International Journal of Contemporary Research in Computer Science and Technology (IJCRCST) e-ISSN: 2395-5325
Volume2, Issue 8 (August '2016)

TOWARDS TO GETTING BETTER ACTIVE PACKET LOSS DIMENSION ON FIREWALL IN CLOUD COMPUTING

Dr.G.Kesavaraj,
Professor and Head,
Department of Computer Science,
Vivekanandha College of Arts and Sciences for Women (Autonomous),
Elayampalayam, Namakkal, Tamilnadu.

N.Jeevitha,
M.Phil Research Scholar,
Department of Computer Science,
Vivekanandha College of Arts and Sciences for Women (Autonomous),
Elayampalayam, Namakkal, Tamilnadu

Abstract: The firewall is one amongst the central technologies permitting high-level access management to organization cloud knowledge networks. Cloud knowledge matching in firewalls involves matching on several fields from the info header. A minimum of 5 fields (protocol variety, supply and destination information processing addresses, and ports) are concerned within the call that rule applies to a given cloud knowledge. Since firewalls ought to filter all the traffic crossing the cloud sharing network perimeter, they must be able to sustain a really high outturn, or risk changing into a bottleneck. Thus, algorithms from process pure mathematics are applied. During this paper we have a tendency to contemplate a classical rule that we have a tendency to tailor to the firewall domain. We have a tendency to decision the ensuing rule "Geometric economical Matching" (GEM). The GEM rule enjoys an exponent matching time performance. Cloud computing could be a new versatile approach for providing higher process power in shared medium. It provides the distributed model supported self-evaluating techniques to enhance the process capabilities of the system with lesser social control issues. This computing model delivers computation capabilities as a calculated service from on top of parts to finish users. Although a good style of devices and their integration are involved, priority of handling security can go down. Implementing firewall for cloud suffers from varied network homeward-bound challenges like load equalization, scheduling, traffic divergence, filtering, dominant the speed of arrival, instance management, attack detection. It additionally aims toward achieving the resource optimizing based mostly provisions and rules to lower the worth related to its possession and operations.

Keywords: Cloud computing, firewall, GEM.

I. INTRODUCTION

Cloud computing is the delivery of computing services over the Internet. Cloud services allow individuals and businesses to use software and hardware that are managed by third parties at remote locations. Examples of cloud services include online file storage, social networking sites, webmail, and online business applications. Cloud computing provides a shared pool of resources, including data storage space, networks, computer processing power, and specialized corporate and user applications. When you store your photos online instead of on your home computer, or use webmail or a social networking site, you are using a "cloud computing" service. If you are an organization, and you want to use, for example, an online invoicing service instead of updating the in-house one you have been using for many years, that online invoicing service is a "cloud computing" service. Cloud computing refers to the delivery of computing resources over the Internet. Instead of keeping data on your own hard drive or updating applications for your needs, you use a service over the Internet, at another location, to store your information or use its applications. Doing so may give rise to certain privacy implications.

For that reason the Office of the Privacy Commissioner of Canada (OPC) has prepared some responses to Frequently Asked Questions (FAQs). To have also developed a Fact Sheet that provides detailed information on cloud computing and the privacy challenges it presents. The following definition of cloud computing has been developed by the U.S. National Institute of Standards and Technology (NIST):

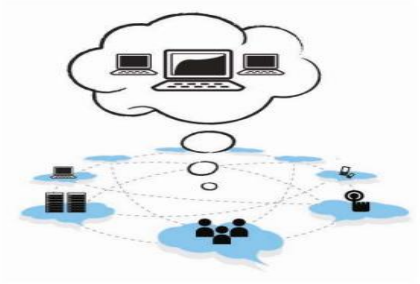

Figure 1: Cloud Computing

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable

IJCRCST © 2016 | All Rights Reserved

computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. In Figure 1 This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.

Service models: The cloud computing service models are Software as a Service (Saas), Platform as a Service (Paas) and Infrastructure as a Service (Iaas). In a Software as a Service model, a pre-made application, along with any required software, operating system, hardware, and network are provided. In Paas, an operating system, hardware, and network are provided, and the customer installs or develops its own software and applications. The Iaas model provides just the hardware and network; the customer installs or develops its own operating systems, software and applications.

Deployment of cloud services: Cloud services are typically made available via a private cloud, community cloud, public cloud or hybrid cloud. Generally speaking, services provided by a public cloud are offered over the Internet and are owned and operated by a cloud provider. Some examples include services aimed at the general public, such as online photo storage services, e-mail services, or social networkinsites. However, services for enterprises can also be offered in a public cloud.

In a **private cloud**, the cloud infrastructure is operated solely for a specific organization, and is managed by the organization or a third party. In a **community cloud**, the service is shared by several organizations and made available only to those groups. A **hybrid cloud** is a combination of different methods of resource pooling (for example, combining public and community clouds).

Privacy is not a barrier but it must be taken into consideration

The Personal Information Protection and Electronic Documents Act (PIPEDA) do not prevent an organization from transferring personal information to an organization in another jurisdiction for processing. However, PIPEDA establishes rules governing those transfers particularly with respect to obtaining consent for the collection, use and disclosure of personal information, securing the data, and ensuring accountability for the information and transparency in terms of practices. For more information on the views of the Office of the Privacy Commissioner of Canada with respect to outsourcing of personal data processing across borders, please see our Guidelines for Processing Personal Data across Borders. These considerations apply whether moving data in the cloud or otherwise. It is important to note that many non-Canadian based cloud providers may also be subject to PIPEDA. To the extent that a cloud provider has a real and substantial connection to Canada, and collects, uses or discloses personal information in the course of a commercial activity, the provider is expected to protect personal information, in keeping with PIPEDA. If you are considering a cloud service, you should think about how your personal information, and that of your customers, can best be protected. Carefully review the terms of service or contracts, and challenge the provider to meet your needs.

- Ruby on Rails (Heroku)
- .NET (Azure Services Platform)
- Web hosting (Mosso)
- Proprietary (Force.com)

Service: A cloud service includes "products, services and solutions that are delivered and consumed in real-time over the Internet". For example, Web Services ("software system[s] designed to support interoperable machine-to-machine interaction over a network") which may be accessed by other cloud computing components, software, e.g., Software plus service, or end users directly.
Specific examples include:

- Identity (OAuth, OpenID)
- Integration (Amazon Simple Queue Service)
- Payments (Amazon Flexible Payments Service, Google Checkout, PayPal)
- Mapping (Google Maps, Yahoo! Maps)
- Search (Alexa, Google Custom Search, Yahoo! BOSS)
- Others (Amazon Mechanical Turk)

Storage : Cloud storage involves the delivery of data storage as a service, including database-like services, often billed on a utility computing basis, e.g., per gigabyte per month.

For Example:

- Database (Amazon SimpleDB, Google App Engine's BigTable datastore)
- Network attached storage (MobileMe iDisk, Nirvanix CloudNAS)
- Synchronization (Live Mesh Live Desktop component, MobileMe push functions)
- Web service (Amazon Simple Storage Service, Nirvanix SDN)

II. RELATED WORK

Above the Clouds: A View of Cloud Computing: Cloud Computing, the long-held dream of computing as a utility, has the potential to transform a large part of the IT industry, making software even more attractive as a service and shaping the way IT hardware is designed and purchased. Developers with innovative ideas for new Internet services no longer require the large capital outlays in hardware to deploy their service or the human expense to operate it. They need not be concerned about over-provisioning for a service whose popularity does not meet their predictions, thus wasting costly resources, or under-provisioning for one that becomes wildly popular, thus missing potential customers and revenue. Moreover, companies with large batch-oriented tasks can get results as quickly as their programs can scale, since using 1000 servers for one hour costs no more than using one server for

1000 hours. This elasticity of resources, without paying a premium for large scale, is unprecedented in the history of IT.

Hey, You, Get Off Of My Cloud: Exploring Information Leakage In Third-Party Compute Clouds: Third-party cloud computing represents the promise of outsourcing as applied to computation. Services, such as Microsoft's Azure and Amazon's EC2, allow users to instantiate virtual machines (VMs) on demand and thus purchase precisely the capacity they require when they require it. In turn, the use of virtualization allows third-party cloud providers to maximize the utilization of their sunk capital costs by multiplexing many customer VMs across a shared physical infrastructure. However, in this paper, we show that this approach can also introduce new vulnerabilities. Using the Amazon EC2 service as a case study, we show that it is possible to map the internal cloud infrastructure, identify where a particular target VM is likely to reside, and then instantiate new VMs until one is placed co-resident with the target. We explore how such placement can then be used to mount cross-VM side-channel attacks to extract information from a target VM on the same machine.

Performance modeling and analysis of network firewalls: Network firewalls act as the first line of defense against unwanted and malicious traffic targeting Internet servers. Predicting the overall firewall performance is crucial to network security engineers and designers in assessing the effectiveness and resiliency of network firewalls against DDoS (Distributed Denial of Service) attacks as those commonly launched by today's Botnets. In this paper, we present an analytical queueing model based on the embedded Markov chain to study and analyze the performance of rule-based firewalls when subjected to normal traffic flows as well as DoS attack flows targeting different rule positions. We derive equations for key features and performance measures of engineering and design significance. These features and measures include throughput, packet loss, packet delay, and firewall's CPU utilization. In addition, we verify and validate our analytical model using simulation and real experimental measurements.

III. THE GEOMETRIC EFFICIENT MATCHING ALGORITHM FOR FIREWALLS:

Firewall packet matching can be viewed as a point location problem: Each packet (point) has 5 fields (dimensions) which need to be checked against every firewall rule in order to find the first matching rule. In this paper we consider a packet matching algorithm, which we call the Geometric Efficient Matching (GEM) algorithm. The GEM algorithm enjoys a logarithmic matching time performance, easily beating the linear time required by the naive matching algorithm. However, the algorithm's theoretical worst-case space complexity is O(n4) for a rule-base with n rules. Based on statistics from real firewall rule-bases, we created a model that generates random, but non-uniform, rulebases. We evaluated GEM via extensive simulation using this rule-base generator. Subsequently, we integrated GEM into the code of the Linux iptables open-source firewall. Our GEM-iptables implementation supports a throughput which is at least 5-10 times higher than that of the unoptimized iptables. Our implementation was able to match over 30,000 packets-per-second even with 10 thousand rules.

Collaborative enforcement of firewall policies in virtual private networks: The widely deployed Virtual Private Network (VPN) tech- nology allows roaming users to build an encrypted tunnel to a VPN server, which henceforth allows roaming users to access some resources as if that computer is residing on their home organization's network. Although the VPN technology is very useful, it imposes security threats to the remote network because their firewall does not know what traffic is flowing inside the VPN tunnel. To address this issue, we pro- pose VGuard, a framework that allows a policy owner and a request owner to collaboratively determine whether the re- quest satisfies the policy without the policy owner knowing the request and the request owner knowing the policy. We first present an efficient protocol, called Xhash, for oblivious comparison, which allows two parties, where each party has a number, to compare whether they have the same number, without disclosing their numbers to each other. Then, we present the VGuard framework that uses Xhash as the basic building block. The basic idea of VGuard is to first convert a firewall policy to non-overlapping numerical rules and then use Xhash to check whether a request matches a rule. Comparing with the Cross-Domain Cooperative Fire- wall (CDCF) framework, which represents the state-of-the- art, VGuard is not only more secure but also orders of magnitude more efficient. On real-life firewall policies, for processing packets, our experimental results show that VGuard is 552 times faster than CDCF on one party and 5035 times faster than CDCF on the other party.

Change-Impact Analysis of Firewall Policies:
Firewalls are the mainstay of enterprise security and the most widely adopted technology for protecting private networks. The quality of protection provided by a firewall directly depends on the quality of its policy (i.e., configuration). Due to the lack of tools for analyzing firewall policies, most firewalls on the Internet have been plagued with policy errors. A firewall policy error either creates security holes that will allow malicious traffic to sneak into a private network or blocks legitimate traffic and disrupts normal business processes, which in turn could lead to irreparable, if not tragic, consequences. A major source of policy errors stem from policy changes. Firewall policies often need to be changed as networks evolve and new threats emerge. In this paper, we first present the theory and algorithms for firewall policy change-impact analysis. Our algorithms take as input a firewall policy and a proposed change, then output the accurate impact of the change. Thus, a firewall administrator can verify a proposed change before committing it.

First Step toward Cloud-Based Firewalling: With the explosive growth of network-based services and attacks, the complexity and cost of firewall deployment and management have been increasing rapidly. Yet, each private network, no matter big or small, has to deploy and manage its own firewall, which is the critical first line of defense. To reduce the complexity and cost in deploying and managing firewalls, businesses have started to outsource the firewall service to their Internet Service Providers (ISPs), such as AT&T, which provide cloud-based firewal service. Such firewalling model saves businesses in managing, deploying, and upgrading firewalls. The current firewall service outsourcing model requires businesses fully trust their ISPs and give ISPs their firewall policies. However, businesses typically need to keep their firewall policies confidential. In this paper, we propose the first privacy preserving firewall outsourcing approach where businesses outsource their firewall services to ISPs without revealing their firewall policies to the ISPs. The basic idea is that businesses first anonymize their firewall policies and send the anonymized policies to their ISP; then the ISP performs packet filtering based on the anonymized firewall policies. For anonymizing firewall policies, we use Firewall Decision Diagrams to cope with the multi-dimensionality of policies and Bloom Filters for the anonymization purpose. This paper deals with a hard problem. By no means that we claim our scheme is perfect; however, this effort represents the first step towards privacy preserving outsourcing of firewall services. We implemented our scheme and conducted extensive experiments. Our experimental results show that our scheme is efficient in terms of both memory usage and packet lookup time. The firewall throughput of our scheme running at ISPs is comparable to that of software firewalls running at businesses themselves.

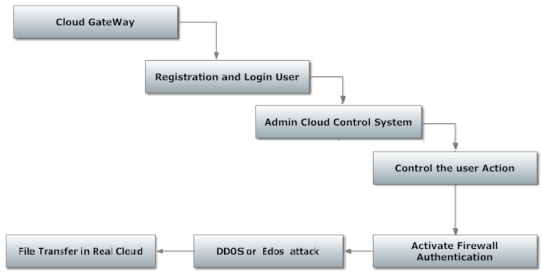

Figure 2: ***System Architecture***

Can We Beat DDoS Attacks in Clouds? (Supplementary Material): In Fig. 2 DDoS attacks aim to exhaust the resources of victims, such as network bandwidth, computing power and operating system data structures. Early DDoS attacks emerged around the year 2000, and well-known web sites, such as CNN, Amazon and Yahoo, have been the targets of hackers since then. The purpose of early attacks was mainly for fun and curiosity about the technique. However, recently we have witnessed an explosive increase in cyber attacks due to the huge financial or political rewards available to cyber attackers [1]. Botnets are the engines behind major DDoS attacks. Hackers exploit the vulnerability of computers connected to the Internet, and establish an overlay network of compromised computers to commit malicious activities, such as DDoS attacks or information phishing. This kind of malicious network is what we call a botnet [2], [3]. A DDoS attack can be carried out in various forms, such as flooding packets or synchronization attacks [2]. Flooding packets is the most common and effective DDoS attack strategy amongst all the available attack weapons. It is critical for defenders to understand the size of botnets, which helps us to estimate the possible attack volume. There has been plenty of work completed on this issue, such as [1] and [4]. Rajab et al. [5] found the number of active bots a botmaster could manipulate was usually at the hundreds or a few thousands level. This means the resources a botnet owner can use is limited. Based on this fact, Yu et al. [6] proposed a similarity based DDoS detection method to beat flash crowd mimicking attacks. Traditionally, a potential victim would be vulnerable if they were left to deal with a DDoS flooding attack by themselves. As nature of the Internet is anarchical, potential victims, such as popular web sites, are usually left.

Inferring Internet Denial-of-Service Activity: In this paper, we seek to answer a simple question: "How prevalent are denial-of-service attacks in the Internet today?". Our motivation is to understand quantitatively the nature of the current threat as well as to enable longerterm analyses of trends and recurring patterns of attacks. We present a new technique, called "backscatter analysis", that provides an estimate of worldwide denial-ofservice activity. We use this approach on three week-long datasets to assess the number, duration and focus of attacks, and to characterize their behavior. During this period, we observe more than 12,000 attacks against more than 5,000 distinct targets, ranging from well known ecommerce companies such as Amazon and Hotmail to small foreign ISPs and dial-up connections. We believe that our work is the only publically available data quantifying denial-of-service activity in the Internet.

A Cost-aware Elasticity Provisioning System for the Cloud: In this paper we present Kingfisher, a cost-aware system that provides efficient support for elasticity in the cloud by (i) leveraging multiple mechanisms to reduce the time to transition to new configurations, and (ii) optimizing the selection of a virtual server configuration that minimizes the cost. We have implemented a prototype of Kingfisher and have evaluated its efficacy on a laboratory cloud platform. Our experiments with varying application workloads demonstrate that Kingfisher is able to (i) decrease the cost of virtual server resources by as much as 24% compared to the current cost unaware approach,

(ii) reduce by an order of magnitude the time to transition to a new configuration through multiple elasticity mechanisms in the cloud, and (iii), illustrate the opportunity for design alternatives which trade-off the cost of server resources with the time required to scale the application.

Adaptive Resource Provisioning for the Cloud Using Online Bin Packing: Data center applications present significant opportunities for multiplexing server resources. Virtualization technology makes it easy to move running application across physical machines. In this paper, we present an approach that uses virtualization technology to allocate data center resources dynamically based on application demands and support green computing by optimizing the number of servers actively used. We abstract this as a variant of the relaxed on-line bin packing problem and develop a practical, efficient algorithm that works well in a real system. We adjust the resources available to each VM both within and across physical servers. Extensive simulation and experiment results demonstrate that our system achieves good performance compared to the existing work.

On Incentive of Customer-Provided Resource Sharing in Cloud: The state-of-the-art cloud computing service has attracted significant interests from the Internet users. How-ever, in the existing cloud platforms, the cloud users are pure consumers; their local resources, though abundant, have been largely ignored. In this paper, we for the first time explore the resource pricing as well as the incentive issues in SpotCloud, a real- world system that enables customer-provided cloud computing service on the Internet. In this system, the resource providers are largely heterogeneous and are not forced to contribute their resources. A working business model is therefore impor- tant to offer them enough sharing incentive. Instead of setting a standardized pricing rule for unit resource, we suggest a distributed market that allows the sellers to decide the quality, quantity, and pricing of their own resources. We demonstrate the efficiency of this business model through a repeated seller competition game. The trace-analysis further indicates that the proposed business model can successfully motivate the resource sharing in our Spot cloud system.

Wide-Area Traffic the Failure of Poisson Modeling: Network arrivals are often modeled as Poisson processes for analytic simplicity, even though a number of traffic studies have shown that packet interarrivals are not exponentially distributed. We evaluate 24 wide-area traces, investigating a number of wide-area TCP arrival processes (session and connection arrivals, FTP data connection arrivals within FTP sessions, and TELNET packet arrivals) to determine the error introduced by modeling them using Poisson processes. We find that user-initiated TCP session arrivals, such as remote login and file-transfer, are well-modeled as Poisson processes with fixed hourly rates, but that other connection arrivals deviate considerably from Poisson; that modeling TELNET packet inter arrivals as exponential grievously underestimates the burstiness of TELNET traffic, but using the empirical Tcplib [Danzig et al, 1992] interarrivals preserves business over many time scales; and that FTP data connection arrivals within FTP sessions come bunched into "connection bursts," the largest of which are so large that they completely dominate FTP data traffic. Finally, we offer some results regarding how our findings relate to the possible self-similarity of wide area traffic.

Queue Length Asymptotic for Generalized Max-Weight Scheduling in the presence of Heavy-Tailed Traffic: We investigate the asymptotic behavior of the steady-state queue length distribution under generalized max weight scheduling in the presence of heavy-tailed traffic. We consider a system consisting of two parallel queues, served by a single server. One of the queues receives heavy-tailed traffic, and the other receives light-tailed traffic. We study the class of throughput optimal max-weight-α scheduling policies, and derive an exact asymptotic characterization of the steady-state queue length distributions. In particular, we show that the tail of the light queue distribution is heavier than a power-law curve, whose tail coefficient we obtain explicitly. Our asymptotic characterization also shows that the celebrated max-weight scheduling policy leads to the worst possible tail of the light queue distribution, among all non-idling policies. Motivated by the above 'negative' result regarding the max weight- α policy, we analyze a log-max-weight (LMW) scheduling policy. We show that the LMW policy guarantees an exponentially decaying light queue tail, while still being throughput optimal.

IV. EXISTING SYSTEM

Cloud computing is becoming popular as the next infrastructure of computing platform. Despite the promising model and hype surrounding, security has become the major concern that people hesitate to transfer their applications to clouds. Concretely, cloud platform is under numerous attacks. As a result, it is definitely expected to establish a firewall to protect cloud from these attacks. However, setting up a centralized firewall for a whole cloud data center is infeasible from both performance and financial aspects. In this paper, we propose a decentralized cloud firewall framework for individual cloud customers. We investigate how to dynamically allocate resources to optimize resources provisioning cost, while satisfying QoS requirement specified by individual customers simultaneously. Moreover, we establish novel queuing theory based model M/Geo/1 and M/Geo/m for quantitative system analysis, where the service times follow a geometric distribution. By employing Z-transform and embedded Markov chain techniques, we obtain a closed-form expression of mean packet response time. Through extensive simulations and experiments, we conclude that an M/Geo/1 model reflects the cloud firewall real system much better than a traditional M/M/1 model. Our numerical results also indicate that we are able to set up cloud firewall with affordable cost to cloud customers.

Disadvantages of existing system

- There is no secure when the packet sending in cloud.
- High configuration of Firewall not used before.
- Time consuming is high.

V. PROPOSED SYSTEM

With available bandwidth increasing rapidly, very efficient matching algorithms need to be deployed in modern firewalls to ensure that the firewall does not become a bottleneck Since firewalls need to filter all the traffic crossing the cloud sharing network perimeter, they should be able to sustain a very high throughput, or risk becoming a bottleneck. Thus, algorithms from computational geometry can be applied. In this paper we consider a classical algorithm that we adapted to the firewall domain. We call the resulting algorithm "Geometric Efficient Matching" (GEM). The GEM algorithm enjoys a logarithmic matching time performance. Cloud computing is a new flexible approach for providing higher computational power in shared medium. It provides the distributed model based on self-evaluating techniques to improve the processing capabilities of the system with lesser managerial concerns

Advantages of proposed system:

- Cloud data filter firewall supports high speed.
- Cloud data filter firewall over configurations of simple network works with more speed.
- The thing behind this is that cloud data filter firewall has the directly connection within external hosts & internal users.
- cloud data filters take decisions on the basis of the each cloud data, it doesn't take decision on the basis of the traffic context.
- It used to implement and enforce a security policy for communication between cloud sharing.

VI. CONCLUSION

We have seen that the GEM algorithm is an efficient and practical algorithm for firewall packet matching. We implemented it successfully in the Linux kernel, and tested its packet-matching speeds on live traffic with realistic large rule bases. GEM's matching speed is far better than the naive linear search, and it is able to increase the throughput of iptables by an order of magnitude. On rule-bases generated according to realistic statistics, GEM's space complexity is well within the capabilities of modern hardware. Thus we believe that GEM may be a good candidate for use in firewall matching engines. We believe it should be quite interesting to implement all of these algorithms and to test them on equal footing, using the same hardware, rule-bases, and traffic load. Furthermore, it would be interesting to do this comparison with real rule-bases, in addition to synthetic Perimeter-model rules. We leave such a "bake-off" for behavior when using more than 4 fields, e.g., matching on the TCP flags, meta data, interfaces, etc.

VII. REFERENCES

[1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," Commun. ACM, vol. 53, no. 4, pp. 50–58, 2010.

[2] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds," in Proc. 16th ACM Conf. Comput. Commun. Security, 2009, pp. 199–212.

[3]K. Salah, K. Elbadawi, and R. Boutaba, "Performance modeling and analysis of network firewalls," IEEE Trans. Netw. Serv. Manage., vol. 9, no. 1, pp. 12–21, Mar. 2012.

[4]. D. Rovniagin and A. Wool, "The geometric efficient matching algorithm for firewalls," IEEE Trans. Dependable Secure Comput., vol. 8, no. 1, pp. 147–159, Jan./Feb. 2011.

[5]. A. X. Liu and F. Chen, "Privacy preserving collaborative enforcement of firewall policies in virtual private networks," IEEE Trans. Parallel Distrib. Syst., vol. 22, no. 5, pp. 887–895, May 2011.

[6]. A. X. Liu, "Firewall policy change-impact analysis," ACM Trans. Internet Technol., vol. 11, no. 4, p. 15, 2012.

[7]. A. R. Khakpour and A. X. Liu, "First step toward cloud-based firewalling," in Proc. IEEE 31st Symp. Reliable Distrib. Syst., 2012, pp. 41–50.

[8]. S. Yu, Y. Tian, S. Guo, and D. Wu, "Can we beat ddos attacks in clouds?" IEEE Trans. Parallel Distrib. Syst., in press, 2014.

[9]. D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage, "Inferring internet denial-of-service activity," ACM Trans. Comput. Syst., vol. 24, no. 2, pp. 115–139, 2006.

[10]. U. Sharma, P. Shenoy, S. Sahu, and A. Shaikh, "A cost-aware elasticity provisioning system for the cloud," in Proc. IEEE 31st Int. Conf. Distrib. Comput. Syst, 2011, pp. 559–570.

[11]. J. Zhu, Z. Jiang, and Z. Xiao, "Twinkle: A fast resource provisioning mechanism for internet services," in Proc. IEEE INFOCOM, 2011, pp. 802–810.

[12]. L. Kleinrock, Queueing systems: Theory, vol. 1. New York, NY, USA: Wiley-interscience, 1975.

[13]. H. Wang, F. Wang, J. Liu, and J. Groen, "Measurement and utilization of customer-provided resources for cloud computing," in Proc. IEEE INFOCOM, 2012, pp. 442–450.

[14]. V. Paxson and S. Floyd, "Wide area traffic: The failure of poisson modeling," IEEE/ACM Trans. Netw., vol. 3, no. 3, pp. 226–244, Jun. 1995.

[15]. K. Jagannathan, M. Markakis, E. Modiano, and J. N. Tsitsiklis, "Queue-length asymptotics for generalized max-weight scheduling in the presence of heavy-tailed traffic," IEEE/ACM Trans. Netw., vol. 20, no. 4, pp. 1096–1111, Aug. 2012.