

Exorcising the Ghost in the Machine



Debunking Myths Around Supply Chain Intrusions

JOE SLOWIK,

PRINCIPAL SECURITY ENGINEER, GIGAMON APPLIED THREAT RESEARCH

Executive Summary

Supply chain intrusions represent a concerning trend in information security. Such incidents attract headlines and researcher attention, but without first investigating what a supply chain event really encompasses or requires. Exploring this idea, we find that there are multiple potential routes to supply chain incidents: via hardware or software products, or through service or contractor relationships. Central to each is a subversion of fundamental trust between the supply chain vector used by the adversary and the ultimate victim.

In investigating this worrisome concept, though, we begin to find multiple points of friction within supply chain intrusions. Essentially, a supply chain event represents two incidents: first identifying and compromising the vector entity, then identifying a way to use that vector to reach ultimate victims for the intended purpose. Throughout this process, adversaries face various challenges and choices in how to shape their operations to maintain stealth, avoid detection, and ensure objectives are finally achieved.

Given this enhanced understanding of supply chain intrusions, defenders can leverage adversary pain points as opportunities for detection and mitigation. Through communication monitoring, network design and implementation, and planning for resilience, system owners and defenders can make an already difficult task even more so for would-be adversaries. In adopting this nuanced view of supply chain intrusions, we arrive at a place where the popular specter of the supply chain attack as a nearly invisible, almost impossible to defend against scenario transforms into just another intrusion vector that can be defeated.

Table of Contents

Executive Summary	2
Introduction	3
Defining a “Supply Chain Attack”	4
Hardware and Software Supply Chain Attacks	4
Trusted Third-Party Intrusions	6
Executing and Controlling a Supply Chain Intrusion	7
Intrusion Vector	8
Delivering Capabilities to Victims	8
Capability Execution and Control	9
Difficulties and Challenges	10
Issues Concerning Targeting and Specificity	11
Concerns on Attack Command and Control	11
Defender Advantages	13
Exploiting Communication and Control	13
Architecture, Visibility, and Monitoring	13
Resilience, Response, and Defense in Depth	15
Conclusion	16
Works Cited	17

Introduction

Supply chain attacks raise the prospect of stealthy, nearly impossible-to-detect intrusions by subverting fundamental trusts between network operators and their suppliers, contractors, and related parties. Examples of breathless reporting on such events range from mainstream reporting on alleged modification of computer hardware such as the largely debunked “The Big Hack” to industry-specific rumors of supply chain manipulation for electric utility components.^{1,2} While concrete proof or direct evidence for any of these alleged incidents is circumstantial at best and typically nonexistent, the nature of the problem makes proving (or disproving) such events difficult or impossible. Once fundamental system trust is questioned, discussion quickly shifts such that one must prove that a device is not compromised, which is a near-impossible task.

As a result of the above dilemma and the supposed power of such potential events, a host of initiatives exist to “solve” the issue of supply chain security. These range from government efforts, such as various initiatives launched by the United States and emerging standards in the European Union,^{3,4,5} to private certification and assessment services.^{6,7} Perhaps most alarming are actions seeking to sever or otherwise interrupt global supply chains, such as the since-revoked US Executive Order relating to non-US components in the electric utility sector.⁸ Ultimately, a near-panic over supply chain concerns and possible attacks has prompted a sequence of potentially costly and disruptive initiatives to address the issue.

Absent from most (if not all) of these discussions is any realistic assessment of just what sort of risk supply chain attacks pose to network operators and critical infrastructure providers. While potential impacts can certainly be quite high depending on circumstances, few commentators bother to investigate the likelihood or feasibility of conducting such attacks. Thus, the traditional risk equation – probability multiplied by impact – has been abandoned, with all emphasis placed on potential impacts while probability is left unexplored and unexamined.

This paper seeks to rectify the above concern by exploring, in detail, just what constitutes a supply chain attack. In identifying the necessary steps and efforts required to successfully carry out such an attack, defenders, policymakers, and other stakeholders can gain greater appreciation for the difficulty in executing such attacks. From this discussion, we can also identify defensive measures and recommendations to address residual risk left over by the potential for supply chain subterfuge – without having to resort to ripping out all equipment and starting from some new, yet still questionable, baseline.

Defining a “Supply Chain Attack”

Before proceeding further, understanding just what constitutes a “supply chain attack” is necessary for focus, as many items seem to enter conversations on this topic. In a traditional business sense “supply chain” designates “a network between a company and its suppliers to produce and distribute a specific product to the final buyer.”⁹ Thus a supply chain, formally defined, can include products, services, equipment, capabilities, and transit nodes from raw materials through final delivery to consumers – multiple potential touch (or manipulation) points for an attacker.

Yet if taken broadly, supply chain events would encompass all events from internet service provider (ISP) intrusions through telecommunication manipulation to events that impact physical logistics and delivery. While academically correct and serious for business continuity planning purposes, such items are superfluous to the current discussion focused on cyber-specific events and only serve to muddy the waters. Additionally, we must differentiate between inadvertent vulnerabilities in software or hardware that may be leveraged by adversaries and the deliberate, active insertion of such vulnerabilities by adversaries for future use. The former represents an issue for coding quality and product testing; the latter stands as a security issue that end users must consider while defending their networks.

This paper will identify two primary types of supply chain attacks with respect to cyber impacts, and primarily focus on one of them given the differences between the two in terms of risk, actualization, and defense: product-focused supply chain intrusions, and services-oriented supply chain incidents.

Hardware and Software Supply Chain Attacks

The most typical example used for a cyber-related supply chain attack are instances – such as the Bloomberg “The Big Hack” article¹ – where an adversary compromises hardware or software to enable insertion of malicious functionality into a product prior to its receipt by a consumer. An apocryphal, largely (although not definitively) debunked example would be the alleged CIA operation against the Soviet gas industry in 1982. In this case, identification of technology desired by the Soviet Union supposedly enabled intelligence operatives to seed a malfunctioning device into the Soviet supply chain, which later led to an alleged explosion on the Trans-Siberian gas pipeline.^{10,11} Yet in this case, subsequent analysis and review of events cast many doubts as to whether this even happened.¹²

In any event, the idea of inserting a difficult or nearly impossible to identify flaw, logic bomb, or related defect into hardware prior to receipt presents an alluring story for how attackers could potentially circumvent multiple layers of defenses at a stroke to wreak future or immediate havoc. In this scenario, attackers would need access to either manufacturing or, depending on the equipment and circumstances, devices in transit to insert, modify, or otherwise alter functionality for malicious purposes.



Figure 1: Product Supply Chain Progression

More recently, through both the increasing adoption of open-source software (OSS) and interconnectivity and reliance of organizations on multiple programs or applications for business, software has taken a near equivalent, if not greater, role to hardware as a supply chain vector. Although not proved at the time of this writing, the attempted commit to the Linux kernel by a Huawei engineer in May 2020 – which included a trivially easy to exploit vulnerability – echoes this fear as vulnerabilities may be deliberately introduced into software for later exploitation.^{13,14} Similar concerns apply to externally sourced software vital for day-to-day business functionality, from industrial control system management software through word processing and web content management applications.

HARDWARE AND SOFTWARE INTRUSIONS

	Hardware	Software
Access	Include functionality or mechanisms within hardware that allow for future access, such as hard-coded user/password combinations or remote access links.	Include or introduce user/password combinations into software or remote access features enabling a party to communicate with the application outside of normal routines.
Destruction	Modify equipment to fail or otherwise perform undesirably in various circumstances to induce disruption or destruction.	Modify software to fail to respond or appropriately behave in extreme instances enabling physical events to propagate to potentially destructive scenarios.

Table 1: Mapping Types and Impacts of Product Supply Chain Intrusions.

As summarized in Table 1 above, access – whether facilitated through hardware manipulation or software injection – may be only one goal of operations. Facilitating an attacker’s ability to communicate with an otherwise difficult to access or denied environment is significant, but other goals exist too. For example, the pipeline story above highlights the possibility of introducing a potentially destructive payload into a sensitive environment through trusted channels.

In this destructive scenario, an attacker modifies either software or hardware to introduce or enable a fault. The critical question, dealt with in greater detail below, is whether such functionality is autonomous (requires no external intervention or attacker control), or merely facilitates an operation by a remote attacker communicating with an implant. The implications for these are quite serious and influence the entire discussion on the feasibility and desirability (from the attacker’s perspective) of hardware and software supply chain attacks.

Trusted Third-Party Intrusions

Another type of supply chain attack of increasing prominence is the trusted third-party attack. In this scenario, an adversary breaches an intermediary – such as a contractor, service provider, or vendor – as a means to gain access to that intermediary’s customers or clients. Where trusted links, especially remote network access, exist between the trusted third party and the ultimate victim, this attack can be a powerful mechanism to breach an otherwise well-secured environment through abuse of legitimate links.

Perhaps the most famous example of trusted third-party access used to facilitate a much larger breach is the Target incident from 2013–2014. In this event, the adversary breached a mechanical contractor with links to refrigeration and HVAC systems inside Target’s network. The intruder used this access to pivot to Target’s point of sale (PoS) systems to facilitate widespread credit card theft.^{15,16} The resulting loss of customer data and subsequent penalties (to say nothing of loss of reputation) cost Target over \$200 million, including legal costs, as of 2017 – all due to an unmonitored, insecure connection with a facilities maintenance contractor.¹⁷




Figure 2: Summary Timeline of Trusted Party Supply Chain Intrusions.

However costly, events such as the Target breach have since been dwarfed by intrusions targeting various managed service providers (MSPs). By nature of their operations, MSPs have privileged, near-unfettered access to client networks to perform their duties – often with little or no security monitoring in place from clients into MSP activities. As a result, an MSP breach can facilitate the follow-on compromise of multiple client entities with little or no indication that something is amiss. Such a situation describes Operation Cloud Hopper as well as other activity attributed to the People’s Republic of China (PRC)-linked group APT10.^{18,19,20}

Seemingly less serious but financially more impactful, in 2017 (and possibly as early as 2015) phishing campaigns initially targeting IT outsourcing giant Wipro subsequently led to the breach of multiple additional entities from 2017 through 2019.^{21,22} More seriously, entities tied to the Russian government targeted multiple contractors to breach several electric utilities in North America and potentially Europe as part of the Palmetto Fusion attacks and linked campaigns from 2017 through 2020.^{23,24,25} Using initial compromises at contractor organizations, the adversary leveraged this access to produce follow-on access mechanisms, such as phishing campaigns from the compromised organizations, to breach ultimate victims in critical infrastructure.

More recently, the actor responsible for the SUNBURST campaign, also linked to Russian government entities, utilized two forms of supply chain compromise from 2019 through 2020. The first, a very well reported compromise of the software build environment for the SolarWinds Orion network visibility software,^{26,27} reflects a software supply chain intrusion such as those discussed in the previous section. In parallel, the entity also utilized compromise of a Microsoft reseller to attempt to breach the security vendor CrowdStrike.^{28,29} While all available evidence suggests this effort failed, we nonetheless gain insight into an entity attempting multiple supply chain intrusions across both product and services vectors within the same timeframe.

Based on the previous sections, there appear to be numerous examples of cyber-focused supply chain intrusions over the past decade. While these are concerning, and have led to some spectacular results such as the previously mentioned Cloud Hopper and SUNBURST campaigns, critical examination of the precise methodology behind supply chain intrusions is insightful to see just how useful and likely such intrusions may be moving forward.

Executing and Controlling a Supply Chain Intrusion

Supply chain intrusions are not “bolt from the blue” events, but rather represent multiple, interdependent operations culminating in some ultimate activity (espionage, theft, or potentially even attack) for final-stage victims. While this superficially appears to map to concepts such as the Cyber Kill Chain,³⁰ we must recognize that there are actually two iterations of the kill chain for adversaries conducting a supply chain intrusion: initial compromise of the supply chain vector or entity, and follow-on compromise of the ultimate victims of the effort.

As illustrated in Figure 3 below, an adversary executing a supply chain intrusion has multiple decision points on methodology and targeting, as well as multiple potential failure points. A supply chain intrusion can, if successful, “leapfrog” several other security controls as part of its deployment – but may also be subject to discovery, interdiction, or outright failure if any one of a number of critical dependencies is disturbed. Adopting this view of supply chain events does not diminish the power of and concern generated by such intrusions when successful, but does usefully highlight the difficulties adversaries face in successful execution.

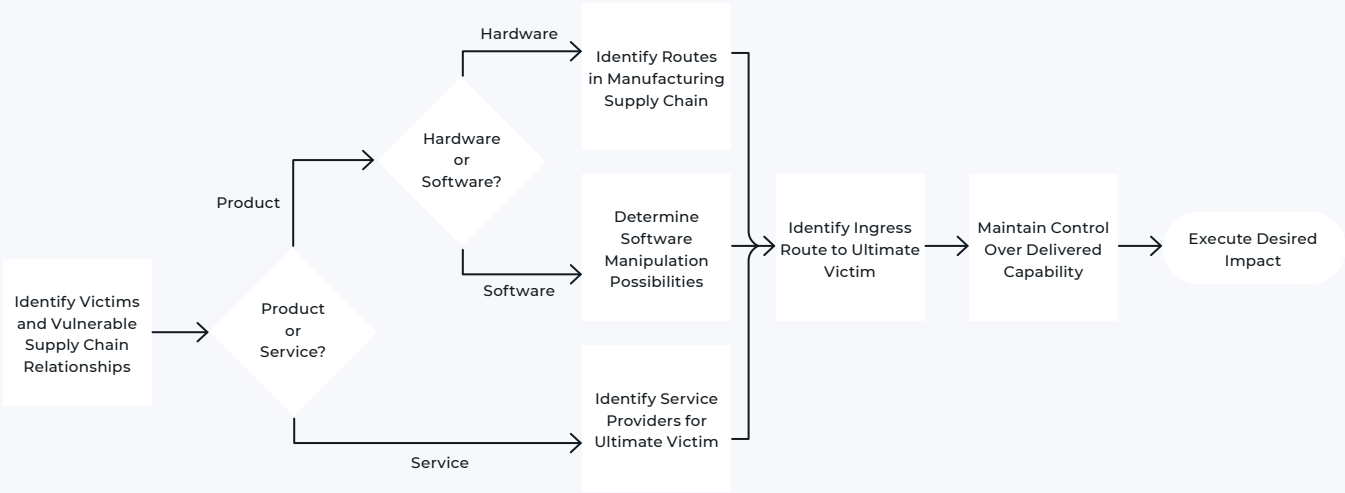


Figure 3: Mapping Supply Chain Intrusion Dependencies.

Before analyzing adversary “pain points” in supply chain intrusions, a review of the complex nature of supply chain intrusion execution is useful for purposes of background.

Intrusion Vector

First, an adversary must identify and weaponize an intrusion vector – how to get into position for follow-on delivery to or infection of ultimate victims. This phase of operation represents the first of two intrusions, where the adversary compromises the vector organization as a route to final intended victims. An adversary faces various choices at this stage in determining whether to pursue a product- or services-focused route, but selection should be influenced by intended victims. The previous point is non-trivial, as supply chain intrusions are typically not opportunistic (obtain a vector, and then see what follow-on targets are possible), but rather deliberate (considering desired end goals, determine appropriate routes to achieve them).

The adversary must perform significant target development to ensure the right vectors exist for the intended final targets. In some cases, such as the SUNBURST instance or the Cloud Hopper campaigns, intermediate victims (the supply chain vector) may be so ubiquitous or widely adopted that the majority of desired, intended follow-on victims will be susceptible to the chosen vector. In other cases, greater discretion and research are required, such as the software update compromise activity executed in the original Dragonfly campaign.³¹ In this example, three relatively small software suppliers focusing on European industrial markets were compromised – a successful intrusion, but one with a far smaller footprint than compromising a widely used product such as SolarWinds Orion.

This last observation is significant and leads to some assumptions and potential difficulties. While not always true, an adversary can likely assume that larger, critical providers of software and services (such as Microsoft, Siemens, or similar) recognize their importance (and the potential embarrassment of being the vector for follow-on compromise). As such these environments should be robustly defended – yet their widespread use and application nonetheless make them tantalizing targets as a supply chain vector. Conversely, smaller, “boutique” suppliers or service providers such as the vendors in the Dragonfly campaign or contractors in the Palmetto Fusion operation likely have significantly weaker security

profiles than larger, more ubiquitous firms – but with the drawback that their customer base is significantly smaller, thus offering reduced scope for weaponization.

The would-be supply chain compromiser thus faces a choice: attempt to breach a high-profile, large customer base entity with potentially greater security or ability to identify and defeat a compromise; or target far smaller entities with weaker defenses but more limited scope for follow-on action. Depending on the goal (e.g., access to very specific entities), a narrow approach may be ideal and appears reflected in activities such as the Palmetto Fusion intrusions. Either way, adversaries must weigh the costs and benefits of these actions, and determine what vectors are relevant and applicable to their ultimate victims, before even initiating operations.

Delivering Capabilities to Victims

Once an adversary identifies a viable “chain,” they now must determine how to leverage that vector to deliver capabilities to victims. This seems an obvious point, but capability delivery is also non-trivial depending on the route chosen and complexity of the vector. For example, an attacker may identify a product susceptible to abuse, but modifying the product in question to include malicious functionality or a backdoor requires inserting the desired functionality while not breaking or noticeably altering the underlying product. In some cases, adversaries may circumvent this requirement by targeting how items are packaged rather than their core functionality – a methodology observed in the original Dragonfly campaign, where software updates were repackaged to include malicious functionality as opposed to directly modifying (or even gaining access to) source code to introduce attacker capabilities.

Significant risk also resides in where the capability is delivered and how to limit exposure or spread. For example, looking at the ShadowHammer incident, where an entity compromised software updates for electronics manufacturer ASUS to deliver a backdoor, the capability was spread quite widely (potentially a half-million victims), but through MAC address filtering, only a few hundred

were actively, meaningfully targeted.^{32,33} The SUNBURST campaign using the SolarWinds Orion intrusion vector also showed victim “filtering” through initial C2 checks, meaning initial spread that was subsequently narrowed through limiting logic.³⁴ In both cases, while capabilities were widely distributed, adversary “controls” prevented “overspread” that could lead to detection.

When capabilities are distributed less carefully, as seen in the NotPetya event, where a compromise of MeDoc accounting software in Ukraine rapidly expanded into a global destructive event,³⁵ impacts and artifacts may spread beyond adversary control resulting in earlier capability detection or unintended consequences. Adversaries thus face a decision of seeking very focused, narrow targeting that may miss potential, desired victims, or widespread campaigns that are relatively indiscriminate in operation with the intended risk of discovery or collateral damage, explored further below.

Capability Execution and Control

Finally, once a capability is in place, adversaries must consider how to execute or control that capability to deliver the desired effect. In the case of backdoors or other access mechanisms designed to facilitate further intrusions, the adversary must possess some means of communicating with the implant, or the implant must be able to beacon to the adversary. Either case results in network traffic and related activity that can be identified by defenders or operators. Looking again at the SUNBURST example, capability deployment required several stages of follow-on C2 activity before providing the NOBELIUM actor with sustained access to victim environments.^{36,37} While the mechanism used to deploy the capabilities referenced may be very difficult-to-detect, subsequent C2 activity remains exposed to network security monitoring and analysis with the possibility of disclosing the breach.

If an adversary requires an extreme degree of stealth and direct, controlled access is not an issue, a sufficiently well-resourced threat actor can develop and deploy autonomous capabilities. In this case, adversaries establish and encode logic, purpose, and capability into a tool pre-deployment so that once within the desired victim environment, the tool can perform its task without further instruction. While there are many commentaries about “AI-powered malware” or similar speculative thinking, actual examples of such capabilities in the wild are vanishingly small. Such items include simplistic wormable malware types such as still-persistent Conficker infections or Sality.^{38,39} For sophisticated, truly targeted capabilities, no real examples exist outside of Stuxnet’s autonomous attack logic,⁴⁰ a capability now over ten years old but with no real public rivals.

Absent highly complex (and very expensive) autonomous operational logic as seen in Stuxnet, the possibilities for intruders are reduced in scope to either being fairly indiscriminate operations (seen in wormable ransomware operations such as WannaCry and NotPetya) or serving as tools for follow-on interactive access with inherent risks of detection. Resources, desire (or need) for stealth, and latent capabilities can all enforce limitations on precisely how an adversary can satisfy the concerns related to capability execution and control following a successful supply chain intrusion.

Difficulties and Challenges

Based on the above, supply chain attacks are exposed as something less than “silver bullets” allowing for rapid and stealthy compromise of numerous victims, and instead revealed as the result of a lengthy, potentially arduous process of multiple intrusions and anguished decision-making. The last point is most important to emphasize, because in seeking a supply chain compromise and its attendant benefits in circumventing security controls and potentially visibility, adversaries are faced with several constraining choices.

Illustrated in Figure 4 below, adversaries must balance targeting specificity with capability control. While these continuums of adversary operational choices are not explicit tradeoffs, as seen by capabilities such as Stuxnet that can combine high degrees of autonomous capability with extremely limited distribution, adversaries must nonetheless determine what balance of possibilities are appropriate for their desired operation - or possess the resources, patience, and capabilities required to defeat any potential trade-offs.

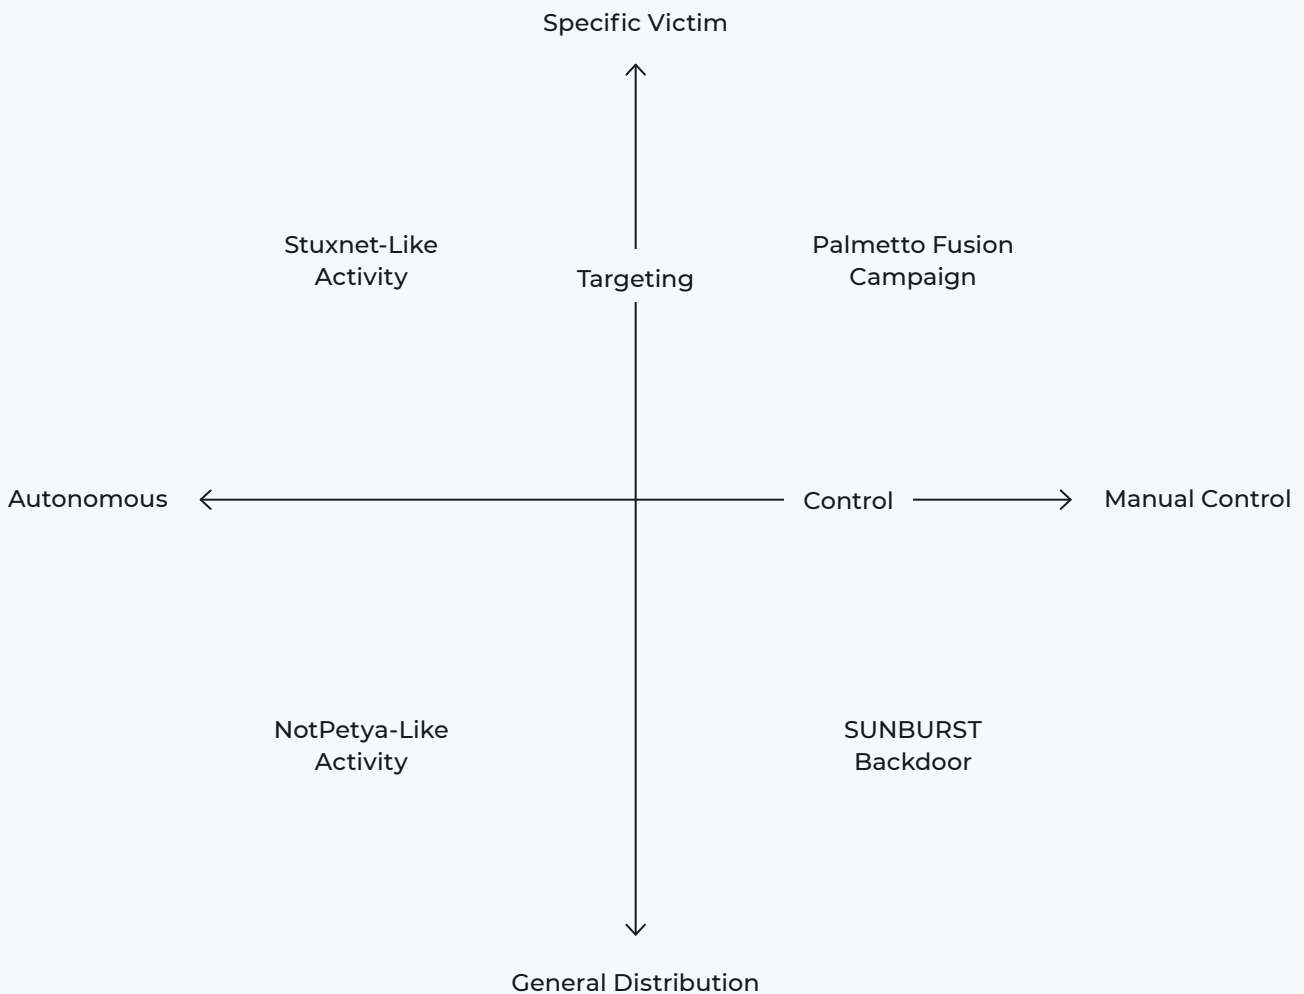


Figure 4: Supply Chain Intrusion Adversary Choices

Issues Concerning Targeting and Specificity

Targeting concerns represent a fundamental aspect of supply chain intrusions, and adversaries must work to ensure subsequent capabilities meet their desired end state – or are willing to accept “collateral damage” when such intrusions spin out of control. To illustrate the risks of overspread, NotPetya provides an especially useful example. The general narrative surrounding the incident focuses on its widespread, global disruptive impact.⁴¹ Yet as discussed previously, the event itself started from far narrower beginnings: as a supply chain intrusion focused on accounting software very specific to Ukraine.^{35,42}

More detailed overviews of the NotPetya incident by Andy Greenberg and Thomas Rid show the event beginning as a primarily Ukrainian affair.^{43,44} Executed on 27 June 2017, the day before Ukrainian Constitution Day, NotPetya obliterated numerous aspects of Ukrainian digital life due to the ubiquity of MeDoc’s software in commercial organizations. As such, this intrusion appeared to map neatly onto a series of other disruptive cyber incidents impacting Ukrainian institutions: the 2015 and 2016 electric power events, various rounds of disruptive phishing activities, and attempted interference in Ukrainian elections. But whereas these incidents remained localized to Ukraine, NotPetya quickly ripped across the global internet, impacting entities far beyond the likely area of focus for the Russian authorities tasked with its deployment.^{45,46,47}

NotPetya therefore appears to have spread dramatically beyond its likely intended area of focus, including significant impacts in Western Europe and North America. NotPetya thus represents an especially risky maneuver by its sponsors given the costly, destructive effects in the United States and elsewhere – entities that represent a far greater threat for retaliation and response than Ukraine. Ultimately, NotPetya’s attribution to Russia by U.S., U.K., and other authorities led to sanctions⁴⁸ – seemingly minor, but still costly and contributing to the increasing economic isolation of Russia.

Based on these impacts, it appears that NotPetya’s controllers either did not anticipate the original

MeDoc-related capability would result in meaningful extension beyond Ukraine or theorized that any potential costs would be acceptable. While we cannot determine the specific calculus behind Russian national command authorities in NotPetya’s deployment, we can arrive at a clear conclusion with respect to such events: they result in at minimum risk of retaliation and at worst realized cost (even if “merely” economic in nature). This extended example thus shows the dangers of a capability – in this case, one able to act autonomously – deployed widely or indiscriminately.

Replacing the MeDoc vector in NotPetya, an obscure company limited to operations in Ukraine but with unexpectedly significant links, with a more widely deployed software or service thus leads to sobering thoughts. Looking at items such as previously discussed SUNBURST and ShadowHammer events, possibilities for widespread distribution in supply chain events are numerous – yet we also see a significant degree of restraint exercised by perpetrators. In just these examples, among others, adversaries introduced “checks” or other controls to prevent “overspread.” The rationale for such action could be operational security to avoid detection, or may align with a desire to avoid the NotPetya situation where a deployed intrusion (and especially follow-on capability) mutates out of control.

At the same time, introducing such barriers limits propagation, which may subsequently inhibit the action and ultimate success of a given intrusion. Placing checks and controls for supply chain-focused spread could achieve greater stealth (or at least reduce the likelihood of discovery given reduced exposure), but at the cost of a more limited number of follow-on victims. Adversaries therefore face a choice aligning with a continuum that, on one end, shows indiscriminate spread (such as NotPetya) while on the other, propagation is more circumscribed and limited.

Concerns on Attack Command and Control

Another consideration for adversaries is that simply deploying a capability or similar may be inherently meaningless unless that actor can

leverage that capability to do something. In this case, a supply chain intrusion is a means to an end, rather than an end in itself, in that it facilitates follow-on access that happens to circumvent various controls and potential monitoring points. While capability delivery may take place via a difficult-to-detect mechanism, subsequent control of that capability requires the same types of communication links and activity seen in less exciting malware samples.

As discussed previously, fully autonomous capabilities are relatively rare, representing something simplistic like a Sality, an item unconstrained such as NotPetya, or a truly well-crafted and designed software such as Stuxnet. Irrespective of relative sophistication, such items are related by a common capability set allowing them to perform their task - whether desired or otherwise - absent any sort of positive operator control. As such, these capabilities can be very powerful and effective in worming their way through networks without waiting for an operator's command or instruction. But such effectiveness comes at the cost of control. Looking at the NotPetya example, as well as the later, more virulent examples of Stuxnet which resulted in the malware being caught, "overspread" and acting well outside intended boundaries are significant risks, bringing on the possibility of unintended consequences.

Conversely, more direct control over a deployed capability requires communication. In this case, the adversary may succeed in evading many layers of security controls and defender visibility, but with the subsequent cost that utilizing the deployed capability will require interaction. Such interaction will require traffic and communication flowing through and then outside the victim network, thus presenting opportunities for detection. Thus, a desire to keep capabilities firmly in line with desired functionality and intention raises the possibility of detection as a result of such interactive control and management.

Adversaries again face a choice: attempt to create an autonomous or independent capability that, once deployed, can no longer be effectively controlled, or ensure continued control and interaction with a deployed capability through active communication. Both contain risks and benefits that an adversary must consider and ultimately balance in deploying a capability. Autonomous capabilities can produce unintended consequences or, if not well-designed, outright fail in their objective, while supply chain intrusions requiring command and control (C2) eliminate significant degrees of stealth and demand a mechanism to enable communication.

Defender Advantages

The previous section highlights several items that adversaries must contend with to successfully execute a supply chain intrusion. Aside from being matters of inconvenience for would-be adversaries, such items also present defenders with multiple touchpoints for identifying and even mitigating adversary actions. While the possibilities for supply chain intrusion defense are numerous, a few particularly valuable items will be considered in the following sections.

Exploiting Communication and Control

As described previously, deployed supply chain capabilities typically require a degree of adversary control to be effective. While autonomous options exist, these are either very narrowly tailored (and developed at significant expense) or open to loss of adversary control, leading to unintended impacts and consequences. For entities unable or unwilling to invest in a Stuxnet-like capability who also wish to avoid events spiraling out of control like a NotPetya, active command and control is not just desirable but necessary.

From this adversary requirement, defenders now have an opportunity. If an implant or modification requires adversary interaction, defenders can see or even potentially prevent such communication from taking place. Since truly “air-gapped” networks are exceedingly rare,⁴⁹ network security monitoring (NSM) and traffic inspection become relevant items for detecting C2 behaviors. Outside of simple indicator monitoring, NSM and intrusion detection technologies enable defenders to spot unusual traffic patterns, destinations, or communicating pairs. Leveraging this knowledge, a beaconing implant added via supply chain compromise can reveal itself, while active C2 over implants may also be flagged.

For example, the SUNBURST campaign appears at first glance to be highly complex with degrees of control in deployment to minimize potential exposure through overspread.⁵⁰ Yet following these checks, the modified SolarWinds Orion application performs a two-stage command and control

sequence, moving from DNS to follow-on HTTP C2 behaviors. NSM tradecraft becomes operational at this stage, with several opportunities to identify activity for further investigation:

1. The initial, lengthy DNS query with encoded information for the first C2 domain
2. Follow-on CNAME response to the query with the second-stage C2 infrastructure information
3. Subsequent HTTP communication to the second-stage C2 infrastructure

While not easy, a well-architected, well-understood network enables a defender to differentiate suspicious from normal traffic, especially if such communications can be traced back to the initiating device. Having a critical service, such as a SolarWinds Orion server, initiating communication with one or several unknown, outside hosts should be an indicator that something is amiss. Such defensive measures can be taken even further by profiling known necessary external communications for vital or valuable services and their hosting machines.⁵¹

Anytime a tool, implant, or other capability deployed via a supply chain intrusion attempts to communicate, defenders must adopt the mindset that such operations represent an opportunity. While adversaries may develop and deploy a variety of mechanisms to obfuscate or otherwise hide their activity to evade detection, given the sensitivity of some systems (such as high-value IT assets or industrial systems), any unusual communication should face scrutiny by defenders or system operators. By focusing on the likely (if not certain) dependency that adversaries will need to C2 over a deployed capability, defenders can identify a potential intrusion for further analysis.

Architecture, Visibility, and Monitoring

Building off the previous section on network communication dependencies for adversaries, defenders and network owners retain significant “first mover” advantage in that they own, manage, and (ideally) can design the landscape on which intruders must operate. While seemingly obvious, such control and ownership represent an

incredible opportunity that can amplify defender advantages with respect to adversary dependencies such as command and control.

One mechanism for applying the above emerges through implementing Zero Trust security architecture.⁵² While commonly abused as a marketing buzzword, an understanding of prerequisites to achieve Zero Trust network posture can enable robust defense for many circumstances, including supply chain intrusions. The core principles of Zero Trust security architecture, shown in Figure 5, are persistent verification, enforcement of least-privileged access, and a posture of “assume breach” within the network. Applied and achieved, network owners implementing a Zero Trust approach are well placed to at minimum severely disadvantage an adversary pursuing a supply chain intrusion.

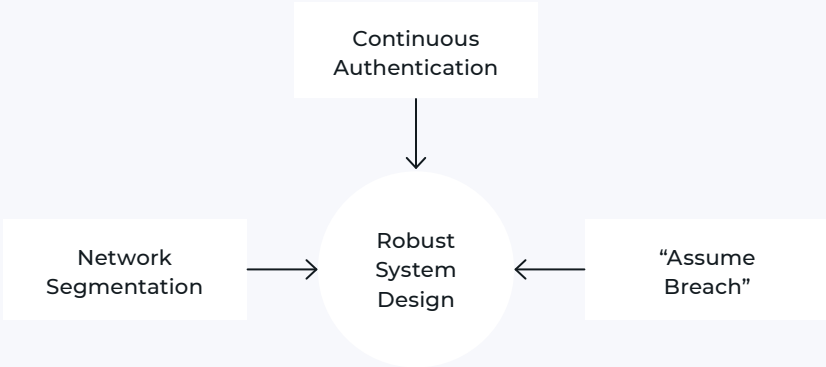


Figure 5: Zero Trust Core Principles.

One of the core mechanisms to achieve and maintain Zero Trust principles is rigorous network segmentation through physical and virtual mechanisms. System owners can reduce direct connectivity between devices and establish authentication or rigorous trust boundaries between segments. Adversary lateral movement then becomes significantly more difficult even if the initial breach takes place via a supply chain mechanism circumventing other controls. Thorough segmentation becomes especially valuable when paired with monitoring and visibility. System owners and network defenders gain insight into internal network traffic flows between discrete zones as opposed to just internal-external communications. Combined with a robust approach to C2 traffic monitoring described in the previous section, defenders gain layered visibility into adversary operations throughout multiple phases of operations.

In the case of a supply chain intrusion, an adversary may be able to deliver a capability deep within the network by compromising a vendor – but once in the intended victim’s network, opportunities for subsequent “breakout” or taking active control of that capability can be severely limited or quickly discovered. In the case of equipment, initial communications may be limited through security and segmentation controls to the network segment on which the given system operates. While difficult to implement in many enterprise IT environments, such an action is not impossible and may be quite relevant and accessible for more sensitive networks such as industrial control system environments.

For vendors and contractors requiring remote access, architecture and policy can limit this access to a dedicated contractor pathway. Applicable design principles include “jump boxes” for enforcing monitoring and limiting communication beyond required assets, dedicated network and access pathways, and temporary credentials for internal systems. If successfully deployed, an adversary that compromises a contractor or vendor gains very little in follow-on capability and access when trying to extend a breach into ultimate victim environments.

Resilience, Response, and Defense in Depth

While this paper has consistently argued that supply chain intrusions are difficult, this argument cannot and should not be distorted to imply they are impossible. Defenders should therefore anticipate supply chain intrusions as a very real, if likely rare, risk to their networks that can, under the right circumstances, evade many defensive controls. In these situations, network defense ceases to be a preventative discipline and transforms into an exercise in limiting adversary impact while minimizing impacts to the organization.

When other controls fail, the first necessary step within the overall incident response process is root cause analysis. Defenders must be in a position to ask and answer questions concerning the origin and nature of an incident to adequately understand its scope. In the case of supply chain intrusions, identifying odd behavior in otherwise isolated or sensitive areas of the network should stand out as immediately interesting and spark the immediate question of how an adversary reached the asset. When standard intrusion pathways appear absent or simply do not align with available evidence, then defenders can proceed to investigating potential supply chain or other non-standard ingress routes to explain available evidence and observations.

As part of root cause analysis, defenders must also have or know where to source the capabilities necessary to answer questions that emerge during the investigation. In the case of supply chain intrusions, analysis and forensics may be difficult given specific systems and software involved. For example, diagnosing the backdoor present in the SolarWinds Orion software (injected at compilation time) is an incredibly difficult item to identify for a typical security team. In such cases, after an investigation reveals the Orion server as the root cause of intrusion by mapping communication and access links within the network, defenders must be able to identify parties for follow-on analysis or to answer questions. Relevant parties include the actual vendor and vendor resources, or third-party entities with specialized forensic and investigative capabilities. Alternatively, organizations must invest in their security people to ensure they

have the skills and tools necessary to answer questions related to compromised or modified software or hardware – a non-trivial ask, but a needed capability given adversary interest in such operations.

Finally, network owners must ensure systems are resilient in the face of cyber intrusions generally, including potential supply chain vectors. When intrusions are successful and adversaries can achieve desired impacts, system owners can still mitigate worst-case scenarios through sound design and planning. Examples of resilience include:

- + Maintain a repository of known-good configurations for critical assets and appliances, including the ability to perform change analysis to identify modifications if necessary, to allow for quick restoration or reimaging
- + Identify critical system dependencies in advance and formulate recovery and restoration plans to rapidly return to known-good operational status
- + In industrial and cyber-physical settings, ensure process and engineering controls are in place and tested to mitigate potential system modifications made in associated information systems

Ultimately the goal for organizations is to maintain known-good, known-safe operations and to continue whatever the organization's primary function may be. Information systems are just one part of this overall operational picture, but at times a critical component. In the case of supply chain attacks, which if successful can be difficult to first diagnose and then repair or remediate, organizations need to plan for scenarios where fundamental aspects of the network suddenly become untrusted. Leveraging not just cyber incident response plans but also incorporating these scenarios into business continuity and disaster recovery planning can ensure continuation of organizational mission even in the face of very complex intrusion scenarios.

Conclusion

Supply chain intrusions are very real and an object of concern for network owners, operators, and defenders. However, stakeholders must recognize the nature and implicit difficulties within supply chain intrusions, both to properly assess the risk of such events and to understand the numerous defensive options available to defeat or mitigate these incidents.

When media, researchers, analysts, and others hype supply chain intrusions as extremely stealthy, near impossible-to-detect events, the defensive community loses sight of just how hard a successful supply chain intrusion is to execute. By understanding the prerequisites to success, we can not only demystify such events but realize that defenders, in many respects, hold the overall advantage in these engagements.

Embracing defender advantages, we can begin implementing security controls to make intruder lives even more difficult. While supply chain intrusion vectors are too valuable for adversaries to ever give up on such operations completely, defenders can work to make such items immensely costly or highly unlikely to succeed. In this way, we can eliminate the assumption of supply chain attacks as somehow inherently impossible to defend against, and move on to the next pressing security concern.

Works Cited

- ¹ J. Robertson and M. Riley, "The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies," Bloomberg, 4 October 2018. [Online]. Available: <https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies>. [Accessed 21 September 2021].
- ² J. Weiss, "Emergency Executive Order 13920 - Response to a Real Nation-State Cyberattack Against the US Grid," ControlGlobal, 4 May 2020. [Online]. Available: <https://www.controlglobal.com/blogs/unfettered/emergency-executive-order-13920-response-to-a-real-nation-state-cyberattack-against-the-us-grid/>. [Accessed 21 September 2021].
- ³ U.S. Department of Homeland Security, "National Strategy for Global Supply Chain Security," 23 January 2012. [Online]. Available: <https://www.dhs.gov/national-strategy-global-supply-chain-security>. [Accessed 21 September 2021].
- ⁴ The National Counterintelligence and Security Center, "Supply Chain Risk Management," Office of the Director of National Intelligence, [Online]. Available: <https://www.dni.gov/index.php/ncsc-what-we-do/ncsc-supply-chain-threats>. [Accessed 21 September 2021].
- ⁵ European Union, "Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA," 17 April 2019. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2019/881/oj>. [Accessed 21 September 2021].
- ⁶ Security Magazine, "UL Announces Supply Chain Cybersecurity Solution," 20 May 2020. [Online]. Available: <https://www.securitymagazine.com/articles/92423-ul-announces-supply-chain-cybersecurity-solution>. [Accessed 22 September 2021].
- ⁷ UL, "Supplier Cyber Trust Level," [Online]. Available: <https://ims.ul.com/supplier-cyber-trust-level>. [Accessed 22 September 2021].
- ⁸ U.S. Office of the President, "Executive Order on Securing the United States Bulk-Power System," 1 May 2020. [Online]. Available: <https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-securing-united-states-bulk-power-system/>. [Accessed 22 September 2021].
- ⁹ W. Kenton, "Supply Chain," Investopedia, 29 August 2021. [Online]. Available: <https://www.investopedia.com/terms/s/supplychain.asp>. [Accessed 22 September 2021].
- ¹⁰ A. Russell, "CIA plot led to huge blast in Siberian gas pipeline," The Telegraph, 28 February 2004. [Online]. Available: <https://www.telegraph.co.uk/news/worldnews/northamerica/usa/1455559/CIA-plot-led-to-huge-blast-in-Siberian-gas-pipeline.html>. [Accessed 22 September 2021].
- ¹¹ Wired Staff, "Soviets Burned by CIA Hackers?," Wired, 26 March 2004. [Online]. Available: <https://www.wired.com/2004/03/soviets-burned-by-cia-hackers/>. [Accessed 22 September 2021].
- ¹² T. Rid, "Think Again: Cyberwar," Foreign Policy, 27 February 2012. [Online]. Available: <https://foreignpolicy.com/2012/02/27/think-again-cyberwar/>. [Accessed 22 September 2021].
- ¹³ R. Jennings, "Best of 2020: Was This Huawei's Failed Attempt at a Linux Backdoor?," Security Boulevard, 30 December 2020. [Online]. Available: <https://securityboulevard.com/2020/12/was-this-huaweis-failed-attempt-at-a-linux-backdoor/>. [Accessed 22 September 2021].
- ¹⁴ C. Cimpanu, "Huawei denies involvement in buggy Linux kernel patch proposal," ZDNet, 12 May 2020. [Online]. Available: <https://www.zdnet.com/article/huawei-denies-involvement-in-buggy-linux-kernel-patch-proposal/>. [Accessed 22 September 2021].
- ¹⁵ B. Krebs, "Target Hackers Broke in Via HVAC Company," KrebsOnSecurity, 5 February 2014. [Online]. Available: <https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company>. [Accessed 22 September 2021].
- ¹⁶ M. Kassner, "Anatomy of the Target data breach: Missed opportunities and lessons learned," ZDNet, 2 February 2015. [Online]. Available: <https://www.zdnet.com/article/anatomy-of-the-target-data-breach-missed-opportunities-and-lessons-learned/>. [Accessed 22 September 2021].
- ¹⁷ Reuters Staff, "Target in \$18.5 mln multi-state settlement over data breach," Reuters, 23 May 2017. [Online]. Available: <https://www.reuters.com/article/target-cyber-settlement-idUSL4N1IP4SR>. [Accessed 22 September 2021].
- ¹⁸ PWC, "Operation Cloud Hopper," April 2017. [Online]. Available: <https://www.pwc.co.uk/cyber-security/pdf/cloud-hopper-report-final-v4.pdf>. [Accessed 22 September 2021].
- ¹⁹ Stilgherrian, "At least nine global MSPs hit in APT10 attacks: ACSC," ZDNet, 20 December 2018. [Online]. Available: <https://www.zdnet.com/article/at-least-nine-global-msps-hit-in-apt10-attacks-acsc/>. [Accessed 22 September 2021].
- ²⁰ B. Barrett, "How China's Elite Hackers Stole the World's Most Valuable Secrets," Wired, 20 December 2018. [Online]. Available: <https://www.wired.com/story/doj-indictment-chinese-hackers-apt10/>. [Accessed 21 September 2021].
- ²¹ Y. Klijnsma and S. Ginty, "Gift Cardsharks: The Massive Threat Campaigns Circling Beneath the Surface," RiskIQ, June 2019. [Online]. Available: <https://cdn.riskiq.com/wp-content/uploads/2019/06/Gift-Cardsharks-Intelligence-Report-2019-RiskIQ.pdf>. [Accessed 21 September 2021].
- ²² J. Reaves, J. Platt and A. Nixon, "Wipro Threat Actors Active Since 2015," Flashpoint, 1 May 2019. [Online]. Available: <https://www.flashpoint-intel.com/blog/wipro-threat-actors-active-since-2015/>. [Accessed 21 September 2021].

- ²³ US Cybersecurity & Infrastructure Security Agency, "Alert (TA17-293A) Advanced Persistent Threat Activity Targeting Energy and Other Critical Infrastructure Sectors," US Department of Homeland Security, 15 March 2018. [Online]. Available: <https://us-cert.cisa.gov/ncas/alerts/TA17-293A>. [Accessed 21 September 2021].
- ²⁴ E. Nakashima, "U.S. officials say Russian government hackers have penetrated energy and nuclear company business networks," The Washington Post, 8 July 2017. [Online]. Available: https://www.washingtonpost.com/world/national-security/us-officials-say-russian-government-hackers-have-penetrated-energy-and-nuclear-company-business-networks/2017/07/08/bbfde9a2-638b-11e7-8adc-fea80e32bf47_story.html. [Accessed 22 September 2021].
- ²⁵ R. Smith and R. Barry, "America's Electric Grid Has a Vulnerable Back Door - and Russia Walked Through It," The Wall Street Journal, 10 January 2019. [Online]. Available: <https://www.wsj.com/articles/americas-electric-grid-has-a-vulnerable-back-door-and-russia-walked-through-it-11547137112>. [Accessed 22 September 2021].
- ²⁶ US Cybersecurity & Infrastructure Security Agency, "Alert (AA20-352A) Advanced Persistent Threat Compromise of Government Agencies, Critical Infrastructure, and Private Sector Organizations," US Department of Homeland Security, 15 April 2021. [Online]. Available: <https://us-cert.cisa.gov/ncas/alerts/aa20-352a>. [Accessed 22 September 2021].
- ²⁷ FireEye, "Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims With SUNBURST Backdoor," FireEye, 13 December 2020. [Online]. Available: <https://www.fireeye.com/blog/threat-research/2020/12/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor.html>. [Accessed 22 September 2021].
- ²⁸ M. Sentonas, "CrowdStrike Launches Free Tool to Identify and Help Mitigate Risks in Azure Active Directory," CrowdStrike, 23 December 2020. [Online]. Available: <https://www.crowdstrike.com/blog/crowdstrike-launches-free-tool-to-identify-and-help-mitigate-risks-in-azure-active-directory/>. [Accessed 22 September 2021].
- ²⁹ S. Vavra, "Microsoft alerts CrowdStrike of hackers' attempted break-in," CyberScoop, 24 December 2020. [Online]. Available: <https://www.cyberscoop.com/crowdstrike-solarwinds-targeted-microsoft/>. [Accessed 22 September 2021].
- ³⁰ E. M. Hutchins, M. J. Cloppert and R. M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," [Online]. Available: <https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf>. [Accessed 22 September 2021].
- ³¹ J. Slowik, "The baffling Berserk Bear: a decade's activity targeting critical infrastructure," 7 October 2021. [Online]. Available: <https://vblocalhost.com/presentations/the-baffling-bersek-bear-a-decades-activity-targeting-critical-infrastructure/>. [Accessed 8 October 2021].
- ³² GREAT, "Operation ShadowHammer," Kaspersky, 25 March 2019. [Online]. Available: <https://securelist.com/operation-shadowhammer/89992/>. [Accessed 22 September 2021].
- ³³ K. Zetter, "Hackers Hijacked ASUS Software Updates to Install Backdoors on Thousands of Computers," Motherboard, 25 March 2019. [Online]. Available: <https://www.vice.com/en/article/pan9wn/hackers-hijacked-asus-software-updates-to-install-backdoors-on-thousands-of-computers>. [Accessed 22 September 2021].
- ³⁴ E. Hjeltnvik, "Finding Targeted SUNBURST Victims with pDNS," Netresec, 4 January 2021. [Online]. Available: <https://www.netresec.com/?page=Blog&month=2021-01&post=Finding-Targeted-SUNBURST-Victims-with-pDNS>. [Accessed 22 September 2021].
- ³⁵ A. Cherepanov, "Analysis of TeleBots' Cunning Backdoor," ESET, 4 July 2017. [Online]. Available: <https://www.welivesecurity.com/2017/07/04/analysis-of-telebots-cunning-backdoor/>. [Accessed 21 September 2021].
- ³⁶ R. Nafisi, A. Lelli, MSTIC and Microsoft 365 Defender Threat Intelligence Team, "GoldMax, GoldFinder, and Sibot: Analyzing NOBELIUM's layered persistence," Microsoft, 4 March 2021. [Online]. Available: <https://www.microsoft.com/security/blog/2021/03/04/goldmax-goldfinder-sibot-analyzing-nobelium-malware/>. [Accessed 22 September 2021].
- ³⁷ L. Smith, J. Leathery and B. Read, "New SUNSHUTTLE Second-Stage Backdoor Uncovered Targeting U.S.-Based Entity; Possible Connection to UNC245," FireEye, 4 March 2021. [Online]. Available: <https://www.fireeye.com/blog/threat-research/2021/03/sunshuttle-second-stage-backdoor-targeting-us-based-entity.html>. [Accessed 22 September 2021].
- ³⁸ P. Howell O'Neill, "Conficker worm still spreading despite being nearly 10 years old," CyberScoop, 8 December 2017. [Online]. Available: <https://www.cyberscoop.com/conficker-trend-micro-2017/>. [Accessed 22 September 2021].
- ³⁹ N. Falliere, "Sality: Story of a Peer-to-Peer Viral Network," July 2011. [Online]. Available: <https://vx-underground.org/archive/Symantec/sality-story-of-peer-to-peer-1-en.pdf>. [Accessed 22 September 2021].
- ⁴⁰ N. Falliere, L. O Murchu and E. Chien, "W32.Stuxnet Dossier," November 2010. [Online]. Available: https://www.wired.com/images_blogs/threatlevel/2010/11/w32_stuxnet_dossier.pdf. [Accessed 22 September 2021].
- ⁴¹ A. Greenberg, "The Untold Story of NotPetya, the Most Devastating Cyberattack in History," Wired, 22 August 2018. [Online]. Available: <https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>. [Accessed 22 September 2021].
- ⁴² D. Maynor, A. Nikolic, M. Olney and Y. Younan, "The MeDoc Connection," Cisco Talos, 5 July 2017. [Online]. Available: <https://blog.talosintelligence.com/2017/07/the-medoc-connection.html>. [Accessed 22 September 2021].
- ⁴³ A. Greenberg, Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers, New York: Penguin Random House LLC, 2019.
- ⁴⁴ T. Rid, Active Measures: The Secret History of Disinformation and Political Warfare, New York: Farrar, Straus, and Giroux, 2020.

- ⁴⁵ US Office of the President, "Statement from the Press Secretary," 15 February 2018. [Online]. Available: <https://trumpwhitehouse.archives.gov/briefings-statements/statement-press-secretary-25/>. [Accessed 22 September 2021].
- ⁴⁶ UK National Cyber Security Centre, "Russian military 'almost certainly' responsible for destructive 2017 cyber attack," 14 February 2018. [Online]. Available: <https://www.ncsc.gov.uk/news/russian-military-almost-certainly-responsible-destructive-2017-cyber-attack>. [Accessed 22 September 2021].
- ⁴⁷ US Department of Justice, "Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace," 19 October 2020. [Online]. Available: <https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and>. [Accessed 22 September 2021].
- ⁴⁸ US Department of the Treasury, "Treasury Sanctions Russia with Sweeping New Sanctions Authority," 15 April 2021. [Online]. Available: <https://home.treasury.gov/news/press-releases/jy0127>. [Accessed 22 September 2021].
- ⁴⁹ J. Slowik, "Mind the (Air) Gap," Stranded on Pylos, 13 May 2021. [Online]. Available: <https://pylos.co/2021/05/13/mind-the-air-gap/>. [Accessed 22 September 2021].
- ⁵⁰ S. Eckels, J. Smith and W. Ballenthin, "SUNBURST Additional Technical Details," FireEye, 24 December 2020. [Online]. Available: <https://www.fireeye.com/blog/threat-research/2020/12/sunburst-additional-technical-details.html>. [Accessed 22 September 2021].
- ⁵¹ C. Anderson, "Protecting Against Supply Chain Attacks by Profiling Suppliers," DomainTools, 12 January 2021. [Online]. Available: <https://www.domaintools.com/resources/blog/protecting-against-supply-chain-attacks-by-profiling-suppliers>. [Accessed 22 September 2021].
- ⁵² S. Rose, O. Borchert, S. Mitchell and S. Connelly, "NIST Special Publication 800-207 Zero Trust Architecture," National Institute of Standards and Technology, August 2020. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>. [Accessed 23 September 2021].