An Analytical Framework to Address the Data Exfiltration of Advanced Persistent Threats

Kamil Nar and S. Shankar Sastry

Abstract— **Detecting and preventing the data exfiltration of advanced persistent threats is a challenging problem. These attacks can remain in their target system for several years while retrieving information at a very slow rate, possibly after reformatting and encrypting the data they have accessed. Tainting and tracking some of the files in the system and deploying honeypots are two of the potentially effective measures against advanced persistent threats. In this paper, we introduce an analytical framework to study the effect of these measures on the amount of files that an attacker can exfiltrate. In particular, we obtain upper bounds on the expected amount of files at risk given a certain ratio of tainted and honey files in the system by using dynamic programming and Pontryagin's maximum principle. In addition, we show that in some cases tainting more of the files does not necessarily improve the security of the system. The results highlight the effectiveness and the necessity of deception for combatting advanced persistent threats.**

I. Introduction

Advanced Persistent Threats (APT) are long-term cyber-attacks that primarily target political organizations, government agencies and facilities, defense contractors, and industries with large influence on global markets [1]. As a type of targeted attacks, APTs intend to affect only their targets and the tools they employ for their attacks are developed specifically for their campaign, which is either sabotage or espionage. The attacks usually leverage the vulnerabilities in their target system which are unknown to the system administrator, and consequently, many APTs can go undetected until the attack reaches its ultimate goal.

Collecting confidential information about the target is either the main goal of an APT, or it is necessary to craft an impactful attack. For example, Flame was designed to collect information about its targets, and it was capable of logging keystrokes to capture passwords, taking screenshots, and recording voice using the internal microphone of the computers [2], [3]. On the other hand, Stuxnet was meant to sabotage nuclear plants in specific countries, but its design required detailed knowledge of what operating system and what type of controllers were used in the target plants and how the communication between machines and sensors could be intercepted [4].

APT attacks are known for their *low and slow* characteristics [1]. Since they aim to gather information about the target system, APTs make great effort to avoid detection and stay in the system as long as possible, and the attack progresses very slowly. For instance, Stuxnex was detected in 2010 even though it was believed to be developed in 2005 [2]. Similarly, Flame was detected in 2012, but some files particular to this APT were first observed in Europe as early as December 2007 [5].

After acquiring access to the information they are looking for, APT attacks can exfiltrate the collected data to their command and control servers at a very small transmission rate, unless the attacker is able to send them at once and sees no benefit in staying in the target system any longer. Before exfiltrating the data, the files could be reformatted, encrypted, or attached onto other files in order to avoid the detection mechanisms of the system. For example, the collected data were exfiltrated as JPEG files in the case of Duqu [6].

Transmission in a form different than the original files, along with very small rates of transmission, renders the detection of data exfiltration of APTs challenging. One of the potentially effective methods to address this problem is tainting, or watermarking, the classified files in the system. If a program in the system attempts to send a tainted file with sensitive information out of the network, the system could detect this attempt and prevent the transmission. However, APTs could copy the original files, recreate them in different formats and possibly add encryption. Therefore, the system needs to keep track of every program that accesses any tainted file and produces some other file anytime later since that file could potentially contain information obtained from the original file. Due to this dependence, taint tracking might easily become burdensome and could lead to frequent false alarms as many of the files produced by benign processes could also be seen as threats. To prevent these false alarms from interfering with the authorized use of the files, transmission of tainted files might be allowed at that instant and transmission logs could be screened later periodically. This, however, puts some of the files with sensitive information at risk of being exfiltrated between two screening times.

Another effective mechanism for detecting APTs is implementing honeypots, which are resources such as computers, account names, passwords, or files placed in an information system whose unauthorized use indicates the existence of an intruder in the system [7]. Required complexity of the honeypots varies depending on its purpose. For instance, if the system administrator wants to understand the motives of a potential intruder, they can place a virtual machine in the system which is isolated from the critical parts of the system but is still capable of interacting with an unauthorized user. If, however, the administrator wants to merely detect the intruder, then they can place fake files in the system which

This research was supported by the U.S. Office of Naval Research (ONR) MURI grant N00014-16-1-2710. Kamil Nar and S. Shankar Sastry are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA. Email: {nar,sastry}@eecs.berkeley.edu

contain no valuable information for the authorized users and should normally not be needed. An attempt to access any of these files, which are referred to as honey files, reveals the presence of an intruder in the system. Honey files do not cause false alarms unless an authorized user accidentally tries to access them, but they occupy some of the memory space in the system and the users need to have the necessary information to distinguish them from the real files.

Even though developing efficient algorithms for dynamic taint tracking is an active research area [8], [9] and the use of honeypots is known to be effective against APTs [10]–[12], there does not exist an analytical model in the literature to evaluate the security provided by these measures quantitatively while taking the dynamic nature of APTs into account. In this paper, we introduce an analytical framework to evaluate the effect of tainted and honey files on data exfiltration. In particular, we obtain upper bounds on the expected amount of files that an intruder can exfiltrate over a long time horizon given a certain ratio of tainted files and honey files in the system. We use dynamic programming and Pontryagin's Maximum Principle from optimal control theory [13], [14] to obtain the upper bounds on the amount of files at risk.

Organization of the rest of the paper is as follows. We introduce a discrete time model to analyze the data exfiltration in Section II. In Section III, we relax some of the conditions of the model, provide a continuous time approximation and obtain an upper bound on the amount of files at risk. Section IV provides a numerical example of the results obtained and Section V concludes the paper. All proofs are given in the Appendix.

II. Data Exfiltration In Discrete Time

Let N_r and N_h be the number of real files and honey files in the system, respectively, and let $N_{r\wedge t}$ of the real files be tainted. If the attacker attempts to access any of the honey files, the attack is detected by the system and any further data exfiltration is disabled. We assume that the attacker has no knowledge about which files are honey or tainted since this is the main reason for their deployment. Therefore, we assume that all files are equally likely to be accessed. Given this assumption, if the attacker tries to access $k \in \mathbb{N}$ files, it is not detected with probability $\binom{N_r}{k} / \binom{N_r+N_h}{k}$. If $k \ll N_r$, accessing files does not cause a significant change in the proportion of files that are not accessed, and we can approximate this probability by $(N_r/(N_r+N_h))^k =: e^{-\alpha k}$, where $\alpha = \log(1+(N_h/N_r))$. Similarly, given that the attacker has accessed only real files, if it tries to send $l \in \mathbb{N}$ files out of the network, the transmission is not detected if none of those files is tainted, which has probability $\binom{N_r - N_{r\wedge t}}{l} / \binom{N_r}{l}$. If $l \ll (N_r - N_{r\wedge t})$, the proportion of the files not transmitted remains almost the same and we can approximate this probability by $((N_r - N_{r\wedge t})/N_r)^l =: e^{-\beta l}$, where $\beta = \log(N_r/(N_r - N_{r\wedge t}))$.

We assume that the attacker decides to access $u_t \in [0,\infty)$ files and transmit $v_t \in [0, M]$ files at times $t \in \mathcal{T} := \{0, 1, \ldots, T-1\}$ and tries to maximize the expected number of files exfiltrated subject to the constraint

$$0 \le \sum_{t=0}^{k} v_t \le \sum_{t=0}^{k} u_t \quad \text{for all } k \in \mathcal{T}, \tag{1}$$

where M is a bound on the amount of files that could be transmitted in each time interval. This bound could be imposed by the transmission capacity of the network, or it could be the maximum transmission rate the attacker can use to avoid getting detected by the detection mechanisms that monitor the data flow rate in the system.

Since we assume that the data exfiltration will be terminated once the attacker is detected, the decision of the attacker at any time is relevant only if it has not been detected until that time. Therefore, the attacker will determine its decision at each time assuming that it has not been detected, and consequently, it will have an open-loop policy.

Our goal is to find an upper bound on the expected number of files that could be exfiltrated given a certain ratio of N_h/N_r and $N_{r\wedge t}/N_r$. Since the objective function and the strategy of the attacker will depend on whether the system detects the exfiltration of tainted files before or after the transmission takes place, we will consider these two cases separately.

A. Detection Before Transmission

Let x_k and y_k denote the total number of accessed and exfiltrated files, respectively, until time k:

$$x_k = \sum_{t=0}^{k} u_t, \quad y_k = \sum_{t=0}^{k} v_t \quad \text{for all } k \in \mathcal{T},$$

and let $d_k = x_k - y_k$ be the total number files that have been accessed but not transmitted until time k. Note that due to constraint (1), we have $d_k \ge 0$ for all $k \in \mathcal{T}$.

Now assume that an attempt to transmit tainted files is detected before the transmission takes place and exfiltration is disabled at that instant. At time $k \in \mathcal{T}$, exfiltration of v_k files is accomplished only if the attacker has not been detected until then, which has probability

$$e^{-\alpha u_0 - \beta v_0} e^{-\alpha u_1 - \beta v_1} \ldots e^{-\alpha u_k - \beta v_k} = e^{-\alpha x_k - \beta y_k}.$$

Therefore, the expected number of files exfiltrated at time k is $v_k e^{-\alpha x_k - \beta y_k}$, and the utility of the attacker can be written as

$$\sum_{t=0}^{T-1} v_t e^{-\alpha x_t - \beta y_t} = \sum_{t=0}^{T-1} v_t e^{-\alpha d_t - \gamma y_t},$$

where $\gamma = \alpha + \beta$. As a result, the goal of the attacker is to solve

$$\max_{\{d_t, v_t | t \in \mathcal{T}\}} \sum_{t=0}^{T-1} v_t e^{-\alpha d_t - \gamma y_t}. \tag{2}$$

Theorem 1. *If*

$$\frac{1}{\gamma} \ge M \frac{1 - e^{-\gamma M T}}{1 - e^{-\gamma M}}, \tag{3}$$

then the optimal strategy for the attacker is given by $v_t^* = M$ for all $t \in \mathcal{T}$. Otherwise, there exists some $k^* \in \mathcal{T}$ such that

$$M \frac{1 - e^{-\gamma M (t+1)}}{1 - e^{-\gamma M}} \le \frac{1}{\gamma} \quad \text{for all } t \le k^* - 1, \tag{4a}$$

$$M\frac{1-e^{-\gamma M(k^*+1)}}{1-e^{-\gamma M}} > \frac{1}{\gamma}, \tag{4b}$$

and the optimal policy is given by

$$v_t^* = \begin{cases} M & \text{if } t > T-k^*-1, \\ \frac{1}{\gamma}\left(1-e^{-\gamma v_{t+1}^*}\right) & \text{if } t \leq T-k^*-1. \end{cases}$$

Corollary 1. *Expected number of files that can be exfiltrated is bounded by*

$$\min\left\{\frac{1}{\gamma},\ M\frac{1-e^{-\gamma MT}}{1-e^{-\gamma M}}\right\}.$$

Theorem 1 shows that small values of γ, which corresponds to small ratio of tainted files and honey files, have no effect on the optimal policy of the attacker. That is, the attacker attempts to exfiltrate files at the highest rate possible as if there is no tainted files or honey files. However, if γ is large enough to influence the policy of the attacker, then the attacker starts with a slow rate of transmission and increases its rate towards the end of the horizon.

Corollary 1 shows that given a certain ratio of tainted files and honey files, the expected number of files that the attacker could exfiltrate can not exceed $\frac{1}{\gamma}$. This implies that integrating these files into the system provides a protection mechanism that is immune to the transmission rate and the length of the horizon, which is typically very long for most APTs.

B. Detection After Transmission

Now assume that the transmission of tainted files is detected only after the transmission has completed. If the attacker accesses and exfiltrates u_t and v_t files at time $t \in \mathcal{T}$, then v_k files are transmitted at time k with probability

$$e^{-\alpha u_0}e^{-\beta v_0-\alpha u_1}\ldots e^{-\beta v_{k-1}-\alpha u_k} = e^{-\alpha x_k-\beta y_k+\beta v_k}.$$

Then the objective of the attacker can be written as

$$\max_{\{u_t,v_t|t\in\mathcal{T}\}}\sum_{t=0}^{T-1} v_t e^{-\alpha x_t-\beta y_t+\beta v_t} + (u_T+d_{T-1})e^{-\alpha x_T-\beta y_{T-1}}. \tag{5}$$

We assume that at the end of the horizon, the attacker transmits all the files that have been accessed but not transmitted, which is represented by d_{T-1} in the second term in the objective function. Then, we have the following theorem, which follows from Lemma 1 and Lemma 2 given in Appendix.

Theorem 2. *Let* $c_T = \frac{1}{\alpha e}$ *and*

$$c_t = \max_{v\in[0,M]}\left\{ve^{-\alpha v}+e^{-\gamma v}c_{t+1}\right\} \quad \forall t \in \mathcal{T},$$

$$v_t^* = \operatorname*{argmax}_{v\in[0,M]}\ \left\{ve^{-\alpha v}+e^{-\gamma v}c_{t+1}\right\} \quad \forall t \in \mathcal{T}.$$

Then an attack policy which has

$$u_t = v_t = v_t^* \quad \forall t \in \mathcal{T}$$

and $u_T = \frac{1}{\alpha}$ *is optimal, and* c_0 *is the expected amount of files that could be exfiltrated.*

Corollary 2. *If*

$$0 = \operatorname*{argmax}_{v\in[0,M]}\ \left\{ve^{-\alpha v}+\frac{1}{\alpha e}e^{-\gamma v}\right\}, \tag{6}$$

then the optimal policy is to collect $\frac{1}{\alpha}$ *files and transmit no files until the end of the horizon, and the expected amount of files that could be exfiltrated is* $\frac{1}{\alpha e}$.

Corollary 2 shows that as long as γ is large or M is small enough to satisfy condition (6), the exact value of γ influences neither the optimal attacker policy nor the expected number of files that could be exfiltrated. In other words, tainting more of the real files does not improve the protection of the system. Note also that condition (6) holds only if

$$\frac{\gamma}{\alpha e} > 1,$$

which will reappear in Theorem 4 in Section III.*B*.

III. Continuous Time Approximation

In Section II, we assumed $u_t \in [0,\infty)$, i.e., there was no upper bound on the amount of files that the attacker can access at any time. Consequently, there was no benefit in accessing a file but not transmitting it, and the optimal policies required accessing only the files that are going to be transmitted at that time instant. In this section, we impose a bound on u_t, and provide a solution by introducing a continuous time approximation to the problem in Section II.

Let $u(t)$ and $v(t)$ denote the rate at which the files are accessed and transmitted at time $t \in [0,T]$, and let $x(t)$ and $y(t)$ be the amount of files that has been accessed and transmitted until time t, respectively:

$$\dot{x} = u, \quad \dot{y} = v,$$

with $x(0) = y(0) = 0$, $u(t) \in [0, M_u]$ and $v(t) \in [0, M_v]$ for all $t \in [0,T]$, where $M_u \geq M_v > 0$. Note that we have the condition

$$0 \leq y(t) \leq x(t) \quad \forall t \in [0,T].$$

A. Detection Before Transmission

The attacker tries to maximize

$$\int_0^T v(t)e^{-\alpha x(t)-\beta y(t)}dt.$$

Similar to the discrete time case, the optimal strategy requires $x(t) = y(t)$ for all $t \in [0,T]$. Defining $\gamma = \alpha+\beta$, we can rewrite the objective of the attacker as maximizing

$$\int_0^T e^{-\gamma y(t)}v(t)dt = \int_{y(0)}^{y(T)} e^{-\gamma y}dy = \frac{1}{\gamma}\left(1-e^{-\gamma y(T)}\right).$$

We observe that the attacker tries to maximize $y(T)$, and therefore, it keeps the transmission rate at the maximum value possible for all time. Similar to discrete time case, $\frac{1}{\gamma}$ is an upper bound on the expected number of files that can be exfiltrated, and this bound becomes tight as the horizon length T increases.

What we observe in the continuous time model is that the optimal strategy is to send the files one by one at the maximum rate, without aggregating or accumulating the accessed files.

Note that the result of the continuous time model is consistent with Theorem 1. As the time intervals in the discrete model become smaller, the bound M also diminishes while the product MT stays the same. As a result, the condition (3) in Theorem 1 holds and the optimal policy requires $v_t^* = M$ for all $t \in \mathcal{T}$.

B. Detection After Transmission

In this case, the objective of the attacker is to maximize

$$\int_0^T v(t)e^{-\alpha x(t)-\beta y(t)}dt + [x(T) - y(T)]e^{-\alpha x(T)-\beta y(T)}.$$

To obtain the optimal strategy, we use Pontryagin's maximum principle with state constraints [15], [16]. First define $d(\cdot) = x(\cdot) - y(\cdot)$, and let

$$\dot{y} = v, \quad v(t) \in [0, M_v] \quad \forall t \in [0,T],$$
$$\dot{d} = r, \quad r(t) \in [-M_v, M_u] \quad \forall t \in [0,T],$$

with the constraint

$$v(t) + r(t) \geq 0 \quad \forall t \in [0,T],$$

which means that the rate of decrease in the number of collected files cannot be more than the transmission rate. Note that we also have the constraint $d(t) \geq 0$ for all $t \in [0,T]$. Then, we can rewrite the objective function in terms of d, y, r and v as

$$\int_0^T v(t)e^{-\alpha d(t)-\gamma y(t)}dt \; + d(T)e^{-\alpha d(T)-\gamma y(T)}.$$

The Hamiltonian and the Lagrangian for this problem are

$$\mathcal{H}(d, y, \lambda_1, \lambda_2, r, v) = ve^{-\alpha d-\gamma y} + \lambda_1 r + \lambda_2 v,$$
$$\mathcal{L}(d, y, \lambda_1, \lambda_2, r, v, \mu) = ve^{-\alpha d-\gamma y} + \lambda_1 r + \lambda_2 v + \mu d,$$

where (λ_1, λ_2) is the costate, and μ is the Lagrange multiplier which satisfies

$$\mu(t) \geq 0, \quad \mu(t)d(t) = 0 \;\; \forall t \in [0,T]. \tag{7}$$

The dynamics of the costate vector is given as

$$\dot{\lambda}_1 = -\frac{\partial \mathcal{L}}{\partial d} = \alpha ve^{-\alpha d-\gamma y} - \mu,$$
$$\dot{\lambda}_2 = -\frac{\partial \mathcal{L}}{\partial y} = \gamma ve^{-\alpha d-\gamma y},$$

with the terminal values

$$\lambda_1(T) = (1 - \alpha d(T))e^{-\alpha d(T)-\gamma y(T)},$$
$$\lambda_2(T) = -\gamma d(T)e^{-\alpha d(T)-\gamma y(T)}.$$

The optimal actions maximize the Hamiltonian over the optimal trajectory:

$$(r^*, v^*) = \underset{\{(r,v)|r+v\geq 0\}}{\arg\max} \; \mathcal{H}(d^*, y^*, \lambda_1^*, \lambda_2^*, r, v).$$

Therefore,

$$r^* = \begin{cases} M_u & \text{if } \lambda_1 > 0, \\ -v^* & \text{if } \lambda_1 < 0, \\ \text{any value in } [-v^*, M_u] & \text{otherwise,} \end{cases}$$

and

$$v^* = \begin{cases} \underset{v\in[0,M_v]}{\arg\max} \;\; v\left(e^{-\alpha d-\gamma y} + \lambda_2\right) & \text{if } \lambda_1 \geq 0, \\ \underset{v\in[0,M_v]}{\arg\max} \;\; v\left(e^{-\alpha d-\gamma y} + \lambda_2 - \lambda_1\right) & \text{if } \lambda_1 < 0, \end{cases}$$

which can be written as

$$v^* = \begin{cases} M_v & \text{if } \lambda_1 \geq 0, \;\, s > 0 \\ & \text{or } \lambda_1 < 0, \;\, s - \lambda_1 > 0, \\ 0 & \text{if } \lambda_1 \geq 0, \;\, s < 0 \\ & \text{or } \lambda_1 < 0, \;\, s - \lambda_1 < 0, \\ \text{any value in } [0, M_v] & \text{otherwise,} \end{cases}$$

where $s = \lambda_2 + e^{-\alpha d-\gamma y}$. Note that

$$\dot{s} = -\alpha re^{-\alpha d-\gamma y},$$
$$s(T) = (1 - \gamma d(T))e^{-\alpha d(T)-\gamma y(T)}.$$

By analyzing the relationship between r, v, λ_1 and λ_2, we can obtain the form of the optimal policies, which is given in Theorem 3.

Theorem 3. *An optimal policy must have a time instant $t_0 \in [0,T]$ such that*

$$v^*(t) \in [0, M_v], \;\, r^*(t) = 0 \quad \forall t \in [0, t_0),$$
$$v^*(t) = 0, \;\, r^*(t) \in [0, M_u] \quad \forall t \in [t_0, T].$$

To obtain an upper bound on the amount of files at risk, we can consider the policies with the specific form given in Corollary 3.

Corollary 3. *There exists an optimal policy with $0 \leq \tau_0 \leq \tau_1 \leq T$ such that*

$$v^*(t) = M_v, \;\, r^*(t) = 0 \quad \forall t \in [0, \tau_0),$$
$$v^*(t) = 0, \;\, r^*(t) = 0 \quad \forall t \in [\tau_0, \tau_1),$$
$$v^*(t) = 0, \;\, r^*(t) = M_u \quad \forall t \in [\tau_1, T].$$

Theorem 4. *Given $T \geq \frac{1}{\alpha M_u}$, the optimal policy of the attacker is the following. If $\gamma > \alpha e$, then*

$$v^*(t) = 0, \; r^*(t) \in [0, M_u] \, \forall \, t \in [0,T] \text{ such that } d(T) = \frac{1}{\alpha}.$$

If $\gamma < \alpha e$,

$$v^*(t) = M_v, \;\, r^*(t) = 0 \;\forall\; t \in [0, T - t^*)$$
$$v^*(t) = 0, \;\, r^*(t) = M_u \;\forall\; t \in [T - t^*, T],$$

where t^* is the unique solution of

$$t^* = \underset{t\in[0,1/\alpha M_u]}{\operatorname{argmin}} \; e^{-\gamma M_v(T-t)}\left[1 - \gamma M_u te^{-\alpha M_u t}\right].$$

Finally, if $\gamma = \alpha e$, then any policy obeying Theorem 3 with $d(T) = \frac{1}{\alpha}$ is optimal.

Corollary 4. *The amount of files that the attacker can exfiltrate is bounded by* $\max\left\{\frac{1}{\gamma}, \frac{1}{\alpha e}\right\}$.

IV. A Numerical Example

Remember that by definition

$$\alpha = \log\left(\frac{N_r + N_h}{N_r}\right),\ \beta = \log\left(\frac{N_r}{N_r - N_{r\wedge t}}\right),\ \gamma = \alpha + \beta,$$

where N_r and N_h are the number of real files and honey files, respectively, and $N_{r\wedge t}$ is the number of real files with taint. Consider a system with $N_r = 1000$ files which detects the exfiltration of the tainted files before the transmission takes place. The expected number of files that an intruder can exfiltrate is given in Table I for different values of N_h and $N_{r\wedge t}$. Note that the values correspond to $\frac{1}{\gamma}$, which was obtained in Corollary 1 and in Section III.*A*. We observe that the number of files at risk strictly decreases as N_h or $N_{r\wedge t}$ increases.

TABLE I

Expected number of files that can be exfiltrated with detection before transmission

N_h \ $N_{r\wedge t}$	20	50	100	200	400
20	25.0	14.1	8.0	4.1	1.9
50	14.5	10.0	6.5	3.7	1.8
100	8.7	6.8	5.0	3.1	1.6
200	4.9	4.3	3.5	2.5	1.4
400	2.8	2.6	2.3	1.8	1.2

Table II displays the expected number of files that an intruder can exfiltrate if the system detects the transmission of the tainted files only after the transmission is completed. The values in the table correspond to the upper bound obtained in Corollary 4: $\max\{\frac{1}{\gamma},\ \frac{1}{\alpha e}\}$.

TABLE II

Expected number of files that can be exfiltrated with detection after transmission

N_h \ $N_{r\wedge t}$	20	50	100	200	400
20	25.0	18.6	18.6	18.6	18.6
50	14.5	10.0	7.5	7.5	7.5
100	8.7	6.8	5.0	3.9	3.9
200	4.9	4.3	3.5	2.5	2.0
400	2.8	2.6	2.3	1.8	1.2

In contrast to Table I, in Table II we observe that increasing the number of tainted files does not strictly improve the security. This is a consequence of the optimal policy of the intruder when exfiltration of tainted files is detected after the transmission. When there are relatively few honey files in the system, and hence, detection due to transmission is more likely than detection due to access, the attacker chooses to collect the files but not to transmit them until the end of the horizon. Since there is no transmission taking place, tainting does not help to detect the intrusion until all the files are exfiltrated at the end of the horizon. Note that existence of honey files becomes crucial to limit the data exfiltration in this case.

V. Conclusion

We introduced an analytical model to quantitatively evaluate the security provided by two of the potentially effective measures against APTs: tainting and tracking the sensitive files and adding honey files into the system. We showed that both of these measures limit the amount of files that an intruder can exfiltrate, and the defense acquired with them is immune to small rates of transmission, which is typical for most APTs. We obtained upper bounds on the amount of files that an intruder can exfiltrate given a certain ratio of tainted files and honey files in the system. We also showed that tainting more of the real files does not necessarily improve the security of the system if the system does not stop the transmission of tainted files at the instant of transmission. In this case, presence of honey files becomes crucial for early detection of an intruder and prevention of data exfiltration.

During our analysis, we assumed that the intruder had the perfect knowledge of the ratio of tainted files and honey files in the system, and we observed that the optimal attack strategy of the intruder changes based on this information. However, in a real scenario, an intruder may not have this knowledge, and it may need to estimate it by observing the system. Then, the system administrator has an incentive to deceive the intruder. That is, the system can benefit from leading the intruder to think that a larger or smaller portion of the files in the system is tainted or honey files. This results in a game of deception [17], which is a future direction for research.

The results obtained show that honey files are essential in some cases, and in others they can be used to decrease the ratio of tainted files needed for a certain level of security, thereby easing the taint tracking. More research and software development need to be put into automatic creation of honey files and methods to help the authorized users distinguish between honey files and real files.

Appendix

Proof of Theorem 1. Note that a policy with $d_t = 0$ for all $t \in \mathcal{T}$ is at least as good as another policy which differs only at time t' with $d_{t'} > 0$. In other words, it is advantageous not to access more files than that is going to be transmitted in each time interval, and there exists an optimal policy with $u_t = v_t$ for all $t \in \mathcal{T}$. Therefore, we can simplify the objective (2) of the attacker as

$$\max_{\{v_0, v_1, \ldots, v_{T-1}\}} \sum_{t=0}^{T-1} v_t e^{-\gamma \sum_{k=0}^{t} v_k}.$$

Let V_k denote the value function at time $k \in \{1, 2, \ldots, T-1\}$:

$$V_k = \max_{\{v_k, v_{k+1}, \ldots, v_{T-1}\}} \sum_{t=k}^{T-1} v_t e^{-\gamma \sum_{i=k}^{t} v_i}. \qquad (8)$$

Then, we have

$$V_{k-1} = \max_{v_{k-1}} \max_{v_k,\ldots,v_{T-1}} \sum\nolimits_{t=k-1}^{T-1} v_t e^{-\gamma \sum_{i=k-1}^{t} v_i}$$
$$= \max_{v_{k-1}} \max_{v_k,\ldots,v_{T-1}} e^{-\gamma v_{k-1}} \Big(v_{k-1} + \sum_{t=k}^{T-1} v_t e^{-\gamma \sum_{i=k}^{t} v_i} \Big)$$
$$= \max_{v_{k-1}} \; e^{-\gamma v_{k-1}} \left(v_{k-1} + V_k \right). \tag{9}$$

The expression to be maximized is quasiconcave in v_{k-1} since its partial derivative with respect to v_{k-1}

$$(1 - \gamma v_{k-1} - \gamma V_k) e^{-\gamma v_{k-1}}$$

changes sign only once. Therefore, the optimal action satisfies

$$v^*_{k-1} = \min \left\{ \frac{1}{\gamma} - V_k, \; M \right\}. \tag{10}$$

Let k^* satisfy the conditions (4a-4b). Then,

$$v^*_t = M, \quad V_t = \sum_{i=0}^{T-1-t} M e^{-\gamma M (i+1)} \quad \text{for all } t > T - k^* - 1$$

due to the equalities (8) and (10). On the other hand,

$$\begin{aligned} v^*_{T-1-k^*} &= \min \left\{ \frac{1}{\gamma} - V_{T-k^*}, \; M \right\} \\ &= -V_{T-k^*} + \min \left\{ \frac{1}{\gamma}, \; M \frac{1 - e^{-\gamma M (k^*+1)}}{1 - e^{-\gamma M}} \right\} \\ &= \frac{1}{\gamma} - V_{T-k^*} < M \end{aligned}$$

where the last inequality follows from (4b).

Now assume $v^*_t < M$, and hence, $v^*_t = \frac{1}{\gamma} - V_{t+1}$ for some $t \in \mathcal{T}$. From equation (9), we have $V_t = e^{-\gamma v^*_t} (v^*_t + V_{t+1})$, and consequently, $V_t = \frac{1}{\gamma} e^{-\gamma v^*_t}$. In addition,

$$\begin{aligned} v^*_{t-1} &= \min \left\{ \frac{1}{\gamma} - V_t, \; M \right\} \\ &= \min \left\{ \frac{1}{\gamma}(1 - e^{\gamma v^*_t}), \; M \right\} \\ &= \frac{1}{\gamma}(1 - e^{-\gamma v^*_t}) < v^*_t < M \end{aligned}$$

where the last equality follows from the assumption $v^*_t < M$ and the inequality

$$(1 - e^{-\gamma z}) < \gamma z \quad \forall z > 0.$$

An induction argument completes the proof. ■

Proof of Corollary 1. Note that V_0 is the maximum value of the utility of the attacker, and it is given by

$$V_0 = \begin{cases} M \frac{1 - e^{-\gamma M T}}{1 - e^{-\gamma M}} & \text{if } M \frac{1 - e^{-\gamma M T}}{1 - e^{-\gamma M}} \le \frac{1}{\gamma}, \\ \frac{1}{\gamma} e^{-\gamma v^*_0} & \text{otherwise.} \end{cases}$$

■

Lemma 1. *Let $V_k(d_{k-1})$ denote the value function at time $k \in \mathcal{T}$ for the problem (5):*

$$V_k(d_{k-1}) = \max_{\{v_t, u_t | t \ge k\}} \left\{ \sum_{t=k}^{T-1} v_t e^{\beta v_t - \alpha \sum_{i=k}^{t} u_i - \beta \sum_{i=k}^{t} v_i} \right.$$
$$\left. + \Big(u_T + d_{k-1} + \sum_{t=k}^{T-1} (u_t - v_t) \Big) e^{-\alpha \sum_{t=k}^{T} u_t - \beta \sum_{t=k}^{T-1} v_t} \right\}.$$

Let S_t denote the set $\{(v_t, d_t) : v_t + d_t \ge d_{t-1}\}$. Then, the value function satisfies

$$V_t(d_{t-1}) = \max_{(v_t, d_t) \in S_t} e^{-\alpha (v_t + d_t - d_{t-1})} \left[v_t + e^{-\beta v_t} V_{t+1}(d_t) \right]$$

for all $t \in \mathcal{T}$ with the terminal value

$$V_T(d_{T-1}) = \begin{cases} \frac{1}{\alpha e} e^{\alpha d_{T-1}} & \text{if } d_{T-1} \le \frac{1}{\alpha}, \\ d_{T-1} & \text{otherwise,} \end{cases}$$

and $V_0(0)$ denotes the expected amount of files that could be exfiltrated.

Proof of Lemma 1. The value function at the end of the horizon is

$$\begin{aligned} V_T(d_{T-1}) &= \max_{u_T} (u_T + d_{T-1}) e^{-\alpha u_T} \\ &= (u_T + d_{T-1}) e^{-\alpha u_T} \big|_{u_T = \max\{0, \; \frac{1}{\alpha} - d_{T-1}\}} \\ &= \begin{cases} \frac{1}{\alpha e} e^{\alpha d_{T-1}} & \text{if } d_{T-1} \le \frac{1}{\alpha}, \\ d_{T-1} & \text{otherwise.} \end{cases} \end{aligned}$$

For every $t \in \mathcal{T}$,

$$V_t(d_{t-1}) = \max \left\{ v_t e^{-\alpha u_t} + e^{-\alpha u_t - \beta v_t} \sum_{k=t+1}^{T-1} v_t e^{\beta v_k - \mu(t,k)} \right.$$
$$\left. + e^{-\alpha u_t - \beta v_t} \Big(u_T + d_{t-1} + u_t - v_t + \sum_{k=t+1}^{T-1} (u_k - v_k) \Big) e^{-\nu(t)} \right\}$$
$$= \max_{u_t \ge 0, v_t \in [0,M]} \left\{ v_t e^{-\alpha u_t} + e^{-\alpha u_t - \beta v_t} V_{t+1}(d_{t-1} + u_t - v_t) \right\}$$

where

$$\mu(t,k) = \sum\nolimits_{i=t+1}^{k} (\alpha u_i + \beta v_i),$$
$$\nu(t) = \alpha \sum\nolimits_{i=t+1}^{T} u_i + \beta \sum\nolimits_{i=t+1}^{T-1} v_i.$$

Note that $d_{t-1} + u_t - v_t = d_t$, and if we write the maximization over (d_t, v_t) instead of (u_t, v_t), the feasible set $\{(u_t, v_t) : u_t \ge 0, \; v_t \in [0, M]\}$ becomes

$$\{(d_t, v_t) | v_t + d_t \ge d_{t-1}, \; v_t \in [0, M]\}.$$

This completes the proof. ■

Lemma 2. *Consider the problem (5). There exists an optimal policy with $d^*_t = 0$ for all $t \in \mathcal{T}$. Furthermore, if $d^*_{t'} > 0$ for some $t' \in \mathcal{T}$, then $v^*_t = 0$ for all $t \ge t'$.*

Proof of Lemma 2. Given two policies $\{(v_t, u_t, d_t)\}_{t \in \mathcal{T}}$ and $\{(\tilde{v}_t, \tilde{u}_t, \tilde{d}_t)\}_{t \in \mathcal{T}}$ which satisfy

$$v_t = \tilde{v}_t, \quad \tilde{d}_t = 0 \quad \forall t \in \mathcal{T},$$
$$\tilde{u}_T = u_T + d_{T-1},$$

we have the inequality

$$\sum_{t=0}^{T-1} v_t e^{-\alpha d_t - \gamma y_t + \beta v_t} + (u_T + d_{T-1}) e^{-\alpha(u_T + d_{T-1}) - \gamma y_{T-1}}$$
$$\le \sum_{t=0}^{T-1} v_t e^{-\gamma y_t + \beta v_t} + \tilde{u}_T e^{-\alpha \tilde{u}_T - \gamma y_{T-1}}. \tag{11}$$

Therefore, expected utility given by a policy can only increase or stay the same if d_t is set to zero for all $t \in \mathcal{T}$. This implies that there exists an optimal policy with $d_t^* = 0$ for all $t \in \mathcal{T}$. The second statement in the lemma is a result of the fact that (11) becomes a strict inequality if $d_{t'} > 0$ and $v_t > 0$ for any $t \geq t' \geq 0$. ■

Lemma 3. *An optimal policy must have two time instants t_0 and t_1 such that $0 \le t_0 \le t_1 \le T$ and*

$$\begin{aligned} v^*(t) &\in [0, M_v],\ r^*(t) = 0 && \forall t \in [0, t_0), \\ v^*(t) &= 0,\ r^*(t) \in [0, M_u] && \forall t \in [t_0, t_1), \\ v^*(t) &\in [0, M_v],\ r^*(t) = M_u && \forall t \in [t_1, T]. \end{aligned}$$

Proof of Lemma 3. Let $\tilde{\tau} \in (0, T]$ be such that $d^*(\tilde{\tau}) > 0$. Then there exists some instant $\tau \in (0, \tilde{\tau}]$ such that $d^*(\tau) > 0$ and $\dot{d}^*(\tau) = r^*(\tau) > 0$. Note that $\mu^*(\tau) = 0$ due to the constraint (7). Since $r^*(\tau) > 0$, we have either $\lambda_1^*(\tau) > 0$ or $\lambda_1^*(\tau) = 0$.

1) $\lambda_1^*(\tau) > 0$:
   $r^*(\tau) = M_u$ and $\dot{\lambda}_1^*(\tau) \ge 0$ since $\mu^*(\tau) = 0$, which implies $\lambda_1^*(t) > 0$ for all $t \ge \tau$, and consequently, $r^*(t) = M_u$ for all $t \in [\tau, T]$.
2) $\lambda_1^*(\tau) = 0$:
   2i. $v^*(\tau) > 0$:
   $\dot{\lambda}_1^*(\tau) > 0 \implies \lambda_1^*(\tau^+) > 0$, which leads to the first case: $r^*(t) = M_u$ and $\lambda_1^*(t) > 0$ for all $t \in (\tau, T]$.
   2ii. $v^*(\tau) = 0$:
   $\dot{\lambda}_1(\tau) = 0$ and $r^*(\tau) \in [0, M_u]$ since $r^*(\tau) \ge -v^*(\tau)$.

Note that we obtained $r^*(t) \ge 0$ for all $t > \tau$, and since $d^*(\tau) > 0$ and $\dot{d}^*(t) = r^*(t)$, we have $d^*(t) > 0$ and $\mu^*(t) = 0$ for all $t \in (\tau, T]$. Consequently, $\lambda_1^*(t)$ is nondecreasing for all $t \in (\tau, T]$. Then the optimal policy on the time interval $(\tau, T]$ must first go through the case 2ii, then the case 2i, and then the case 1. However, $\tilde{\tau}$ is an arbitrary time instant with $d^*(\tilde{\tau}) > 0$; therefore, we can consider τ to be the first instant with $\dot{d}^*(\tau) > 0$, which corresponds to t_0 in the lemma, and this completes the proof. ■

Proof of Theorem 3. We need to show that $v^*(t) = 0$ for all $t \in [t_1, T]$ in Lemma 3. From the proof of Lemma 3, note that $\lambda_1^*(t) \ge 0$ for all $t \in [t_0, T]$, and therefore, $v^*(t)$ is determined based on the sign of $s^*(t)$ in the interval $[t_0, T]$. Since $v^*(t) = 0$ for all $t \in [t_0, t_1)$, we know either $s^*(t_1) < 0$ or $s^*(t_1) = 0$.

1) $s^*(t_1) < 0$:
   Since $r^*(t) \ge 0$ for all $t \in [0, T]$, we have $s^*(t) < 0$ and $v^*(t) = 0$ for all $t \ge t_1$.
2) $s^*(t_1) = 0$:
   Since $r^*(t) = M_u$ for all $t \ge t_1$, we have $\dot{s}^*(t) < 0$ for $t \ge t_1$, $s^*(t) < 0$ for $t > t_1$, and consequently, $v^*(t) = 0$ for all $t \in (t_1, T]$.

■

Proof of Corollary 3. Consider an optimal policy that satisfies the conditions in Theorem 3. With this policy, the objective function of the attacker attains

$$\int_0^{t_0} v^*(t) e^{-\gamma y^*(t)} dt + d^*(T) e^{-\alpha d^*(T) - \gamma y^*(t_0)}$$
$$= \frac{\left(1 - e^{-\gamma y^*(t_0)}\right)}{\gamma} + d^*(T) e^{-\alpha d^*(T) - \gamma y^*(t_0)},$$

which only depends on $y^*(t_0)$ and $d^*(T)$. If we choose τ_0 and τ_1 as

$$\tau_0 = \frac{y^*(t_0)}{M_v},\ \tau_1 = T - \frac{d^*(T)}{M_u},$$

then the policy given in Corollary 3 is optimal. ■

Proof of Theorem 4. We can write the utility of a policy given in Corollary 3 as

$$\int_0^{\tau_0} M_v e^{-\gamma M_v t} dt + M_u (T - \tau_1) e^{-\alpha M_u (T - \tau_1) - \gamma M_v \tau_0}$$
$$= \frac{1}{\gamma} - \frac{e^{-\gamma M_v \tau_0}}{\gamma} \left[1 - \gamma M_u (T - \tau_1) e^{-\alpha M_u (T - \tau_1)}\right].$$

Then, we can solve

$$\min_{0 \le \tau_0 \le \tau_1 \le T} e^{-\gamma M_v \tau_0} \underbrace{\left[1 - \gamma M_u (T - \tau_1) e^{-\alpha M_u (T - \tau_1)}\right]}_{h(\tau_1)} \tag{12}$$

to obtain the optimal policy. Note that if the term in the brackets, $h(\tau_1)$, can be negative, which corresponds to the condition $\gamma > \alpha e$, then we minimize it and set $\tau_0 = 0$. This means that the attacker will collect $1/\alpha$ files and not transmit anything until the end of the horizon.

If the minimum value $h(\tau_1)$ can achieve is zero, which corresponds to $\gamma = \alpha e$, then the optimal policy requires collecting $\frac{1}{\alpha}$ files to make $h(\tau_1) = 0$. The attacker can possibly transmit some files before starting to collect them in congruence with Theorem 3; this does not change the utility of the attacker.

If $h(\tau_1)$ is positive for all $\tau_1 \in [0, T]$, which corresponds to the condition $\gamma < \alpha e$, then τ_0 gets the largest possible value, which is τ_1, to minimize the expression in (12). If we make the change of variable $t = T - \tau_1$, the problem (12) becomes

$$\min_{t \in [0, T]} e^{-\gamma M_v (T - t)} \left[1 - \gamma M_u t e^{-\alpha M_u t}\right].$$

This expression is strictly increasing over $[1/\alpha M_u, T]$; therefore, the minimum is achieved on $[0, 1/\alpha M_u]$. In addition, the function is quasiconvex over $[0, 1/\alpha M_u]$ with a unique minimizer as long as $M_u \ge M_v$. ■

REFERENCES

[1] Symantec, “Advanced persistent threats: a Symantec perspective,” 2011.

[2] N. Virvilis, D. Gritzalis, “The big four – What we did wrong in advanced persistent threat detection?”, in Proc. of the International Conference on Availability, Reliability and Security, 2013.

[3] N. Virvilis, D. Gritzalis, T. Apostolopoulos, “Trusted computing vs. advanced persistent threats: Can a defender win this game?”, in Proc. of the 10th IEEE International Conference on Ubiquitous Intelligence and Computing, 2013.

[4] N. Falliere et al., *W32.Stuxnet Dossier: Symantec Security Response.* Mountain View, CA: Symantec, 2011.

[5] B. Bencsath et al., “The Cousins of Stuxnet: Duqu, Flame, and Gauss,” Future Internet, Vol. 4, pp. 971–1003, 2012.

[6] E. Chian, L. O. Murchu, and N. Falliere, “W32. Duqu: the precursor to the next Stuxnet,” in Proc. of the 5th USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET), 2012.

[7] L. Spitzner, “Honeypots: catching the insider threat,” in Proc. of the 19th Annual Computer Security Applications Conference, 2003.

[8] D. Zhu et al., “Tainteraser: protecting sensitive data leaks using application-level taint tracking,” SIGOPS Operating Systems Review, Vol. 45, No.1, pp. 142–154, 2011.

[9] S. Ma et al., “Protracer: towards practical provenance tracing by alternating between logging and tainting,” in Proc. of Network and Distributed System Security Symposium, 2016.

[10] N. Virvilis et al., “Changing the game: the art of deceiving sophisticated attackers,” in Proc. of the 6th International Conference on Cyber Conflict, 2014.

[11] D. Fronimos et al., “Evaluating low interaction honeypots and on their use against advanced persistent threats,” in Proc. of the 18th Panhellenic Conference on Informatics, 2014.

[12] Z. Saud et al., “Towards proactive detection of advanced persistent threat (apt) attacks using honeypots,” in Proc. of the 8th International Conference on Security of Information and Networks, 2015.

[13] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. I,* 3rd ed. Belmont, MA: Athena Scientific, 2005.

[14] D. Liberzon, *Calculus of Variations and Optimal Control Theory: A Concise Introduction.* Princeton, NJ: Princeton University Press, 2012.

[15] R. F. Hartl et al., “A Survey of the Maximum Principles for Optimal Control Problems with State Constraints”, SIAM Review, Vol. 37, No. 2, pp. 181–218, 1995.

[16] R. Vinter, *Optimal Control. Modern Birkhauser Classics.* Boston, MA: Birkhauser, 2010.

[17] V. J. Baston, F. A. Bostock, “Deception games,” International Journal of Game Theory, Vol. 17, No. 2, pp. 129–134, 1988.