

Forensic Analysis for Epidemic Attacks in Federated Networks

Yinglian Xie

Vyas Sekar

Michael K. Reiter

Hui Zhang

Carnegie Mellon University

{ylxie, vyass, reiter, hzhang}@cs.cmu.edu

Abstract— We present the design of a Network Forensic Alliance (NFA), to allow multiple administrative domains (ADs) to jointly locate the origin of epidemic spreading attacks. ADs in the NFA collaborate in a distributed protocol for post-mortem analysis of worm-like attacks. Information exchange between any two participating ADs is limited to traffic records that are known to both sides, maintaining the privacy of participants. Such an architecture is incentive-compatible – participants benefit by gaining better local investigative capabilities, even with partial deployment. Further, we also show that by sharing local investigation results, ADs can achieve global investigative capabilities that are comparable to a centralized implementation with access to global traffic records. Our evaluations demonstrate that it is feasible for large-scale attack investigation to be incrementally deployed in an Internet-like federation.

I. INTRODUCTION

Attackers today can launch epidemic attacks (worms) of dramatic intensity and reach. Much effort has been spent on designing more effective detection mechanisms (e.g., [1]), building better defense measures (e.g., [2]), and generating vulnerability-specific remedies (e.g., [3]). While these approaches provide immediate benefit to operators and end users, there is little ability to perform post-mortem diagnostic and forensic investigation of these incidents. Specifically, a service provider or enterprise needs a means to determine the entry point of a worm to its network so that the vulnerability it initially exploited can be identified. In addition, the ability to determine the true origin of a worm on a large network such as the Internet could help to prosecute and, ultimately, deter these attacks.

In this paper, we address the problem of providing forensic capabilities for investigating large-scale attacks in distributed, federated networks. A federated network is composed of multiple independent Administrative Domains (ADs). Conceptually, an AD is a network under a single administrative authority, and a single AD may be composed of one or more Autonomous Systems (ASes) in a BGP context. The most popular instance of a federated network is the Internet as we know it today, which is composed of multiple (possibly competing) ISPs, with diverse economic and peering relationships.

We propose the concept of a Network Forensic Alliance (NFA)¹—a collaborative effort involving multiple ADs to provide network-wide and localized forensic capabilities, as

depicted in Figure I. Using the techniques we describe in this paper, an NFA allows multiple ADs in a federated network to jointly perform local attack investigation, with the additional ability to diagnose attacks globally.




Fig. 1. NFA: Architecture for federated forensics

Each participating AD in the NFA possesses an independent traffic monitoring infrastructure. Traffic data (in the form of flow records) are collected and saved for a period of time during which an investigation can be launched using them. Investigative capabilities are realized using a distributed protocol which builds on, and enhances our earlier work [5] on worm origin identification. To launch an investigation, each AD feeds its local traffic records to its analysis engine running the distributed operations (described in Section IV). In the course of running the protocol, the analysis engine may communicate intermediate results to other participating ADs. Each analysis engine subsequently post-processes its results to obtain a local view of the attack entry within its domain, or shares the results with other domains to obtain a global view of the attack.

To our knowledge, we are the first to formulate the problem of forensic analysis for epidemic attacks in the context of federated networks. Our design is not only attack-agnostic (within the class of epidemic-spreading attacks), but also addresses the real world deployment concerns of privacy, participation incentives, and incremental deployment. The NFA: (i) is privacy-preserving, in that each AD reveals to another AD only records for traffic that each has already seen; (ii) provides participation incentives for individual ADs to collaborate with other ADs, i.e., each AD improves its own local diagnosis by cooperating with others; and (iii) offers strong incremental deployment properties, in that participants

¹The name is inspired by existing efforts at a Fingerprint Sharing Alliance [4], for generating and sharing attack signatures.

benefit even in partial-deployment scenarios. We consider (i)-(iii) to be essential to a viable federated approach, given the reticence of competing ADs to cooperate except when it is in their interests to do so.

We also present a methodology to address the fundamental challenges of evaluating the effectiveness of worm forensics in such a distributed setting, in the absence of multi-AD traffic datasets. Our framework permits us to specify well-defined measures of incentives that any solution in this domain need to address. Our evaluation includes two sets of experiments: one uses simulation on a network topology defined by the Internet AS-level topology, and another uses traffic collected from the *Internet2* [6] educational backbone. These experiments show that our approach realizes investigative capabilities comparable to an ideal unified network administrative model, while operating in a federated setting. The evaluations also validate that our approach greatly increases the local investigative capabilities of participating ADs; enables ADs to generate a network-wide reconstruction of the epidemic attack; and degrades gracefully with only partial deployment.

II. BACKGROUND

Early work in network forensics focused on stepping-stone detection and IP traceback. Stepping-stone detection attempts to bridge sources of attack indirection, in attacks that are launched through a series of intermediate machines. The existing techniques (e.g., [7], [8]) typically require packet-level analysis, and in some cases require the analysis to be applied close to the compromised machines. Both of these would be obstacles in applying these techniques to trace the origin of a large-scale worm after the fact, and our approach requires neither. IP traceback focuses on identifying the true sources of packets with spoofed source addresses (e.g., [9], [10], [11]). With respect to large-scale epidemic attacks, traceback techniques typically are not needed since source addresses are seldom spoofed. However, even in epidemic attacks with spoofing, these techniques would simply identify another victim in the vast majority of cases. It is this last obstacle that we seek to overcome here.

Network telescopes have been suggested as an alternative approach to reconstruct worm attacks [12], [13]. Kumar et al. [12] present an approach for reconstructing the spread of the Witty worm by using scan traffic collected by network telescopes to reverse engineer the pseudorandom number sequences used. Rajab et al. [13] use the telescope-observed scan traffic to infer infection times by studying the inter-arrival times between successful scans. Both approaches tackle the problem of worm forensics, but are targeted toward specific types of random-scanning attacks. Furthermore, the robustness of using network telescopes for forensic analysis is a topic of ongoing research [14].

Our earlier work has proposed a random moonwalk algorithm to identify the origin of an epidemic attack [5]. The algorithm takes traffic flow-records as inputs, and generates as output a set of “high-ranking” flows. The algorithm works by repeatedly performing random “moonwalks” on the graph and

correlating the set of flows traversed by the walks. Each moonwalk starts from an arbitrarily chosen flow. The algorithm then randomly picks a next step backward in time from the set of flows that arrived at the source host of the current flow, within the previous Δt seconds. For each next step, the above process repeats until there are no candidate flows to continue the path, or the walk has traversed a maximum of d hops. Our earlier analysis and empirical evaluation [5] suggest techniques for selecting the sampling window size Δt and the maximum path length d , using a trace-driven adaptive approach. After many such walks, the algorithm returns a set of flows that are most frequently traversed by the walks. Due to the tree-structured nature of epidemic attacks, the initial causal flows of the attack emerge among the highly ranked flows returned.

The random moonwalk algorithm assumes a unified network administrative model. In large federations such as the Internet, establishing a centralized traffic repository is hardly feasible due to privacy and economic considerations of multiple competing ADs. The approach we propose here builds from the random moonwalk algorithm, but overcomes the shortcomings for deploying this algorithm in Internet-like federated environments. While each AD can perform *isolated* attack investigation using the centralized algorithm, this does not achieve global forensics goals. In addition, as we will show, isolated forensics provides weaker local investigation capabilities than the distributed protocol we propose here, since each AD observes only a small portion of a large-scale, distributed attack.

III. SYSTEM DESIGN

In this section, we motivate the key component of our design, and outline the three stages that comprise the distributed operation of an NFA.

A. Why Random Moonwalks

The key insight behind our design decision of adopting the random moonwalk algorithm is an alternative interpretation of it as a Monte-Carlo simulation method to compute the largest eigenvector of a network flow graph. In this light, known results suggest that this algorithm should be relatively resilient to missing data, and so could form the basis of a solution offering good partial deployment properties.

To make this alternative interpretation more concrete, consider a directed *flow-graph* G_f , where each node represents a flow. Given a sampling window size Δt for random moonwalks, we insert a directed edge from node n_2 to node n_1 , if the flow n_1 arrived at the source of the flow n_2 at most Δt time units before that machine initiated n_2 . That is, there is a directed edge from node n_2 to node n_1 if a random moonwalk could (single-) step from the flow n_2 to the flow n_1 . After inserting all such edges, we then insert an edge from each sink n_2 (i.e., with zero out-degree) to every other node in G_f . A random moonwalk then corresponds to a Markov random walk on G_f , where at each node n_2 , the distribution of the next node is uniform over all nodes to which n_2 has an edge. With this Markov random walk interpretation, the normalized

adjacency matrix A_f defined by G_f is a stochastic matrix whose row sums are all 1. We consider the spectral analysis of matrix A_f , namely (non-trivial) solutions to the equation:

$$A_f x = \lambda x$$

The stationary distribution π_0 of A_f , computed as the largest eigenvector of A_f , represents the probability of reaching every node in a global stable state. The highest probability flows being traversed in a random moonwalk therefore correspond to those nodes with largest numerical values in the largest eigenvector. These flows, with high probability, are the initial causal flows. We can therefore view random moonwalks as a Monte Carlo sampling method to compute the largest eigenvector of the flow graph G_f .

Such a spectral analysis view suggests that random moonwalks can be robust to missing traffic, as previous studies have shown that spectral techniques are effective in identifying the underlying structure of graphs even with partial data available [15], [16]. For this reason, we choose to build the NFA from this random moonwalk algorithm, extending it into a distributed version that can operate in a federated environment.

B. Federated Forensics

In the NFA, multiple ADs independently collect traffic data, and jointly perform attack investigation to identify both the local attack entry points and the global attack source. The high-level idea is that multiple domains can perform loosely coordinated random moonwalks inside their own domains, so that when these walks are viewed together, they achieve the same effect as performing random moonwalks on a unified network. More precisely, this process involves the following three stages:

1. *Distributed traffic monitoring*: Each participating AD logs flow-level traffic records. Prior work on hash-based traceback has demonstrated that it is feasible to build such fine-grained traffic logging capabilities [11]. As shown in Figure I, each AD deploys traffic collectors (e.g., routers) to log network flow records, which will then be accessed or queried by an analysis engine.

2. *Collaborative random moonwalks*: Federated domains perform joint attack investigation via a distributed protocol. The analysis engines interact with each other to launch coordinated local random moonwalks on the traffic collected within their own domains. The shared intermediate results will be used to guide new local random moonwalks in an iterative fashion. These intermediate results should not release data that are proprietary to an AD or that should otherwise be protected.

3. *Post-processing for attacker identification*: After performing the collaborative random moonwalks, each participating domain obtains a set of suspicious flows. Each participant can either choose to use the flows independently for local attack investigation, or to further share information to achieve global investigative capabilities.

Since the NFA would be deployed incrementally, one major challenge throughout the above process is to gracefully handle missing information, so that the NFA utilizes all available traffic records even when certain ADs do not participate. Next, we first describe how multiple ADs jointly perform collaborative random moonwalks and perform post-processing. We then discuss strategies to deal with missing traffic records in partial deployment scenarios.

IV. PROTOCOL FOR DISTRIBUTED RANDOM MOONWALKS

In this section, we present a distributed protocol for multiple domains to jointly perform attack investigation in a federated network. In the distributed protocol, each AD participating in the NFA performs two operations: (1) it runs local random moonwalks on flow-records it has collected, and (2) it exchanges intermediate results with other collaborating ADs. These intermediate results will ensure that moonwalks continue even when they cross the network boundaries of ADs which originated them. When these local moonwalks are “stitched” together (via the inter-AD exchanges), they approximate random moonwalks on a global traffic trace (i.e., the union of the traffic observed at all the participating ADs). Throughout the protocol, each AD independently keeps a count of the number of times each flow has been traversed in the local random moonwalks. The output of the protocol at each domain will be a set of Z top frequency flows, which are returned as candidate flows for further investigation. To simplify the description of the protocol, we assume that all ADs in the network collaborate in the distributed operation, and defer a discussion of partial deployment until Section VI.

The protocol begins with a relatively lightweight bootstrap phase to initialize parameters. Participating ADs first decide on the time boundary for investigating the attack. Next, they select H , the maximum number of AD hops each distributed moonwalk can traverse, and the sampling window size Δt . Δt specifies the maximum look-back window within which candidate flows can be selected for continuing moonwalks. Each domain could also independently specify a maximum local walk-length $d(i)$ for walks it performs inside its own domain. ADs also agree on a traffic normalization factor, η , which determines the number of moonwalks that each AD launches initially.

After the initialization, each AD_i participates in the distributed random moonwalk protocol, and its operation in the protocol is as defined by Figure 2. Using the notation in Table I, the protocol works as follows:

1. *Initial launch*: Each AD_i initially launches random moonwalks from within its domain. The number of walks $W(i)$ launched by AD_i , is calculated as $W(i) = \eta \times |\text{FlowSet}(i)|$. AD_i then selects $W(i)$ flows from $\text{FlowSet}(i)$, uniformly at random with replacement. For each selected flow F , AD_i sets its AD hop count $h = 0$, constructs the tuple $\langle F, h \rangle$, and inserts it into $\text{InitSet}(i)$ ².

²We use multisets for keeping tuples of flow records and their AD hop counts. Each tuple can occur multiple times in a multiset.

Notation	Description
$\text{Flow}(i, j)$	A flow record with the source-host in domain AD_i and the destination-host in domain AD_j .
$\text{FlowSet}(i)$	The set of network flow records whose sources belong to AD_i .
$\text{InitSet}(i)$	A multiset of flow records and their AD hop counts that AD_i uses to start local moonwalks.
$W(i)$	The number of walks initially started by AD_i .
$d(i)$	The maximum length of a local random moonwalk within AD_i .
H	The maximum number of AD hops a distributed moonwalk can traverse.
Δt	The sampling window size, defined as the maximum time window to look back for continuing a moonwalk.

TABLE I

NOTATION USED IN THE DISTRIBUTED OPERATIONS. $W(i)$ AND $d(i)$ ARE LOCAL PARAMETERS DEFINED BY AD_i . H AND Δt ARE GLOBAL PARAMETERS AGREED BY ALL PARTICIPATING ADS BEFORE THE OPERATION.

```

// delta_t denotes the sampling window size,
// H is the maximum AD hop count
// TSet is a data-structure implementing a set
TSet distributed_moonwalk(TSet FlowSet(i), int W(i),
                        int d(i), int H, int delta_t){
    // Select W(i) flows to initialize InitSet(i)
    for (int j = 1 to W(i)){
        Select a flow F randomly from FlowSet(i)
        Insert <F,0> into InitSet(i);
    }
    // Perform local moonwalks
    while (InitSet(i) not empty){
        foreach tuple <F,h> in InitSet(i) {
            // Perform a random moonwalk starting at F
            // Let endflow be the last step of a local walk
            endflow = random_moonwalk(F, d(i), delta_t);

            delete <F,h> from InitSet(i)

            // Check if the walk reached the max. AD hopcount
            h = h + 1;
            if (h >= H)
                continue;

            // If the source host of endflow is in AD(j),
            // Send the tuple <endflow,h> to AD(j) so that
            // AD(j) will continue the walk from endflow
            int j = find_source_domain(endflow);
            if (j != i){
                send <endflow,h> to AD(j);
            }
        } // end foreach
    } // end while
    Return a set of flows with highest traversal counts;
}

// tuple <F,h> is the message received from another AD
void received_flow_record(Tuple <F,h>){
    insert <F,h> into InitSet(i)
}

```

Fig. 2. Pseudocode for the operations each AD_i executes in the protocol. Each analysis engine performs two operations of running distributed random moonwalks and updating its $\text{InitSet}(i)$.

In our two-domain example (Figure 3), both AD_1 and AD_2 use $W(i) = 2$. During this initial step, AD_1 selects flow $A(1,1)$ twice and inserts two entries into $\text{InitSet}(1)$, both with hop count set to zero. AD_2 similarly inserts two entries into $\text{InitSet}(2)$.

2. *Local random moonwalks*: For every entry $\langle F, h \rangle$ in $\text{InitSet}(i)$, AD_i will launch a random moonwalk, with F as the initial step. After performing the local moonwalk from F , AD_i deletes the corresponding tuple $\langle F, h \rangle$ from $\text{InitSet}(i)$. A walk will stop within domain AD_i if it meets any of the following criteria.

- There is no flow to continue the walk, within the previous

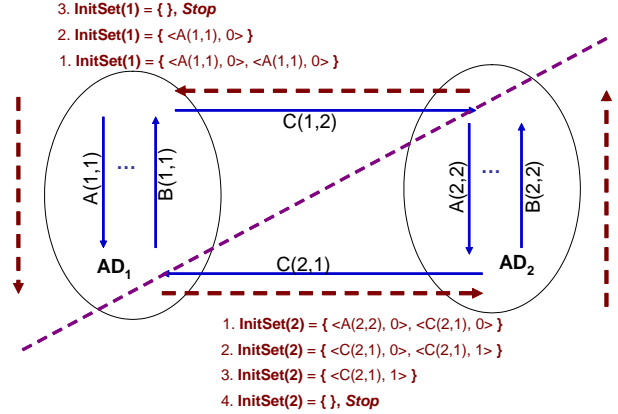


Fig. 3. A simple example with two domains collaborating in the distributed protocol. Solid arrows represent the direction of network flows, and dashed arrows represent the direction of random moonwalks (reversed with respect to the flow directionality).

Δt seconds.

- The walk has traversed the maximum number of local hops $d(i)$, within AD_i .
- The walk has reached a cross-domain flow whose source host is located in a different AD.

The first two stopping criteria are the same as in the centralized algorithm. The third condition arises due to the distributed setting. Since each AD has only a local view of the global traffic, it lacks sufficient information to select a next step to continue the moonwalk. Specifically, for flows originating from foreign hosts that belong to other ADs, AD_i cannot observe all possible incoming flows at such hosts. In the example, a random moonwalk starting from flow $A(1,1)$ stopped within AD_1 because it reached flow $C(2,1)$, whose source is inside AD_2 .

3. *Cross-domain exchanges*: If a local random moonwalk starting from F reaches a cross-domain flow $\text{Flow}(j, i)$ ³, AD_i increases the corresponding AD hop count h by one. If the new h is smaller than the maximum AD hop count H , AD_i sends the tuple $\langle \text{Flow}(j, i), h \rangle$ to AD_j .

On reception of the tuple $\langle \text{Flow}(j, i), h \rangle$, AD_j inserts it into $\text{InitSet}(j)$. AD_j will subsequently launch a local random moonwalk starting from $\text{Flow}(j, i)$, to continue the

³We assume that the source-AD information for each flow can be obtained from whois lookups or data from public routerservers, or using specialized tools [17]

distributed random moonwalk.

For example, AD_1 sends $\langle C(2,1), 1 \rangle$ to AD_2 , which then inserts the tuple into $InitSet(2)$. AD_2 does not send any flow record to AD_1 because the local moonwalk starting from flow $A(2,2)$ does not reach a cross domain flow.

4: *Termination condition*: If $InitSet(i)$ becomes empty, AD_i terminates its operation, and returns the set of Z flows with the highest counts (i.e., the sum of counts across all of its local random moonwalks).

The inter-AD exchanges have two important properties: (1) minimal information disclosure, and (2) low overhead. The flow records that can be exchanged between a pair of ADs can only be among the set of inter-domain flows between them. These cross-domain flows can be independently logged by both ADs, and thus exchanging these flow records does not reveal internal traffic information that is not already available to either AD.

The worst case communication overhead occurs when every step along the walk triggers a message to be sent across a pair of ADs, resulting $\sum_1^n W(i) \times H$ flow records to be exchanged in the federation. $\sum_1^n W(i)$ will typically be a small fraction of traffic, depending on the traffic normalization factor η (typically of the order of 0.1-1%). So the number of records that need to be exchanged will be a very small fraction of total unique flow records. However each flow-record that needs to be exchanged, incurs a per-record communication overhead between the pair of ADs. To reduce the number of communication events, ADs can choose to aggregate the set of cross-domain flows at which local walks terminate, and exchange these flow-records in a single exchange with another collaborating AD.

V. POST-PROCESSING FOR INVESTIGATION

After performing the distributed random moonwalks, each participating domain obtains a set of top frequency flows, to serve as starting points of further attack investigation. Participating ADs can either choose to use these flows for independent local attack investigation, or further share information to achieve global investigative capabilities.

A. Local View

ADs can use the flow counts obtained after the distributed moonwalks on the set of flows collected within their network boundaries, and return the set of Z top frequency flows for local diagnosis. The locally observed top frequency flows can then be used to diagnose the initial attack entry points, or initial infected hosts within a domain. Identifying the attack entry points can help detect configuration or firewall rule errors that may have allowed the attack to enter the network perimeter in the first place.

With such a *Local View* option, ADs realize benefits of collaboration by just participating in the distributed protocol. No additional information will be shared after the distributed protocol has been executed. Section VII-A.2 will quantitatively evaluate the participation benefit in terms of improved quality

of local attack investigation compared with isolated forensics using various measures.

B. Global Merge

ADs can choose to share selected suspicious flow records from within each domain for the NFA as a whole to realize global benefits. This functionality is comparable to operation under a unified network model to reconstruct the initial stages of the attack, providing diagnostic aids to pinpoint the origin of the attack.

In such a *Global Merge* strategy, each AD provides a set of locally identified top frequency flows that originate from its own domain, together with their counts, after executing the distributed protocol. Each AD_i identifies Z top frequency flows from $FlowSet(i)$, and broadcast the counts of local flows to every other participating domain. On reception of all other domains' top frequency flow counts, each AD can identify the Z flows with the highest global counts.

We note that to compute global flow ranks, each AD_i contributes Z flows from $FlowSet(i)$, instead of all collected flows that are used in Local View. Since for $i \neq j$, $FlowSet(i)$ and $FlowSet(j)$ do not overlap, flows will not be double counted by different domains. For each flow in $FlowSet(i)$, AD_i is the only domain to select its next step in the distributed random moonwalks; thus the frequency count that AD_i maintains for that flow is the flow's global frequency count. Therefore, the computed global flow ranks will be equivalent to the ranks computed by an ideal centralized moonwalk algorithm, assuming the entire federated network is a large unified network.

ADs need not reveal internal traffic records that are identified, and can share only high frequency cross-domain flow records by broadcasting these flow records and their counts to other participants. This can provide sufficient information to understand how the attack propagated across domains. Participating ADs can also choose to reveal these selected flow records using different degrees of source-address anonymization. The extent to which the globally reconstructed initial attack graph aids further diagnosis, depends on the degree of anonymization that the participating ADs engage in. At one end of the spectrum, ADs can choose to construct the graph at the granularity of actual end-hosts, providing the most fine-grained diagnosis possible. At the other end, ADs can choose to construct the graph at the AD granularity. In the intermediate space, ADs can reconstruct the attack graph at the granularity of coarse-grained network prefix ranges (e.g., using /8 or /16 masks).

VI. HANDLING PARTIAL DEPLOYMENT

So far, we assumed that all the ADs in a federated network participate in the NFA. In reality, such an NFA will be deployed in an incremental fashion. Even though the NFA may continue to observe a significant fraction of traffic, we need to extend the protocol for the ADs to utilize all available traffic records in partial deployment scenarios. Specifically, during the distributed random moonwalks, it may not be possible




Fig. 4. An example partial deployment scenario. W , X , Y are three ADs that are in the federation, and AD Z does not participate. $F1$, $F2$, and $F3$ represent three network flows to or from host u in AD Z .

to continue the walk from flows whose source-ADs do not participate in the NFA.

Figure 4 shows an example partial deployment scenario, where ADs W , X , Y participate in the federated network, while AD Z does not. In the distributed moonwalks, suppose X performs a local moonwalk that reaches flow $F2$, whose source host u is located inside domain Z , then none of the participating ADs will be able to observe all the incoming flows to u to continue this walk across domains.

A naive *discard* solution is to assume a closed world environment where ADs simply discard flows originating at non-participating domains. This solution is sub-optimal, as it does not utilize all the observed traffic data among all collaborating ADs, and interesting traffic flow records may end up being discarded in the process. A better solution is to try and perform a best-effort next step selection based on the data available to the participants in the federation. We consider the following strategies to enhance the distributed random moonwalks for continuing a walk once some participating AD reaches a flow originated from a non-participating domain.

Random selection: The simplest strategy is to randomly pick a participating AD and request it to continue the walk. In our example, X could randomly pick W or Y to continue the walk from $F2$, based on their local views of the incoming flows to host u .

Routing-path based selection: The second option for the AD is to use routing information and select the participating AD that is closest to the source host on the AD-path. The assumption here is that this selected AD will have the best available partial view to continue the walk. In our example scenario, flow $F2$ is routed through AD Y to X . Hence with routing-based selection, X would send a message to Y , requesting it to continue the walk, instead of choosing AD W .

Routing-horizon based selection: A generalization of the routing-based selection is to identify a set of participating ADs, defined as a *routing horizon*, that provide the greatest coverage over routing paths from the non-participating AD. Given the routing horizon, the AD that reaches a cross-domain flow in the random moonwalks can broadcast the flow record among the participating ADs on the routing horizon, for selecting candidate next step flows. Recall that given the current step f of the walk, candidate flows are those whose destination is the source of f , and which occur within previous Δt seconds of f . From the set of candidate flows, the AD can

randomly select a flow to continue the walk, and hands off the current flow to the source AD of the selected flow. In our example, when AD X reaches flow $F2$, it could query both W and Y for their observed incoming flows to host u within the previous Δt seconds to $F2$. Suppose both $F1$ and $F3$ are such candidate flows, X could then randomly pick one, say $F1$ as the next moonwalk step.

VII. EVALUATION

Our primary results are drawn from a simulation study, with synthetically generated traffic traces mapped into an Internet AS-level topology with over 16,000 ASes (Section VII-A). We also independently validate our findings using data derived from the *Internet2* educational backbone (Section VII-B). Broadly, our evaluate attempts to answer the following questions:

- How does the forensics performance achieved in a federated setting compare with similar analysis in a unified network model? What is the overhead of message exchanges in the distributed protocol? (Section VII-A.1)
- Is the architecture incentive-compatible, i.e, what are the benefits an AD can gain by collaborating with other ADs? (Section VII-A.2)
- How does the performance degrade under different partial deployment scenarios? Is the architecture incrementally deployable, and in particular do participants realize benefits even under partial deployment? (Section VII-A.3)

A. Simulation Study

We use a simulated trace-driven study to evaluate the deployability of an NFA in an Internet-scale large federation. The network topology for our study is defined by the Internet AS-level topology, obtained from Routeviews [18]. Each AS in the AS-level topology is mapped into a single AD in our federated network. We model the address-space ownership distribution across different ADs using prefix-ownership advertisements observed from Routeviews, proportionally scaling them down to a smaller address space of 10^7 hosts. Routing in our federated network is modeled using hop-count based shortest-path routing.

We use an event-driven simulator to generate synthetic flow-level traffic traces in our simulated network topology. Each of the 10^7 end-hosts in the address space is associated with a *working-set* of destination end-hosts that it talks to. The working-set size distribution across the hosts is captured using a discretized power-law distribution, with working-set sizes ranging between 10 and 1000 (the power-law ensures that most of the hosts have small working-sets and a small number of hosts have large working-sets). To initiate a flow, a host picks a destination from its designated working set using a preferential selection policy (also modeled with a discretized power-law distribution over the working-set size).

Attack traffic is specified using a worm scanning model, a scanning rate, and the fraction of hosts vulnerable. In our evaluation, the set of vulnerable hosts are selected uniformly at random from within the address-space of 10^7 hosts with

a total of 10% of hosts vulnerable to the attacks⁴. We use two different worm scanning models: random and local-preferential scanning. The random-scanning worm selects destinations uniformly at random from the entire address-space. The local-preferential scanning worm selects destinations to infect based on topological locality. For example, an infected host will try and infect hosts within the same source AD with a higher probability, and this probability decreases as a function of topological proximity (measured in terms of hop-count).

For each of the following experiments, we perform 5 independent runs and report the mean. The standard deviations observed across the different runs were small and are not reported for brevity. We set the traffic normalization factor η in the distributed random moonwalks to 10^{-3} . For each experiment, we use the methodology in [5] to select the optimal sampling window size, and set the maximum path length $d(i) = 20$ for each local moonwalk inside a domain.

1) *Accuracy and Overhead*: We first compare the performance of the distributed random moonwalks with a centralized approach, by varying the worm propagation rate with a random scanning attack. The centralized approach [5] assumes access to global traffic records from all ADs in the network. The collaborative, distributed protocol is as described in Section IV. The performance of the centralized and distributed approaches is measured in terms of the detection accuracy, defined as the fraction of the top Z returned flows that are actually causal flows. With the distributed operations, these Z flows are identified using the *Global Merge* procedure presented in Section V-B.

Figure 5 shows that the detection performance achieved by the distributed random moonwalks closely approximates that of the centralized algorithm. The worm scanning rates are presented relative to the mean of the normal per-host traffic flow rates. Both the centralized algorithm and the collaborative approach achieve high detection accuracy regardless of the worm rates. For the rest of our evaluations, we only present results from a single worm rate, with a scanning rate equal to six times the mean normal traffic rate.

We show in Figure 6, the detection accuracy and communication overhead as a function of H , the maximum number of AD-hops that moonwalks may traverse. Figure 6 (a) shows how the performance approaches that achieved by the centralized algorithm, as H increases. With a small H , the distributed version is less accurate in identifying causal flows. The reason is that the initial steps of the moonwalks start at random flows, and non-causal attack flows are inherently more numerous than causal flows. As ADs exchange intermediate results over time, the majority of the returned flows are causal flows, and the detection accuracy converges after five or six hops. The communication overhead (Figure 6 (b)), in terms of the total number of flow records exchanged in the network, is on the order of 10^{-3} of the total traffic records in the network. While the overhead is a monotonically increasing function

⁴The random moonwalk algorithm is robust to the fraction of vulnerable host population, as suggested in [5]. We do not duplicate the results here due to space.

of the accuracy, it does not increase significantly beyond the convergence limit of five or six hops.




Fig. 5. Global causal flow detection accuracy with the distributed protocol in the NFA compared with the centralized implementation in a unified network model. We vary the worm rate (relative to the mean normal traffic rate), and select the top $Z = 100$.

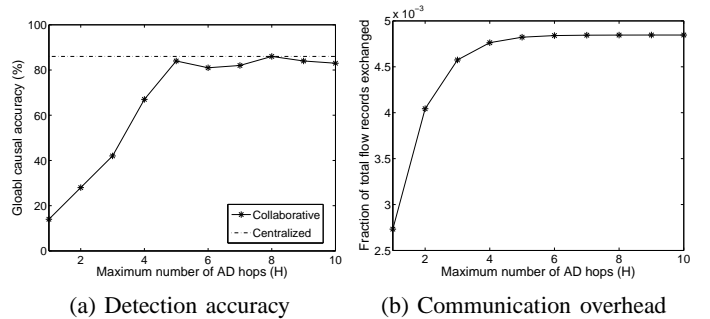


Fig. 6. Accuracy and communication overhead increase as a function of the maximum AD-hops H a moonwalk can traverse, ($Z = 100$).




Fig. 7. Distribution of the communication overhead in terms of the fraction of total exchanged flow records received and the number of ADs contacted.

We further examine how the communication overhead distributes across the participating ADs in this case. Figure 7 (a) shows the fraction of the total exchanged flow records that each participating AD receives in the course of the distributed protocol. The distribution is highly skewed, with a small number of ADs receiving a large fraction of the total flow records exchanged in the network. Figure 7 (b) shows a similar trend in terms of the number of ADs that each participating AD needs to communicate with in the protocol. Further investigation shows that the ADs that receive a high volume of flow-records and that need to communicate with a large number of fellow participants are ADs that own large

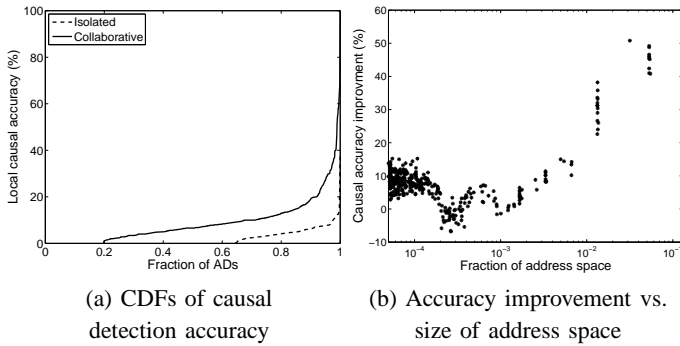


Fig. 8. Causal flow detection accuracy with and without collaboration for each AS ($Z = 100$).

chunks of the overall address space (and also contain a large fraction of infected hosts). Thus a lot of flows traversed by random moonwalks originated from these domains.

2) *Benefits of Participation*: For a federated architecture to be viable, each participating AD must benefit by collaborating with other ADs. Here, we consider the case where ADs collaborate in the distributed random moonwalks, but subsequently choose to perform only the *Local View* (Section V-A) for post-processing. This is compared against an isolated investigation scenario, where each AD runs a centralized random moonwalk algorithm on locally available traffic data, with optimally chosen parameters.

We consider three measures to quantify the perceived local benefit. The first metric is the (local) causal detection accuracy. Figure 8 (a) compares the CDFs of the causal flow detection accuracy across all the ADs with and without collaboration, where each AD returns the top $Z = 100$ after the protocol. With isolated forensics, only 40% of the ADs are able to identify causal flows, while through collaboration more than 80% of ASes can successfully identify causal flows. Overall, we find that there is a substantial improvement in the local detection accuracy through participation over an isolated execution. Figure 8(b) correlates the improvement in the local causal flow detection accuracy (defined as $Accuracy_{LocalView} - Accuracy_{Isolated}$) with the size of the address-space owned by each AD. We observe the larger ADs gain a lot more by collaboration than the smaller ADs.

The other two metrics quantify the ability of each AD to identify the initial infected hosts and the initial attack entry point within its network. Figure 9 (a) depicts the distribution of the accuracy in identifying initial infected hosts. Each AD selects the first 50 internal hosts from among the set of top-frequency flows it has available to it after the distributed protocol⁵. The accuracy of initial host identification is then defined as the fraction of these returned hosts that are among the 50 hosts that are actually infected earliest in time within the AD. And such detection accuracy is improved across ADs in general when they participate the NFA.

ADs may also be interested in identifying the first few

⁵For many of the smaller ADs, there may not be 50 hosts available to return. In such cases, these ADs simply return all internal hosts that appear along moonwalks

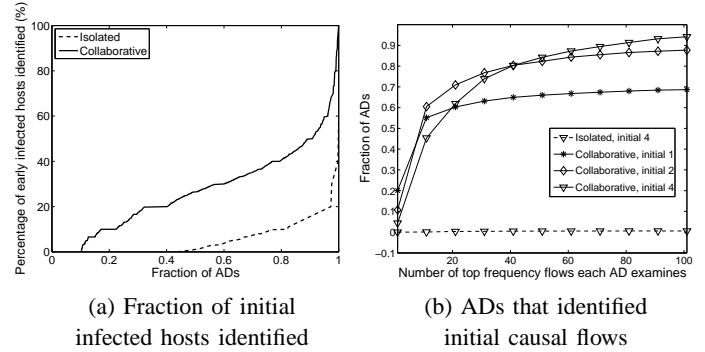


Fig. 9. The ability to identify the initial infected hosts and the attack entry point with and without collaboration for each AD.

causal flows where the infection first entered its network. Figure 9 plots the fraction of ADs that successfully identified at least one of the first several entry points, by examining various number of top frequency flows returned with *Local View*. The vast majority of ADs (more than 60%) need to examine only a small number of flows (fewer than 20) before they can detect the first one or two causal flows. These results further corroborate the observations in Figure 8 (a)—collaboration significantly boosts detection performance, and this serves as a strong participation incentive for ADs.

3) *Partial Deployment*: There are two natural concerns with incremental deployability: performance and participation incentives. We would like the causal detection accuracy under partial deployment to be comparable to that under complete deployment. Similarly, ADs should be able to observe participation benefits (as defined by the metrics discussed above) even with partial deployment. We consider two partial deployment scenarios, by varying the set of ADs that do not participate based on AD degree and address space size. In each scenario, we consider the case where the first k large domains collaborate and the case where the first k large domains are missing from the NFA, and vary the parameter k .

Global Accuracy: Figure 10 plots the global causal flow

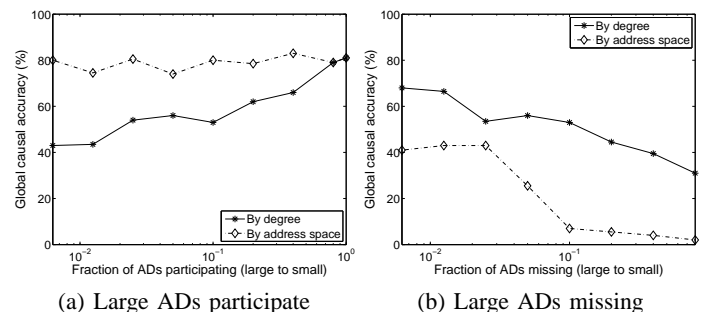


Fig. 10. Global detection accuracy under different missing scenarios, with the random scanning worm. We use the best recovery strategy, i.e., *routing horizon*-based recovery.

detection accuracy by varying the number of large ADs participating/missing, with the random scanning worm model. We find that when large ADs collaborate, the global causal

accuracy is high. In particular, when ADs with largest address space participate in the NFA, the performance is comparable to full deployment. However, when these large ADs are missing from the NFA, the performance degrades significantly, with the global accuracy going down from 80% to 30-40%, even when the rest majority of small ADs collaborate.

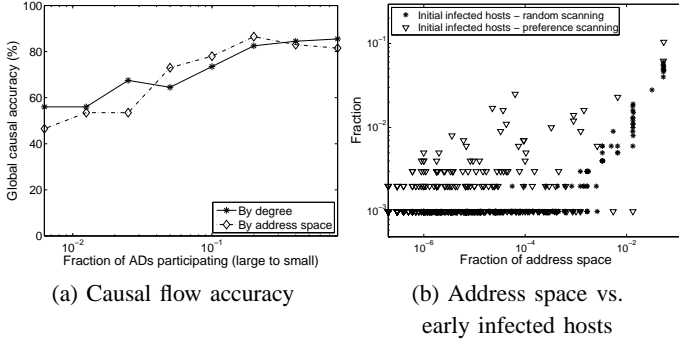


Fig. 11. Impact of worm scanning model on partial deployment performance. We use the best recovery strategy, i.e., *routing horizon*-based recovery.

To investigate the reason, we vary the worm model to local preferential scanning attack, and show how the global accuracy increases with the number of large ADs collaborating (Figure 11 (a)). Under this worm model, the performance is not as good as the random scanning attack, when the ADs with largest address space participate. The moonwalk algorithm suggests that the global accuracy is tightly coupled with the initial causal flows in the attack. Hence, we examine the correlation between the AD address-space size and the fraction of initial infected hosts (defined as the 200 hosts that are infected earliest in time in the network) that are actually present within each AD (Figure 11 (b)). For the random scanning worm, ADs with larger address space also observe larger fractions of early infected hosts within their domains, and thus a larger fraction of initial causal flows for the random moonwalk to return successfully. With a local-preferential attack, larger ADs do not necessarily observe many initial infected hosts, hence the corresponding decrease in the number of initial causal flows available to them.

Figure 12 studies the impact of different recovery strategies (discussed in Section VI) on the global detection accuracy, using the random scanning worm. Intuitively, we expect the *Discard* strategy to have the worst performance, and the *Routing-horizon* strategy to have the best performance. However, we find in Figure 12 (a) that when the ADs with large address-space ownership collaborate, the recovery strategies have negligible impact on the overall performance, as these ADs already observe the most number of initial infected hosts and a large amount of attack traffic with the random scanning worm model. With only a small number of large degree ADs participating, the recovery strategies do have significant impact on performance (Figure 12 (b))⁶.

The above results suggest that both attack models and

recovery strategies may have impact on performance with partial deployment. Further understanding their implications under more diverse partial deployment scenarios is a topic of ongoing work.

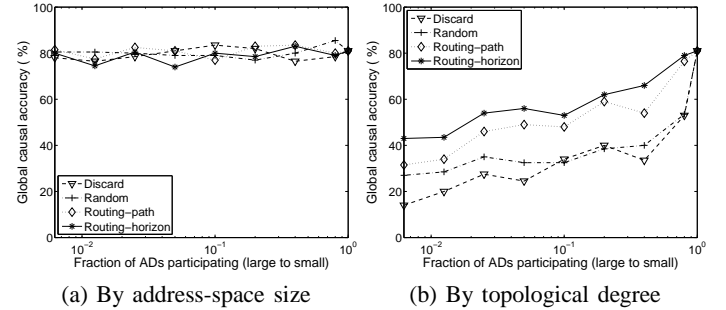


Fig. 12. Impact of recovery strategies on performance, with large ADs participating

Participation Benefits: Next, we investigate the participation benefits of individual ADs under partial deployment, once again using the random-scanning worm with normalized attack-rate of six. Figures 13 (a) and (b) show how different partial deployment scenarios affect the participation benefits. We sort ADs based on address-space ownership and degree (from large to small), respectively, and categorize ADs based on their ranks in the sorted order. We then compute the average local causal flow detection accuracy ($Z = 100$) of ADs within each category, and compare the performance achieved by varying the number of large ADs participating. Overall, participation yields benefit for all categories of ADs. Specifically, we find that the large ADs perceive similar benefits to the full collaboration case, even when the rest of the 90% of the ADs do not participate in the NFA. We also observed earlier (Figure 8 (b)) that the perceived performance benefit is greatest for the large ADs under the NFA. These observations bode well for the deployability of the NFA, since it suggests that, for the vast majority of attacks known today which use random scanning for destination selection, domains that have the greatest impact on global performance also have very strong incentives to participate in the NFA.

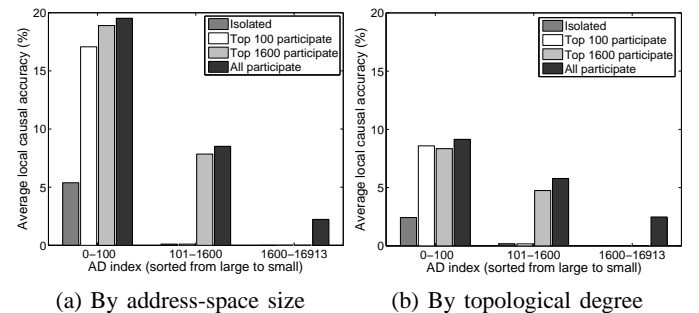


Fig. 13. Benefits of participation under partial deployment, assuming the large ADs participate

⁶Our analysis using the Routeviews dataset suggests that the correlation between the address-space ownership and AS-degree is not very pronounced

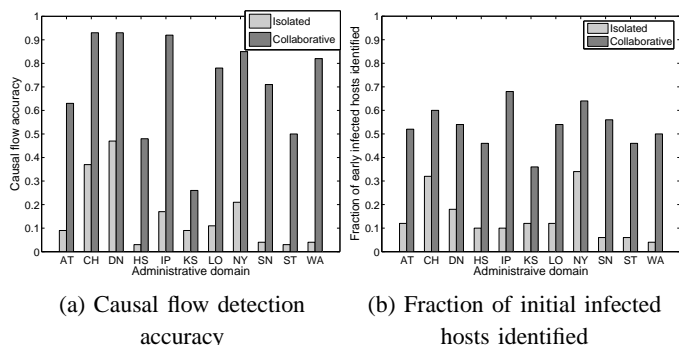


Fig. 14. Internet-2 study results. Each node in the federation corresponds to a router in the Abilene network, denoted by a two-letter city code. We compare the performance of a collaborative effort against investigation in isolation.

B. Internet-2 Trace Study

We also independently validated our results over a one-hour traffic snapshot from each of the 11 routers in the Abilene backbone. The data set consists of sampled unidirectional flow records (with 1 in 100 sampling). IP addresses in the dataset were anonymized by zeroing out the last 11 bits of the source and destination IP address fields. In our evaluation, each of the 11 routers serves as an independent domain for our 11-AD federation over the Abilene topology [6]. We map each observed /21 prefix range to a unique end-host, and subsequently assign the hosts as being “owned” by the router that most frequently observes traffic to/from that prefix range.

We generate synthetic attack using a random-scanning worm, with each infected host scanning at 0.5 scans per second and 10% of the hosts being vulnerable. Given that our goal is not to perform attack analysis over Internet2, but rather to validate our results using a more realistic workload of background traffic, we believe this dataset can serve this purpose, and our results are not biased by these transformations.

We first examine the local causal flow detection accuracy, shown in Figure 14 (a). For all the domains, we find that there is a substantial improvement in the detection accuracy through participation over an isolated execution. For example, the performance of domain *IP* improves from around to 20% to 80%. Even though the absolute accuracy is relatively low for domain *KS* compared with other domains after collaboration, the relative performance increase is at least 80%.

The second measure we have discussed is the ability to identify the set of hosts that are infected earliest in time. Each domain identifies the first 50 internal hosts by examining the top frequency flows in order. Figure 14 (b) shows the fraction of these 50 hosts which are among the 50 internal hosts that are actually infected earliest in time for each AD. We observe that compared with isolated execution of random moonwalks, the ability to find the initial infected hosts increases significantly with collaboration.

Our evaluations with both the simulation study, and the Internet2 data suggest that each AD’s local investigation capability can benefit significantly by collaborating with other ADs in the network.

VIII. CONCLUSIONS

We proposed a protocol for performing distributed random moonwalks, which can enable participating ADs to identify local attack entry points, and can provide additional functionality to pinpoint global attack origins. Our design is suitable for deployment in an Internet-like federation for three reasons. First, the protocol operates with limited information disclosure, without requiring participants to reveal traffic records to other participants that would otherwise not be available. Second, participating domains realize enhanced local attack investigation, and thus receive substantial incentives for cooperation. Third, the framework is incrementally deployable and can handle non-participation and missing data gracefully. We believe that our design and results provide a technical basis of a Network Forensic Alliance (NFA), a collaborative effort involving multiple ADs to provide network-wide and localized forensic capabilities.

REFERENCES

- [1] J. Jung, V. Paxson, A. W. Berger, and H. Balakrishnan, “Fast Portscan Detection Using Sequential Hypothesis Testing,” in *Proc. of IEEE Symposium on Security and Privacy*, 2004.
- [2] D. Moore, C. Shannon, G. Voelker, and S. Savage, “Internet Quarantine: Requirements for Containing Self-Propagating Code,” in *Proc. of IEEE INFOCOM*, Apr. 2003.
- [3] H. A. Kim and B. Karp, “Autograph: Toward Automated, Distributed Worm Signature Detection,” in *Proc. of 12th USENIX Security Symposium*, 2004.
- [4] “Fingerprint Sharing Alliance,” <http://www.arbor.net/fingerprint-sharing-alliance.php>.
- [5] Y. Xie, V. Sekar, D. Maltz, M. K. Reiter, and H. Zhang, “Worm Origin Identification Using Random Moonwalks,” in *Proc. of IEEE Symposium on Security and Privacy*, 2005.
- [6] “The Internet2 Abilene Network,” <http://abilene.internet2.edu>.
- [7] Y. Zhang and V. Paxson, “Detecting Stepping Stones,” in *Proc. of 9th USENIX Security Symposium*, 2001.
- [8] D. L. Donoho, A. G. Flesia, U. Shankar, V. Paxson, J. Coit, and S. Staniford-Chen, “Multiscale Stepping-Stone Detection: Detecting Pairs of Jittered Interactive Streams by Exploiting Maximum Tolerable Delay,” in *Proc. of The 5th International Symposium on Recent Advances in Intrusion Detection (RAID)*, 2002.
- [9] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, “Practical Network Support for IP Traceback,” in *Proc. of ACM SIGCOMM*, 2000.
- [10] H. Burch and B. Cheswick, “Tracing Anonymous Packets to Their Approximate Source,” in *Proc. of USENIX LISA Systems Administration Conference*, 2000.
- [11] A. C. Snoeren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, S. T. Kent, and W. T. Strayer, “Hash-Based IP Traceback,” in *Proc. of ACM SIGCOMM*, 2001.
- [12] A. Kumar, V. Paxson, and N. Weaver, “Exploiting Underlying Structure for Detailed Reconstruction of an Internet-scale Event,” *Proc. of Internet Measurement Conference*, 2005.
- [13] M. A. Rajab, F. Monrose, and A. Terzis, “Worm Evolution Tracking via Timing Analysis,” in *Proc. of Workshop on Rapid Malcode (WORM)*, 2005.
- [14] J. Bethencourt, J. Franklin, and M. Vernon, “Mapping Internet Sensors with Probe Response Attacks,” in *Proc. of 14th USENIX Security Symposium*, 2005.
- [15] Y. Azar, A. Fiat, A. Karlin, F. McSherry, and J. Saia, “Spectral Analysis of Data,” in *Symposium on Theory of Computing*, 2001.
- [16] S. Acharyya and J. Ghosh, “Outlink estimation for pagerank computation under missing data,” in *Proc. of the 13th international World Wide Web conference*, 2004.
- [17] Z. Mao, J. Rexford, J. Wang, and R. Katz, “Towards an Accurate AS-Level Traceroute Tool,” in *Proc. of ACM SIGCOMM*, 2003.
- [18] “The Route Views Project,” <http://www.routeviews.org>.