Volume 1, No. 3, Sept-Oct 2010
ISSN No. 0976-5697
International Journal of Advanced Research in Computer Science
RESEARCH PAPER
Available Online at www.ijarcs.info

Crosstalk: A Scalable Crossprotocol Monitoring System For Anomaly Detection

Iamuna S, Hari Kishore
School of IT and Engg.
VIT University,
Vellore, India
Iamunas66@yahoo.co.in

Jeyanthi. N.*
School of IT and Engg.
VIT University,
Vellore, India
njeyanthi@vit.ac.in

Abstract: As there is a relentless growth in IP, Monitoring is important both in the case of operations of a network and the services that run on it. Operators on the network perform monitoring on various purposes such as traffic engineering, Quality of Service, security and detection of faults and misconfigurations. Monitoring and detection of anomalies in a network is a very challenging problem. The most common forms are botnet detection and Denial of Service attacks. Many of these anomalies use several protocols to carry out their works. Botnets uses IRC to control and SMTP to send out spam. Another example is VoIP where calls tend to be split into signalling and media traffic as in the case with SIP and RTP. These have to be detected using cross-protocol correlation. In addition to cross-protocol correlation, monitoring needs to be done in a distributed fashion, since traffic from a particular attack or misconfigurations may cross different monitoring points in the network. While previous work has looked into the area of cross-protocol detection [2] , it has focused on single-point solutions, and so did not scale nor could it correlate attack traffic traversing more than one monitoring point. This raises a serious scalability issues while designing, which needs to monitor large quantities of traffic and also to aggregate results to provide network wide anomaly detection. In this paper we introduce Crosstalk - a scalable and efficient protocol to detect anomalies using cross-protocol correlation in a distributed fashion. Evaluation of detecting anomalies in distributed system does not show how it would scale under heavy load. For the purposes of evaluating Crosstalk's scalability and performance, we focused on SIP-based VoIP attacks. Here, we used network simulator and evaluate the performance. Based on CDR (Call Detail Record) and on the traffic the results simulated and anomaly is detected. The probes monitor the network and collect the data in bloom filters and export the measurements to the mediators and collector in DAT tree structure. Based on the application the results were simulated.

Keywords: Cross talk, CDR, Denial of Service attack, SIP, RTP, Voice over IP, DAT.

I. INTRODUCTION

Monitoring large networks in order to detect anomalies is inherently difficult for several reasons. Many of these anomalies require cross-protocol correlation in order to be detected. Botnet, for example, often use several protocols to coordinate activities and to carry out attacks (e.g., IRC for control and SMTP to send out spam). An important place where cross-protocol correlation needed is VoIP in which calls tend to be split into signaling and media traffic, as is the case with SIP and RTP. In addition to cross-protocol correlation, monitoring needs to be done in a distributed fashion, since traffic from a particular attack or misconfigurations may cross different monitoring points in the network. Making matters more difficult is the relentless growth of IP traffic volume, nearly doubling every two years.

Crosstalk, a scalable architecture that gathers data from a potentially large set of distributed monitoring probes, and performs cross-protocol correlation to detect network anomalies. While previous work has looked into the area of cross-protocol detection it has focused on single-point solutions, and so did not scale nor could it correlate attack traffic traversing more than one monitoring point.

II. CROSS TALK'S ARCHITECTURE

Crosstalk's architecture consists of three main features that allow it to perform distributed detection in a scalable way: leveraging Distributed Aggregation Trees (DATs), taking advantage of probabilistic data structures (e.g., Bloom filters), and using a novel mechanism called *backtracking* (BT).

A. Distributed Aggregation Trees

The simple approach of exporting data from several monitoring probes to a centralized location clearly does not scale. In order to cope with this scalability issue, efforts both in the research and standardization communities have focused on creating tree-based hierarchies, whereby monitoring *probes* export measurements to intermediate nodes called *mediators*. These in turn perform some sort of data reduction operation (e.g., aggregating packet counts) and export the results up the tree hierarchy. In the final step the root, which is a special mediator called a *collector*, stores the aggregated results. Ideally we would like to have a way of deriving such a tree-based topology dynamically in order to adapt to traffic conditions.

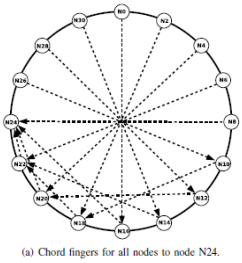

(a) Chord fingers for all nodes to node N24.

Figure 1. Chord fingers

The basic insight behind a DAT is that Chord's fingers already provide a tree structure. In order to illustrate this,figure 1(a) shows a regular Chord network with dotted lines representing the path from each node to node N24 (perhaps the responsible node for a particular key); figure 1(b) then shows

these same connections but this time drawn as a tree. As canbe seen, for any given key, Chord naturally builds a tree rooted at the node responsible for that key. In this way, each key has its own DAT, with all the DATs sharing the same peer-to-peer infrastructure. Within each DAT, intermediate nodes (i.e., all nodes except the leaves of the tree) can aggregate data as it travels towards the root, thus providing scalability.

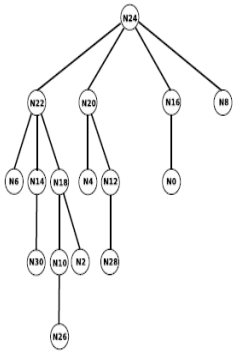

(b) Same Chord fingers, this time shown as a tree.

.

B. Probabilistic Data Structures

Clearly nodes in the DAT will need to export information about the monitored data, and this will consume bandwidth. Another consideration for real-time monitoring and detection is being able to perform the cross-protocol correlation quickly. To achieve both of these goals we rely on probabilistic data structures, and more specifically Bloom filters (BFs).

Our attack detection application works as follows: each of the probes monitors all the on-going SIP and RTP traffic. The SIP messages are parsed and, for each call, the two end-points of the media traffic are located by examining the SDP data; as for the RTP traffic, the two end-points simply correspond to the source and destination addresses of the messages.

Our method assumes the probes to synchronously and periodically export their probabilistic summaries of the monitored traffic.

The monitored data is used to fill two Bloom filters:

• A Bloom filter for keeping track of the end-points of the media traffic corresponding to SIP calls that have been terminated (or redirected) within the last measurement period.

• A Bloom filter for keeping track of the end-points of the RTP sessions that have been terminated (or redirected) within the last measurement period.

Clearly, the hash functions associated with these BFs must be the same so that the RTP and SIP endpoints associated with the same call are hashed into the same bit positions. Once these BFs are created, they are exported to the nearest mediator (i.e., the parent of the probe in the DAT). The mediator then joins all of the SIP BFs and all the RTP BFs received from its sons by performing a bitwise "OR", thus obtaining two summarized BFs that it forwards to its own mediator.

C. Backtracking

While some applications might be content to only receive the summarized data from a DAT's collector, others will use such summarized data as a trigger for retrieving more detailed information (e.g., packet headers) at the monitoring probes, perhaps to determine the cause of the trigger. In addition, Bloom filters carry a low but non-negligible probability of false positives, and so we need a way to verify whether a result is valid or just a false positive. In order to accomplish these goals we introduce a mechanism called *backtracking*. The idea behind it is simple: when exporting Bloom filters to nodes in the DAT, keep a copy of them locally so that the system can track back to the original probes that monitored the traffic.

The detection of the anomalous behavior is achieved by a node in the DAT performing a bit-wise "XOR" of the RTP and SIP BFs: if two bits in the same position are different, that means that either the data stream or the control stream has not been terminated. In that case, all of the node's children receive a backtracking request which includes the indices of the unmatched bits (i.e., the set bits that appeared in one BF but not the other). Each intermediate node then checks such bits against its cached aggregated BFs, and, if at least one among those is set, it propagates the BT request to its children. Such a procedure is repeated recursively until all the probes which have logged relevant information are reached.

Figure 2 shows the process in greater detail. Probes P0 and P3 monitor traffic and export data about it in the form of Bloom filters, depicted as a set of squares with each square representing a bit in the filter (note that the figure is simplified for explanatory purposes: normally an entry in the Bloom filter would use up several bits, and more than one Bloom filter would be used to represent the protocols to be correlated). In addition, probes, as well as mediators, keep a local copy of exported Bloom filters, shown in the figure in grey. As the exported filters travel up the tree, mediators perform a bitwise operation to combine the filters, which eventually reach the collector C. Upon receiving all the combined data, C correlates Bloom filters from different protocols and, depending on the application, triggers a backtracking request to all its immediate mediators, in this case M4 and M5. The request includes the collector's Bloom filter (shown in white), which the mediators use to compare it with their locally stored state by performing a bitwise operation: if the number of set bits in the resulting filter is higher than a user-defined threshold, the backtracking request is propagated to all of the mediator's children; otherwise, no relevant probes exist in this area of the DAT and the backtracking process finishes. Eventually the back tracking message arrives at the probes, in this case P0and P3. In section IV we provide an evaluation of the costs associated with this mechanism and show its applicability even in large networks.

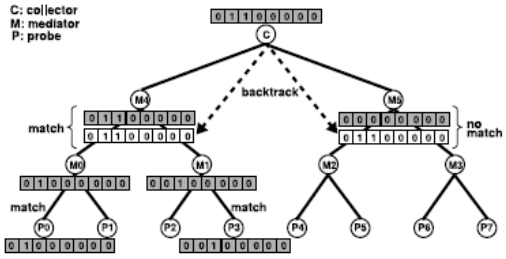

Figure 2. Network moniotr

III. EVALUATION

In this section we provide extensive simulation results to show the performance of Crosstalk, and in particular that of the VoIP attack detection application. Please note that throughout this section we use the term report to mean the Bloom filters exported between probes and mediators as a result of the monitoring and aggregation process.

A. Setup

In order to assess the performance of our solution, we evaluated several performance parameters through extensive simulations. In greater detail, we extended the Oversim overlay network simulator [14] by implementing a new application module with Crosstalk's basic functionality and which runs on top of the Chord. The input to the simulation consists of Call Data Records (CDRs), in order to match the format used by our VoIP data set (a CDR is a short record of a VoIP communication, including fields like caller, callee, and call duration). To have control over their distribution, CDRs are fed to the simulated monitoring probes by a centralized CDR dispatcher module: each node is assigned a given range of the overall hash ID space and each CDR is handed over to the responsible node (based on its source address).

In order to simulate the fact that RTP and SIP traffic for the same call may traverse different paths, two separate copies of the same CDR, representing in turn the RTP and SIP traffic associated with a given call, are assigned to two distinct probes by using two independent hash functions. Such a choice is rather conservative, since in a significant fraction of the real cases, the two traffic streams are likely to follow the same path, but this approach is still useful to show that our system can cope with even this extreme case. It is worth noting that the simulated probes are actually nodes on the DATs, meaning that they can act as probes for one key but as mediators (and even collector) for others simultaneously. Regarding CDR generation, we took two approaches. First, we generated CDRs randomly by setting the timestamps and the call durations according to a Poisson process (such a simple model has been extensively used in the field of telephone traffic measurement). The purpose here was to be able to effectively tune and change the simulation parameters to show the performance of the system. In the second approach we relied on an extensive data set gathered from a large VoIP operator in order to demonstrate Crosstalk's applicability to a real world scenario. In both cases we modified the CDRs at a certain rate (set as a percentage of the total CDRs) in order to simulate malicious calls.

IV. PERFROMANCE ANALYSIS

In this section we present simulation results based on generated CDRs in order to assess the system's performance and scalability. Crosstalk's VoIP application depends on a number of different parameters:

- **Bloom filter size**, which affects several factors such as the missed detection rate, the bandwidth consumed and how much state nodes in the DAT keep.
- **The call rate**, in other words, how much traffic the system needs to monitor, export, and correlate.
- **The measurement interval**, which determines how long the probes keep data locally before exporting (longer intervals result in lower overheads but increase the detection delay).
- **The anomaly rate**, or percentage of malicious calls, which increases the costs associated with backtracking requests.
- **The number of probes**, equal in our case to the number of nodes in the p2p system, affecting the DAT's topology and therefore the messaging overhead, the amount of aggregation, and the detection delay.

A.Bloom filter size and call rate:

For the first experiment we took a look at the first two parameters and their relationship to false negatives and positives. In other words, given a certain call rate, how would an operator deploying Crosstalk dimension the Bloom filter size (which affects things like bandwidth consumption) so that the false negative and positive rates are relatively low? To this end, consider that missed detections (i.e., false negatives) happen when a collision in one BF causes a match with a "true" set bit in the other BF, resulting in the "XOR" matching operation to return 0 (the misdetection). Because the cause is the collision within a BF, all the well-known results about BF performance evaluation and dimensioning apply to our system. In particular, as the number of keys in the BFs equals the call rate λ times the measurement period *T*, the missed detection (md) probability can be expressed as:

$$P(md) = (1-(1- 1M)\ K\ \lambda\ T\ )K \sim (1-e-K\lambda TM\ )K$$

Where *M* stands for the BF size (in bits) and *K* for the number of hash functions (which is set to an optimal value depending on the other parameters).

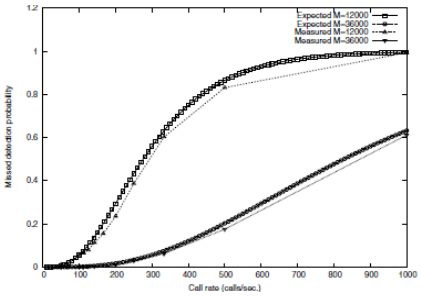

To get a feel for the system's performance we rely on this model and on our CDR database, which shows a peak call rate well below 100 calls per second. Even if we assume that since more and more users are migrating from PSTN to VoIP such a figure will increase in the future by an order of magnitude, our system can handle the resulting traffic volume (1,000 calls/sec.): exporting data every 10 seconds and using 17KB-wide BFs yields a target missed detection probability of 10−3, while using 34KB-wide BFs yields a target missed detection probability of 10−6. For a measuring infrastructure made up of 1,000 probes the *total* reporting traffic adds up to only few MB/sec.

B.Measurement interval

If the BF size does not vary, a longer measurement period implies a larger number of keys in the BF, and, in turn, an increased missed detection probability. On the other hand, of course, this involves lower bandwidth consumption, as data summaries are exported less frequently. Depending on the operational constraints, the previously presented mathematical model allows finding out a good trade-off; we do not present more extensive results here due to space constraints.

C.Anomaly rate

The next parameter we looked at was the anomaly rate, and in particular how it affects the costs related to backtracking (BT). Backtracking is triggered by either a detected anomaly or a false positive. Assuming a well dimensioned system with a low false positive rate (e.g., less than 1%) and no anomalies, simulation results show about an order of magnitude difference between export messages and BT messages. Arriving at more precise figures is difficult since the actual number of backtracking messages generated depends on the number of probes which have to be reached by a BT request and on the topology of the tree. Having said that, we ran a simulation to get a feel for the effects of the anomaly rate on the system, and in particular the cost of backtracking.

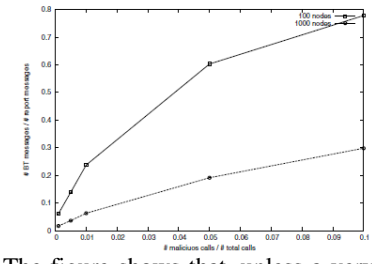

The figure shows that, unless a very unrealistic scenario is assumed (a network where one in ten calls is malicious), the fraction of BT messages is small (usually an order of magnitude smaller) with respect to the number of reports, which proves that the BT mechanism can locate the relevant probes without flooding the DAT with messages. Further, the behavior of the system improves as the number of nodes increases. These figures can be even further improved by reducing the size of each BT message. Observe that the Bloom filter obtained through a bit-wise "XOR" of the aggregated SIP and RTP reports must have a very limited number of set bits (in fact, the number of such bits should be lower than the number of malicious calls times the number of hash functions), which lends itself to compression. In order to effectively compress such a "sparse" bitmap, it is sufficient to include the indices of the set bits within the BT message: the resulting message size would be roughly some dozens of bytes, which is negligible when compared to the bandwidth consumed by the reporting messages.

D.Number of nodes

The number of probes does not affect the accuracy of our system (that, in fact, depends on the overall number of monitored calls) but rather the aggregation and backtracking delay and the overall bandwidth consumption. The former depends on the depth of the tree, which, in turn, grows logarithmically with the number of probes. On the other hand, the overall bandwidth consumption due to the report messages grows linearly with the number of probes (each additional probe corresponds to an additional edge on the tree, which, in turn, corresponds to an additional report being transmitted). As for the BT requests, their amount depends on several variables, but we already showed their bandwidth consumption to be negligible with respect to that of the report messages.

V. CONCLUSIONS AND FUTURE WORK

We have presented Crosstalk, a scalable and distributed monitoring system for detecting cross-protocol anomalies. We have implemented a VoIP attack detection application over it and presented extensive simulation results on a large VoIP data set. In addition, we used a mathematical model to show that Crosstalk performs well even when presented with much higher loads than those conveyed by current VoIP infrastructures.

The results confirm that Crosstalk can scale to a very large number of monitoring probes, deal with a large call rate of 1,000 calls/sec and a high percentage of anomalous calls, all while using small Bloom filter sizes of only dozens of KB. The system can be easily tuned to achieve arbitrarily small missed detection rates with a limited increase in terms of overhead. Moreover, in case missing an anomaly is not acceptable, a slight change in the system layout allows Crosstalk to fulfill such a requirement. One of the topics we did not discuss due to space constraints is tree topologies. The DATs we used relied on Chord's normal routing algorithm, which can result in unbalanced trees, especially when the DAT contains a large number of nodes. Towards a solution, previous work [16] modified Chord to provide (almost) balanced binary trees. While certainly an improvement, what we would like is not only to have a mostly balanced tree, but also the ability to control its depth; in other words, controlling the trade-off between scalability through aggregation (achieved with deeper trees) and aggregation delay and messaging overheads (reduced by using shallower trees). We are currently working on an algorithm to achieve this.

VI. REFERENCES

[1] Symantec Corporation, "Internet Security Threat Report Volume XI," http://www.symantec.com/enterprise/threatreport/index.jsp, March 2007.

[2] Y.-S. Wu, S. Bagchi, S. Garg, N. Singh, and T. Tsai, "Scidive: A stateful and cross protocol intrusion detection architecture for voice-over-ip environments," in DSN '04: Proceedings of the 2004 InternationalConference on Dependable Systems and Networks.

[3] B. Barry and A. Chan, "Towards intelligent cross protocol intrusion detection in the next generation networks based on protocol anomaly detection," in The 9th International Conference on Advanced CommunicationTechnology, 2007, pp. 1505–1510.

[4] P. Yalagandula and M. Dahlin, "A scalable distributed information management system," in SIGCOMM '04: Proceedings of the 2004conference on Applications, technologies, architectures, and protocolsfor computer communications. New York, NY, USA: ACM, 2004, pp. 379–390.

[5] R. Zhang, X. Wang, X. Yang, and X. Jiang, "Billing attacks on sipbasevoipsystems," in WOOT '07: Proceedings of the first USENIXworkshop on Offensive Technologies. Berkeley, CA, USA: USENIX Association, 2007, pp. 1–8.

[6] M. Cai and K. Hwang, "Distributed Aggregation Algorithms with Load-Balancing for Scalable Grid Resource Monitoring," Parallel andDistributed Processing Symposium, International, vol. 0, p. 123, 2007.

[7] Sven Ehlert, Dimitris Geneiatakis, Thomas Magedanz, "Survey of network security systems to counter SIP-based denial of service attacks," University of the Aegean, Greece.

[8] Chung-Hsin Liu, Chun-Lin Lo, "The simulation for the VoIP DDoS attacks," 2008 , International Conference on MultiMedia and Information Technology

