Google

Google Workspace Networking Best Practices for Large Deployments

Administrator guide

Google Workspace

Google, Inc.
1600 Amphitheatre
Parkway Mountain View, CA
94043
www.google.com

Part number: NETBP_GAPPS_3.8

December 2020

Table of Contents

Chapter 1: Introduction

About This Guide

This document discusses best practices for optimizing your large-scale IP network for Google Workspace.

The recommendations and information in this guide have been gathered through our work with a variety of customers in both the public and private sectors and partners in many network environments. We thank our customers and partners for sharing their insight and experience.

Target Audience

This document is intended for Google Workspace customers with complex networks, especially those that are spread across a large geographical area. Administrators with smaller networks or networks in a single location may find some of this information useful and may find answers to specific questions, but some of the major network routing, capacity, and testing issues may not apply.

Benefits

Optimizing your network configuration will help you to improve your Google Workspace implementation in the following ways.

- Improve the responsiveness of Google Workspace by reducing latency in your network.
- Reduce bandwidth consumption by optimizing network routing and network services.
- Predict network performance and capacity needs by collecting baseline metrics for latency, packet loss, and network availability.
- Reduce upload and download times with Google Workspace for large files, such as internal videos and attachments.
- Efficiently migrate data from existing legacy servers into Google Workspace.

Level of Effort

The level of effort needed to implement the recommendations in this guide will depend on your requirements, your current infrastructure, and the skills of your network team. The design principles and implementation best practices in this document are not industry- or technology- specific. The principles in this document do not require specific technical expertise outside of an industry-standard network architecture and network engineering skill set.

Getting the most out of this guide

This guide includes testing and planning methodologies, answers to common questions about the impact of Google Workspace on IP networks, and results of field studies on best practices for integrating your network with Google Workspace.

This guide is designed to be used in the following ways:

- As a reference guide of network best practices and recommended network tools. The checklist in the next chapter provides a reference to each network topic, with links to further information.
- As an in-depth discussion of a variety of network best practices and related topics for Google Workspace services. You can read this document in its entirety to gain a detailed understanding of all topics related to networks and Google Workspace.
- As a reference guide to answer questions on specific topics about network best practices.

Life Cycle of your Google Workspace Implementation

The information presented in this guide is associated with multiple milestones of your Google Workspace deployment:

1. **Network Evaluation**: This section contains information on evaluating your current network before you deploy Google Workspace. While this information may be helpful after you have deployed Google Workspace, you will see the best results if you run these evaluations and tests before any other steps.
2. **Network Configuration**: This section contains notes and information on how to setup your network to work best with Google Workspace. This section includes network routing information, IP addresses, protocols and port numbers, proxy server configuration, DNS configuration, firewall setup, and mail server setup.
3. **Client Configuration**: This section provides advice on setting up the environment for your users. This includes client information, mobile network expectations, and migration.
4. **Network Monitoring**: This section includes notes on maintenance of your network, and troubleshooting if problems occur after a complete deployment.

Disclaimer for Third-Party Product Configurations

This guide describes how Google Workspace products work with common servers and the configurations that Google recommends. These instructions are designed to work with the most common scenarios. Any changes to your configuration should be made at the

discretion of your administrators.

Google does not provide technical support for configuring third-party products. In the event of a third-party issue, you should consult your network administrator. GOOGLE ACCEPTS NO RESPONSIBILITY FOR THIRD-PARTY PRODUCTS. You may also contact Google Solutions Providers for consulting services. Links to third-party Web sites are provided for your convenience. The links and their content may change without notice. Please consult the appropriate products' websites for the latest configuration and support information.

Chapter 2: Network Action Checklist

About This Checklist

This section contains a summary checklist of all action items in this guide. If you don't have the time to review this guide end-to-end, we suggest you start by reviewing this Network Action Checklist.

Each topic is described in detail later in this guide.

Network Evaluation

Evaluate your current network and plan for capacity needs. To achieve the best results during testing, use the following methodology:

- Conduct an inventory of all your network locations, including location name, Internet access type (e.g., T1, VPN, DSL), and available Internet bandwidth.
- Test DNS resolution from all network locations to Google Workspace, to ensure that clients in your network can resolve Google Workspace hostnames - and do so as close to the user as possible.
- Test TCP/UDP reliability from all locations to Google Workspace, to ensure that clients in your network can reliably establish and maintain a connection to Google Workspace.
- Assess WAN bandwidth between your Internet egress location and network locations which use that egress point.
- Identify any requirements for limiting access to Google services (including SSL inspection) and how those policies will be put in place.
- If you intend for your users to connect to Google Workspace through a proxy server, create a test environment and measure how many connections to expect per user, so you can calculate the expected number of outbound connections on your proxy server.

For more information on evaluating your network, see Test Your Network Environment and Proxy Server Evaluation and Sizing in later sections of this guide.

Network Configuration

The following recommendations describe network configurations that will help provide users the best experience with Google Workspace. These recommendations increase network availability and performance, and can reduce costs by simplifying the network equipment required to reach Google Workspace.

Network Routing

To achieve the best performance with network connections to Google Workspace:

- Use Google's network by egressing traffic to the Internet as quickly as possible. Route network traffic to the Internet as close to the end user as possible, in terms of both geography and topology.
- Focus on addressing latency issues over bandwidth requirements. Above a minimum level, bandwidth considerations are generally less significant for Google Workspace.
- Open your firewalls to the ports that Google Workspace services require. For details, see Google Protocols.
- Avoid using specific IP addresses to permit access to Google Workspace. See Google IP Addresses.
- If you are using a hub-and-spoke network topology—or if your network has multiple locations with a single network egress point spanning multiple continents—Google recommends identifying potential regional egress points.

For more information about network routing, see Network Addressing and Protocols and Network Routing later in this document.

Proxy Servers

- Avoid routing Google Workspace data through proxy infrastructure that inspects the content of HTTP traffic. This will reduce performance and be largely ineffective.
- While not recommended, if you are using a proxy server that supports SSL Terminations, set up your proxy server to inspect Google Workspace content while relaying the secure connection.
- When required, keep your proxy servers in a location that is close to your users and their Internet egress point, in terms of both geography and network topology.
- If you need to filter web traffic by URI, consider using a PAC configuration file on the client's desktop, since URIs in encrypted HTTP traffic are not visible to the proxy.

For more information about setting up proxy servers, see Proxy Servers later in this guide.

Other Services

- Use a DNS resolver in a location as close to the user as possible, in terms of both geography and network topology. Using DNS resolvers located in remote network locations will greatly slow down connections to Google Workspace.
- If it's not feasible to use a DNS resolver local to users, use a DNS server that supports the edns-client-subnet extension —such as Google's DNS server or OpenDNS—which allows the resolver to pass part of the client's IP address.
- Adhere to the advertised TTL value for all DNS record types.
- Set up firewall rules to allow unrestricted outbound HTTPS traffic to Google Workspace. You do not need to set up special rules for inbound traffic; Google Workspace does not

initiate inbound traffic to users unless specifically noted.

- If possible, avoid routing inbound and outbound mail through a gateway inside your network. If inbound and outbound mail is routed to a gateway inside your network, mail traffic will consume unnecessary network resources.

For more information on network services, see Other Network Services below.

Client Configuration

After your network is configured, prepare your user environment to work with Google Workspace. This can include setting up clients, SSO authentication, and preparing for data migration.

Client Access

When planning for clients that will connect to Google Workspace, consider the following:

- Google highly recommends utilizing Google Chrome as your default browser for Google Workspace services in order to provide the best possible user experience.
  .
- Other supported browsers include Mozilla Firefox, Microsoft Internet Explorer, Microsoft Edge and Apple Safari. However, some Google Workspace features and functionalities are not supported in these browsers. More information on browser support can be found here.

For more information on setting up client environments, see Client Access.

Authentication

If you plan to set up Single Sign-On (SSO) authentication, consider the following:

- Set up SSO servers in distributed network locations, rather than a central location.
- Implement SSO servers in conjunction with VPN servers (if required) to avoid routing authentication traffic of VPN connected users to a different location.
- Set up internal DNS servers to redirect SSO traffic to the nearest SSO server, and ensure that alternate SSO servers are in place for redundant service.

For more information on SSO Authentication, see Authentication.

Migration

Google Workspace deployments often involve migration traffic, either from local clients like Google Workspace Migration for Microsoft Outlook, or from server-side clients like Google Workspace Migration for IBM Notes and Google Workspace Migration for Microsoft Exchange.

Migration of legacy data into Google Workspace is typically a resource-intensive activity. If you plan to migrate user data, consider the following:

- In the case of a large migration, it is highly recommended to choose a server-side approach.
- Ensure that your migration servers are in the same location as your legacy data servers, or that the connectivity between servers has low latency and high bandwidth.
- Avoid routing traffic from the migration servers to Google through proxy servers.
- Assess your network capacity before migration to determine the maximum amount of data that you can migrate concurrently. Adjust your migration plan accordingly.
- During migration, some of the connections established to Google servers can stay open for extended periods of time depending on the migration tool. To avoid possible disruption, and to reduce the need to re-migrate data, it is important to keep these sessions open and not close them prematurely with any in-line network infrastructure.

For more information about data migration, see Migration below.

Network Monitoring

Use monitoring tools to maintain and administer an existing IP network that is already working with Google Workspace.

- There are a variety of network monitoring tools available that are well-suited to monitor Google Workspace traffic. For a list of recommended network monitoring tools, see the Monitoring Tools section later in this guide.
- Network Packet Captures can assist in identifying possible performance issues during troubleshooting - or when working with your network provider / Google support. For more information, see Network Packet Captures.

Bandwidth Measurement

Bandwidth needs vary widely between customer environments and different usage patterns; no single measurement will fit all use Google Workspace customer profiles.

Some recommended practices include:

- Design and execute a plan to monitor bandwidth usage in the early phases of any Google Workspace deployment.
- While deploying Google Workspace, monitor bandwidth usage for a variety of user types at multiple locations to ensure data diversity.

Chapter 3: Network Evaluation

Summary

When planning for implementation, you will achieve better results if you first understand your current network capacity and the expected amount of network load from Google Workspace. The best way to predict load is to benchmark bandwidth usage in your network, and create a test environment to simulate how much capacity an average user will require. This section discusses approaches to test your network environment.

If you have already deployed Google Workspace without running environment tests and benchmarking, it may still be valuable to do so. This may give benchmarks for future planning and capacity requirements, and may help to identify potential issues before they impact user experience.

Test Your Network Environment

Network testing prior to a Google Workspace implementation is primarily focused on assessing capacity and identification of network bottlenecks, Internet proxies, firewalls, and any other network components involved in routing or monitoring Internet-based traffic.

Below are the recommended steps to assess and test your network prior to deployment.

- Conduct an inventory of all of your network locations, including location name, Internet access type (e.g., T1, VPN, DSL), and available Internet bandwidth.
- Test DNS resolution from all network locations to Google Workspace, to ensure that clients in your network can resolve Google Workspace hostnames.
- Test ICMP connectivity from all network locations to Google Workspace, to ensure that clients in your network can reach Google servers.
- Test TCP/UDP reliability from all locations to Google Workspace, to ensure that clients in your network can reliably establish and maintain a connections..
- Assess WAN bandwidth between your Internet egress location and network locations which use that egress point.

Inventory of Network Locations

When planning for a Google Workspace implementation, it is important to create an inventory of all locations from which users will access Google Workspace. The goal of this inventory is to gather information about the Internet connectivity and capacity from each network location.

When conducting an inventory, include the following information about each network location:

- The location's name and a description of its Internet access. Example: "Headquarters, DS3."
- Internet bandwidth average and peak usage. Example: "50% average usage, 70% peak usage."
- Number of proxy servers, and current average and peak usage.
- Number of firewall appliances, and current average and peak usage.
- Number of DNS servers, and current average and peak usage.

Once you have collected this information for every network location, use the data to assess current capacity, and whether upgrades are needed.

Network Testing

Use the information you have gathered during your network inventory to test each network route, DNS server, and proxy server. Run the following tests for all relevant network connections in each location.

Note: The third-party testing software described in this section is available for various operating systems, including Linux, Unix, Mac OS X, and Windows.

DNS Resolution Test

Ensure that clients in your network can resolve Google Workspace hostnames and URIs by testing DNS Resolution from all of your network locations to Google Workspace hostnames, as follows:

1. Open the sample list **text file hosted on GitHub**. (Note that this is a sample only, not a complete list.)
2. Save the sample .txt file in the directory where you will be using the test commands:
   a. Click **View raw file**.
   b. Right-click the page, and click **Save As**.
3. Run the following command to test DNS resolution:

```
% dig +all +trace -f GoogleAppsDomains.txt
```

ICMP Connectivity Test

Ensure that clients in your network can reach the hostname mail.google.com, by testing ICMP connectivity from all your network locations to Google Workspace. Test that your users can reach
Google Workspace, especially from all users' VLANs.

```
% ping -s 512 -c 400 -n mail.google.com
```

If you see slow or failed connections on your ping requests, this may indicate a loss of connectivity. Investigate each step of your connection to identify the source of the problem.

TCP/UDP Reliability Test

Ensure that clients in your network can reliably establish and maintain a connection to Google Workspace servers. Use the Hping tool to test link reliability over a period of time. Run the following command:

```
% time hping3 -S mail.google.com -p 443 --fast -c 1000
```

Run this test for each domain listed in the GoogleAppsDomains.txt file mentioned above.

Note: TCP/UDP reliability tests are intrusive and can affect network performance. Run these tests during off-hours to gather data while causing the minimum possible the impact to your network.

Available WAN Bandwidth Assessment

Use the iperf tool to assess the amount of bandwidth available from each location to its network egress point. This test is run on both the client and the network egress point.

This test is intended to assess bandwidth within your WAN network. It is not suited for testing bandwidth between your network and Google Workspace servers.

On each remote location that is connected over a WAN network, run this command:

```
% iperf -c CLIENT IP ADDRESS -d
```

On the network egress location, run this command:

```
% iperf -s
```

Note: WAN Bandwidth tests are intrusive and can affect network performance. Run these tests during off-hours to gather data while causing the minimum possible the impact to your network. If you need to run these tests during business hours, be careful of possible effects this test may have on your network performance.

Network Testing Tools

You can obtain the tools discussed above from the following online sources:

- Download the Hping packet analyzer tool from **hping.org**.
- Download the iperf bandwidth performance measuring tool from **SourceForge**.

Proxy Server Evaluation and Sizing

In a cloud computing environment, there are typically more outbound requests for external hosts than are generated in a legacy environment. The increase in outbound requests may impact the number of proxy servers required in your network.

If you intend for your users to connect to Google Workspace through a proxy server, you can determine what level of proxy server load to expect by running these tests beforehand. Use this information to estimate whether you need to increase your proxy server capacity.

Google Services largely leverage asynchronous calls within a browser session. For this reason, try to avoid implementing connection limiting parameters for user sessions if possible.

Follow these steps to evaluate your proxy server needs:

1. Create a test environment with each platform and browser that you plan to use in your user environment.

2. For each browser, measure the number of connections that occur during testing, including minimum and maximum concurrent connections, both for idle use and active use.

3. Based on this information and the number of users you expect on your system, calculate the expected number of connections for your proxy server.

4. Use these calculations to plan for any proxy server capacity changes needed. See below for more information on these steps.

Benchmark Proxy Load Per User

To benchmark the amount of proxy resources used by a typical user, establish a testing environment where you can test the various platforms and browsers that you support. Your testing environment should include machines on your network that can connect to Google Workspace using the same routes that you plan to use for your users. Once the testing environment is ready, direct traffic to a test proxy where you can measure the number of connections.

Collect the following data for each environment, while using Google Workspace services available in your domain. For instance, open Gmail, Google Hangouts, Google Docs, and Google Calendar.

- Average connections/sec
- Peak connections/sec
- Non-peak connections/sec

Additionally, Google Workspace, like many web-based applications that run in the cloud, keeps several connections open to the remote server to poll for new data. To evaluate the load caused by these open connections, measure the following in your test environment.

- Minimum amount of connections an idle user has with your browser platform
- Maximum amount of connections an idle user has with your browser platform

Once you have gathered these numbers, you can compile this information to estimate the load you might experience given your unique environment.

Estimate Expected Proxy Resources Needed

To estimate the amount of load you can expect during a Google Workspace rollout, multiply the number of connections for each test environment by the number of users you expect for that environment.

Use the following calculations.

Estimated average load = Sum (average load of each test machine environment X estimated number of users who will use that environment)

Estimated peak load = Sum (peak load of each test machine environment X estimated number of users who will use that environment)

Estimated idle load = Sum (idle load of each test machine environment X estimated number of users who will use that environment)

If the estimated average load, plus any additional traffic your proxies handle, exceeds your current capacity, make plans to expand your proxy server capacity, or change your proxy server implementation so that your proxy servers do not handle the requests that your users will make to Google Workspace.

Example

In the following example, a large enterprise plans to deploy the following:

- 5,000 users running Chrome on Windows 10.
- 3,000 users running Firefox on Windows 10.

Note: Google recommends all users run Google Chrome for best performance with Google Workspace.

During benchmarking, tests show the following sample numbers of concurrent connections through the proxy server. (Note: These are for example only. Your environment will vary.)

- **Chrome on Windows 10**

  Connections when entering URI: 1
  Connections during initial load: 3
  Connections during login: 6
  Connections after a few minutes idle: 4
  Connections when opening Calendar and Docs: 4
  Connections when loading a document: 6

  Average load: 3.6 connections
  Peak load: 6 connections
  Idle load: 3.1 connections

- **Firefox on Windows 10**

  Connections when entering URI: 1
  Connections during initial load: 4
  Connections during login: 9 Connections
  after a few minutes idle: 3
  Connections when opening Calendar and Docs: 11
  Connections when loading a document: 17

  Average load: 4.1 connections
  Peak load: 17 connections
  Idle load: 3.8 connections

Based on this, the expected load is:

- **Average**: (5000 x 3.6) + (3000 x 4.1) = 30,300 connections.
- **Peak**: (5000 x 6) + (3000 x 17) = 81,000 connections.
- **Idle**: (5000 x 3.1) + (3000 x 3.8) = 26,900 connections.

Based on this estimate, the proxy environment needs to be able to support at least 30,000 connections, possibly more to avoid problems during peak periods, or if growth is expected. If the current proxy server environment is running at 50% capacity with 20,000 connections, this is a sign that it will be necessary to deploy significantly more proxy servers, or to route Google Workspace traffic so that it bypasses the proxy server.

Chapter 4: Network Configuration

Summary

This section includes details on how to optimize your network for Google Workspace. This includes information on Google's IP addresses, protocols used, routing suggestions, proxy server configuration options, and DNS configuration. Use this information as a guide when configuring your network, and as a reference for what types of requests Google Workspace clients will make to Google servers.

Google Workspace traffic routing best practice

If required, the recommended and more robust way to allowlist Google Workspace traffic for traffic redirection and prioritization is to use IP ranges in conjunction with wildcarded hostnames.
Using IP ranges only or wildcarded hostnames only is not recommended. Refer to the following sections for more details.

Network Addressing and Protocols

Google IPv4 Addresses

Google Workspace services can be accessed both via IPv4 and IPv6 - we'll use IPv4 in the following examples but the same approach can be applied to IPv6.

Google Workspace exists in a multi-tenant server environment that includes both Google Workspace and consumer services. Therefore, Google Workspace shares the same IP address space as Google's consumer services. It also means that different services can be run from the same IP range. For example, Google Docs servers could use the same IP address space as Google Photos - and service both enterprise and consumer users. In addition, a specific IP address for a Google hostname, such as mail.google.com or drive.google.com, might be serving *both* Google Workspace and consumer users at the same time. This allows for unparalleled reliability for all users of all services.

Since Google Workspace uses the same IP address space as other Google products (including consumer products), it is also not always possible to distinguish traffic to different services using IP addresses.

For any Google hostname, such as mail.google.com or docs.google.com, IP addresses are not static and is valid only for its time-to-live (TTL) value returned in the DNS lookup of the hostname.

For example, if we query the A record for mail.google.com, several results are returned (note that output is not meant to be authoritative):

```
% dig a mail.google.com +ttl

;; ANSWER SECTION:

mail.google.com.          60    IN    CNAME    googlemail.l.google.com.
googlemail.l.google.com.  60    IN    A        216.58.198.229
```

The second column in the result set is the TTL for the records in seconds. Based on these sample results, we can determine that the IP addresses are valid for only a minute.

Google IP addresses for specific hostnames are not static. For example, do not assume mail.google.com will always be 216.58.198.229 - or any other result you receive in testing. If you need to configure your environment to accept mail from Google for a mail gateway, include all of the subnets from '_spf.google.com' record per this **Help Center** article.

It is not recommended to use Google's IP address space to permit access to Google (see Google Global Cache below); however, IP addresses can be used to implement traffic redirection and prioritization to the Internet knowing the implications of Google Global Cache (a recommendation stated throughout this document).

A more robust option to implement these prioritizations can be Google's hostnames (see Google Host Names, below) in conjunction with Google's IP space.

Google IP ranges can be obtained following this Help Center article.

We highly encourage customers to monitor the updates and implement a tracking script (see the Monitoring Tools section below).

Google Host Names

Google owns and operates a large amount of domains to serve our applications. To efficiently serve and operate such a large, global presence requires advanced network engineering and optimizations. We do not recommend using Google's hostnames as a means to allow access.

Rather, hostnames should be used to implement traffic redirection or prioritization to the Internet; a recommendation stated throughout this document.

However, a list of Google Workspace hostnames can be found in this Help Center article.

Note: The information contained in the Help Center article is subject to change without notice

Google Global Cache

Many of Google's services and applications participate in the Google Global Cache (GGC) content delivery system. The goal of this system is to provide the best service to all users by locating termination points as close to users as possible.

The GGC system involves Network Operators and Internet Service Providers in the distribution of commonly accessed resources - mostly static content. The participants in GGC have deployed a number of Google owned and operated servers inside their network to serve popular Google content. This results in IP addresses being used with Google services and applications that are owned by these host operators. Therefore, any use of Google's IP addresses to allow access should not be used. Rather, IP addresses may be used to implement traffic redirection or prioritization knowing that there may be some Google related traffic going to IP addresses not listed.

Google's use of GGC for content delivery is most effective for users with a large "network-distance" from Google (see Google's data center locations). Google's use of GGC is dynamic in both the services and client networks it applies to. Refer to the frequently asked questions at peering.google.com for more information related to GGC and its use.

Google Protocols

The table below shows common Google Workspace services, and the protocol used for each. As shown in the table, Google Workspace services are always SSL based except in the case of Google Meet.

Application	Protocol	Port
Mail, Calendar, Docs, Sites	TCP	443
Google Workspace Sync for Microsoft Office	TCP	443
Google Chat	TCP	443
Google Meet	UDP UDP TCP TCP	19302 - 19309 19302 - 19309 80 443
Google Workspace Migration for Microsoft Exchange	TCP (API)	443
Google Workspace Migration for Lotus Notes	TCP (API)	443

Google Meet

To provide users best performance with Google Meet refer to the Admin Help Center articles Network connectivity requirements and Optimize your Network for Google Meet.

Some networking considerations Google Meet include:

- Proxying traffic adds latency and may cause Meet to automatically reduce the video and audio quality, therefore it is not recommended to use proxy servers for Meet traffic.
- Avoid using packet inspection or protocol analyzers for Meet traffic; they introduce latency that may cause the Meet infrastructure to automatically reduce video meeting quality.
- Google recommends not utilizing QoS for Meet in your network. Meet automatically adapts to network conditions. If you have compelling reasons to use QoS for meet, please refer to the Meet QoS best practices guide.
- Make sure that network latency is low and consistent so that Meet traffic takes the shortest path between the client and Google.
- Make sure your network has enough bandwidth to handle all concurrent video meetings in a location. Refer to the Help Center article for the recommended bandwidth.

- Open the outbound ports to allow UDP and TCP traffic to flow to and from your networks. If you don't want to allow UDP out from clients on your network, at a minimum, permit TCP out from clients on your network to Google (see the Help Center article Optimize your Network for Meet for more detail). Forcing a TCP connection for services such as voice and video may create a poor experience for your users; therefore, we recommend allowing the use of UDP out from your network.

Network Routing

When routing to Google Workspace, the simplest network routing generally provides the best performance. Reduce complexity and unnecessary network routing from users' locations to Google data centers. A primary goal for your network design should be to reduce the total round trip time from your network to Google. If you see performance issues, address any latency problems before increasing bandwidth as this will most often yield greater results.

To achieve the best performance with connections to Google Workspace:

- Egress network traffic to the Internet as close to the end user as possible, in terms of geography and network topology.
- Focus on addressing latency issues over bandwidth requirements. Above a minimum bandwidth level, bandwidth considerations are generally less significant for Google Workspace.
- Ensure global firewalls are open to all ports Google Workspace services utilize.
- Consider traffic prioritization if you are using a hub-and-spoke network topology - or if your network has multiple locations with a single network egress point.

WAN Optimization

When planning your network cloud strategy, try to reduce latency and round-trip time. Users in remote offices will experience reduced performance if WAN traffic must traverse large geographic areas to reach the Internet. Implement network egress points as geographically close as possible to the user to reduce traffic across your bit-expensive links. Part of this optimization can be accomplished through DNS resolution changes.

For more information, see DNS Resolution on page 26.

Traffic Prioritization

You may be able to improve Google Workspace performance with traffic prioritization. This is accomplished by giving Google Workspace traffic priority over other network traffic to reduce latency during congestion. Traffic prioritization is possible on the data link and network layers; see the sections below for more information.

You may wish to consider traffic prioritization to reduce potential latency if you have any of the following environments:

- Hub and spoke network topologies.

- Multiple locations with a single network egress point.

Network Layer (Layer 3) Prioritization

As noted earlier in this document, Google Workspace uses the same set of IP addresses that other Google products use, including consumer products like Gmail and Google Photos. It is not possible to distinguish traffic to different products.

If you require network-layer prioritization, we suggest you do one or more of the following:

- Create a proxy PAC file that directs Google Workspace traffic to a proxy that routes only Google Workspace traffic. For more information, see Proxy PAC file configuration on page 23.
- Configure your networking equipment to prioritize your proxy network interface.
- Distribute proxies to avoid the creation of a hub and spoke proxy topology.

For information on the Google IP addresses and TCP Port usage, see Google IP Addresses on page 17.

Peering

Peering is the direct interconnection of your network to Google's network. This reduces latency and improves the reliability of the connection between your network and Google.

For most Google Workspace customers, the best way to do this is to choose an ISP or network provider that already peers with Google. Google peers with many Internet Service Providers in most geographies across the globe. This is the easiest and fastest way to realize the benefits of peering to Google. Contact your ISP to find out if they have peering established with Google.

For larger corporate networks, it may be possible to peer with Google directly. There are a number of requirements to peer with Google. In general, if you are not peering with other networks already, then it is more appropriate to let your upstream network provider handle peering relationships.

For Google's peering requirements, which apply to ISPs, network operators, and corporate networks, see the Google entry on **PeeringDB**. PeeringDB also contains the list of Internet Exchanges and other locations where Google is capable of peering.

If you or your Internet Service Provider qualifies for peering based on Google's peering requirements, discuss a peering relationship with your Google deployment or support representative.

Network Routing Tools

- A variety of useful tools are available to generate detailed data regarding your Internet connection performance on the external website **Measurement Lab**. You can use these tools to measure your overall Internet access performance.

Proxy Servers

When planning your proxy infrastructure for Google Workspace, keep in mind the following best practices:

- Avoid routing Google Workspace data through a proxy that inspects the content of HTTPS traffic as this will reduce performance.
- Keep your proxy servers in a location that is close to your users and their Internet egress point, in terms of both geography and network topology.
- If you need to filter web traffic by URI, consider using a PAC configuration file on the client's desktop, since URIs in encrypted HTTP traffic are not visible to the proxy.
- If you are using a proxy server that supports SSL Terminations, you can set up your proxy server to inspect Google Workspace content while relaying the secure connection.

Proxy Server Configuration

We recommend that you do not route Google Workspace traffic through a proxy server unless you have compelling reason to do so. If you decide to send Google Workspace traffic through your proxy, look for settings on your proxy server that might disrupt Google Workspace traffic.

Look for configurations and settings that include the following conditions:

- Content filters that might mark Google-related traffic as prohibited
- Settings that can lower the total amount of possible concurrent connections/sec per client
- Exceptionally long or short SSL time-outs (The default setting is recommended)
- Outdated firmware versions
- SSL Inspection without hardware acceleration

Filtering Google Workspace traffic through a Proxy

The vast majority of traffic originating from your users to Google Workspace servers consists of HTTPS transactions. This type of traffic is preferred because it is secure and reliable. Although interruption of traffic to Google Workspace for filtering is possible, it can drastically reduce the overall experience for your users.

In browsers and protocols that support the Server Name Identifier (SNI) extension to TLS, you will see the request for the hostname in the initial HELLO from the client in your proxy logs. A list of those browsers is available on the following page in **Wikipedia**. Consult your browser documentation to learn about SNI support.

After the initial HELLO request between the client/server and once the TLS connection is established, all traffic is encrypted including the URI path after the hostname.

If you need to filter your users' traffic, there are two recommended ways to accomplish this:

- Filter your users traffic with a proxy PAC file at the browser level prior to encryption is easier and less costly to implement. See Proxy PAC file configuration below.
- Perform SSL interception and inspection after the encryption is more secure but is more difficult and costly to implement. See SSL Inspection on page 24.

Proxy PAC file configuration

A Proxy PAC file is a cost-effective way to filter traffic because URI and IP evaluation is performed on the client machine prior to encryption.

A proxy PAC file is a set of JavaScript commands that the browser uses to evaluate against the URI requests received from the user.

The following sample script includes code to test

- If a URI is a plain hostname
- If a URI matches one of the Google Workspace wildcarded hostnames.
  Note that it is recommended to use two different rules to catch traffic for a top-level domain and potential subdomains (for example https://*.google.com/* and https://google.com/*)
- If the IP is a private address
- If the IP is in Google Netblocks

In all the above cases, the request will follow the "DIRECT" route, otherwise it will be routed through the default proxy server.

Note: This PAC file is meant to be provided only as an example. IP addresses and URL lists should be checked by a network administrator with the recommendations provided throughout this document and should be updated, owned and maintained by the Enterprise.

```
function FindProxyForURL(url, host) {

       // Plain hostnames. ( e.g. http://server )
       if (isPlainHostName(host)) {
              return "DIRECT";
       }

       // Private address classes.
       if (isInNet(dnsResolve(host), "10.0.0.0", "255.0.0.0") ||
           isInNet(dnsResolve(host), "172.16.0.0",  "255.240.0.0") ||
           isInNet(dnsResolve(host), "192.168.0.0", "255.255.0.0") ||
           isInNet(dnsResolve(host), "127.0.0.0", "255.255.255.0")) {
              return "DIRECT";
       }

       // Google Netblocks (_netblocks.google.com)

       if (isInNet(dnsResolve(host), '216.239.32.0', '255.255.224.0') ||
           isInNet(dnsResolve(host), '64.233.160.0', '255.255.224.0') ||
           isInNet(dnsResolve(host), '66.249.80.0', '255.255.240.0') ||
           isInNet(dnsResolve(host), '72.14.192.0', '255.255.192.0') ||
           isInNet(dnsResolve(host), '209.85.128.0', '255.255.128.0') ||
           isInNet(dnsResolve(host), '66.102.0.0', '255.255.240.0') ||
           isInNet(dnsResolve(host), '74.125.0.0', '255.255.0.0') ||
           isInNet(dnsResolve(host), '64.18.0.0', '255.255.240.0') ||
           isInNet(dnsResolve(host), '207.126.144.0', '255.255.240.0') ||
```

```
            isInNet(dnsResolve(host), '108.177.8.0', '255.255.248.0') ||
            isInNet(dnsResolve(host), '216.58.192.0', '255.255.224.0') ||
            isInNet(dnsResolve(host), '172.217.0.0', '255.255.224.0') ||
            isInNet(dnsResolve(host), '173.194.0.0', '255.255.0.0')) {
               return "DIRECT";
        }

        // Catch any wildcard Google domains that have fallen through.
        if (shExpMatch(url, "https://*.google.com/*") ||
            shExpMatch(url, "https://doubleclick.net/*") ||
            shExpMatch(url, "https://*.doubleclick.net/*") ||
             shExpMatch(url, "https://googleadservice.net/*") ||
            shExpMatch(url, "https://*.googleadservice.net/*") ||
             shExpMatch(url, "https://googledrive.com/*") ||
            shExpMatch(url, "https://gmail.com/*") ||
            shExpMatch(url, "https://ssl.google-analytics.com/*") ||
            shExpMatch(url, "https://*.googlegroups.com/*") ||
             shExpMatch(url, "https://googlegroups.com/*") ||
            shExpMatch(url, "https://googleapis.com/*") ||
            shExpMatch(url, "https://*.googleusercontent.com/*") ||
            shExpMatch(url, "https://*.gstatic.com/*") ||
            shExpMatch(url, "https://*.ggpht.com/*") ||
            shExpMatch(url, "https://*.googleapis.com/*") ||
            shExpMatch(url, "https://s.ytimg.com/*")) {
               return "DIRECT";
        }


        // Default rule falls back to the proxy servers.
        return "PROXY myproxyserver.corp.mycompany.com:3128; PROXY
myproxyserver2.corp.mycompany.com:3128";
}
```

isInNet() and *shExpMatch()* functions are used to evaluate the host to check if it is in one of Google's publicized IP network address blocks or to a wildcarded list of Google hostnames.

When the browser requests a page from a host that makes the function FindProxyForURL() evaluate to true, the configuration file will instruct the browser to use a direct connection to that host.
If the function returns false it will send everything to the proxy defined, e.g. "myproxyserver.corp.mycompany.com:3128".

More examples for developing a proxy PAC file can be found on the external website **FindProxyForURL**.

Proxy PAC file testing

Implementing a functional proxy PAC file requires careful testing. Use a PAC file testing tool like pactester to test different JavaScript functions. A PAC file tester will allow you to pass a hostname and URI and see which path the browser will take given your PAC file. Download pactester from the **GitHub pactester project site**.

SSL Inspection

Avoid SSL inspection if possible. SSL inspection is effectively an SSL "man in the middle attack" on your own users to examine the contents of HTTPS traffic. With SSL terminations, your users connect to a proxy as an end point. The proxy then terminates the SSL connection and inspects traffic, then establishes a new connection to the destination server

forwarding the traffic. This can cause a significant increase of load on traditional proxies that perform these operations in software, rather than a network appliance.

There are many commercial appliance vendors as well as many software proxy servers that can perform SSL inspection. Typically this requires additional proxy configuration.

Each proxy server SSL Inspection setup is different, but the typical steps are as follows:

1. Self-sign an SSL Certificate with an internal hostname, such as mail.example.com.
2. Install the mail.example.com certificate on the proxy server.
3. Write custom proxy rules. For instance, rewrite connections from https://mail.example.com/ to https://mail.google.com/a/example.com/.
4. Reject connections with a Host header that contains mail.google.com.

Note: Some proxies will allow you to keep the hostname the same, and use a built-in certificate. This requires that the user's browser trust the certificate, or users will receive a certificate error. For information on how to resolve these problems related to SSL inspection, consult your proxy server vendor and documentation.

Blocking Access to Google Consumer Services

As an administrator, you might want to prevent users on your network from signing in to a Google service using a consumer account instead of the Google Workspace account you provided. For example, you may not want them to use their personal Gmail accounts. In addition, you might also want to prevent users from signing in to a Google Workspace account from *another* domain.

Note that implementing this requires the use of web proxy and SSL interception, a practice that is not recommended as stated throughout this document.

A common means of blocking access to web services is using a web proxy server to filter traffic directed at particular URIs or hostnames. This approach is ineffective in this case because all the URIs accessed between consumer and Google Workspace accounts are the same.

To only allow users to access Google services using specific Google accounts from your domain, you need the web proxy to add an HTTP header to all traffic directed to *google.com. The header identifies domains whose users can access Google services. Since most Google Workspace traffic is encrypted, your proxy server also needs to support SSL interception. (See Block access to consumer accounts in the Admin Help Center for a list of proxy servers known to support both SSL interception and HTTP header insertion.)

To prevent users from signing in to Google services using Google accounts other than those you explicitly specify:

1. Route all traffic outbound to google.com through your web proxy server(s).
2. Enable SSL interception on the proxy server.

   Since you will be intercepting SSL requests, you will probably want to manage client certificates on every device using the proxy, so that the user's browser does not issue warnings for the requests.
3. For each google.com request:

a. Intercept the request.

b. Add the HTTP header X-GoogApps-Allowed-Domains, whose value is a comma-separated list with allowed domain name(s). Include the domain you registered with Google Workspace and any secondary domains you might have added.

For example, to allow users to sign in using accounts ending @altostrat.com and tenorstrat.com, create the following header with the domain names you want to allow:

X-GoogApps-Allowed-Domains = altostrat.com,tenorstrat.com

4. Optionally, create a proxy policy to prevent users from inserting their own headers.

Monitor URI Filtering

Avoid URI filtering with SSL inspection if possible. If you are using URI filtering, set up a policy to monitor URIs in proxy logs. Look for any URIs that were incorrectly blocked or allowed.

These changes in the accessed URIs can cause Google Workspace to load partially, slowly, or not at all. To avoid problems with URI filtering, if you are filtering your proxy servers, set up a policy for constant monitoring of your proxy load, and be prepared to adjust the rules if necessary.

To help discover what these new URIs might be, test new Google Workspace features or services in a test environment before allowing their use in production. To help with this you can install a tool like **HttpWatch** or **HttpFox**.

Proxy Configuration Tools

Download the following tools which may be helpful when configuring Proxy Servers:

- Use pactester or a similar tool to validate PAC files for different URIs. Download pactester from the **Github project site**.
- Download **HttpWatch** or **HttpFox** (Firefox extension) to help you see what URIs are being requested by the browser prior to encryption.

Countries or regions with restricted access

Google restricts access to some of its business services in certain countries or regions. Certain Google services might be available in these countries or regions for personal use, but not for business or education use. Refer to the Help Center for the list of countries or regions.

Other Network Services

Google runs a sophisticated load-balancing system to ensure the best experience for users. One factor in Google's load-balancing systems is the way Google answers DNS requests for services. Google attempts to determine the geographical location of a user partly through the location of the DNS resolver's IP address.

To ensure the best experience for your users:

- Use a DNS resolver in a location that is close to the user, in terms of both geography and network topology. Using DNS resolvers located in remote network locations will greatly slow down connections to Google Workspace.
- If it's not feasible to use a DNS resolver that's close to the user, use a DNS server that supports the edns-client-subnet extension (**Draft Proposal 2671**)—such as **Google's DNS server** or **OpenDNS**—which allows the resolver to pass part of the client's IP address.
- Adhere to the advertised TTL value for all DNS record types.
- Set up firewall rules to allow unrestricted outbound HTTPS traffic to Google Workspace. You do not need to set up special rules for inbound traffic; Google Workspace does not generally initiate inbound traffic to users.
- Avoid routing inbound and outbound mail through a gateway inside your network. If inbound and outbound mail is routed to a gateway inside your network it will consume unnecessary network resources.

DNS Resolution

The diagram below shows typical DNS resolution for a Google Workspace user on an enterprise network.

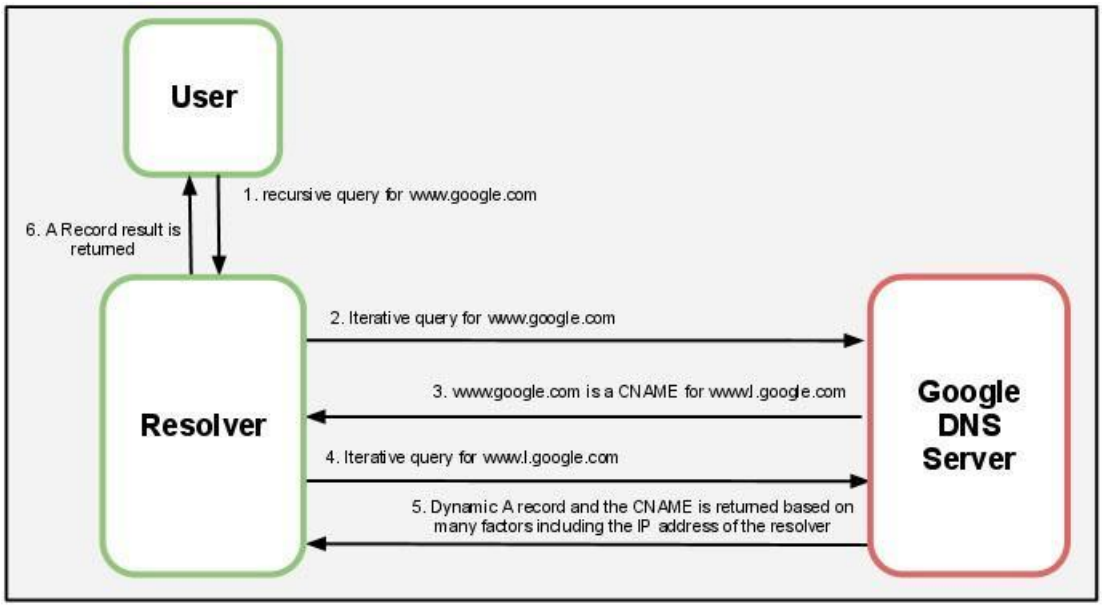

Google serves DNS A record queries dynamically to ensure users receive the best experience at the time they make their request. To ensure that this occurs properly, configure your DNS caching resolvers to adhere to the TTL values specified with each record. Using the cached result beyond the TTL value on the DNS record can lead to a poor experience for the user, because the cached DNS record might direct users to a suboptimal IP address.

Below is an example of the TTL values for www.l.google.com:

```
%dig +ttl www.l.google.com
```

For this query, you might see the following results:

```
; <<>> DiG 9.4.3-P3 <<>> +ttl www.l.google.com
;; global options:                printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 54488
;; flags: qr rd ra; QUERY: 1, ANSWER: 6, AUTHORITY: 4, ADDITIONAL: 4
;; QUESTION SECTION:
;www.l.google.com. IN A
;; ANSWER SECTION:
www.l.google.com.      184  IN   A  209.85.225.104
www.l.google.com.      184  IN   A  209.85.225.99
www.l.google.com.      184  IN   A  209.85.225.103
www.l.google.com.      184  IN   A  209.85.225.105
www.l.google.com.      184  IN   A  209.85.225.147
www.l.google.com.      184  IN   A  209.85.225.106
```

In this example, the TTL value is 184 seconds, which equates to 3 minutes. Be sure your DNS servers adhere to this value when caching results.

Using a centralized DNS server architecture will obscure the user making the request from Google's DNS servers. If DNS queries are routed through a central server to resolve Internet hosts, users may not connect to the closest Google Workspace servers. In extreme cases, this architecture can cause users in one continent to connect to servers in another continent.

The ideal solution is to place local DNS resolvers close to users and have the remote DNS resolvers send all DNS traffic through an Internet connection that's also local to users. For internal-only addresses, forward the requests to the appropriate internal corporate DNS server.

Alternatively, you can use a DNS service that supports the edns-client-subnet extension (**Draft Proposal 2671**), such as **Google's DNS server** or **OpenDNS**.

Note: Clients and DNS servers using the edns-client-subnet extension require more data to be sent with the request, causing the traditional 512-byte limit to be exceeded. It's common for poorly implemented or configured services between the client and the authoritative DNS server to incorrectly handle the request. For more information, including instructions on how to test your infrastructure, see the **DNS-OARC site**.

Firewall Configuration

With Google Workspace and other cloud applications, users reach outside your network for resources. This causes a shift of HTTP connections, from internal to external resources.

Because of this change, outbound firewalls that were previously properly sized in your network might become overwhelmed. Be aware of this possible increased footprint on your outbound firewall.

The average, peak, and idle connections from your benchmarking of proxy server load is a good estimate of the connection load to expect on your outbound firewall. The only connections you will not see on your outbound firewall are those that your proxy server does not allow through. For more information on gathering and using this data, see Proxy Server

Evaluation and Sizing on page 13.

Outbound Firewall Rules

To ensure the best possible experience for users of Google Workspace, and to provide a low-latency connection to our systems, we recommend leaving outbound firewall rules as open as possible on ports 80/443 for TCP/IP traffic.

Inbound Firewall Rules

Google Workspace does not initiate connections from Google data centers into your network. All traffic is initiated by clients inside of your network to Google.

Mail Routing

After you change your MX records to route mail traffic to Google Workspace, your email is no longer delivered to your servers. Instead, inbound email is directed to the Google Workspace servers. This essentially eliminates inbound SMTP mail traffic on your network.

Outbound Mail Connections

Depending on your needs, you may have some outbound mail traffic that you wish to send from your own network, such as automated or batched communications from applications in your system. You can use Google's **SMTP Relay Service** to route and filter your outbound mail securely. Ensure there is sufficient planning and estimation of how the expected SMTP relay volume will affect your overall network needs.

Chapter 5: Client Configuration

Summary

It is important to understand the effects that different clients can have on the performance of Google Workspace. This section discusses elements of the user environment that can impact Google Workspace performance, suggestions for setting up authentication for use with Google Workspace, and advice for migrating your users' data from an existing server into Google Workspace.

Client Access

When planning for the clients that your users will use to access Google Workspace, consider the following:

- In order to provide users with the best performance - and to experience all Google Workspace features - Google recommends Google Chrome.
- Google supports the latest version of Google Chrome (which automatically updates whenever it detects that a new version) and one major revision backward.
- Other supported browsers include the current and one previous major releases of Mozilla Firefox, Microsoft Internet Explorer, Microsoft Edge and Apple Safari.
  Note that not all Google Workspace features and functionalities are supported on these browsers..

Mobile

Android devices (which use the Google Sync protocol) and Windows Phone and Apple iOS devices (which use the ActiveSync protocol) communicate directly to Google servers without using your network resources.

See the chart below for an illustration.

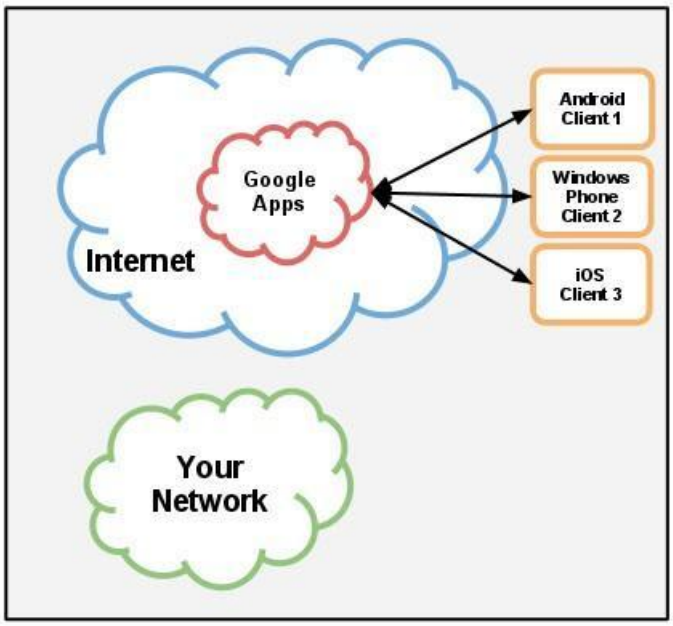

These devices do not access your network when using Google Workspace. With ActiveSync or Google Sync, Google Workspace delivers this mail directly to the user's device.

Authentication

Users can authenticate to the Google Workspace service in two ways:

- Single Sign-On service
- Google Authentication

Large enterprise organizations often use a Single Sign-On system to authorize users. There are also options for cloud based Single Sign-On systems for smaller organizations.

Single Sign-On

Google Workspace supports SAML 2.0 based authentication for all Google Workspace services. Client-side applications like Google Workspace Sync for Microsoft Outlook also support Single Sign-On.

If you plan to set up Single Sign-On authentication, consider the following suggestions:

- Set up SSO servers in distributed network locations, rather than a central location.
- Set up internal DNS servers to redirect SSO traffic to the nearest SSO server, and ensure that alternate SSO servers are in place for redundant service in case of disruption that prevents users from accessing the SSO server in a particular location.

Single Sign-On Process

When an unauthenticated user signs into Google Workspace, and an SSO URI is configured

for the domain, authentication takes several steps. See the chart below.

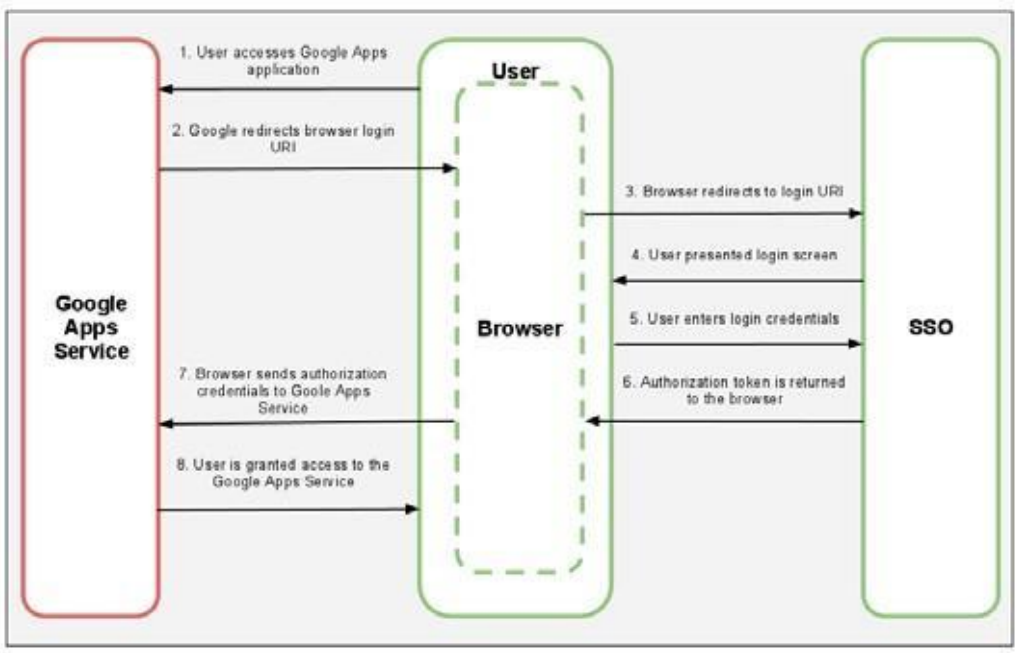

This is the process of SSO Authentication:

1. The user makes a request for a Google Workspace service.
2. The Google Workspace Authentication System redirects the user's browser to the configured URI for the SSO System. If the SSO/SAML server is not available, the user is unable to authenticate to the service.
3. The browser redirects to the login URI.
4. The SSO server displays a login screen.
5. The user enters login credentials and authenticates to the SSO System.
6. The SSO System passes an authorization token to the user's browser.
7. The user's browser sends the authorization credentials to the Google Workspace Service.
8. The user is granted access to the Google Workspace service.

Authentication Tools

A helpful tool to resolve any SAML-related errors during the authentication process is a SAML 2.0 debugger, such as **SAML 2.0 Debugger**.

Migration

Google Workspace deployments often involve traffic from migrating user data, either through local clients like **Google Workspace Migration for Microsoft Outlook,** or server-side clients like **Google Workspace Migration for IBM Notes** or **Google Workspace Migration for Microsoft Exchange**.

If you are migrating user data as part of your Google Workspace deployment, you can expect substantial data load depending on the amount of data you choose to migrate. To limit the impact to your network, we recommend following these best practices:

- If not cloud based, ensure that your migration servers are in the same location as your legacy data servers, or at least that the connectivity between servers has low latency and high bandwidth.
- Avoid routing traffic from the migration servers to Google through proxy servers.
- Assess your network capacity before migration to determine the maximum amount of data that you can migrate concurrently. Adjust your migration plan accordingly.
- During migration, some of the connections established to Google servers can stay open for a long period of time - depending on the migration tool. To avoid any possible migration errors, and to reduce the need to re-migrate data, it is important to keep these sessions open and not close them prematurely with any proxy or firewall timeouts.
- Consider scheduling the migration outside business hours, in order to reduce the network load during operating hours. This will obviously extend the total duration time for migration.

Chapter 6: Network Monitoring

Summary

After your network is setup to work with Google Workspace and your users are enabled, you can maintain the quality of your users' experience by monitoring the health of your network. To ensure the best user experience, follow these suggestions for monitoring tools and network traces.

Monitoring Tools

There are many commercial and open source tools to monitor various aspects of your network. A comprehensive directory of network monitoring tools is available on the **SLAC Network Monitoring Tools** site.

Specific recommended tools are listed in the table below.

Type of Monitoring	Tool	Description
Device Monitoring	**mrtg**	Monitors and graphs various aspects of network devices.
DNS changes	**Netblock Monitor**	Monitors Google's netblocks and alerts you to changes.
Host Monitoring	**smokeping**	Monitors and plots round-trip times to many destinations. Highly configurable.
Looking glass server	**Example list**	A looking glass server provides a read-only view of a network operator's routing information -- including connections, latency, and other factors -- at a remote point on another network.
Network	**pingplotter**	Helps monitor network latency, uptime, and route changes.
Network	**multiping**	Helps monitor network latency, uptime, and route changes.
Packet Capture	**Wireshark**	Performs packet captures.
RTT latency	**wbox**	Attempts to measure RTT of web application latency using HTTP/TCP latency.
Trace	**tcptrace**	Similar to traceroute but uses TCP packets rather than ICMP packets

Network Packet Captures

A network packet capture can help you to discover problems that may negatively affect the round-trip time or overall latency for Google Workspace users, such as:

- Different types of network flooding problems (ARP, TCP, UDP, IP, etc.)
- MTU mis-matches for Ethernet
- Malicious traffic on your network

Packet captures are helpful even though Google Workspace typically uses HTTPS connections. Packet captures will still show dropped packets, retransmits, window resizing, and evidence of saturated links.

One way to gather this type of data is to enable port mirroring, which allows you to capture traffic for a certain port or VLAN and divert it to another port where a service listens and logs all the traffic. Another approach is to use technologies such as **Wireshark** to capture data on a machine for later analysis.

An easy way to capture interactions between the browser and a website is to use a HAR (HTTP ARchive) file. It is basically a JSON object which contains detailed information about network requests.
HAR files can be generated by Google Chrome, Firefox, Internet Explorer, Microsoft Edge. The HAR file can also be analyzed using a utility available in Google Workspace Toolbox.

It is also possible to use chrome://net-export to capture network logs in Google Chrome and save them to disk.