A Campus-Level View of Outlook Email Traffic

Zhengping Zhang
Nanjing University of Aeronautics and Astronautics
University of Calgary
zhang.zhengping@ucalgary.ca

Carey Williamson
University of Calgary
cwill@ucalgary.ca

ABSTRACT

This paper presents a workload characterization study of Outlook email traffic, as viewed at a campus scale. Using a combination of passive and active approaches to network traffic measurement, we identify the key components in the email delivery infrastructure, and then characterize the email traffic on our campus network, using a month of empirical network measurement data. The main observations from our study include the complexity of modern email services, the strong diurnal patterns for human-driven email activities, and the low throughput achieved for large email attachments.

CCS Concepts

•**Networks** → **Network Performance Evaluation;**

Keywords

Network Traffic Measurement, Outlook Email, Workload Characterization

1. INTRODUCTION

Electronic mail (email) has grown and changed dramatically since its first inception on the Internet in 1969. From humble beginnings with SMTP (Simple Mail Transfer Protocol), the `maild` daemon, and the `sendmail` program on Unix-based systems, email has since used a variety of mail access protocols, including POP (Post Office Protocol), IMAP (Internet Mail Access Protocol), and HTTP (Hyper-Text Transfer Protocol). Among these protocols, Web-based email using HTTPS is prevalent today, and has helped make email one of the most popular services on the Internet.

In the last two decades, numerous personal, business, and commercial email services have emerged, including Gmail, Outlook, Yahoo!, and Zoho. Based on worldwide email service statistics, there are over 3.7 billion email users around the world, and this number is projected to reach 4.2 billion by the end of 2022 [25].

DOI: 10.1145/1235

A decade ago, most academic institutions provided their own email service [9]. However, because of limited IT budgets, security concerns, and the growing volume and complexity of email, small email service providers rarely have the resources required to maintain a high-quality email service. In particular, they are often slow in adopting new standards or fixing security vulnerabilities [10].

In recent years, many universities have migrated their email to large cloud-based service providers [27], such as Google's Gmail or Microsoft's Outlook (a.k.a. Office 365). Our own university is one example of this. In 2014, the University of Calgary started to transfer its email service to Outlook, which is one of the most popular email services. By the end of 2016, all email users at the University of Calgary were switched to this service provider. This cloud-based email service is now used on a daily basis by 32,000 students and 3,000 faculty and staff.

In this paper, we conduct an in-depth measurement study of how this cloud-based email service is being used by our campus community, and how well it performs. We first investigate how the Outlook email service is provided, including its key components. Then we present a high-level overview of our campus email traffic. Finally, we present a detailed workload characterization study, and identify some performance-related issues for Outlook email.

There are two major technical challenges in our work. First, Outlook email relies on a number of other sub-services, including authentication, spam filtering, and content delivery. Each sub-service in turn has distinct usage, configuration, and workload characteristics. It is difficult to analyze email service as a whole without adequate understanding of its underlying sub-services. Second, for security reasons, all email traffic is encrypted using either SSL or TLS [13, 22]. Having all application-layer data obfuscated makes it difficult to conduct our analysis, for which only TCP/IP packet header information is available.

To address these challenges, we employ both passive and active measurement approaches in our workload study. Via passive monitoring at the campus edge router, we have collected more than two years of network traffic data, a month of which is used for our analysis in this paper. By using active measurement approaches, we are able to better understand the workflow of Outlook, including the application layer. For this purpose, we also rely upon several Outlook technical documents [17, 18, 19].

The results of our measurement study show that Outlook email usage has strong diurnal and weekly patterns on our campus network. The connection duration, as well as vol-

ume of data exchanged, vary with the type of service being used. We also identify several performance-related issues for Outlook email service.

The rest of this paper is organized as follows. Section 2 reviews prior related work on network traffic measurement. Section 3 introduces the methodology used for our study. Section 4 discusses the five main components involved in email delivery. In Section 5, we present a high-level overview of Outlook related traffic, while Section 6 provides an in-depth analysis. Finally, Section 7 concludes the paper.

2. RELATED WORK

Network traffic measurement is a widely-adopted approach for analyzing network usage and understanding network application performance. Such studies can provide not only a statistical characterization of the network ecosystem, but also help identify anomalies, misconfigurations, or performance limitations within the network. A tutorial introduction to network traffic measurement is provided in [26].

From the 1990's to the present, researchers have characterized many kinds of network applications, including Web traffic [1, 4], peer-to-peer applications [3], video streaming services [5, 14], and email [10]. For example, Arlitt *et al.* [1] characterized the workload of Internet Web servers, and identified ten common properties that are still prominent in today's Web traffic. Crovella *et al.* [8] discovered self-similarity in Web traffic, based on traces collected from both clients and servers.

Video traffic has also been a main subject of study. For example, Cha *et al.* [5] studied user-generated videos on YouTube, while Gill *et al.* [14] studied YouTube usage on a campus network. More recently, Laterman *et al.* [16] provided a campus-level characterization study of two popular video streaming services (Netflix and Twitch), and identified the similarities and differences between these two services.

In recent years, several email-related measurement studies have been conducted. Dominik *et al.* [24] presented a flow-based classification of Webmail traffic. By extracting features from HTTPS flows, their method can detect Webmail traffic with high accuracy, even when it is encrypted. Elieb *et al.* [10] presented a comprehensive measurement study of SMTP security extensions. This research indicated that top mail providers are more proactive in adopting new security configurations. Ramachandran *et al.* [21] characterized the properties of spammers, and developed an approach to detect spam based on network-level footprints. In [11], Steven *et al.* discussed privacy issues in email transfers. They pointed out that third-party tracking is a problem in commercial email. In [15], Trinabh *et al.* addressed the tradeoffs between email privacy and essential functions such as spam filtering. They showed that spam filtering is still achievable, even with end-to-end encryption.

These recent prior works focus on the privacy and security of email. However, to the best of our knowledge, the workload of email traffic itself has not been studied, which provides the primary motivation for our work.

3. METHODOLOGY

In general, network traffic measurement approaches can be divided into *passive* and *active* measurement, based on the data collection mechanism. For a fuller picture of how Outlook email works, we use both approaches in our work. Besides, This approach is commonly used for traffic measurement, which can be generalized to study other Email systems as well.

3.1 Passive Measurement

In passive measurement, data is gathered by listening to the Internet traffic passively. Specifically, it does not generate any additional traffic of its own. This approach is used for collecting network statistics or evaluating network performance, especially using long-term measurements.

In our work, we use the network traffic collected by an Endace DAG network monitor, which is installed at the edge router of our campus network. The functionality of our monitor is similar to that of Wireshark [6], but uses specialized hardware to capture packets at multi-Gigabit rates. The monitor is equipped with two Intel Xeon E5-2690 CPUs, 64 GB RAM, and 5.5 TB of hard disk for data storage.

We use the Bro [20] intrusion detection system to process our dataset. Bro takes the headers of all captured Internet traffic as input, and produces hourly connection-level summaries as output. Since the vast majority of email-related traffic uses HTTPS, our primary interest is the SSL logs, as well as connection logs [12]. Collectively, these logs describe the TCP connection information, including transport-layer endpoints, number of packets transferred, and number of payload bytes sent/received.

3.2 Active Measurement

We also employ active measurement approaches for understanding network activities. This approach allows us to test the behaviour of Outlook email services under different scenarios (e.g., login, email sending, email reading, logout).

The main tool we use is the Man-In-The-Middle (MITM) proxy [7], as illustrated in Figure 1. This setup allows us to decrypt HTTPS messages in our own test email sessions. These traces are used to discern essential information missed in passive measurement of encrypted email traffic.

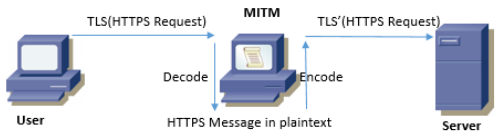

Figure 1: Example of MITM proxy configuration

We also use other common tools to conduct active measurements, such as MaxMind for IP address geolocation, `traceroute` for routing, and `ping` for network latency.

3.3 Server Classification

In our logs, we identify Outlook email-related traffic using specific IP address ranges. However, Outlook email itself consists of several different components, and each has different IP ranges. Therefore, we introduce these components here, along with the IP ranges for each.

The servers supporting Outlook email can be divided into three main categories [17]: Main Server (Portal and Exchange Online), Content Distribution Network (Exchange Online CDN), and Spam Filtering (Exchange Online Protection). Table 1 summarizes the IP ranges for these services, along with the underlying SMTP, POP, and IMAP services.

Table 1: IP ranges for different services involved in Outlook email.

Service	Domain Name	IP Address	Location	Port
Portal and Exchange Online	portal.office365.com	13.107.6.156 52.109.2.136	Redmond, WA San Francisco, CA	443
	*.office365.com	40.97.0.0/16 40.100.0.0/16 40.101.0.0/16	Redmond, WA	
Exchange Online CDN	r1.res.office365.com	23.41.182.156 104.125.231.63 184.27.135.86 ...	Seattle, WA	443
	r4.res.office365.com	23.36.176.113	Vancouver, BC	
Exchange Online Protection	protection.outlook.com	23.103.157.42 23.103.157.10	Toronto, ON	443
SMTP server	smtp.office365.com	40.97.0.0/16	Redmond, WA	587
POP3 server	outlook.office365.com	40.97.0.0/16	Redmond, WA	995
IMAP server	outlook.office365.com	40.97.0.0/16	Redmond, WA	993

3.3.1 Main Servers

The Main Servers in Outlook have several responsibilities, including the sending and receiving of email messages, delivering shared resources, and maintaining the connection between clients and the Outlook server. The Main Server uses a Class B IP address range: 40.97.0.0/16. In addition, two other Class B IP ranges are used for additional servers: 40.100.0.0/16 and 40.101.0.0/16.

3.3.2 CDN Nodes

The Exchange Online CDN is used to deliver shared resources, such as icons and scripts. Microsoft indicates that both Azure and Akamai nodes can be used to assist with content delivery. However, only two CDN nodes are observed in email delivery for our campus, namely r1.res.office365.com and r4.res.office365.com. Based on our understanding, `r1.res` is used to support mobile devices, while `r4.res` handles requests from desktop PCs. In both cases, the resources are stored publicly on these nodes. That is, as long as you know the corresponding URL, you can access these shared resources without authentication (i.e., using HTTP).

3.3.3 Protection Server

Outlook Protection is a spam filtering server to maintain email security. If an email is potentially malicious, it is forwarded to the Protection server, which decides whether to deliver the email or tag it as spam. Generally, the Protection server is invisible to users, since it is a hidden step in the email delivery process. However, the SMTP server within our campus can talk to it directly. Based on our analysis, the traffic between these two specific servers is dominated by spam-related activities initiated from compromised computers within our campus network.

4. EMAIL SESSION STRUCTURE

In this section, we present a structural analysis of the network communication involved in the Outlook email service. We divide the Outlook email services into five major steps: Login, Authentication, Email Sending, Email Receiving, and Logout. Other less common interactions are ignored. We present a brief introduction of the mechanisms in each step, to better understand how the email services are provided.

4.1 Login

There are three ways to access the Outlook email service: Outlook PC application, Outlook Mobile application, and the Outlook Web site. The Mobile application approach is very similar to the PC application, therefore, we will focus on the other two approaches.

Using Outlook email via the Outlook application is the most common approach. In general, the login process is straightforward. The application first leads the users to outlook.office365.com to enter their email address. When a valid `ucalgary.ca` email address is typed, the login process is redirected to the Central Authentication Service (CAS) at the University of Calgary for authentication. Once authenticated, users are issued an authentication token, which can be used to access all Office365 services, including Outlook, Word, and PowerPoint. TCP connections are established with several other servers, such as `*.aria.microsoft.com` and `nexus.officeappps.live.com`. However, these servers are shared with other Office365 services, and are not directly related to the Outlook email service. Therefore, we simply ignore them in our discussion.

Compared with the application-based approach, the login process through a Web browser is rather complex. Several different servers are involved in the initial step of email service. The most frequently observed ones include the Login Portal (`login.microsoftonline.com`), the Delve service (`loki.delve.office.com`), the Aria service for Web browsers (`browser.pipe.aria.microsoft.com`) and the Web shell (`webshell.suite.office.com`). In a Web-based scenario, more steps are inherently required to support the service. However, there are many superfluous TCP connections, such as Delve and Skype, which are not essential to the email service.

4.2 Authentication

The authentication traffic occurs between the CAS server within our campus network and the authentication server at Microsoft's data center. When a user attempts to login to Outlook from the University of Calgary domain, Microsoft's authentication server asks the CAS server to check the user's identity, since all the authentication information is stored locally at the University of Calgary. If a user modifies his or her password, a new authentication process is required.

Figure 2 shows the Authentication step for Outlook email.

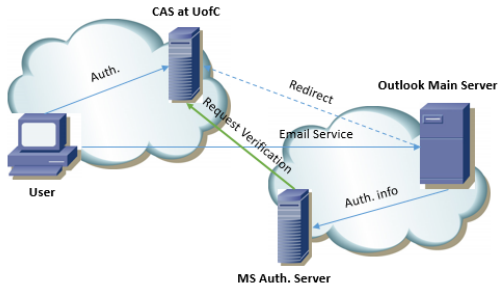

Figure 2: Authentication by Outlook and CAS (green line)

4.3 Sending Email

The sending of email is relatively straightforward. The email content is always sent directly to the Outlook main server, using the HTTP POST method. If there are email attachments, these get handled in one of two different ways, depending on their size. Small attachments are sent directly to `attachment.outlook.office.com`. Large attachments (20 MB or more) are redirected to OneDrive.

4.4 Receiving Email

Incoming email is received through a periodic detection procedure. After the login step, several parallel connections (typically 3 or 6) are maintained between the client and the Outlook servers. The client periodically sends (every minute) an HTTP POST request to the Outlook server to check for new messages. Interestingly, the update interval for the Web-based approach is only 10 seconds. Hence, when new email comes in, the Web-based Outlook often alerts the user much sooner than the application-based version.

4.5 Logout

There are several ways to end an email session, such as selecting "Sign Out" in the PC or browser Outlook application, or manually closing the browser window. In general, TCP connections should be terminated through a four-way handshake using FIN and ACK flags. However, Outlook does not use a normal FIN to terminate the session. Instead, a RST flag is commonly witnessed in the logout step, regardless of how the session was terminated. This phenomenon is a long-standing issue with Microsoft's TCP implementation [2].

Another interesting observation from our measurements is that connections with CDN nodes are terminated with *both* FIN and RST flags. To be specific, the client sends a FIN, followed shortly by an RST, to force connection termination. As a result, there is no ACK for the original FIN from the server. In general, about 35% of the connections with the Outlook server are terminated with RST, with the remainder ending properly. This seems to be caused by non-standard TCP configurations from Microsoft.

5. TRAFFIC OVERVIEW

In this section, we provide an overview of the Outlook traffic on our campus network in February 2018 (Figure 3). We investigate the traffic for each sub-service separately to show the distinctive patterns within different services.

5.1 Authentication Server

Figure 3a shows the number of TCP connections launched between CAS and the Outlook Authentication server. Ignoring the 3-day monitor outage early in the month, there are strong diurnal and weekly patterns seen. There are 8,000-10,000 TCP connections every hour during the busy part of each work day, with a slight decline on weekends.

The peak hour often appears early in the morning, reflecting the typical workday schedule for most people. However, Authentication traffic continues throughout the day and into the evenings. The explanation is that no matter where the user is when doing authentication (i.e., at work, at home, or on travels), it is captured by the monitor, since the Authentication server is physically on our campus.

5.2 Main Server

The main server (Portal and Exchange Online) handles most of the Outlook email traffic. In total, we recorded over 100 million individual TCP connections during our month-long observational period. These connections were all launched by clients within our campus, targeting the Outlook Main Server. This workload indicates how heavily the Outlook email service is used within our campus.

Figure 3b shows the number of TCP connections initiated for Outlook email in every hour of each day during the month of February 2018. A strong diurnal pattern is evident in this data.

On typical weekdays, the peak hour usually appears in late morning or early afternoon, with a slight downward dip during the noon-hour lunch break (see Figure 4). The peak hour of a day often has more than 500,000 individual TCP connections. However, this number declines rapidly in mid-to-late afternoon, as people start leaving campus for the day. There is often a small peak in the evenings as well, indicating that some people check their email just before bedtime.

Interestingly, the average email traffic early in the work week (e.g., Monday to Wednesday) is always slightly higher than that later in the week (e.g., Thursday and Friday). This indicates either more people on campus on those days (consistent with our earlier measurements of the Learning Management System on our campus network [23]), or that people are more active in email activities at the beginning of each week. The weekends have a lower volume of email traffic, with the peak hour below 100,000 TCP connections. Also, the peak hour during weekends tends to be in the evenings, rather than mid-day.

Our university's Reading Week break (February 19-23, 2018) is also evident in Figure 3b. Traffic volumes are lower because of the statutory holiday (February 19), for which the traffic is similar to that of a weekend. For the remaining four days of that week, the average volume of traffic is a bit lower than normal workdays, but the diurnal pattern remains the same. Although many students are away during the break, there are still many other faculty, staff, and students working on campus.

5.3 CDN Nodes

As mentioned earlier, there are two CDN nodes providing the Exchange Online CDN service. Figure 3c shows the number of CDN-related TCP connections in each individual hour of our data. The diurnal and weekly patterns are still apparent, though the number of connections is much lower. The peak hour on a weekday typically has 6,000 to 8,000

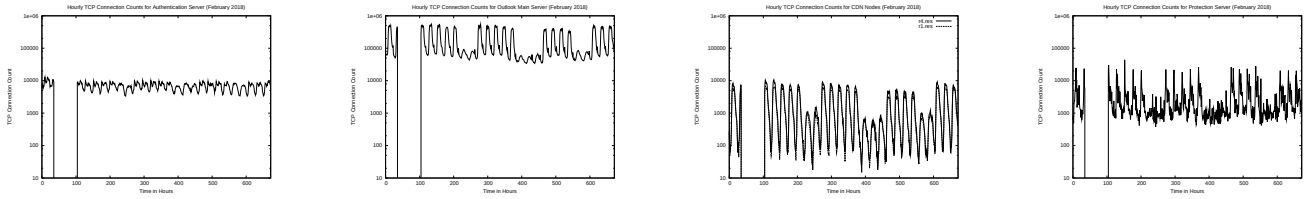

(a) Authentication server (b) Main server (c) CDN server (d) Protection server

Figure 3: Hourly Counts of TCP Connections Initiated for Email-related Traffic (February 2018)

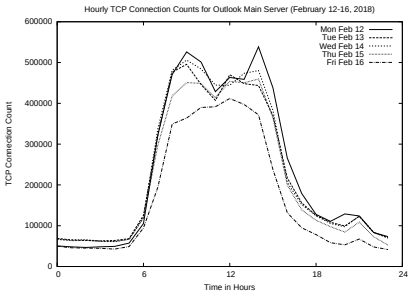

Figure 4: Detailed Hourly Count of TCP Connections with Outlook Main Server

TCP connections for `r1.res` and `r4.res`, respectively.

These connection counts are a lot lower than for the Main Server, since CDN nodes are not directly involved for sending/receiving emails. The nightly peak does not show up in the CDN workload either. This may indicate that although people check their email at night, CDN nodes are rarely involved. Similarly, the volume of traffic to CDN nodes remains low during Reading Week.

Another interesting observation is that the workloads of the two CDN nodes are not balanced. Specifically, `r4.res` has about 20% more traffic than `r1.res`. This may indicate that there are more PC Outlook users than mobile users. We explore these differences further in Section 6.

5.4 Protection Servers

The Online Protection server is responsible for spam filtering and other related services. Hence the number of Online Protection connections represents effort spent processing (potentially) suspicious email traffic.

In our campus network, all incoming email messages are first sent to the Protection server for a spam filtering check. This is true regardless of the originating domain (i.e., internal or external to the U. of Calgary) of the email message. However, externally-generated messages receive the most rigorous scrutiny.

Figure 3d plots the Online Protection traffic in each individual hour. The diurnal and weekly patterns remain, but the peak hour often appears at around 9:00 AM, and drops rapidly afterwards. Also, the count of TCP connections in the peak hour fluctuates a lot. Surprisingly, the workload during Reading Week is not reduced, compared with other weekdays. This suggests some machine-initiated email activities, rather than human-driven activities. We believe that these patterns reflect spam-related activity between our campus SMTP server and the Protection server.

6. DETAILED ANALYSIS

In this section, we focus on two important features of Outlook email traffic, namely the duration and data volume of each individual TCP connection. These two features reflect the network-level workload of each connection. It also sheds some light on how the Outlook email service performs.

We use a similar structure to the previous section to present the results, with Authentication, Protection, Main Server, and CDN traffic analysis. However, we supplement our analysis with two new features. First, the communication with the Main Server is divided into two groups, based on how connections are terminated (i.e., FIN or RST). Second, we explore the reasons for the load imbalance between the two CDN nodes. Therefore, in this section, there are six (rather than four) graphs for each analysis.

6.1 Connection Duration

We study the duration of each TCP connection between our campus network and the Outlook email-related services. This duration is the elapsed time between the start (i.e., SYN) and the end (i.e., FIN or RST) of each TCP connection. In our analysis, we classify TCP connections into one of three categories, namely *short* (less than 10 seconds), *medium* (between 10 seconds and 5 minutes), and *long* duration connections (over 5 minutes). The latter (long) are typically used to maintain state between the client and the Outlook email server.

6.1.1 Authentication Server

Figure 5a shows the TCP connection durations with the Authentication server. The communication patterns have very consistent durations, with most connections lasting 1 to 5 seconds. This result makes sense, since the initial login procedure often finishes within a few seconds.

6.1.2 Protection Server

Figure 5b shows the TCP connection duration with the Protection server. These results show that the Protection connections last 120 to 200 seconds. Considering that the default timeout for persistent connections in Microsoft's Internet Information Server (IIS) is 2 minutes, it is clear that these sessions use persistent TCP connections.

6.1.3 Main Server

(a) Authentication (b) Protection

(c) Main Server (FIN) (d) Main Server (RST)

(e) CDN node `r1.res` (f) CDN node `r4.res`

Figure 5: CDF of TCP Connection Duration for Email-related Traffic (February 2018)

Figure 5c and Figure 5d show the durations of connections with the Outlook Main Server. These connections have a wide range of durations, no matter how the sessions ended. For FIN connections, 50% are short, 45% are medium duration, and only 5% are long. For RST connections, less than 20% are short, about 60% are medium duration, and 20% are long. The long sessions maintain state between the client and the Outlook email servers. Recalling that 35% of the connections are terminated with RST, we can infer that more than 60% of the long connections are not terminated properly.

6.1.4 *CDN Nodes*

The duration of connections with `r1.res` and `r4.res` are shown in Figure 5e and Figure 5f, respectively. In general, all connections with CDN nodes have a medium duration. This implies connections with CDN nodes are not terminated immediately after the data transfer. No long durations are witnessed in these connections.

There are also some perceptible differences in the CDFs for the two CDN nodes. For node `r1.res`, around 30% of the connections lasts 10 seconds. On the other hand, for connections with `r4.res`, a one-minute duration is typical, with about 35% of the sessions having this duration. This discrepancy suggests that mobile-based Outlook has smaller data transfers with CDN nodes than PC Outlook. We explore this further in the next section.

6.2 Data Traffic Volume

Our next analysis focuses on the number of data bytes exchanged on the TCP connections. Since TCP allows full-duplex communication, we do a separate analysis of inbound and outbound data traffic. Specifically, traffic that goes from the Internet to the University of Calgary campus is considered inbound, and that in the opposite direction is considered outbound.

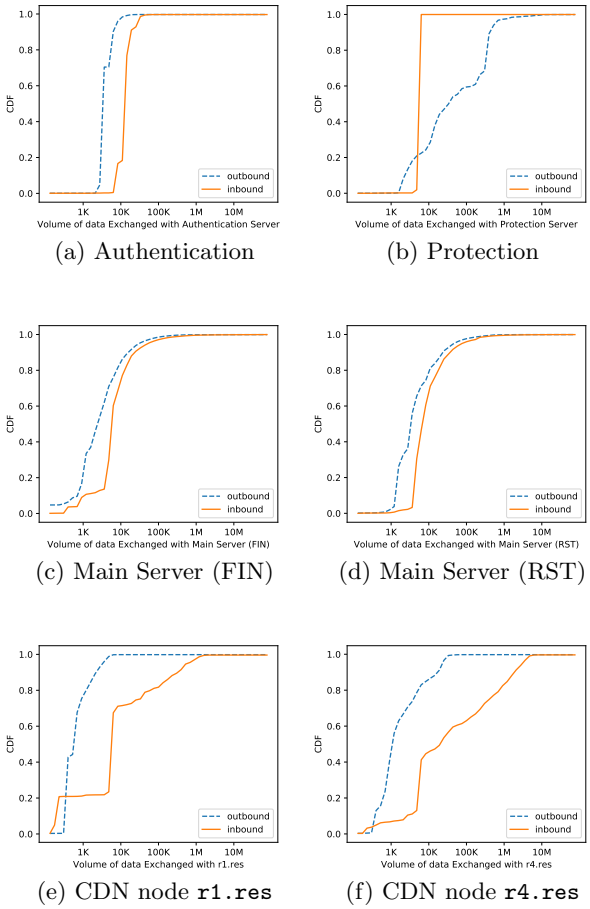

(a) Authentication (b) Protection

(c) Main Server (FIN) (d) Main Server (RST)

(e) CDN node `r1.res` (f) CDN node `r4.res`

Figure 6: CDF of TCP Connection Data Volume for Email-related Traffic (February 2018)

6.2.1 *Authentication Server*

Figure 6a shows the data volumes exchanged with the Authentication server. These results show consistent volumes for both inbound and outbound traffic. The typical volumes are 3 KB for inbound, and 12 KB for outbound. The data exchanged with the Authentication server involves user information, such as user credentials and (perhaps) some type of certificate. Our results show that these have a small and consistent size for all users.

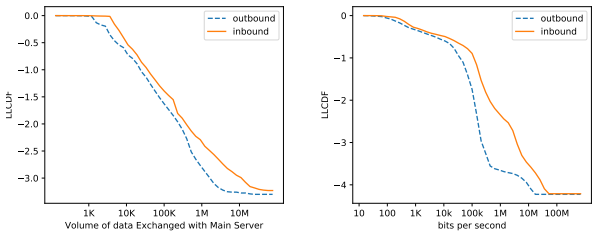

(a) Transfer size (bytes) (b) Throughput (bps)

Figure 7: LLCD plots for Main Server TCP Connections

6.2.2 *Protection Server*

Figure 6b shows the volume of data exchanged with the Protection server. The sizes for outbound traffic vary a lot (from 1 KB to 1 MB), reflecting the typical size distribution for email messages. On the other hand, the inbound traffic is extremely consistent in its size, perhaps reflecting a standard set of email headers for the spam filtering report. Our results suggest that the Protection server typically receives a single email message as input. No matter what is received, the Protection server returns about 4,800 bytes as a reply.

6.2.3 *Main Server*

Figure 6c and Figure 6d show the data volumes exchanged between clients and the Main Server. The two distributions are similar, regardless of how the session is terminated (FIN or RST). The only difference appears in connections with low data volume. For FIN connections, about 15% of them have less than 1 KB of inbound data, and less than 2 KB of outbound data. For RST connections, however, only 4% of them have these structural characteristics. This result implies that (short) sessions with limited data exchange are more likely to be terminated normally, while medium or long duration connections are more likely to have a RST.

Most inbound and outbound connections have less than 1 MB of data volume. For example, 90% of them exchange less than 100 KB of data. However, a few connections have very large data transfers. The largest inbound data volume seen was 1.7 GB. Figure 7a shows a Log-Log Complementary Distribution (LLCD) plot of the transfer sizes. This graph has power-law structure, showing evidence of a heavy-tailed distribution.

We also calculate the average throughput (in bits per second) for TCP connections with the Main Server (see Figure 7b). Most connections have very low throughput. This is not unexpected, since users maintain sessions regardless of whether there is new email or not. However, the largest outbound throughput achieved by any TCP connection is only about 12 Mbps, which is disappointing for large data transfers, especially over a Gigabit network. Furthermore, the throughput achieved by outbound connections is markedly lower than that for inbound connections. The highest throughput we witnessed reaches 50 Mbps. Detailed investigations using Wireshark show that inbound connections use TCP window scaling (1 MB) to improve throughput, while outbound connections are limited to a maximum TCP window size of 64 KB. This constrains the effective throughput, especially when uploading large email attachments.

6.2.4 *CDN Nodes*

Figure 6e and Figure 6f show the data volume for the two CDN nodes. As evident from the graphs, the workloads for `r1.res` and `r4.res` differ. For both nodes, the volume of inbound traffic varies widely (e.g., from 30 bytes to 2 MB), while outbound traffic does not (e.g., from 100 bytes to 5 KB). There are also some distinct step-like patterns in the CDFs. For `r1.res`, about 40% of the inbound sizes are about 8 KB, and 20% are about 40 KB. For `r4.res`, about 15% of the inbound traffic is near 8 KB, with most of the transfers larger than this. The frequently observed sizes indicate popular shared resources, such as icons or scripts. Upon login, users are required to download dozens of scripts, some of which are relatively large (e.g., 200 KB). Also, since the scripts are updated frequently, the inbound data volumes change over time.

7. CONCLUSION

In this paper, we presented a detailed measurement study of Outlook email service at the University of Calgary. Outlook email is a cloud-based service, and is heavily used by faculty, staff, and students at our university. However, the service itself has a very complex implementation.

The main contributions in our work are a better understanding of the operation of Outlook email, and a workload characterization of email usage on our campus network. We first identified four email-related services supporting Outlook email, namely Authentication, Protection, Main Server, and CDN. All these servers show strong daily and weekly patterns, which reflects the human-driven activities in email usage. However, based on the different roles for these servers, the characteristics of the workloads for these different servers are distinct from each other. Second, we introduced the five major structural components in an email session. Finally, we presented our overall and detailed measurement of network traffic with distinct email-related servers. The results reflect the usage of Outlook email service at the campus scale.

We identified several performance-related issues related to Outlook email. First, in the initialization step, Outlook builds many parallel TCP connections with extraneous servers, such as Delve and Skype, which are never used. These increase the delay for email session initialization. Second, although the TCP implementation at the server supports TCP window scaling, the maximum advertised receiver window on the server is only 64 KB. This setting results in unduly low throughput when uploading large attachments. Finally, Outlook uses TCP resets to terminate many of its TCP connections. While this does not in itself cause particular performance problems, it is non-standard, and limits the effectiveness of TCP/IP traffic analysis using passive network measurements. Since all RST comes from Microsoft, a redesign on both client and server side of Outlook Email Service could alleviate this problem.

Acknowledgements

The authors thank University of Calgary Information Technologies (UCIT) for facilitating the collection of our network traffic measurement data, and the ICNCC 2018 reviewers for their constructive feedback on our work. Financial support for this work was provided in part by Canada's Natural Sciences and Engineering Research Council (NSERC) through

the Discovery Grants program.

8. REFERENCES

[1] M. Arlitt and C. Williamson, "Internet Web Servers: Workload Characterization and Performance Implications", *IEEE/ACM Transactions on Networking*, Vol. 5, No. 5, pp. 631–645, October 1997.

[2] M. Arlitt and C. Williamson, "An Analysis of TCP Reset Behaviour on the Internet", *ACM Computer Communication Review*, Vol. 35, No. 1, pp. 37-44, January 2005.

[3] N. Basher, A. Mahanti, A. Mahanti, C. Williamson, and M. Arlitt, "A Comparative Analysis of Web and Peer-to-peer Traffic", *Proceedings of the 17th World Wide Web Conference*, pp. 287–296, Beijing, China, May 2008.

[4] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web Caching and Zipf-like Distributions: Evidence and Implications", *Proceedings of IEEE INFOCOM*, pp. 126-134, New York, NY, 1999.

[5] M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon, "I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User-Generated Content Video System", *Proceedings of ACM Internet Measurement Conference (IMC)*, pp. 1-14, San Diego, CA, November 2007.

[6] L. Chappell and G. Combs, *Wireshark Network Analysis: The Official Wireshark Certified Network Analyst Study Guide*, Protocol Analysis Institute, Chappell University, 2010.

[7] A. Cortesi, M. Hils, T. Kriechbaumer, *et al.*, `mitmproxy`: A Free and Open Source Interactive HTTPS Proxy, 2010-2018. https://mitmproxy.org/.

[8] M. Crovella and A. Bestavros, "Self-similarity in World Wide Web Traffic: Evidence and Possible Causes", *IEEE/ACM Transactions on Networking*, Vol. 5, No. 6, pp. 835–846, December 1997.

[9] J. Drucker and H. Goldstein, "Assessing the Regional Economic Development Impacts of Universities: A Review of Current Approaches", *Intl. Regional Science Review*, Vol. 30, No. 1, pp. 20–46, 2007.

[10] Z. Durumeric, D. Adrian, A. Mirian, J. Kasten, E. Bursztein, N. Lidzborski, K. Thomas, V. Eranti, M. Bailey, and J. Halderman, "Neither Snow nor Rain nor mitm...: An Empirical Analysis of Email Delivery Security", *Proceedings of ACM IMC*, pp. 27–39, Tokyo, Japan, October 2015.

[11] S. Englehardt, J. Han, and A. Narayanan, "I Never Signed Up for This! Privacy Implications of Email Tracking", *Proceedings of Privacy Enhancing Technologies Symposium*, pp. 109–126, Barcelona, Spain, July 2018.

[12] I. Foster, J. Larson, M. Masich, A. Snoeren, S. Savage,and K. Levchenko, "Security by Any Other Name: On the Effectiveness of Provider-based Email Security", *Proceedings of ACM SIGSAC Conference on Computer and Communications Security*, pp. 450–464, Denver, CO, October 2015.

[13] T. Fowdur, S. Aumeeruddy, and Y. Beeharry, "Implementation of SSL/TLS-based Security Mechanisms in E-commerce and E-mail Applications using Java", *Journal of Electrical Engineering, Electronics, Control and Computer Science*, Vol. 4, No. 1, pp. 13–26, 2018.

[14] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "YouTube Traffic Characterization: A View from the Edge", *Proceedings of ACM IMC*, pp. 15–28, San Diego, CA, November 2007.

[15] T. Gupta, H. Fingler, L. Alvisi, and M. Walfish, "Pretzel: Email Encryption and Provider-supplied Functions are Compatible", *Proceedings of ACM SIGCOMM Conference*, pp. 169–182, Los Angeles, CA, August 2017.

[16] M. Laterman, M. Arlitt, and C. Williamson, "A Campus-level View of Netflix and Twitch: Characterization and Performance Implications", *Proceedings of SCS SPECTS*, pp. 1–8, Seattle, WA, July 2017.

[17] Microsoft, Office 365 URLs and IP Address Ranges. https://docs.microsoft.com/en-us/microsoftteams/office-365-urls-ip-address-ranges.

[18] Microsoft, Server Settings of Office 365 Email Service. https://support.office.com/en-us/article/server-settings-you-ll-need-from-your-email-provider.

[19] Microsoft, Office 365 Content Delivery Network. https://docs.microsoft.com/en-us/sharepoint/dev/general-development/office-365-cdn.

[20] V. Paxson, "Bro: A System for Detecting Network Intruders in Real-time", *Computer Networks*, Vol. 30, No. 23-24, pp. 2435–2463, 1999.

[21] A. Ramachandran and N. Feamster, "Understanding the Network-level Behavior of Spammers", *Proceedings of ACM SIGCOMM Conference*, pp. 291–302, Pisa, Italy, September 2006.

[22] E. Rescorla, *SSL and TLS: Designing and Building Secure Systems*, Addison-Wesley, 2001.

[23] S. Roy and C. Williamson, "Why is my LMS so slow? A Case Study of D2L Performance Issues", *Proceedings of ISCA Computers and Their Applications (CATA)*, Las Vegas, NV, pp. 34-39, March 2018.

[24] D. Schatzmann, W. Mühlbauer, T. Spyropoulos, and X. Dimitropoulos, "Digging into HTTPS: Flow-based Classification of Webmail Traffic", *Proceedings of ACM IMC*, pp. 322–327, Melbourne, Australia, November 2010.

[25] H. Tschabitscher, "How many email users are there?", https://www.lifewire.com/ accessed Oct. 10, 2018.

[26] C. Williamson, "Internet Traffic Measurement" *IEEE Internet Computing*, Vol. 5, No. 6, pp. 70–74, 2001.

[27] S. Zhang, S. Zhang, X. Chen, and X. Huo, "Cloud Computing Research and Development Trends", *Proceedings of IEEE International Conference on Future Networks*, pp. 93–97, 2010.