International Journal of Communication Networks and Security

Volume 1 | Issue 2 Article 11

2011

Dynamic Detection of Packet Losses by CRDP

N. Sailaja
Narasaraopeta Engineering College Department-CSE, Narasaraopet, AP, sailaja.mtech2011@gmail.com

K. LakshmiNadh
Narasaraopeta Engineering College Department-CSE, Narasaraopet, AP, l_nadh@yahoo.co.in

Follow this and additional works at: https://www.interscience.in/ijcns

 Part of the Computer Engineering Commons, and the Systems and Communications Commons

Recommended Citation

Sailaja, N. and LakshmiNadh, K. (2011) "Dynamic Detection of Packet Losses by CRDP," *International Journal of Communication Networks and Security*: Vol. 1 : Iss. 2 , Article 11.
Available at: https://www.interscience.in/ijcns/vol1/iss2/11

This Article is brought to you for free and open access by Interscience Research Network. It has been accepted for inclusion in International Journal of Communication Networks and Security by an authorized editor of Interscience Research Network. For more information, please contact sritampatnaik@gmail.com.

Dynamic Detection of Packet Losses by CRDP

N.Sailaja, *K. LakshmiNadh*
Narasaraopeta Engineering College
Department-CSE, Narasaraopet, AP
sailaja.mtech2011@gmail.com

Abstract— ***TCP has provided the primary means to transfer data reliably across the Internet, however TCP has imposed limitations on several applications. Measurement and estimation of packet loss characteristics are challenging due to the relatively rare occurrence and typically short duration of packet loss episodes. While active probe tools are commonly used to measure packet loss on end-to-end paths, there has been little analysis of the accuracy of these tools or their impact on the network. The main objective is to understand the problem of detecting whether a compromised router is maliciously manipulating its stream of packets. In particular to this concern a simple yet effective attack in which a router selectively drops packets destined for some Victim. Unfortunately, it is quite challenging to attribute a missing packet to a malicious action because normal network congestion can produce the same effect. Modern networks routinely drop packets when the load temporarily exceeds their buffering capacities. Previous detection protocols have tried to address this problem with a user-defined-threshold: too many dropped packets imply malicious intent. However, this heuristic is fundamentally unsound; setting this threshold is, at best, an art and will certainly create unnecessary false positives or mask highly focused attacks.***

Index Terms—Internet dependability, intrusion detection and tolerance, distributed systems, reliable networks, malicious routers.

1 INTRODUCTION

The Internet is not a safe place. Unsecured hosts can expect to be compromised within minutes of connecting to the Internet and even well-protected hosts may be crippled with denial-of-service (DoS) attacks.

In this paper, we develop a compromised router detection protocol (CRDP) that dynamically infers the precise number of congestive packet losses that will occur. Once the congestion ambiguity is removed, subsequent packet losses can be safely attributed to malicious actions. We believe our protocol is the first to automatically predict congestion in a systematic manner and that it is necessary for making any such network fault detection practical. Once a router has been compromised in such a fashion, an attacker may interpose on the traffic stream and manipulate it maliciously to attack others—selectively dropping, modifying, or rerouting packets.

Several researchers have developed distributed protocols to detect such traffic manipulations, typically by validating that traffic transmitted by one router is received unmodified by another [3], [4]. However, all of these schemes—including our own—struggle in interpreting the absence of traffic. While a packet that has been modified in transit represents clear evidence of tampering, a missing packet is inherently ambiguous: it may have been explicitly blocked by a compromised router or it may have been dropped benignly due to network congestion. In fact, modern routers routinely drop packets due to bursts in traffic that exceed their buffering capacities, and the widely used Transmission Control Protocol (TCP) is designed to cause such losses as part of its normal congestion control behavior. Thus, existing traffic validation systems must inevitably produce false positives for benign events and/or produce false negatives by failing to report real malicious packet dropping.

2 BACKGROUND

There are inherently two threats posed by a compromised router. The attacker may subvert the network control plane (e.g., by manipulating the routing protocol into false route updates) or may subvert the network data plane and forward individual packets incorrectly. The first set of attacks have seen the widest interest and the most activity—largely due to their catastrophic potential. By violating the routing protocol itself, an attacker may cause large portions of the network to become inoperable.

While groundbreaking, Perlman's work required significant commitments of router resources and high levels of network participation to detect anomalies. However, we also assumed that the problem of congestion ambiguity could be solved, without providing a solution. This paper presents a protocol that removes this assumption.

3 INFERRING CONGESTIVE LOSS

In building a traffic validation protocol, it is necessary to explicitly resolve the ambiguity around packet losses. Should the absence of a given packet be seen as malicious or benign? In practice, there are some approaches for addressing this issue:

- Static Threshold: Low rates of packet loss are assumed to be congestive, while rates above some predefined threshold are deemed malicious.
- Traffic measurement. Individual packet losses are predicted as a function of measured traffic load and

router buffer capacity. Deviations from these predictions are deemed malicious.

Instead of using a static threshold, if the probability of congestive losses can be modeled, then one could resolve ambiguities by comparing measured loss rates to the rates predicted by the model. A simplified stochastic model of TCP congestion control yields the following famous square root formula:

$$B = \frac{1}{RTT}\sqrt{\frac{3}{2bp}}$$

Where B is the throughput of the connection, RTT is the average round trip time, b is the number of packets that are acknowledged by one ACK, and p is the probability that a TCP packet is lost. The steady-state throughput of long-lived TCP flows can be described by this formula as a
function of RTT and p.

4 SYSTEM MODEL

Our work proceeds from an informed, yet abstracted, model of how the network is constructed, the capabilities of the attacker, and the complexities of the traffic validation problem. In this section, we briefly describe the assumptions underlying our model. We use the same system model as in our earlier work [4].

4.1 Network Model

We consider a network to consist of individual homogeneous routers interconnected via directional point-to point links. This model is an intentional simplification of real networks (e.g., it does not include broadcast channels or independently failing network interfaces) but is sufficiently general to encompass such details if necessary.

Within a network, we presume that packets are forwarded in a hop-by-hop fashion, based on a local forwarding table This is critical, as we depend on the routing protocol to provide each node with a global view of the current network topology. Finally, we assume the administrative ability to assign and distribute cryptographic keys to sets of nearby routers.

4.2 Threat Model

As explained in Section 1, this paper focuses solely on data plane attacks (control plane attacks can be addressed by other protocols the protocol we develop validates traffic whose source and sink routers are uncompromised.

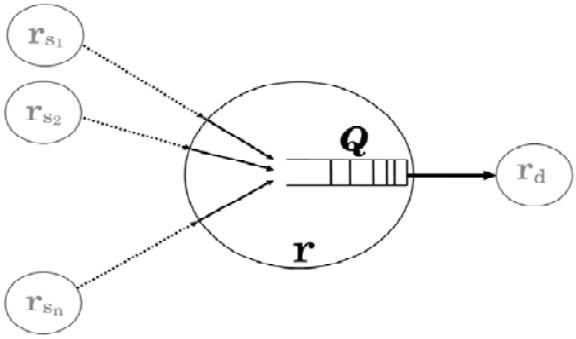

A protocol faulty router can send control messages with arbitrarily faulty information, or it can simply not send some or all of them. A faulty router is one that is traffic faulty, protocol faulty, or both. Attackers can compromise one or more routers in a network. However, for simplicity, we assume in this paper that adjacent routers cannot be faulty. Our work is easily extended to the case of k adjacent faulty routers. Misdetection of legitimate behavior by TV results in a false positive.

$$\begin{aligned}
c_{single} &= Prob(fp \text{ is maliciously dropped}) \\
&= Prob(\text{there is enough space in the queue to buffer } fp) \\
&= Prob(q_{act}(ts) + ps \le q_{limit}) \\
&= Prob(X + q_{pred}(ts) + ps \le q_{limit}) && \text{Random variable } X = q_{act}(ts) - q_{pred}(ts) \\
& && \text{with mean } \mu \text{ and standard deviation } \sigma \\
&= Prob(X \le q_{limit} - q_{pred}(ts) - ps) \\
&= Prob\left(Y \le \frac{q_{limit} - q_{pred}(ts) - ps - \mu}{\sigma}\right) && \text{Random variable } Y = (X - \mu)/\sigma \\
&= Prob(Y \le y_1) && y_1 = \frac{q_{limit} - q_{pred}(ts) - ps - \mu}{\sigma} \\
&= \frac{1 + erf(y_1/\sqrt{2})}{2} && erf \text{ is the error function.}
\end{aligned}$$

5 ANALYSIS OF PROTOCOL

In this section, we consider the properties and overhead of protocol χ.

There are two steps in showing the accuracy and completeness of χ:

- Showing that TV is correct.

in a false negative, and any

Within the given system model of Section 4, the example TV predicate in Section 5.1 is correct. However, the system model is still simplistic. In a real router, packets may be legitimately dropped due to reasons other than congestion: for example, errors in hardware, software or memory, and transient link errors.

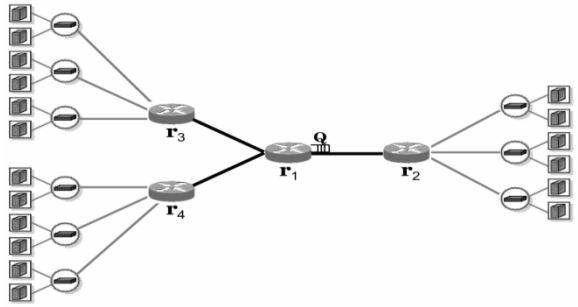

Fig:3 Simple topology

6 EXPERIENCES

We have implemented and experimented with protocol χ in the Emulab [35], [36] testbed. In our experiments, we used the simple topology shown in Fig. 3. The routers were Dell PowerEdge 2850 PC nodes with a single 3.0-GHz 64-bit Xeon processor and 2 Gbytes of RAM, and they were running Redhat-Linux-9.0 OS software. Each router except for r1 was connected to three LANs to which user machines were connected. The links between routers were configured with 3-Mbps bandwidth, 20-ms delay, and 75,000-byte capacity FIFO queue.

Each pair of routers shares secret keys; furthermore, integrity and authenticity against the message tampering is provided by message authentication codes.

In the second experiment, we first ran a training run to measure the mean and standard deviation of q_{error}. We found $\mu=0$ and $\sigma=1{,}750$. We then ran protocol χ under a high traffic load for more than 1 h, which generated more than half a million packets. Approximately 4,000 validation rounds occurred within this run, and approximately 16,000 packets were dropped due to congestion distribution, and the lower false positive rate for the combined packet drop test is because the test is not done on a simple random sample. We are investigating this further. In all of the subsequent experiments, we used the same mean, standard deviation, and two significance levels given here.

7 ISSUES

7.1 Quality of Service

Real routers implement Quality of Service (QoS) providing preferential treatment to specified traffic via several different traffic-handling techniques, such as traffic shaping,
traffic policing, packet filtering, and packet classification. Given the configuration files, our work can be extended to handle these fairly complex real-life functions, even those involving nondeterminism, if the expected
behavior of the function can be modeled.

7.2 Adjacent Faulty Routers

We assume that there exists no adjacent faulty routers in our threat model for simplicity. This assumption eliminates consorting faulty routers that collude together to produce fraudulent traffic information in order to hide their faulty behavior. This is the same approach that we used in [4], and it increases the overhead of detection.

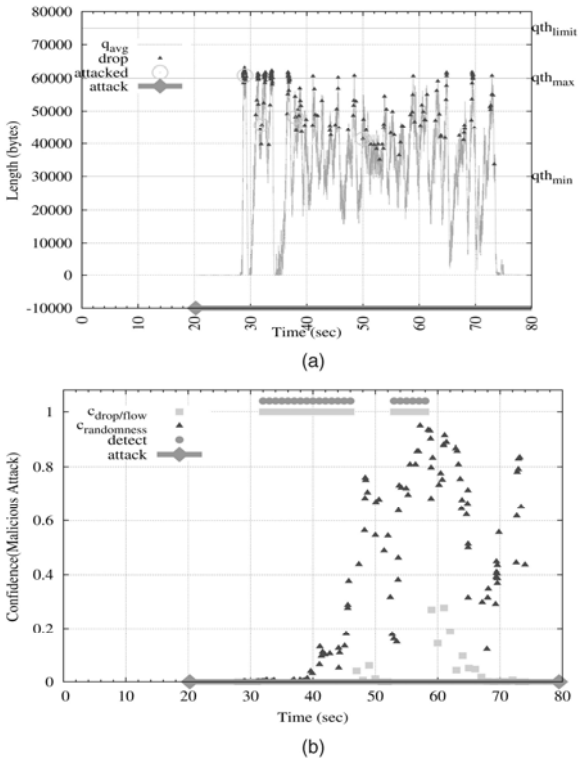

Fig. 4. Attack 5: Target a host trying to open a connection by dropping SYN packets. (a) Average queue length. (b) Statistical test results.

This assumption is necessary, in order to protect against faulty terminal routers that drop packets they receive from an end host or packets they should deliver to an end host. However, it also excludes DoS attacks wherein a faulty router introduces bogus traffic claiming that the traffic originates from a legitimate end host. Yet, none of these protocols explicitly address this

problem. Of course, standard rate-limit scheme can be applied against these kinds of DoS attacks.

8 CONCLUSION

To the best of our knowledge, this paper is the first serious attempt to distinguish between a router dropping packets maliciously and a router dropping packets due to congestion. Previous work has approached this issue using a static user-defined threshold, which is fundamentally limiting. Using the same framework as our earlier work (which is based on a static user-defined threshold) [4], we developed a compromised router detection protocol χ that dynamically infers, based on measured traffic rates and buffer sizes, the number of congestive packet losses that will occur.

We evaluated the effectiveness of protocol χ through an implementation and deployment in a small network. We show that even fine-grained attacks, such as stopping a host from opening a connection by discarding the SYN packet, can be detected.

APPENDIX

Similar to the specification that we have defined in [4], we cast the problem as a failure detector with accuracy and completeness properties.

- a-Accuracy: A failure detector is a-Accurate if whenever a correct router suspects (π, τ), then $|\pi| \leq$ a and some router $r \in \pi$ was faulty in π during τ.

We use the term traffic faulty to indicate a router that drop packets from transit traffic and the term *protocol faulty* to indicate a router that behaves arbitrarily with respect to the detection protocol. The a-Accuracy requirement can result in a detection if a router is either protocol faulty or traffic faulty.

ACKNOWLEDGMENTS

The first author would like to thank Mustafa Arisoylu for the insightful discussions about the dynamics of RED algorithm. This material is based upon Alper T. M1zrak's dissertation research at University of California, San Diego.

REFERENCES

[1] X. Ao, Report on DIMACS Workshop on Large-Scale Internet Attacks, http://dimacs.rutgers.edu/Workshops/Attacks/internet-attack- 9-03.pdf, Sept. 2003.

[2] R. Thomas, ISP Security BOF, NANOG 28, http://www.nanog.org/mtg-306/pdf/thomas. pdf, June 2003.

[3] K.A. Bradley, S. Cheung, N. Puketza, B. Mukherjee, and R.A. Olsson, "Detecting Disruptive Routers: A Distributed Network Monitoring Approach," Proc. IEEE Symp. Security and Privacy (S&P '98), pp. 115-124, May 1998.

[4] A.T. Mizrak, Y.-C. Cheng, K. Marzullo, and S. Savage, "Detecting and Isolating Malicious Routers," IEEE Trans. Dependable and Secure Computing, vol. 3, no. 3, pp. 230-244, July-Sept. 2006.

[5] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz, "Listen and Whisper: Security Mechanisms for BGP," Proc. First Symp. Networked Systems Design and Implementation (NSDI '04), Mar. 2004.

[6] S. Kent, C. Lynn, J. Mikkelson, and K. Seo, "Secure Border Gateway Protocol (Secure-BGP)," IEEE J. Selected Areas in Comm., vol. 18, no. 4, pp. 582-592, Apr. 2000.

[7] Y.-C. Hu, A. Perrig, and D.B. Johnson, "Ariadne: A Secure On- Demand Routing Protocol for Ad Hoc Networks," Proc. ACM MobiCom '02, Sept. 2002.