Hold the Door: Examining Exfiltration Activity and Applying Countermeasures

blog.gigamon.com/2021/06/17/hold-the-door-examining-exfiltration-activity-and-applying-countermeasures

June 17, 2021

Security / June 17, 2021

Joe Slowik

Updated October 14, 2021.

Background

Data exfiltration exists as a cornerstone of malicious cyber operations. Yet the nature of exfiltration and its impact changed significantly over the past few years.

Historically, adversaries linked to state sponsors or similar activities would engage in data theft to further goals of industrial espionage or intellectual property theft. More recently, adversaries have engaged in exfiltration, not based on the value of the data to themselves or

their sponsors, but rather on the desire of victims to not have such data exposed. The result is an ecosystem of exfiltration and disclosure for either monetary or “hacktivist” purposes, where the malicious party seeks to damage (or threaten damage) to the victim. Relevant, recent examples include ransomware-related data loss at meat processor JBS, which prompted to company to pay $11 million to avoid disclosure of data, and an intrusion at gaming company Electronic Arts resulting in the loss and potential sale of nearly 800 gigabytes of sensitive information. Given these trends, network owners, operators, and defenders must work to ensure controls are in place to detect, mitigate, or outright block such activity given potential costs.

Exfiltration Impacts

On its face, exfiltration appears less concerning than other cyber-nexus impacts such as network disruption. While data exfiltration typically has some immediate impact on an organization, long-term hazards may make loss of sensitive data even more destructive than an especially virulent ransomware incident.

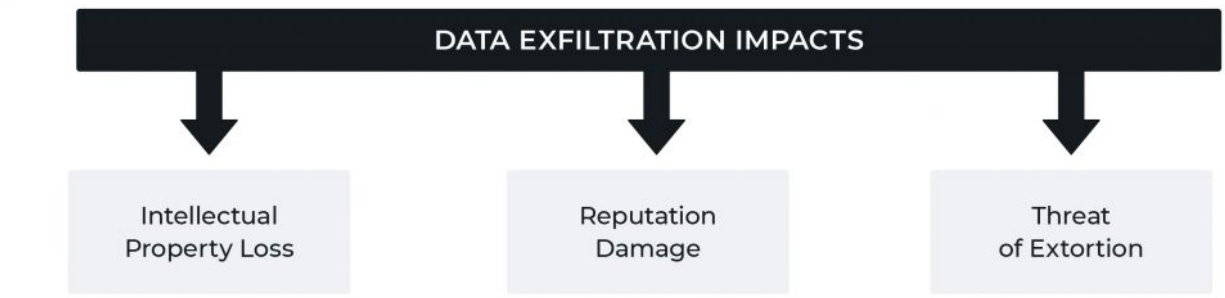

Figure 1: The three primary impacts of data exfiltration.
As illustrated above, three primary risks adhere to data loss or exfiltration events:

- Loss of sensitive data or intellectual property, which can place an organization at a competitive disadvantage or even induce regulatory risk in the case of sensitive customer or client data
- Damage to an organization’s reputation following a publicized incident, including but not limited to public relations work to reduce impacts and long-term association with a major, high-profile data loss event
- Opening up the organization to extortion by actors who extract data, then threaten to publish it absent payment or some other action

The first case constitutes the typical focus of data loss concern, as this relates to incidents ranging from state-sponsored industrial espionage to criminal entities attempting to steal banking or personal information. The second adheres to all (publicly revealed or acknowledged) exfiltration incidents, and forms a difficult-to-quantify cost that may take years

to overcome. The third is a relatively recent development that aligns with emerging trends in "double-extortion" ransomware and similar criminal activity, although such tactics possess a longer history among politically motivated hacktivist entities.

Overall, while risks are significant, their quantification and time to manifestation are typically unclear, making these items hard to assess. Theft of intellectual property may have impacts felt only years (or decades) after the action takes place, while potential reputational losses may be impossible to adequately quantify under any circumstance. Yet given an increase in both hacktivist and especially ransomware operations incorporating data leak or data extortion activities, data exfiltration activity is rapidly becoming a more acute and immediate issue than in years past. Thus, network defenders and operators must align defenses in such a fashion to identify and potentially even prevent such activity where appropriate.

Identifying Exfiltration Behaviors

Exfiltration activity would be expected to be noisy and relatively easy to spot, given the movement of large amounts of data to new or unfamiliar sources. Yet even making this determination requires a combination of network visibility and active network monitoring such that most organizations either lack insight into such behaviors, or such activities blend in sufficiently with legitimate activity so as to produce significant "noise" in detections.

Furthermore, adversaries of various types employ multiple techniques to minimize visibility and evade detection, either through obfuscation or blending in with other actions. Examples include:

- Using legitimate third-party services, such as cloud backup systems or web-based storage, as destinations for leaked data. Examples range from common items such as Google Drive and Dropbox to more specialized products such as an ecosystem of items related to the Mega.io service.
- Tunneling traffic over non-HTTP services or using alternative protocols for large data transfer that might not be monitored with the same level of scrutiny.
- Dividing data into smaller pieces for exfiltration to avoid abnormally large traffic flows leaving the network.

Although all of these are concerning, none is impossible to detect. Rather, defenders can employ a combination of monitoring for more general traffic pattern anomalies along with specific identification of certain techniques or behaviors to detect this type of activity. Even in cases where data is encoded or encrypted, or where visibility into network activity is somewhat limited, access to datasets such as network flow provides a number of possibilities to flag suspicious operations.

At the most general level, identification of suspicious network flows represents a potentially powerful technique for identifying exfiltration activity. Aside from just looking for large data flows, with "large" likely dependent upon the normal operations and expectations within the

monitored network, identification of directionality and upload/download ratios are necessary enrichments to accurately leverage flow monitoring. For example, a “large” flow may just be indicative of a long-running, data-rich connection such as streaming, remote access, or related activity. However, when also looking at directionality and data ratios, identifying large flows where the majority of data (80–90 percent) leaves the network can flag a large upload session. While there may certainly be legitimate use cases for such activity, such as sharing large project files or datasets, these flows may also be indicative of data loss or other unauthorized activity.

The above becomes even more powerful when combined with an analytic approach to network connections. Identifying not just anomalous, outbound flows but connecting these to new, previously unobserved or suspicious network infrastructure enables powerful detection possibilities. For example, linking a suspicious outbound flow to new network infrastructure or a virtual private server (VPS) instance (such as Linode, DigitalOcean, or a similar provider) can identify activity of concern. Extending further, the same logic can also be applied to nonstandard connections (like external FTP or other protocols) to unknown or untrusted destinations to identify potential exfiltration activity.

While these strategies are effective in identifying when an exfiltration event occurs, they are lacking in that such approaches are fundamentally “backward looking.” In other words, these items will identify exfiltration as or shortly after it actually *takes place*. This can be of value to alert organizations that something is wrong, thus enabling response and mitigation operations and reducing time to detection and time to response metrics. Yet ideally defenders and their respective employers would be able to *prevent* exfiltration from taking place at all.

Applying “Whole of Kill Chain” Defense

Preventing exfiltration requires a “whole of Cyber Kill Chain” perspective with respect to network monitoring and defense. In this scenario, defenders look for overall adversary intrusion pathways and dependencies to identify necessary precursors (initial access, lateral movement, data collection, and data staging) prior to data leaving the network.

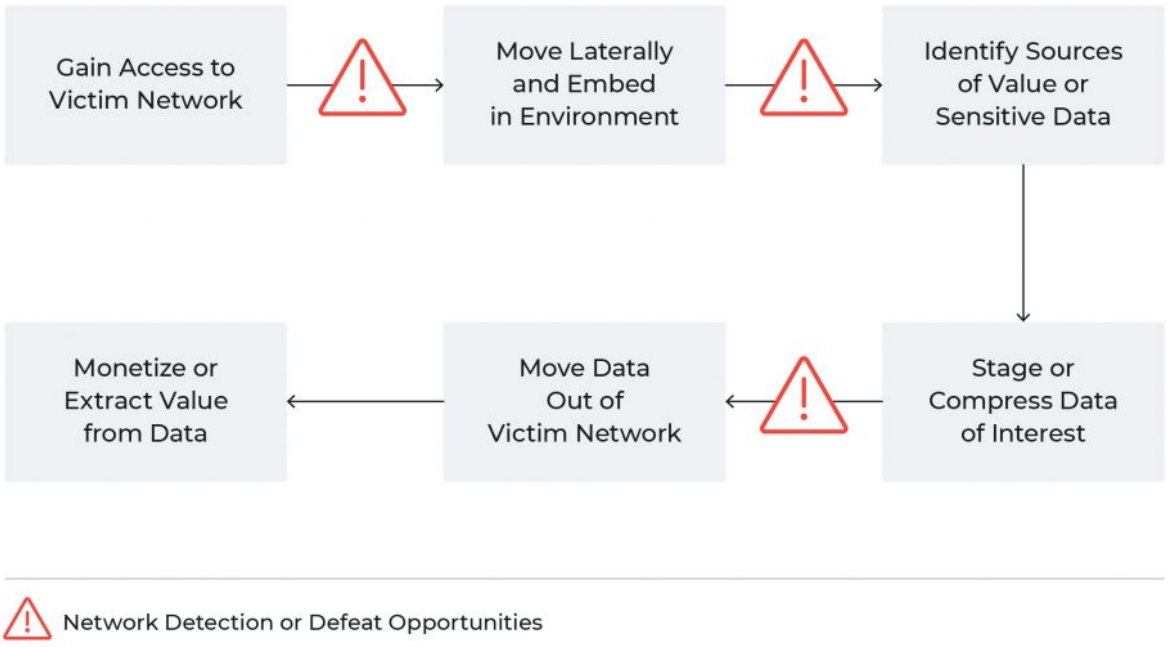

Figure 2: Threat intrusion pathways and points of defender intervention to detect or defeat attacks.

Through a critical examination of what an adversary requires to succeed, we can begin implementing more general controls across both host and network capabilities to defeat not just exfiltration operations but a variety of intrusions. At minimum, identifying likely adversary ingress mechanisms and either monitoring them closely or implementing controls to reduce attack surface can significantly reduce the likelihood of a threat actor gaining access to the environment. Relevant controls include patching external-facing systems, reducing available services from external access, limiting traffic types inbound to the network, and monitoring sensitive activity such as remote administration or access sessions.

Many organizations focus on these steps as a means to reduce risk, but true, layered defense requires going beyond the network boundary to ensure coverage across subsequent adversary actions as well. Given a near-infinite number of possible initial access mechanisms available to adversaries, as well as the possibility of subverting systems or users even when items are patched or otherwise monitored, defense must extend into internal network traffic flows and host items. Defenders must identify lateral movement techniques and their artifacts across both network and host behaviors. While increasing adoption of endpoint defense and response (EDR) products covers host-centric observations, investment in East-West traffic visibility and monitoring is essential to capture adversary traversal of defended networks. Implemented together, these items can ensure coverage of adversary operations ranging from opportunistic criminal actors to focused hacktivists to state-sponsored threats.

Overall, defenders need to combine an understanding of adversary behaviors and tendencies (such as through production or ingestion of cyber threat intelligence) to first determine how adversaries operate and what techniques are relevant in intrusion operations. Once determined, a combination of EDR and network defense and response (NDR) must be employed to ensure layered detection and monitoring and cover any potential gaps in visibility that adversaries may attempt to use to their advantage.

Conclusion

Data exfiltration is increasingly incorporated into the actions of a variety of threats, from criminals to hacktivists to state-directed intrusions. Through the application and monitoring of robust network controls, defenders and asset owners can ensure awareness of such behaviors when they occur, allowing for quicker response and potentially event mitigation.

Defenders cannot focus solely on detecting anomalous outbound data flows. Instead, robust defense requires identifying adversary actions across all phases of intrusion operations. Once understood, defenders can implement network and host controls or visibility to cover each step of an intrusion. Only through this robust, defense-in-depth approach can defenders ensure not just awareness of potential malicious activity (including but not limited to data exfiltration operations), but also enable possible intrusion interdiction or disruption when actions are caught early in the adversary's lifecycle. While neither easy nor inexpensive, defense against modern cyber threats, from ransomware operations to data exfiltration for a variety of purposes, demands such investment to ensure defenders keep pace with a rapidly evolving threat landscape.

GET YOUR COPY >

WEBINAR

Ransomware Best Practices: Agentless Threat Hunting

WATCH ON-DEMAND >

REPORT

2022 TLS Trends Data

DOWNLOAD REPORT >

WEBPAGE

Suddenly, Ransomware Has Nowhere to Hide

TAKE A LOOK >

OLDER ARTICLE
Applied Threat Research and Guided-SaaS NDR
NEWER ARTICLE
Gigamon ThreatINSIGHT Guided-SaaS NDR Releases 2021.4 and 2021.5

TOP