Exploring Potential Vulnerabilities in Data Center Network *

Zhenqian Feng*†, Haitao Wu*, Jinshu Su†
*Microsoft Research Asia, Beijing, P. R. China
†School of Computer, National University of Defense Technology, Changsha, P.R. China

ABSTRACT

In this paper, we study the potential security issues in data center networks (DCN). We focus on TCP, and propose a synchronized denial of service (SDoS) attack model. In SDoS, malicious tenants may employ multiple virtual machines (VM) as attackers and launch a low-rate synchronized DoS attack to VMs of the targeted victims. We give an analysis and our preliminary experiments confirm the feasibility of such SDoS attack.

Categories and Subject Descriptors

C.2.0 [**Security and protection**]: Denial of Service

General Terms

Security, Measurement, Performance

Keywords

Denial of Service, TCP, Data Center Network

1. INTRODUCTION

Recent cloud computing provides scalable service using rental VMs, e.g., Amazon EC2. Although the computing resources on any particular server is well allocated by VMs, the network resource is actually shared directly among different tenants[5, 1]. Lack of efficient bandwidth and traffic isolation makes such service vulnerable to attacks inside DCN.

Malicious tenants may launch arbitrary traffic pattern to attack the VMs of other tenants in the same data center. In this paper, we focus on attacks targeting at TCP, which has been widely used as reliable transport protocol. Previous study shows that pulsing denial of service(PDoS) alike attack[3] may cause excessive timeouts on other victims' TCP connections. Compared with traffic flooding attack, PDoS is hard to detect.

Although on Internet the assumed bottleneck saturated by one PDoS flow may be found, such case is hard to deploy in a well-constructed DCN. A DCN is usually provided with high-bandwidth and low-latency. Considering such low-latency can be leveraged for synchronization, we study a synchronized denial of service (SDoS) attack. In SDoS, multiple PDoS flows are synchronized carefully (at microsecond granularity), and are aggregated to generate large traffic burst in milliseconds time. Both the inter-arrival time of flows and the burst length is different from previous study on Internet at 100 ms[4]. Recent study on TCP incast [2] reflects the impact. In this paper, we build a SDoS model to study the characteristics of such attack.

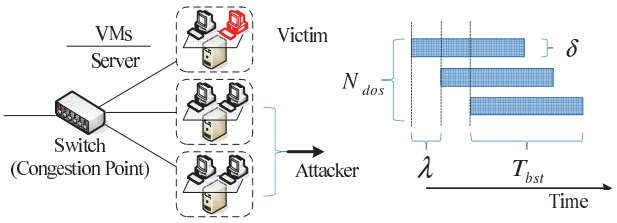

Figure 1: DoS scenario and system model.

2. MODEL OF ATTACK

To illustrate how SDoS attack works, we begin with the simplest case as shown in Fig. 1, where tenant attackers (TAs) and the victim(s) share the same FIFO pattern ToR switch. TAs synchronize multiple attack flows to overflow the buffer periodically, like[3]. We consider a single victim TCP flow to simplify our model. For cases that multiple victim flows under the attack simultaneously, our model accounts for the worst case, namely flow with the largest RTT, as longer bursts

*This work was performed when Zhenqian was an intern at Microsoft Research Asia.

ACM CoNEXT Student Workshop, November 30, Philadelphia, USA.

S_{buf}	The buffer size of switch, counted by packets
τ	The time to transmit a packet
N_{DoS}	The number of DoS attack flows
$\delta \in [0,1]$	The rate of homogeneous attack flows, normalized against the capacity of the link
T_{bst}	The duration of persistent attack, namely the burst length in[3]
$\lambda \in [0, T_{bst}]$	inter-arrival time to characterize the degree of synchronization of the attack flows
R_0	RTT of the victim TCP without queueing
N_{pkt}	Packets of the victim TCP issued per round

Table 1: Notations.

would be required to suppress it.

Table. 1 introduces the notations for our model. We assume the victim TCP already entered the steady state before the attack, and a round corresponds to the duration for the foremost packet of a full TCP window to arrive at the switch successively. For convenience, we further presume all packets are of the same size.

From the TA's point of view, the whole objective is to cause enough packet losses for TCP to timeout while minimizing the attack time T_{bst}.

We decompose T_{bst} into two parts: time to saturate the switch buffer (T_1) and time to drop enough packets (T_2). For the former, as more SDoS packets pile into the buffer, both the RTT and the queue length would increase until a packet loss and evolve as follows:

$$R(k) = R_0 + Q(k) \times \tau \quad (1)$$

$$Q(k) = Q(k-1) + N_{pkt} + (\phi - 1) \times R(k-1)/\tau \quad (2)$$

where $R(k)$ and $Q(k)$ are the RTT and queue length in the k^{th} round respectively, ϕ indicates the aggregate effective rate of SDoS flows and:

$$\phi = \frac{N_{DoS} \times T_{bst} \times \delta}{T_{bst} + (N_{DoS} - 1) \times \lambda} \quad (3)$$

Let $Q(k') = S_{buf}$, one can find the round when the buffer begins to overflow. Using (1)-(3), we then get:

$$T_1 = \sum\nolimits_{i \in [1,k']} R(i), where: k' = Q^{-1}(S_{buf}) \quad (4)$$

For the latter, we simply assume that an additional round after the buffer gets full would cause enough packet losses for TCP to timeout, hence:

$$T_2 = R(k') = R_0 + S_{buf} \times \tau \quad (5)$$

Thus we can conclude the time TAs need to attack efficiently is $T_{bst} \geq T_1 + T_2$ by combining (4) and (5).

3. PRELIMINARY RESULTS

We deployed a testbed with tens of servers and one Quanta LB4G 48-port Gigabit Ethernet switch with 2MB buffer per chip evenly shared by 24 ports, as shown in Fig. 1. We restrict TCP receive window size to control the packets issued per round, and launch each attack flow with $\delta \approx 0.94$. Further, we found λ can be less than 100us thus assumed $\lambda \ll T_{bst}$. Observing that, we take $\phi = \frac{N_{DoS} \times T_{bst} \times \delta}{T_{bst} + (N_{DoS} - 1) \times \lambda} \approx N_{DoS} \times \delta$ for simplicity.

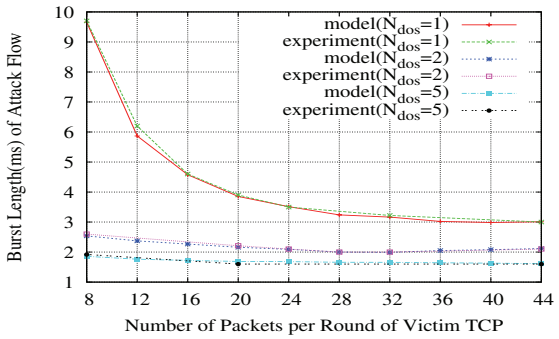

Figure 2: T_{bst} vs num of packets per round.

Besides a good match between our model and the practical experiments, two observations are found in Fig. 2:(1)more packets issued per round makes TCP more vulnerable, as the attacker gains more help from the victim itself, and (2)more participants in an attack lower the risk of exposure of TAs.

4. CONCLUSIONS

Compared with traditional PDoS attack, SDoS attack leverages the low-latency feature of DCN by well synchronizing multiple flows to generate high traffic burst at milliseconds time. As future work, we are considering bottleneck at high-level switches and other packet dropping policy like RED.

As for the counter-measures, performance isolation or guarantee[5, 1] are promising, but to handle attacks at milliseconds scale is challenging. Proposals for mitigating TCP Incast may help, yet elaborate attacks may preserve the congestion for a long time, e.g., using different combinations of flows to attack.

5. ACKNOWLEDGMENTS

This work is supported by the grants from National Grand Fundamental Research 973 Program of China under Grant No. 2009CB320503, the National 863 Development Plan of China under Grant No.2008AA01A325, No.2009AA01Z432, and No. 2009AA01A346.

6. REFERENCES

[1] C. Guo et al. Secondnet: A data center network virtualization architecture with bandwidth guarantees. In *Proc. CoNext*, 2010.
[2] V. Vasudevan et al. Safe and effective fine-grained tcp retransmissions for datacenter communication. In *Proc. Sigcomm*, 2009.
[3] X. Luo et al. On a new class of pulsing denial-of-service attacks and the defense. In *Proc. NDSS*, 2005.
[4] Y. Zhang et al. Low-rate tcp-targeted dos attack disrupts internet routing. In *Proc. ISOC NDSS*, 2007.
[5] A. Shieh. Seawall:performance isolation for cloud datacenter networks. *HotCloud.*, 2010.