage dating back five years before Stuxnet. 2026-04-23.
https://www.sentinelone.com/labs/fast16-mystery-shadowbrokers-reference-reveals-high-precision-software-sabotage-5-years-before-stuxnet/

[2] Symantec. Stuxnet 0.5: The Missing Link [R/OL]. Symantec Security Response, 2013. https://docs.broadcom.com/doc/security-response-w32-stuxnet-0-5-the-missing-link-13-en

[3] Antiy. Comprehensive Analysis Report on the Stuxnet Worm Attack Against Industrial Control Systems [R/OL]. (2010-09-27)
https://www.antiy.com/response/stuxnet/Report_on_the_Worm_Stuxnet_Attack.html

[4] Antiy Labs virusview.Trojan/Win32.Fast16[EB/OL].(2026-04-01).
https://www.virusview.net/malware/Trojan/Win32/Fast16

[5] NAKASHIMA E. NSA Plans for a Supercomputer to Break Encryption Codes [M/OL]. The Washington Post, January 28, 2012. https://www.washingtonpost.com/business/economy/nsa-plans-for-a-supercomputer-to-break-encryption-codes/2012/01/28/gIQAt3NScQ_story.html

[6] From Equations to “Systems of Equations”: An Analysis of the Cross-Platform Capabilities of EQUATION’s Advanced Malware
https://www.antiy.com/response/EQUATIONS/EQUATIONS.html

[7] U.S. Joint Chiefs of Staff. JP 3-12, Cyberspace Operations [S/OL]. 2018-06-08.
https://irp.fas.org/doddir/dod/jp3_12.pdf

[8] Antiy. Recreating the Winter Myth—Analysis of Printer USB Interface Attack Paths [R/OL]. XCon 2008

[9] Antiy. Jamming and Attacks on Long-Wave Time-Signal Broadcasts [R/OL]. XCon 2012, 2012.

[10] Antiy Labs. Research on Instruction Tampering in 3D Printers and Finished-Product Structural Distortion Attacks [R/OL]. XCon 2013

[11] NSA/CSS. NSA/CSS Technical Cyber Threat Framework (NTCTF)[S/OL]. National Security Agency/Central

Security Service, 2018.
https://www.nsa.gov/portals/75/documents/what-we-do/cybersecurity/professional-resources/ctr-nsa-css-technical-cyber-threat-framework.pdf

[12] Global Times. Annual Cybersecurity Industry Report: The United States Remains the Primary Global Threat to Cybersecurity
https://world.huanqiu.com/article/4GJd0ZOyYDe

[13] Global Times. Can China's cybersecurity capabilities not detect U.S. cyberattacks? A new report thoroughly refutes such claims. 2024
https://3w.huanqiu.com/a/5e93e2/4H4KWWI0JtC

[14] An Tian. Review and Prospect of Network Security Threats in 2015 [R/OL]. 2015-12.
http://www.antiy.com/response/2015_Antiy_Annual_Security_Report.html

[15] CLOPPERT M. Why Stuxnet Isn't APT?[EB/OL]. SANS Institute, 2011-03-07.
https://digital-forensics.sans.org/blog/2011/03/07/why-stuxnet-isnt-apt

[16] Mindlab: The "Anti-China Chorus" of U.S. Cybersecurity Think Tanks [EB/OL]. (2025-09-22).
https://www.guancha.cn/xinzhiguanchasuo/2025_09_22_790880.shtml

[17] Hearing on "China's Cyber Capabilities: Warfare, Espionage, and Implications for the United States" "China's Cyberspace Capabilities: Cyber Warfare, Espionage, and Implications for the United States" Hearing, February 17, 2022
https://www.uscc.gov/hearings/chinas-cyber-capabilities-warfare-espionage-and-implications-united-states

[18] Antiy. Technical Analysis: Reflecting on the Stuxnet Attack 15 Years After Its Exposure [R/OL]. (2026-03-06).
https://www.antiy.com/response/Stuxnet_Revisited-15_Years_of_Technical_Insights.html

Appendix III: About Antiy

Antiy is committed to enhancing the network security defense capabilities of its customers and effectively responding to security threats. Through more than 20 years of independent research and development, Antiy has developed technological leadership in areas such as threat detection engines, advanced threat countermeasures, and large-scale threat automation analysis.

Antiy has developed IEP (Intelligent Endpoint Protection System) security product family for PC, server and other system environments, as well as UWP (Unified Workload Protect) security products for cloud hosts, container and

other system environments, providing system security capabilities including endpoint antivirus, endpoint protection (EPP), endpoint detection and response (EDR), and Cloud Workload Protection Platform (CWPP) , etc. Antiy has established a closed-loop product system of threat countermeasures based on its threat intelligence and threat detection capabilities, achieving perception, retardation, blocking and presentation of the advanced threats through products such as the Persistent Threat Detection System (PTD), Persistent Threat Analysis System (PTA), Attack Capture System (ACS), and TDS. For web and business security scenarios, Antiy has launched the PTF Next-generation Web Application and API Protection System (WAAP) and SCS Code Security Detection System to help customers shift their security capabilities to the left in the DevOps process. At the same time, it has developed four major kinds of security service: network attack and defense logic deduction, in-depth threat hunting, security threat inspection, and regular security operations. Through the Threat Confrontation Operation Platform (XDR), multiple security products and services are integrated to effectively support the upgrade of comprehensive threat confrontation capabilities.

Antiy provides comprehensive security solutions for clients with high security requirements, including network and information authorities, military forces, ministries, confidential industries, and critical information infrastructure. Antiy has participated in the security work of major national political and social events since 2005 and has won honors such as the Outstanding Contribution Award and Advanced Security Group. Since 2015, Antiy's products and services have provided security support for major spaceflight missions including manned spaceflight, lunar exploration, and space station docking, as well as significant missions such as the maiden flight of large aircraft, escort of main force ships, and Antarctic scientific research. We have received several thank-you letters from relevant departments.

Antiy is a core enabler of the global fundamental security supply chain. Nearly a hundred of the world's leading security and IT enterprises have chosen Antiy as their partner of detection capability. At present, Antiy's threat detection engine provides security detection capabilities for over 1.3 million network devices and over 3 billion smart terminal devices worldwide, which has become a "national-level" engine. As of now, Antiy has filed 1,877 patents in the field of cybersecurity and obtained 936 patents. It has been awarded the title of National Intellectual Property Advantage Enterprise and the 17th (2015) China Patent Excellence Award.

Antiy is an important enterprise node in China emergency response system and has provided early warning and comprehensive emergency response in major security threats and virus outbreaks such as “Code Red”, “Dvldr”, “Heartbleed”, “Bash Shellcode” and “WannaCry”. Antiy conducts continuous monitoring and in-depth analysis

against dozens of advanced cyberspce threat actors (APT groups) such as “Equation”, “White Elephant”, “Lotus” and “Greenspot” and their attack actions, assisting customers to form effective protection when the enemy situation is accurately predicted.

For more information, please visit:	*http://www.antiy.com (Chinese)* *http://www.antiy.net (English)*
For more information about Antiy Enterprises Security, please visit:	*http://www.antiy.cn*
For more information about Antiy Mobile Security (AVL TEAM), please visit:	*http://www.avlsec.com*