# NGINX Cookbook

Advanced Recipes for High-Performance Load Balancing

Compliments of

NGINX
Part of F5

Derek DeJonghe

NGINX Cookbook

NGINX is one of the most widely used web servers available today, in part because of its capabilities as a load balancer and reverse proxy server for HTTP and other network protocols. This revised cookbook provides easy-to-follow examples of real-world problems in application delivery. Practical recipes help you set up and use either the open source or commercial offering to solve problems in various use cases.

For professionals who understand modern web architectures, such as n-tier or microservice designs and common web protocols such as TCP and HTTP, these recipes provide proven solutions for security and software load balancing and for monitoring and maintaining NGINX's application delivery platform. You'll also explore advanced features of both NGINX and NGINX Plus, the free and licensed versions of this server.

You'll find recipes for:

- High-performance load balancing with HTTP, TCP, and UDP
- Securing access through encrypted traffic, secure links, HTTP authentication subrequests, and more
- Deploying NGINX to Google, AWS, and Azure cloud
- Setting up and configuring NGINX Controller
- Installing and configuring the NGINX App Protect module
- Enabling WAF through Controller ADC
- NGINX Instance Manager, Service Mesh, and the njs module

"NGINX is one of the most powerful and complete tools out there today, and this book is the ultimate instrumentation guide for large architectures. The use cases presented can help anyone resolve almost all difficulties that appear while working in a microservice environment without losing focus on the business."

—**Gonzalo Spina**
Software Engineer, Brubank

Derek DeJonghe's in-depth background and experience in web development, system administration, and networking give him a well-rounded understanding of modern web architecture. He specializes in cloud migrations and operations of all sizes, leading a team of cloud architects and solution engineers and producing self-healing, autoscaling infrastructure for many different applications.

SYSTEM ADMINISTRATION

ISBN: 978-1-098-12626-1

Improve Kubernetes Resilience, Visibility, and Security with F5 NGINX

If you're ready to implement the methods discussed in this eBook, you need tools with the right features to support your use cases. NGINX can help with our suite of production-grade Kubernetes traffic-management tools:

NGINX Ingress Controller – Based on NGINX Plus, with advanced traffic control and shaping, monitoring and visibility, authentication and SSO, and API gateway functions.

NGINX App Protect – Holistic protection for modern apps and APIs, built on F5's market-leading security technologies and integrated with NGINX Ingress Controller and NGINX Plus. Two independent modules give you deployment flexibility and optimal resource utilization:

- **NGINX App Protect WAF** – Strong, lightweight protection against OWASP Top 10 and beyond, plus PCI DDS compliance
- **NGINX App Protect DoS** – Behavior-based, adaptive, and consistent DoS detection and mitigation across clouds and architectures

NGINX Service Mesh – Lightweight, turnkey, and developer-friendly, featuring NGINX Plus as an enterprise sidecar.
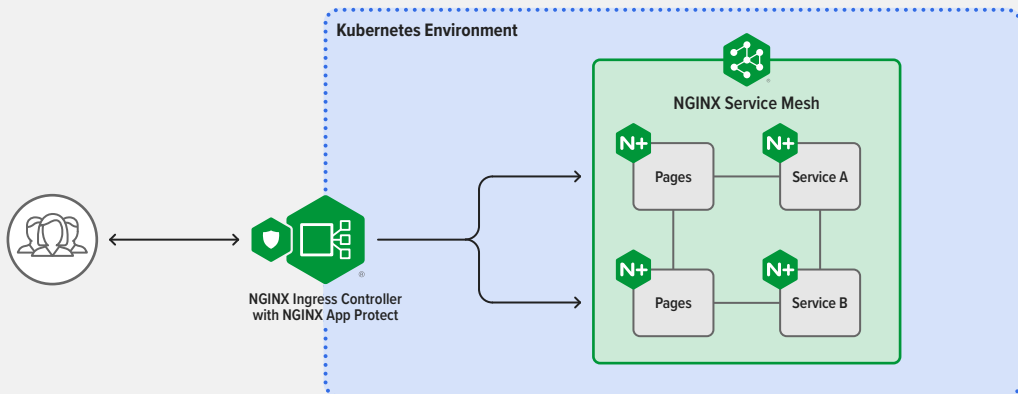
Get started by requesting your **free 30-day trial** of NGINX Ingress Controller with NGINX App Protect WAF and DoS, and **download** the always-free NGINX Service Mesh. Learn more today at **nginx.com**.

NGINX
Part of F5

NGINX Cookbook

Advanced Recipes for High-Performance Load Balancing

Derek DeJonghe

Beijing · Boston · Farnham · Sebastopol · Tokyo   **O'REILLY®**

Table of Contents

Foreword

Welcome to the 2022 edition of the NGINX Cookbook. O'Reilly has been publishing the *NGINX Cookbook* for seven years, and we continue to iterate and improve on the content. NGINX today is the world's most popular web server, according to W3Techs, and runs nearly 50% of the top 1000 busiest sites. Since we released the first version of NGINX in 2002, the product has evolved to become a Swiss Army Knife in the best sense. We architected NGINX for flexibility and scale, which made it possible to extend our capabilities beyond web serving to load balancing, reverse proxy, and API gateway as well. We take it as a testament to the value of NGINX that many of the load-balancing services offered by major public clouds and CDNs are actually based on NGINX code.

NGINX has also continued to expand into new realms; our web application firewall now applies machine learning to adapt to new threats as they arise. NGINX Ingress Controller for Kubernetes is natively integrated with NGINX and provides key capabilities for managing both east-west and north-south traffic, a critical requirement in the world of Kubernetes. None of this would matter if NGINX did not continue to deliver speed, resilience, and agility—all fundamental requirements of modern distributed applications.

The *NGINX Cookbook* is the "go-to" guide to NGINX from the people who know the code best. Whether you are running NGINX Open Source for a small project or NGINX Plus for a mission-critical enterprise deployment across multiple regions, and whether you are running locally or in the cloud, the Cookbook shows you how to get the most out of NGINX. It features dozens of easy-to-follow recipes covering how to install NGINX, configure it for almost any use case, secure and scale it, and troubleshoot common issues.

The 2022 edition updates many sections, covering new functionality in NGINX and adding entirely new sections and chapters focused on cloud native computing deployments and the new tools and deployment patterns for that paradigm shift. The world of technology and applications is changing fast, and NGINX is changing with it to continue contributing to your success.

The original vision of NGINX was highly scalable, reliable, fast, and secure web serving. Everything we do today is built on that original vision and designed to help our community deploy the apps they need, in any environment, at any scale, with confidence and trust. Please enjoy this latest edition and feel free to tell us what you think (in our Slack channel or online). We are listening.

—Christopher Guzikowski
Product Marketing, NGINX, Inc.

Preface

The *NGINX Cookbook* aims to provide easy-to-follow examples of real-world problems in application delivery. Throughout this book, you will explore the many features of NGINX and how to use them. This guide is fairly comprehensive, and touches on most of the main capabilities of NGINX.

Throughout this book, there will be references to both the free and open source NGINX software, as well as the commercial products from NGINX Inc.: NGINX Plus and NGINX Controller. Features and directives that are only available as part of the paid subscription to NGINX Plus will be denoted as such. Because NGINX Plus is an *application delivery controller* (ADC), and provides many advanced features, it's important to highlight these features to gain a full view of the possibilities of the platform.

The book will begin by explaining the installation process of NGINX and NGINX Plus, as well as some basic getting-started steps for readers new to NGINX. From there, the sections will progress to load balancing in all forms, accompanied by chapters about traffic management, caching, and automation. Chapter 6, "Authentication", covers a lot of ground but is important, because NGINX is often the first point of entry for web traffic to your application, and the first line of application-layer defense against web attacks and vulnerabilities. There are a number of chapters that cover cutting-edge topics such as HTTP/2, media streaming, cloud, service mesh, and container environments—wrapping up with more traditional operational topics such as monitoring, debugging, performance, and operational tips. The book will end with an introduction to NGINX Instance Manager and NGINX Controller, an application-centric management platform.

I personally use NGINX as a multitool, and believe this book will enable you to do the same. It's software that I believe in and enjoy working with. I'm happy to share this knowledge with you, and hope that as you read through this book you relate the recipes to your real-world scenarios and will employ these solutions.

Conventions Used in This Book

The following typographical conventions are used in this book:

Italic

    Indicates new terms, URLs, email addresses, filenames, and file extensions.

`Constant width`

    Used for program listings, as well as within paragraphs to refer to program elements such as variable or function names, databases, data types, environment variables, statements, and keywords.

This element signifies a general note.

This element indicates a warning or caution.

O'Reilly Online Learning

For more than 40 years, *O'Reilly Media* has provided technology and business training, knowledge, and insight to help companies succeed.

Our unique network of experts and innovators share their knowledge and expertise through books, articles, and our online learning platform. O'Reilly's online learning platform gives you on-demand access to live training courses, in-depth learning paths, interactive coding environments, and a vast collection of text and video from O'Reilly and 200+ other publishers. For more information, visit *http://oreilly.com*.

How to Contact Us

Please address comments and questions concerning this book to the publisher:

    O'Reilly Media, Inc.
    1005 Gravenstein Highway North
    Sebastopol, CA 95472
    800-998-9938 (in the United States or Canada)

707-829-0515 (international or local)
707-829-0104 (fax)

We have a web page for this book, where we list errata, examples, and any additional information. You can access this page at *https://oreil.ly/NGINX-cookbook-2e*.

Email *bookquestions@oreilly.com* to comment or ask technical questions about this book.

For news and information about our books and courses, visit *https://oreilly.com*.

Find us on LinkedIn: *https://linkedin.com/company/oreilly-media*

Follow us on Twitter: *https://twitter.com/oreillymedia*

Watch us on YouTube: *https://youtube.com/oreillymedia*

Acknowledgments

I would like to thank Hussein Nasser and Gonzalo Josue Spina for their helpful and detailed reviews of this title. I would also like to thank my wife for her support during the writing process.

Basics

1.0 Introduction

To get started with NGINX Open Source or NGINX Plus, you first need to install it on a system and learn some basics. In this chapter, you will learn how to install NGINX, where the main configuration files are, and what the commands are for administration. You will also learn how to verify your installation and make requests to the default server.

1.1 Installing NGINX on Debian/Ubuntu

Problem

You need to install NGINX Open Source on a Debian or Ubuntu machine.

Solution

Update package information for configured sources and install some packages that will assist in configuring the official NGINX package repository:

```
apt-get update
apt install -y curl gnupg2 ca-certificates lsb-release \
    debian-archive-keyring
```

Download and save the NGINX signing key:

```
curl https://nginx.org/keys/nginx_signing.key | gpg --dearmor \
    | tee /usr/share/keyrings/nginx-archive-keyring.gpg >/dev/null
```

Use `lsb_release` to set variables defining the OS and release names, then create an *apt* source file:

```
OS=$(lsb_release -is | tr '[:upper:]' '[:lower:]')
RELEASE=$(lsb_release -cs)
echo "deb [signed-by=/usr/share/keyrings/nginx-archive-keyring.gpg] \
    http://nginx.org/packages/${OS} ${RELEASE} nginx" \
    | tee /etc/apt/sources.list.d/nginx.list
```

Update package information once more, then install NGINX:

```
apt-get update
apt-get install -y nginx
nginx
```

Discussion

The commands provided in this section instruct the advanced package tool (APT) package management system to utilize the official NGINX package repository. The NGINX GPG package signing key was downloaded and saved to a location on the filesystem for use by APT. Providing APT the signing key enables the APT system to validate packages from the repository. The `lsb_release` command was used to automatically determine the OS and release name, so that these instructions can be used across all release versions of Debian or Ubuntu. The `apt-get update` command instructs the APT system to refresh its package listings from its known repositories. After the package list is refreshed, you can install NGINX Open Source from the official NGINX repository. After you install it, the final command starts NGINX.

1.2 Installing NGINX on RedHat/CentOS

Problem

You need to install NGINX Open Source on RedHat or CentOS.

Solution

Create a file named *etc/yum.repos.d/nginx.repo* that contains the following contents:

```
[nginx]
name=nginx repo
baseurl=http://nginx.org/packages/OS/OSRELEASE/$basearch/
gpgcheck=0
enabled=1
```

Alter the file, replacing OS in the middle of the URL with `rhel` or `centos`, depending on your distribution. Replace OSRELEASE with `7` or `8` for version 7.x or 8.x, respectively. Then, run the following commands:

```
yum -y install nginx
systemctl enable nginx
systemctl start nginx
```

```
firewall-cmd --permanent --zone=public --add-port=80/tcp
firewall-cmd --reload
```

Discussion

The file you just created for this solution instructs the YUM package management system to utilize the official NGINX Open Source package repository. The commands that follow install NGINX Open Source from the official repository, instruct systemd to enable NGINX at boot time, and tell it to start it now. The firewall commands open port 80 for the TCP protocol, which is the default port for HTTP. The last command reloads the firewall to commit the changes.

1.3 Installing NGINX Plus

Problem

You need to install NGINX Plus.

Solution

Visit *http://cs.nginx.com/repo_setup*. From the drop-down menu, select the OS you're installing to and then follow the instructions. The instructions are similar to those of the installation of the open source solutions; however, you need to install a certificate in order to authenticate to the NGINX Plus repository.

Discussion

NGINX keeps this repository installation guide up to date with instructions on installing the NGINX Plus. Depending on your OS and version, these instructions vary slightly, but there is one commonality. You must obtain a certificate and key from the NGINX portal, and provide them to your system, in order to authenticate to the NGINX Plus repository.

1.4 Verifying Your Installation

Problem

You want to validate the NGINX installation and check the version.

Solution

You can verify that NGINX is installed and check its version by using the following command:

```
$ nginx -v
nginx version: nginx/1.21.3
```

As this example shows, the response displays the version.

You can confirm that NGINX is running by using the following command:

```
$ ps -ef | grep nginx
root      1738     1  0 19:54 ?  00:00:00 nginx: master process
nginx     1739  1738  0 19:54 ?  00:00:00 nginx: worker process
```

The `ps` command lists running processes. By piping it to `grep`, you can search for specific words in the output. This example uses `grep` to search for `nginx`. The result shows two running processes: a `master` and `worker`. If NGINX is running, you will always see a master and one or more worker processes. Note the master process is running as `root`, as NGINX needs elevated privileges in order to function properly. For instructions on starting NGINX, refer to the next section. To see how to start NGINX as a daemon, use the init.d or systemd methodologies.

To verify that NGINX is returning requests correctly, use your browser to make a request to your machine or use `curl`. When making the request, use the machine's IP address or hostname. If installed locally, you can use `localhost` as follows:

```
curl localhost
```

You will see the NGINX Welcome default HTML site.

Discussion

The `nginx` command allows you to interact with the NGINX binary to check the version, list installed modules, test configurations, and send signals to the master process. NGINX must be running in order for it to serve requests. The `ps` command is a surefire way to determine whether NGINX is running either as a daemon or in the foreground. The configuration provided by default with NGINX runs a static-site HTTP server on port 80. You can test this default site by making an HTTP request to the machine at `localhost` as well as the host's IP and hostname.

1.5 Key Files, Directories, and Commands

Problem

You need to understand the important NGINX directories and commands.

Solution

NGINX files and directories

The following files, directories, and commands are important to know to get started with NGINX.

/etc/nginx/
> The */etc/nginx/* directory is the default configuration root for the NGINX server. Within this directory you will find configuration files that instruct NGINX on how to behave.

/etc/nginx/nginx.conf
> The */etc/nginx/nginx.conf* file is the default configuration entry point used by the NGINX service. This configuration file sets up global settings for things like worker processes, tuning, logging, loading dynamic modules, and references to other NGINX configuration files. In a default configuration, the */etc/nginx/nginx.conf* file includes the top-level `http` block, or context, which includes all configuration files in the directory described next.

/etc/nginx/conf.d/
> The */etc/nginx/conf.d/* directory contains the default HTTP server configuration file. Files in this directory ending in *.conf* are included in the top-level `http` block from within the */etc/nginx/nginx.conf* file. It's best practice to utilize `include` statements and organize your configuration in this way to keep your configuration files concise. In some package repositories, this folder is named *sites-enabled*, and configuration files are linked from a folder named *site-available*; this convention is deprecated.

/var/log/nginx/
> The */var/log/nginx/* directory is the default log location for NGINX. Within this directory you will find an *access.log* file and an *error.log* file. The access log contains an entry for each request NGINX serves. The error logfile contains error events and debug information if the `debug` module is enabled.

NGINX commands

`nginx -h`
> Shows the NGINX help menu.

`nginx -v`
> Shows the NGINX version.

`nginx -V`
> Shows the NGINX version, build information, and configuration arguments, which show the modules built into the NGINX binary.

```
nginx -t
```
Tests the NGINX configuration.

```
nginx -T
```
Tests the NGINX configuration and prints the validated configuration to the screen. This command is useful when seeking support.

```
nginx -s signal
```
The `-s` flag sends a signal to the NGINX master process. You can send signals such as `stop`, `quit`, `reload`, and `reopen`. The `stop` signal discontinues the NGINX process immediately. The `quit` signal stops the NGINX process after it finishes processing in-flight requests. The `reload` signal reloads the configuration. The `reopen` signal instructs NGINX to reopen logfiles.

Discussion

With an understanding of these key files, directories, and commands, you're in a good position to start working with NGINX. Using this knowledge, you can alter the default configuration files and test your changes with the `nginx -t` command. If your test is successful, you also know how to instruct NGINX to reload its configuration using the `nginx -s reload` command.

1.6 Serving Static Content

Problem

You need to serve static content with NGINX.

Solution

Overwrite the default HTTP server configuration located in *etc/nginx/conf.d/default.conf* with the following NGINX configuration example:

```
server {
    listen 80 default_server;
    server_name www.example.com;

    location / {
        root /usr/share/nginx/html;
        # alias /usr/share/nginx/html;
        index index.html index.htm;
    }
}
```

Discussion

This configuration serves static files over HTTP on port 80 from the directory */usr/share/nginx/html/*. The first line in this configuration defines a new `server` block. This defines a new context for NGINX to listen for. Line two instructs NGINX to listen on port 80, and the `default_server` parameter instructs NGINX to use this server as the default context for port 80. The `listen` directive can also take a range of ports. The `server_name` directive defines the hostname or the names of requests that should be directed to this server. If the configuration had not defined this context as the `default_server`, NGINX would direct requests to this server only if the HTTP host header matched the value provided to the `server_name` directive. With the `default_server` context set, you can omit the `server_name` directive if you do not yet have a domain name to use.

The `location` block defines a configuration based on the path in the URL. The path, or portion of the URL after the domain, is referred to as the uniform resource identifier (URI). NGINX will best match the URI requested to a `location` block. The example uses / to match all requests. The `root` directive shows NGINX where to look for static files when serving content for the given context. The URI of the request is appended to the `root` directive's value when looking for the requested file. If we had provided a URI prefix to the `location` directive, this would be included in the appended path, unless we used the `alias` directive rather than `root`. The `location` directive is able to match a wide range of expressions. Visit the link in the Also See section for more information. Lastly, the `index` directive provides NGINX with a default file, or list of files to check, in the event that no further path is provided in the URI.

Also See

NGINX HTTP `location` directive documentation

1.7 Graceful Reload

Problem

You need to reload your configuration without dropping packets.

Solution

Use the `reload` method of NGINX to achieve a graceful reload of the configuration without stopping the server:

```
nginx -s reload
```

This example reloads the NGINX system using the NGINX binary to send a signal to the master process.

Discussion

Reloading the NGINX configuration without stopping the server provides the ability to change configurations on the fly without dropping any packets. In a high-uptime, dynamic environment, you will need to change your load-balancing configuration at some point. NGINX allows you to do this while keeping the load balancer online. This feature enables countless possibilities, such as rerunning configuration management in a live environment, or building an application- and cluster-aware module to dynamically configure and reload NGINX to meet the needs of the environment.

High-Performance Load Balancing

2.0 Introduction

Today's internet user experience demands performance and uptime. To achieve this, multiple copies of the same system are run, and the load is distributed over them. As the load increases, another copy of the system can be brought online. This architecture technique is called *horizontal scaling*. Software-based infrastructure is increasing in popularity because of its flexibility, opening up a vast world of possibilities. Whether the use case is as small as a set of two system copies for high availability, or as large as thousands around the globe, there's a need for a load-balancing solution that is as dynamic as the infrastructure. NGINX fills this need in a number of ways, such as HTTP, TCP, and user datagram protocol (UDP) load balancing, which we cover in this chapter.

When balancing load, it's important that the impact to the client's experience is entirely positive. Many modern web architectures employ stateless application tiers, storing state in shared memory or databases. However, this is not the reality for all. Session state is immensely valuable and vast in interactive applications. This state might be stored locally to the application server for a number of reasons; for example, in applications for which the data being worked is so large that network overhead is too expensive in performance. When state is stored locally to an application server, it is extremely important to the user experience that the subsequent requests continue to be delivered to the same server. Another facet of the situation is that servers should not be released until the session has finished. Working with stateful applications at scale requires an intelligent load balancer. NGINX Plus offers multiple ways to solve this problem by tracking cookies or routing. This chapter covers session persistence as it pertains to load balancing with NGINX and NGINX Plus.

It's important to ensure that the application that NGINX is serving is healthy. For a number of reasons, upstream requests may begin to fail. It could be because of network connectivity, server failure, or application failure, to name a few. Proxies and load balancers must be smart enough to detect failure of upstream servers (servers behind the load balancer or proxy), and stop passing traffic to them; otherwise, the client will be waiting, only to be delivered a timeout. A way to mitigate service degradation when a server fails is to have the proxy check the health of the upstream servers. NGINX offers two different types of health checks: passive, available in NGINX Open Source; and active, available only in NGINX Plus. Active health checks at regular intervals will make a connection or request to the upstream server, and can verify that the response is correct. Passive health checks monitor the connection or responses of the upstream server as clients make the request or connection. You might want to use passive health checks to reduce the load of your upstream servers, and you might want to use active health checks to determine failure of an upstream server before a client is served a failure. The tail end of this chapter examines monitoring the health of the upstream application servers for which you're load balancing.

2.1 HTTP Load Balancing

Problem

You need to distribute load between two or more HTTP servers.

Solution

Use NGINX's HTTP module to load balance over HTTP servers using the upstream block:

```
upstream backend {
    server 10.10.12.45:80      weight=1;
    server app.example.com:80  weight=2;
    server spare.example.com:80  backup;
}
server {
    location / {
        proxy_pass http://backend;
    }
}
```

This configuration balances load across two HTTP servers on port 80, and defines one as a backup, which is used when the primary two are unavailable. The weight parameter instructs NGINX to pass twice as many requests to the second server, and the weight parameter defaults to 1.

Discussion

The HTTP `upstream` module controls the load balancing for HTTP. This module defines a pool of destinations—any combination of Unix sockets, IP addresses, and DNS (Domain Name Service) records, or a mix. The `upstream` module also defines how any individual request is assigned to any of the upstream servers.

Each upstream destination is defined in the upstream pool by the `server` directive. The `server` directive is provided a Unix socket, IP address, or a *fully qualified domain name* (FQDN), along with a number of optional parameters. The optional parameters give more control over the routing of requests. These parameters include the weight of the server in the balancing algorithm; whether the server is in standby mode, available, or unavailable; and how to determine if the server is unavailable. NGINX Plus provides a number of other convenient parameters, like connection limits to the server, advanced DNS resolution control, and the ability to slowly ramp up connections to a server after it starts.

2.2 TCP Load Balancing

Problem

You need to distribute load between two or more TCP servers.

Solution

Use NGINX's `stream` module to load balance over TCP servers using the `upstream` block:

```
stream {
    upstream mysql_read {
        server read1.example.com:3306  weight=5;
        server read2.example.com:3306;
        server 10.10.12.34:3306        backup;
    }

    server {
        listen 3306;
        proxy_pass mysql_read;
    }
}
```

The `server` block in this example instructs NGINX to listen on TCP port `3306` and balance load between two MySQL database read replicas, and lists another server as a backup that will be passed traffic if the primaries are down.

This configuration is not to be added to the *conf.d* folder, as that folder is included within an `http` block; instead, you should create another folder named *stream.conf.d*,

open the `stream` block in the *nginx.conf* file, and include the new folder for stream
configurations. An example follows.

In the */etc/nginx/nginx.conf* configuration file:

```
user nginx;
worker_processes auto;
pid /run/nginx.pid;

stream {
    include /etc/nginx/stream.conf.d/*.conf;
}
```

A file named */etc/nginx/stream.conf.d/mysql_reads.conf* may include the following
configuration:

```
upstream mysql_read {
    server read1.example.com:3306  weight=5;
    server read2.example.com:3306;
    server 10.10.12.34:3306        backup;
}

server {
    listen 3306;
    proxy_pass mysql_read;
}
```

Discussion

The main difference between the `http` and `stream` context is that they operate at
different layers of the OSI model. The `http` context operates at the application layer,
7, and `stream` operates at the transport layer, 4. This does not mean that the `stream`
context cannot become application-aware with some clever scripting, but that the
`http` context is specifically designed to fully understand the HTTP protocol, and the
`stream` context, by default, simply routes and load balances packets.

TCP load balancing is defined by the NGINX `stream` module. The `stream` module,
like the `HTTP` module, allows you to define upstream pools of servers and configure
a listening server. When configuring a server to listen on a given port, you must
define the port it will listen on, or optionally, an address and a port. From there,
a destination must be configured, whether it be a direct reverse proxy to another
address or an upstream pool of resources.

A number of options that alter the properties of the reverse proxy of the TCP
connection are available for configuration. Some of these include SSL/TLS validation
limitations, timeouts, and keepalives. Some of the values of these proxy options can
be (or contain) variables, such as the download rate and the name used to verify an
SSL/TLS certificate.

The upstream for TCP load balancing is much like the upstream for HTTP, in that it defines upstream resources as servers, configured with Unix socket, IP, or FQDN, as well as server weight, maximum number of connections, DNS resolvers, connection ramp-up periods, and if the server is active, down, or in backup mode.

NGINX Plus offers even more features for TCP load balancing. These advanced features can be found throughout this book. Health checks for all load balancing will be covered later in this chapter.

2.3 UDP Load Balancing

Problem

You need to distribute load between two or more UDP servers.

Solution

Use NGINX's `stream` module to load balance over UDP servers using the `upstream` block defined as `udp`:

```
stream {
    upstream ntp {
        server ntp1.example.com:123  weight=2;
        server ntp2.example.com:123;
    }

    server {
        listen 123 udp;
        proxy_pass ntp;
    }
}
```

This section of configuration balances load between two upstream network time protocol (NTP) servers using the UDP protocol. Specifying UDP load balancing is as simple as using the `udp` parameter on the `listen` directive.

If the service over which you're load balancing requires multiple packets to be sent back and forth between client and server, you can specify the `reuseport` parameter. Examples of these types of services are OpenVPN, Voice over Internet Protocol (VoIP), virtual desktop solutions, and Datagram Transport Layer Security (DTLS). The following is an example of using NGINX to handle OpenVPN connections and proxy them to the OpenVPN service running locally:

```
stream {
    server {
        listen 1195 udp reuseport;
        proxy_pass 127.0.0.1:1194;
```

```
        }
    }
```

Discussion

You might ask, "Why do I need a load balancer when I can have multiple hosts in a DNS A or service record (SRV record)?" The answer is that not only are there alternative balancing algorithms with which we can balance, but we can load balance over the DNS servers themselves. UDP services make up a lot of the services that we depend on in networked systems, such as DNS, NTP, QUIC, HTTP/3, and VoIP. UDP load balancing might be less common to some, but it's just as useful in the world of scale.

You can find UDP load balancing in the `stream` module, just like TCP, and configure it mostly in the same way. The main difference is that the `listen` directive specifies that the open socket is for working with datagrams. When working with datagrams, there are some other directives that might apply where they would not in TCP, such as the `proxy_response` directive, which specifies to NGINX how many expected responses can be sent from the upstream server. By default, this is unlimited until the `proxy_timeout` limit is reached. The `proxy_timeout` directive sets the time between two successive read-or-write operations on client or proxied server connections before the connection is closed.

The `reuseport` parameter instructs NGINX to create an individual listening socket for each worker process. This allows the kernel to distribute incoming connections between worker processes to handle multiple packets being sent between client and server. The `reuseport` feature works only on Linux kernels 3.9 and higher, DragonFly BSD, and FreeBSD 12 and higher.

2.4 Load-Balancing Methods

Problem

Round-robin load balancing doesn't fit your use case because you have heterogeneous workloads or server pools.

Solution

Use one of NGINX's load-balancing methods, such as least connections, least time, generic hash, random, or IP hash. This example sets the load-balancing algorithm for the backend upstream pool to be least connections:

```
upstream backend {
    least_conn;
    server backend.example.com;
```

```
    server backend1.example.com;
}
```

All load-balancing algorithms, with the exception of generic hash, random, and least time, are standalone directives, such as the preceding example. The parameters to these directives are explained in the following discussion.

Discussion

Not all requests or packets carry equal weight. Given this, round robin, or even the weighted round robin used in previous examples, will not fit the need of all applications or traffic flow. NGINX provides a number of load-balancing algorithms that you can use to fit particular use cases. In addition to being able to choose these load-balancing algorithms or methods, you can also configure them. The following load-balancing methods are available for upstream HTTP, TCP, and UDP pools:

Round robin
> This is the default load-balancing method, which distributes requests in the order of the list of servers in the upstream pool. You can also take weight into consideration for a weighted round robin, which you can use if the capacity of the upstream servers varies. The higher the integer value for the weight, the more favored the server will be in the round robin. The algorithm behind weight is simply the statistical probability of a weighted average.

Least connections
> This method balances load by proxying the current request to the upstream server with the least number of open connections. Least connections, like round robin, also takes weights into account when deciding which server to send the connection to. The directive name is `least_conn`.

Least time
> Available only in NGINX Plus, least time is akin to least connections in that it proxies to the upstream server with the least number of current connections, but favors the servers with the lowest average response times. This method is one of the most sophisticated load-balancing algorithms and fits the needs of highly performant web applications. This algorithm is a value-add over least connections because a small number of connections does not necessarily mean the quickest response. When using this algorithm, it is important to take into consideration the statistical variance of services' request times. Some requests may naturally take more processing and thus have a longer request time, increasing the range of the statistic. Long request times do not always mean a less performant or overworked server. However, requests that require more processing may be candidates for asynchronous workflows. A parameter of `header` or `last_byte` must be specified for this directive. When `header` is specified, the

time to receive the response header is used. When `last_byte` is specified, the time to receive the full response is used. The directive name is `least_time`.

Generic hash

The administrator defines a hash with the given text, variables of the request or runtime, or both. NGINX distributes the load among the servers by producing a hash for the current request and placing it against the upstream servers. This method is very useful when you need more control over where requests are sent or for determining which upstream server most likely will have the data cached. Note that when a server is added or removed from the pool, the hashed requests will be redistributed. This algorithm has an optional parameter, `consistent`, to minimize the effect of redistribution. The directive name is `hash`.

Random

This method is used to instruct NGINX to select a random server from the group, taking server weights into consideration. The optional `two [method]` parameter directs NGINX to randomly select two servers and then use the provided load-balancing method to balance between those two. By default the `least_conn` method is used if `two` is passed without a method. The directive name for random load balancing is `random`.

IP hash

This method works only for HTTP. IP hash uses the client IP address as the hash. Slightly different from using the remote variable in a generic hash, this algorithm uses the first three octets of an IPv4 address or the entire IPv6 address. This method ensures that clients are proxied to the same upstream server as long as that server is available, which is extremely helpful when the session state is of concern and not handled by shared memory of the application. This method also takes the `weight` parameter into consideration when distributing the hash. The directive name is `ip_hash`.

2.5 Sticky Cookie with NGINX Plus

Problem

You need to bind a downstream client to an upstream server using NGINX Plus.

Solution

Use the `sticky cookie` directive to instruct NGINX Plus to create and track a cookie:

```
upstream backend {
    server backend1.example.com;
    server backend2.example.com;
    sticky cookie
```

```
              affinity
              expires=1h
              domain=.example.com
              httponly
              secure
              path=/;
    }
```

This configuration creates and tracks a cookie that ties a downstream client to
an upstream server. In this example, the cookie is named `affinity`, is set for
example.com, expires in an hour, cannot be consumed client-side, can be sent only
over HTTPS, and is valid for all paths.

Discussion

Using the `cookie` parameter on the `sticky` directive creates a cookie on the first
request that contains information about the upstream server. NGINX Plus tracks
this cookie, enabling it to continue directing subsequent requests to the same server.
The first positional parameter to the `cookie` parameter is the name of the cookie
to be created and tracked. Other parameters offer additional control, informing the
browser of the appropriate usage—like the expiry time, domain, path, and whether
the cookie can be consumed client-side or whether it can be passed over unsecure
protocols.

2.6 Sticky Learn with NGINX Plus

Problem

You need to bind a downstream client to an upstream server by using an existing
cookie with NGINX Plus.

Solution

Use the `sticky learn` directive to discover and track cookies that are created by the
upstream application:

```
upstream backend {
    server backend1.example.com:8080;
    server backend2.example.com:8081;

    sticky learn
          create=$upstream_cookie_cookiename
          lookup=$cookie_cookiename
          zone=client_sessions:2m;
}
```

This example instructs NGINX to look for and track sessions by looking for a cookie named COOKIENAME in response headers, and looking up existing sessions by looking for the same cookie on request headers. This session affinity is stored in a shared memory zone of 2 MB that can track approximately 16,000 sessions. The name of the cookie will always be application-specific. Commonly used cookie names, such as jsessionid or phpsessionid, are typically defaults set within the application or the application server configuration.

Discussion

When applications create their own session-state cookies, NGINX Plus can discover them in request responses and track them. This type of cookie tracking is performed when the sticky directive is provided: the learn parameter. Shared memory for tracking cookies is specified with the zone parameter, with a name and size. NGINX Plus is directed to look for cookies in the response from the upstream server via specification of the create parameter, and searches for prior registered server affinity using the lookup parameter. The values of these parameters are variables exposed by the HTTP module.

2.7 Sticky Routing with NGINX Plus

Problem

You need granular control over how your persistent sessions are routed to the upstream server with NGINX Plus.

Solution

Use the sticky directive with the route parameter to use variables about the request to route:

```
map $cookie_jsessionid $route_cookie {
    ~.+\.(?P<route>\w+)$ $route;
}

map $request_uri $route_uri {
    ~jsessionid=.+\.(?P<route>\w+)$ $route;
}

upstream backend {
    server backend1.example.com route=a;
    server backend2.example.com route=b;

    sticky route $route_cookie $route_uri;
}
```

This example attempts to extract a Java session ID, first from a cookie by mapping the value of the Java session ID cookie to a variable with the first `map` block, and second by looking into the request URI for a parameter called `jsessionid`, mapping the value to a variable using the second `map` block. The `sticky` directive with the `route` parameter is passed any number of variables. The first nonzero or nonempty value is used for the route. If a `jsessionid` cookie is used, the request is routed to `backend1`; if a URI parameter is used, the request is routed to `backend2`. Although this example is based on the Java common session ID, the same applies for other session technology like `phpsessionid`, or any guaranteed unique identifier your application generates for the session ID.

Discussion

Sometimes, utilizing a bit more granular control, you might want to direct traffic to a particular server. The `route` parameter to the `sticky` directive is built to achieve this goal. Sticky route gives you better control, actual tracking, and stickiness, as opposed to the generic hash load-balancing algorithm. The client is first routed to an upstream server based on the route specified, and then subsequent requests will carry the routing information in a cookie or the URI. Sticky route takes a number of positional parameters that are evaluated. The first nonempty variable is used to route to a server. Map blocks can be used to selectively parse variables and save them as other variables to be used in the routing. Essentially, the `sticky route` directive creates a session within the NGINX Plus shared memory zone for tracking any client session identifier you specify to the upstream server, consistently delivering requests with this session identifier to the same upstream server as its original request.

2.8 Connection Draining with NGINX Plus

Problem

You need to gracefully remove servers for maintenance or other reasons, while still serving sessions with NGINX Plus.

Solution

Use the `drain` parameter through the NGINX Plus API, described in more detail in Chapter 5, to instruct NGINX to stop sending new connections that are not already tracked:

```
$ curl -X POST -d '{"drain":true}' \
  'http://nginx.local/api/3/http/upstreams/backend/servers/0'
{
  "id":0,
  "server":"172.17.0.3:80",
```

```
    "weight":1,
    "max_conns":0,
    "max_fails":1,
    "fail_timeout":
    "10s","slow_start":
    "0s",
    "route":"",
    "backup":false,
    "down":false,
    "drain":true
}
```

Discussion

When session state is stored locally to a server, connections and persistent sessions must be drained before the server is removed from the pool. Draining connections is the process of letting sessions to a server expire natively before removing the server from the upstream pool. You can configure draining for a particular server by adding the `drain` parameter to the `server` directive. When the `drain` parameter is set, NGINX Plus stops sending new sessions to this server but allows current sessions to continue being served for the length of their session. You can also toggle this configuration by adding the `drain` parameter to an upstream server directive, then reloading the NGINX Plus configuration.

2.9 Passive Health Checks

Problem

You need to passively check the health of upstream servers.

Solution

Use NGINX health checks with load balancing to ensure that only healthy upstream servers are utilized:

```
upstream backend {
    server backend1.example.com:1234 max_fails=3 fail_timeout=3s;
    server backend2.example.com:1234 max_fails=3 fail_timeout=3s;
}
```

This configuration passively monitors upstream health by monitoring the response of client requests directed to the upstream server. The example sets the `max_fails` directive to three, and `fail_timeout` to 3 seconds. These directive parameters work the same way in both stream and HTTP servers.

Discussion

Passive health checking is available in the open source version of NGINX, and is configured by using the same `server` parameters for HTTP, TCP, and UDP load balancing. Passive monitoring watches for failed or timed-out connections as they pass through NGINX as requested by a client. Passive health checks are enabled by default; the parameters mentioned here allow you to tweak their behavior. The default `max_fails` value is 1, and the default `fail_timeout` value is `10s`. Monitoring for health is important on all types of load balancing, not only from a user-experience standpoint, but also for business continuity. NGINX passively monitors upstream HTTP, TCP, and UDP servers to ensure that they're healthy and performing.

Also See

HTTP Health Checks Admin Guide
TCP Health Checks Admin Guide
UDP Health Checks Admin Guide

2.10 Active Health Checks with NGINX Plus

Problem

You need to actively check your upstream servers for health with NGINX Plus.

Solution

For HTTP, use the `health_check` directive in a location block:

```
http {
    server {
        # ...
        location / {
            proxy_pass http://backend;
            health_check interval=2s
                fails=2
                passes=5
                uri=/
                match=welcome;
        }
    }
    # status is 200, content type is "text/html",
    # and body contains "Welcome to nginx!"
    match welcome {
        status 200;
        header Content-Type = text/html;
        body ~ "Welcome to nginx!";
    }
}
```

This health-check configuration for HTTP servers checks the health of the upstream servers by making an HTTP GET request to the URI "/" every 2 seconds. The HTTP method can't be defined for health checks; only GET requests are performed, as other methods may change the state of backend systems. The upstream servers must pass five consecutive health checks to be considered healthy. They are considered unhealthy if they fail two consecutive checks. The response from the upstream server must match the defined match block, which defines the status code as 200, the header Content-Type value as 'text/html', and the string "Welcome to nginx!" in the response body. The HTTP match block has three directives: status, header, and body. All three of these directives have comparison flags as well.

Stream health checks for TCP/UDP services are very similar:

```
stream {
    # ...
    server {
        listen 1234;
        proxy_pass stream_backend;
        health_check interval=10s
            passes=2
            fails=3;
        health_check_timeout 5s;
    }
    # ...
}
```

In this example, a TCP server is configured to listen on port 1234, and to proxy to an upstream set of servers, for which it actively checks for health. The stream health_check directive takes all the same parameters as in HTTP, with the exception of uri, and the stream version has a parameter to switch the check protocol to udp. In this example, the interval is set to 10 seconds, requires two passes to be considered healthy, and requires three fails to be considered unhealthy. The active-stream health check is also able to verify the response from the upstream server. The match block for stream servers, however, has just two directives: send and expect. The send directive is raw data to be sent, and expect is an exact response or a regular expression to match.

Discussion

In NGINX Plus, passive or active health checks can be used to monitor the source servers. These health checks can measure more than just the response code. In NGINX Plus, active HTTP health checks monitor based on a number of acceptance criteria of the response from the upstream server. You can configure active health-check monitoring for how often upstream servers are checked, how many times a server must pass this check to be considered healthy, how many times it can fail before being deemed unhealthy, and what the expected result should be. For more

complex logic, a `require` directive for the match block enables the use of variables whose value must not be empty or zero. The `match` parameter points to a match block that defines the acceptance criteria for the response. The match block also defines the data to send to the upstream server when used in the stream context for TCP/UDP. These features enable NGINX to ensure that upstream servers are healthy at all times.

Also See

HTTP Health Checks Admin Guide
TCP Health Checks Admin Guide
UDP Health Checks Admin Guide

2.11 Slow Start with NGINX Plus

Problem

Your application needs to ramp up before taking on full production load.

Solution

Use the `slow_start` parameter on the `server` directive to gradually increase the number of connections over a specified time as a server is reintroduced to the upstream load-balancing pool:

```
upstream {
    zone backend 64k;

    server server1.example.com slow_start=20s;
    server server2.example.com slow_start=15s;
}
```

The `server` directive configurations will slowly ramp up traffic to the upstream servers after they're reintroduced to the pool. `server1` will slowly ramp up its number of connections over 20 seconds, and `server2` over 15 seconds.

Discussion

Slow start is the concept of slowly ramping up the number of requests proxied to a server over a period of time. Slow start allows the application to warm up by populating caches, initiating database connections without being overwhelmed by connections as soon as it starts. This feature takes effect when a server that has failed health checks begins to pass again and re-enters the load-balancing pool, and is only available in NGINX Plus. Slow start can't be used with `hash`, `ip_hash`, or `random` load-balancing methods.

Traffic Management

3.0 Introduction

NGINX and NGINX Plus are also classified as web-traffic controllers. You can use NGINX to intelligently route traffic and control flow based on many attributes. This chapter covers NGINX's ability to split client requests based on percentages; utilize the geographical location of the clients; and control the flow of traffic in the form of rate, connection, and bandwidth limiting. As you read through this chapter, keep in mind that you can mix and match these features to enable countless possibilities.

3.1 A/B Testing

Problem

You need to split clients between two or more versions of a file or application to test acceptance or engagement.

Solution

Use the `split_clients` module to direct a percentage of your clients to a different upstream pool:

```
split_clients "${remote_addr}AAA" $variant {
    20.0%    "backendv2";
    *        "backendv1";
}
```

The `split_clients` directive hashes the string provided by you as the first parameter and divides that hash by the percentages provided to map the value of a variable provided as the second parameter. The addition of "AAA" to the first parameter is

to demonstrate that this is a concatenated string that can include many variables, as mentioned in the generic hash load-balancing algorithm. The third parameter is an object containing key-value pairs where the key is the percentage weight and the value is the value to be assigned. The key can be either a percentage or an asterisk. The asterisk denotes the rest of the whole after all percentages are taken. The value of the $variant variable will be backendv2 for 20% of client IP addresses and backendv1 for the remaining 80%.

In this example, backendv1 and backendv2 represent upstream server pools and can be used with the proxy_pass directive as such:

```
location / {
    proxy_pass http://$variant
}
```

Using the variable $variant, our traffic will split between two different application server pools.

To demonstrate the wide variety of uses split_clients can have, the following is an example of splitting between two versions of a static site:

```
http {
    split_clients "${remote_addr}" $site_root_folder {
        33.3%    "/var/www/sitev2/";
        *        "/var/www/sitev1/";
    }
    server {
        listen 80 _;
        root $site_root_folder;
        location / {
            index index.html;
        }
    }
}
```

Discussion

This type of A/B testing is useful when testing different types of marketing and front-end features for conversion rates on ecommerce sites. It's common for applications to use a type of deployment called *canary release*. In this type of deployment, traffic is slowly switched over to the new version by gradually increasing the percentage of users being routed to the new version. Splitting your clients between different versions of your application can be useful when rolling out new versions of code, to limit the blast radius in the event of an error. Even more common is the blue-green deployment style, where users are cut over to a new version and the old version is still available while the deployment is validated. Whatever the reason for splitting clients between two different application sets, NGINX makes this simple because of this split_clients module.

Also See

3.2 Using the GeoIP Module and Database

Problem

You need to install the GeoIP database and enable its embedded variables within NGINX to utilize the physical location of your clients in the NGINX log, proxied requests, or request routing.

Solution

The official NGINX Open Source package repository, configured in Chapter 2 when installing NGINX, provides a package named `nginx-module-geoip`. When using the NGINX Plus package repository, this package is named `nginx-plus-module-geoip`. The following examples show how to install the dynamic NGINX GeoIP module package, as well as how to download the GeoIP country and city databases:

RHEL/CentOS NGINX Open Source:

```
# yum install nginx-module-geoip
```

Debian/Ubuntu NGINX Open Source:

```
# apt-get install nginx-module-geoip
```

RHEL/CentOS NGINX Plus:

```
# yum install nginx-plus-module-geoip
```

Debian/Ubuntu NGINX Plus:

```
# apt-get install nginx-plus-module-geoip
```

Download the GeoIP country and city databases and unzip them:

```
# mkdir /etc/nginx/geoip
# cd /etc/nginx/geoip
# wget "http://geolite.maxmind.com/\
download/geoip/database/GeoLiteCountry/GeoIP.dat.gz"
# gunzip GeoIP.dat.gz
# wget "http://geolite.maxmind.com/\
download/geoip/database/GeoLiteCity.dat.gz"
# gunzip GeoLiteCity.dat.gz
```

This set of commands creates a *geoip* directory in the */etc/nginx* directory, moves to this new directory, and downloads and unzips the packages.

With the GeoIP database for countries and cities on the local disk, you can now instruct the NGINX GeoIP module to use them to expose embedded variables based on the client IP address:

```
load_module "/usr/lib64/nginx/modules/ngx_http_geoip_module.so";

http {
    geoip_country /etc/nginx/geoip/GeoIP.dat;
    geoip_city /etc/nginx/geoip/GeoLiteCity.dat;
# ...
}
```

The `load_module` directive dynamically loads the module from its path on the filesystem. The `load_module` directive is only valid in the main context. The `geoip_country` directive takes a path to the *GeoIP.dat* file containing the database mapping IP addresses to country codes and is valid only in the `http` context.

Discussion

To use this functionality, you must have the NGINX GeoIP module installed, and a local GeoIP country and city database. Installation and retrieval of these prerequisites was demonstrated in this section.

The `geoip_country` and `geoip_city` directives expose a number of embedded variables available in this module. The `geoip_country` directive enables variables that allow you to distinguish the country of origin of your client. These variables include `$geoip_country_code`, `$geoip_country_code3`, and `$geoip_country_name`. The country code variable returns the two-letter country code, and the variable with a `3` at the end returns the three-letter country code. The country name variable returns the full name of the country.

The `geoip_city` directive enables quite a few variables. The `geoip_city` directive enables all the same variables as the `geoip_country` directive, just with different names, such as `$geoip_city_country_code`, `$geoip_city_country_code3`, and `$geoip_city_country_name`. Other variables include `$geoip_city`, `$geoip_latitude`, `$geoip_longitude`, `$geoip_city_continent_code`, and `$geoip_postal_code`, all of which are descriptive of the value they return. `$geoip_region` and `$geoip_region_name` describe the region, territory, state, province, federal land, and the like. Region is the two-letter code, where region name is the full name. `$geoip_area_code`, only valid in the US, returns the three-digit telephone area code.

With these variables, you're able to log information about your client. You could optionally pass this information to your application as a header or variable, or use NGINX to route your traffic in particular ways.

Also See

`geoip` module documentation
GeoIP Update GitHub

3.3 Restricting Access Based on Country

Problem

You need to restrict access from particular countries for contractual or application requirements.

Solution

Map the country codes you want to block or allow to a variable:

```
load_module
  "/usr/lib64/nginx/modules/ngx_http_geoip_module.so";

http {
    map $geoip_country_code $country_access {
        "US"    0;
        default 1;
    }
    # ...
}
```

This mapping will set a new variable `$country_access` to a `1` or a `0`. If the client IP address originates from the US, the variable will be set to a `0`. For any other country, the variable will be set to a `1`.

Now, within our `server` block, we'll use an `if` statement to deny access to anyone not originating from the US:

```
server {
    if ($country_access = '1') {
      return 403;
    }
    # ...
}
```

This `if` statement will evaluate `True` if the `$country_access` variable is set to 1. When `True`, the server will return a 403 unauthorized. Otherwise the server operates as normal. So this `if` block is only there to deny people who are not from the US.

Discussion

This is a short but simple example of how to only allow access from a couple of countries. This example can be expounded on to fit your needs. You can utilize

this same practice to allow or block based on any of the embedded variables made available from the GeoIP module.

3.4 Finding the Original Client

Problem

You need to find the original client IP address because there are proxies in front of the NGINX server.

Solution

Use the `geoip_proxy` directive to define your proxy IP address range and the `geoip_proxy_recursive` directive to look for the original IP:

```
load_module "/usr/lib64/nginx/modules/ngx_http_geoip_module.so";

http {
    geoip_country /etc/nginx/geoip/GeoIP.dat;
    geoip_city /etc/nginx/geoip/GeoLiteCity.dat;
    geoip_proxy 10.0.16.0/26;
    geoip_proxy_recursive on;
# ...
}
```

The `geoip_proxy` directive defines a *classless inter-domain routing* (CIDR) range in which our proxy servers live and instructs NGINX to utilize the `X-Forwarded-For` header to find the client IP address. The `geoip_proxy_recursive` directive instructs NGINX to recursively look through the `X-Forwarded-For` header for the last client IP known.

A header named `Forwarded` has become the standard header for adding proxy information for proxied requests. The header used by the NGINX GeoIP module is `X-Forwarded-For` and cannot be configured otherwise at the time of writing. While `X-Forwarded-For` is not an official standard, it is still very widely used, accepted, and set by most proxies.

Discussion

You may find that if you're using a proxy in front of NGINX, NGINX will pick up the proxy's IP address rather than the client's. For this you can use the `geoip_proxy` directive to instruct NGINX to use the `X-Forwarded-For` header when connections are opened from a given range. The `geoip_proxy` directive takes an address or a CIDR range. When there are multiple proxies passing traffic in front of NGINX,

you can use the `geoip_proxy_recursive` directive to recursively search through `X-Forwarded-For` addresses to find the originating client. You will want to use something like this when utilizing load balancers such as Amazon Web Services Elastic Load Balancing (AWS ELB), Google's load balancer, or Microsoft Azure's load balancer in front of NGINX.

3.5 Limiting Connections

Problem

You need to limit the number of connections based on a predefined key, such as the client's IP address.

Solution

Construct a shared memory zone to hold connection metrics, and use the `limit_conn` directive to limit open connections:

```
http {
    limit_conn_zone $binary_remote_addr zone=limitbyaddr:10m;
    limit_conn_status 429;
    # ...
    server {
        # ...
            limit_conn limitbyaddr 40;
        # ...
    }
}
```

This configuration creates a shared memory zone named `limitbyaddr`. The predefined key used is the client's IP address in binary form. The size of the shared memory zone is set to 10 MB. The `limit_conn` directive takes two parameters: a `limit_conn_zone` name and the number of connections allowed. The `limit_conn_status` sets the response when the connections are limited to a status of 429, indicating too many requests. The `limit_conn` and `limit_conn_status` directives are valid in the `http`, `server`, and `location` context.

Discussion

Limiting the number of connections based on a key can be used to defend against abuse and share your resources fairly across all your clients. It is important to be cautious with your predefined key. Using an IP address, as we are in the previous example, could be dangerous if many users are on the same network that originates from the same IP, such as when behind a Network Address Translation (NAT). The entire group of clients will be limited. The `limit_conn_zone` directive is only valid in

the `http` context. You can utilize any number of variables available to NGINX within the `http` context in order to build a string by which to limit. Utilizing a variable that can identify the user at the application level, such as a session cookie, may be a cleaner solution depending on the use case. The `limit_conn_status` defaults to 503, service unavailable. You may find it preferable to use a 429, as the service is available, and 500-level responses indicate server error whereas 400-level responses indicate client error.

Testing limitations can be tricky. It's often hard to simulate live traffic in an alternate environment for testing. In this case, you can set the `limit_req_dry_run` directive to on, then use the variable `$limit_req_status` in your access log. The `$limit_req_status` variable will evaluate to either PASSED, DELAYED, REJECTED, DELAYED_DRY_RUN, or REJECTED_DRY_RUN. With dry run enabled, you'll be able to analyze the logs of live traffic and tweak your limits as needed before enabling, providing you with assurance that your limit configuration is correct.

3.6 Limiting Rate

Problem

You need to limit the rate of requests by a predefined key, such as the client's IP address.

Solution

Utilize the rate-limiting module to limit the rate of requests:

```
http {
    limit_req_zone $binary_remote_addr
        zone=limitbyaddr:10m rate=3r/s;
    limit_req_status 429;
    # ...
    server {
        # ...
            limit_req zone=limitbyaddr;
        # ...
    }
}
```

This example configuration creates a shared memory zone named `limitbyaddr`. The predefined key used is the client's IP address in binary form. The size of the shared memory zone is set to 10 MB. The zone sets the rate with a keyword argument. The `limit_req` directive takes a required keyword argument: `zone`. `zone` instructs the directive on which shared memory request–limit zone to use. Requests that exceed the expressed rate are returned a 429 HTTP code, as defined by the `limit_req_`

`status` directive. I advise to set a status in the 400-level range, as the default is a `503`, implying a problem with the server, when the issue is actually with the client.

Use optional keyword arguments to the `limit_req` directive to enable two-stage rate limiting:

```
server {
    location / {
        limit_req zone=limitbyaddr burst=12 delay=9;
    }
}
```

In some cases, a client will need to make many requests all at once, and then will reduce its rate for a period of time before making more. You can use the keyword argument `burst` to allow the client to exceed its rate limit but not have requests rejected. The rate-exceeded requests will have a delay in processing to match the rate limit up to the value configured. A set of keyword arguments alter this behavior: `delay` and `nodelay`. The `nodelay` argument does not take a value, and simply allows the client to consume the burstable value all at once; however, all requests will be rejected until enough time has passed to satisfy the rate limit. In this example, if we used `nodelay`, the client could consume 12 requests in the first second, but would have to wait 4 seconds after the initial request to make another. The `delay` keyword argument defines how many requests can be made up front without throttling. In this case, the client can make nine requests up front with no delay, the next three will be throttled, and any more within a 4-second period will be rejected.

Discussion

The rate-limiting module is very powerful for protecting against abusive rapid requests, while still providing a quality service to everyone. There are many reasons to limit rate of request, one being security. You can deny a brute-force attack by putting a very strict limit on your login page. You can set a sane limit on all requests, thereby disabling the plans of malicious users who might try to deny service to your application or to waste resources. The configuration of the rate-limit module is much like the connection-limiting module described in Recipe 3.5, and many of the same concerns apply. You can specify the rate at which requests are limited in requests per second or requests per minute. When the rate limit is reached, the incident is logged. There's also a directive not in the example, `limit_req_log_level`, which defaults to `error`, but can be set to `info`, `notice`, or `warn`. In NGINX Plus, rate limiting is now cluster-aware (see Recipe 12.5 for a zone sync example).

Testing limitations can be tricky. It's often hard to simulate live traffic in an alternative environment for testing. In this case, you can set the `limit_req_dry_run` directive to `on`, then use the variable `$limit_req_status` in your access log. The `$limit_req_status` variable will evaluate to `PASSED`, `REJECTED`, or `REJECTED_DRY_RUN`. With dry run

enabled, you'll be able to analyze the logs of live traffic and tweak your limits as needed before enabling, providing you with assurance that your limit configuration is correct.

3.7 Limiting Bandwidth

Problem

You need to limit download bandwidth per client for your assets.

Solution

Utilize NGINX's `limit_rate` and `limit_rate_after` directives to limit the bandwidth of response to a client:

```
location /download/ {
    limit_rate_after 10m;
    limit_rate 1m;
}
```

The configuration of this location block specifies that for URIs with the prefix *download*, the rate at which the response will be served to the client will be limited after 10 MB to a rate of 1 MB per second. The bandwidth limit is per connection, so you may want to institute a connection limit as well as a bandwidth limit where applicable.

Discussion

Limiting the bandwidth for particular connections enables NGINX to share its upload bandwidth across all of the clients in a manner you specify. These two directives do it all: `limit_rate_after` and `limit_rate`. The `limit_rate_after` directive can be set in almost any context: `http`, `server`, `location`, and `if` when the `if` is within a location. The `limit_rate` directive is applicable in the same contexts as `limit_rate_after`; however, it can alternatively be set by a variable named `$limit_rate`.

The `limit_rate_after` directive specifies that the connection should not be rate limited until after a specified amount of data has been transferred. The `limit_rate` directive specifies the rate limit for a given context in bytes per second by default. However, you can specify `m` for megabytes or `g` for gigabytes. Both directives default to a value of 0. The value 0 means not to limit download rates at all. This module allows you to programmatically change the rate limit of clients.

Massively Scalable Content Caching

4.0 Introduction

Caching accelerates content serving by storing request responses to be served again in the future. Content caching reduces load to upstream servers, caching the full response rather than running computations and queries again for the same request. Caching increases performance and reduces load, meaning you can serve faster with fewer resources. The scaling and distribution of caching servers in strategic locations can have a dramatic effect on user experience. It's optimal to host content close to the consumer for the best performance. You can also cache your content close to your users. This is the pattern of content delivery networks, or CDNs. With NGINX you're able to cache your content wherever you can place an NGINX server, effectively enabling you to create your own CDN. With NGINX caching, you're also able to passively cache and serve cached responses in the event of an upstream failure. Caching features are only available within the `http` context.

4.1 Caching Zones

Problem

You need to cache content and need to define where the cache is stored.

Solution

Use the `proxy_cache_path` directive to define shared memory-cache zones and a location for the content:

```
proxy_cache_path /var/nginx/cache
                 keys_zone=CACHE:60m
                 levels=1:2
```

```
                    inactive=3h
                    max_size=20g;
    proxy_cache CACHE;
```

The cache definition example creates a directory for cached responses on the filesystem at */var/nginx/cache* and creates a shared memory space named CACHE with 60 MB of memory. This example sets the directory structure levels, defines the release of cached responses after they have not been requested in 3 hours, and defines a maximum size of the cache of 20 GB. The proxy_cache directive informs a particular context to use the cache zone. The proxy_cache_path is valid in the http context and the proxy_cache directive is valid in the http, server, and location contexts.

Discussion

To configure caching in NGINX, it's necessary to declare a path and zone to be used. A cache zone in NGINX is created with the directive proxy_cache_path. The proxy_cache_path designates a location to store the cached information and a shared memory space to store active keys and response metadata. Optional parameters to this directive provide more control over how the cache is maintained and accessed. The levels parameter defines how the file structure is created. The value is a colon-separated value that declares the length of subdirectory names, with a maximum of three levels. NGINX caches are based on the cache key, which is a hashed value. NGINX then stores the result in the file structure provided, using the cache key as a file path and breaking up directories based on the levels value. The inactive parameter allows for control over the length of time a cache item will be hosted after its last use. The size of the cache is also configurable with the use of the max_size parameter. Other parameters relate to the cache-loading process, which loads the cache keys into the shared memory zone from the files cached on disk.

4.2 Cache Locking

Problem

You don't want NGINX to proxy requests that are currently being written to cache to an upstream server.

Solution

Use the proxy_cache_lock directive to ensure only one request is able to write to the cache at a time, where subsequent requests will wait for the response to be written:

```
    proxy_cache_lock on;
    proxy_cache_lock_age 10s;
    proxy_cache_lock_timeout 3s;
```

Discussion

The `proxy_cache_lock` directive instructs NGINX to hold requests, destined for a cached element, that are currently being populated. The proxied request that is populating the cache is limited in the amount of time it has before another request attempts to populate the element, defined by the `proxy_cache_lock_age` directive, which defaults to 5 seconds. NGINX can also allow requests that have been waiting a specified amount of time to pass through to the proxied server, which will not attempt to populate the cache by use of the `proxy_cache_lock_timeout` directive, which also defaults to 5 seconds. You can think of `proxy_cache_lock_age` and `proxy_cache_lock_timeout` as "You're taking too long, I'll populate the cache for you," and "You're taking too long for me to wait, I'm going to get what I need and let you populate the cache in your own time," respectively.

4.3 Caching Hash Keys

Problem

You need to control how your content is cached and retrieved.

Solution

Use the `proxy_cache_key` directive along with variables to define what constitutes a cache hit or miss:

```
proxy_cache_key "$host$request_uri $cookie_user";
```

This cache hash key will instruct NGINX to cache pages based on the host and URI being requested, as well as a cookie that defines the user. With this you can cache dynamic pages without serving content that was generated for a different user.

Discussion

The default `proxy_cache_key`, which will fit most use cases, is `"$scheme$proxy_host$request_uri"`. The variables used include the scheme, HTTP or HTTPS, the `proxy_host`, where the request is being sent, and the request URI. Altogether, this reflects the URL that NGINX is proxying the request to. You may find that there are many other factors that define a unique request per application—such as request arguments, headers, session identifiers, and so on—to which you'll want to create your own hash key.[1]

---

1 Any combination of text or variables exposed to NGINX can be used to form a cache key. A list of variables is available in NGINX.

Selecting a good hash key is very important and should be thought through with understanding of the application. Selecting a cache key for static content is typically pretty straightforward; using the hostname and URI will suffice. Selecting a cache key for fairly dynamic content like pages for a dashboard application requires more knowledge around how users interact with the application and the degree of variance between user experiences. Due to security concerns, you may not want to present cached data from one user to another without fully understanding the context. The `proxy_cache_key` directive configures the string to be hashed for the cache key. The `proxy_cache_key` can be set in the context of `http`, `server`, and `location` blocks, providing flexible control on how requests are cached.

4.4 Cache Bypass

Problem

You need the ability to bypass the caching.

Solution

Use the `proxy_cache_bypass` directive with a nonempty or nonzero value. One way to do this dynamically is with a variable set to anything other than an empty string or 0 within a location block that you do not want cached:

```
proxy_cache_bypass $http_cache_bypass;
```

The configuration tells NGINX to bypass the cache if the HTTP request header named `cache_bypass` is set to any value that is not `0`. This example uses a header as the variable to determine if caching should be bypassed—the client would need to specifically set this header for their request.

Discussion

There are a number of scenarios that demand that the request is not cached. For this, NGINX exposes a `proxy_cache_bypass` directive so that when the value is nonempty or nonzero, the request will be sent to an upstream server rather than be pulled from the cache. Different needs and scenarios for bypassing cache will be dictated by your application's use case. Techniques for bypassing cache can be as simple as using a request or response header, or as intricate as multiple map blocks working together.

You may want to bypass the cache for many reasons, such as troubleshooting or debugging. Reproducing issues can be hard if you're consistently pulling cached pages or if your cache key is specific to a user identifier. Having the ability to bypass the cache is vital. Options include, but are not limited to, bypassing the cache when a particular cookie, header, or request argument is set. You can also turn off the cache

completely for a given context, such as a location block, by setting `proxy_cache off;`.

4.5 Cache Performance

Problem

You need to increase performance by caching on the client side.

Solution

Use client-side cache-control headers:

```
location ~* \.(css|js)$ {
  expires 1y;
  add_header Cache-Control "public";
}
```

This location block specifies that the client can cache the content of CSS and Java-Script files. The `expires` directive instructs the client that their cached resource will no longer be valid after one year. The `add_header` directive adds the HTTP response header `Cache-Control` to the response, with a value of `public`, which allows any caching server along the way to cache the resource. If we specify `private`, only the client is allowed to cache the value.

Discussion

Cache performance has many factors, disk speed being high on the list. There are many things within the NGINX configuration that you can do to assist with cache performance. One option is to set headers of the response in such a way that the client actually caches the response and does not make the request to NGINX at all, but simply serves it from its own cache.

4.6 Cache Purging with NGINX Plus

Problem

You need to invalidate an object from the cache.

Solution

Use the purge feature of NGINX Plus, the `proxy_cache_purge` directive, and a non-empty or zero-value variable:

```
map $request_method $purge_method {
    PURGE 1;
```

```
        default 0;
    }
    server {
        # ...
        location / {
            # ...
            proxy_cache_purge $purge_method;
        }
    }
```

In this example, the cache for a particular object will be purged if it's requested with a method of PURGE. The following is a curl example of purging the cache of a file named *main.js*:

```
$ curl -XPURGE localhost/main.js
```

Discussion

A common way to handle static files is to put a hash of the file in the filename. This ensures that as you roll out new code and content, your CDN recognizes it as a new file because the URI has changed. However, this does not exactly work for dynamic content to which you've set cache keys that don't fit this model. In every caching scenario, you must have a way to purge the cache. NGINX Plus has provided a simple method of purging cached responses. The proxy_cache_purge directive, when passed a nonzero or nonempty value, will purge the cached items matching the request. A simple way to set up purging is by mapping the request method for PURGE. However, you may want to use this in conjunction with the geo_ip module or simple authentication to ensure that not just anyone can purge your precious cache items. NGINX has also allowed for the use of *, which will purge cache items that match a common URI prefix. To use wildcards, you will need to configure your proxy_cache_path directive with the purger=on argument.

4.7 Cache Slicing

Problem

You need to increase caching efficiency by segmenting the file into fragments.

Solution

Use the NGINX slice directive and its embedded variables to divide the cache result into fragments:

```
proxy_cache_path /tmp/mycache keys_zone=mycache:10m;
server {
    # ...
    proxy_cache mycache;
```

```
    slice 1m;
    proxy_cache_key $host$uri$is_args$args$slice_range;
    proxy_set_header Range $slice_range;
    proxy_http_version 1.1;
    proxy_cache_valid  200 206 1h;

    location / {
        proxy_pass http://origin:80;
    }
}
```

Discussion

This configuration defines a cache zone and enables it for the server. The `slice` directive is then used to instruct NGINX to slice the response into 1 MB file segments. The cache files are stored according to the `proxy_cache_key` directive. Note the use of the embedded variable named `slice_range`. That same variable is used as a header when making the request to the origin, and that request HTTP version is upgraded to HTTP/1.1 because 1.0 does not support byte-range requests. The cache validity is set for response codes of `200` or `206` for 1 hour, and then the location and origins are defined.

The Cache Slice module was developed for delivery of HTML5 video, which uses byte-range requests to pseudostream content to the browser. By default, NGINX is able to serve byte-range requests from its cache. If a request for a byte range is made for uncached content, NGINX requests the entire file from the origin. When you use the Cache Slice module, NGINX requests only the necessary segments from the origin. Range requests that are larger than the slice size, including the entire file, trigger subrequests for each of the required segments, and then those segments are cached. When all of the segments are cached, the response is assembled and sent to the client, enabling NGINX to more efficiently cache and serve content requested in ranges.

Cache Slice module should be used only on large files that do not change. NGINX validates the ETag each time it receives a segment from the origin. If the ETag on the origin changes, NGINX aborts the cache population of the segment because the cache is no longer valid. If the content does change and the file is smaller, or your origin can handle load spikes during the cache fill process, it's better to use the Cache Lock module described in the blog listed in the following Also See section. This module is not built by default, and needs to be enabled by the `--with-http_slice_module` configuration when building NGINX.

Also See

"Smart and Efficient Byte-Range Caching with NGINX & NGINX Plus"

Programmability and Automation

5.0 Introduction

Programmability refers to the ability to interact with something through program-ming. The API for NGINX Plus provides just that: the ability to interact with the configuration and behavior of NGINX Plus through an HTTP interface. This API provides the ability to reconfigure NGINX Plus by adding or removing upstream servers through HTTP requests. The key-value store feature in NGINX Plus enables another level of dynamic configuration—you can utilize HTTP calls to inject infor-mation that NGINX Plus can use to route or control traffic dynamically. This chapter will touch on the NGINX Plus API and the key-value store module exposed by that same API.

Configuration management tools automate the installation and configuration of servers, which is an invaluable utility in the age of the cloud. Engineers of large-scale web applications no longer need to configure servers by hand; instead, they can use one of the many configuration management tools available. With these tools, engineers only need to write configurations and code once to produce many servers with the same configuration in a repeatable, testable, and modular fashion. This chapter covers a few of the most popular configuration management tools available and how to use them to install NGINX and template a base configuration. These examples are extremely basic but demonstrate how to get an NGINX server started with each platform.

5.1 NGINX Plus API

Problem

You have a dynamic environment and need to reconfigure NGINX Plus on the fly.

Solution

Configure the NGINX Plus API to enable adding and removing servers through API calls:

```
upstream backend {
    zone http_backend 64k;
}
server {
    # ...
    location /api {
        api [write=on];
        # Directives limiting access to the API
        # See chapter 7
    }

    location = /dashboard.html {
        root   /usr/share/nginx/html;
    }
}
```

This NGINX Plus configuration creates an upstream server with a shared memory zone, enables the API in the /api location block, and provides a location for the NGINX Plus dashboard.

You can utilize the API to add servers when they come online:

```
$ curl -X POST -d '{"server":"172.17.0.3"}' \
  'http://nginx.local/api/3/http/upstreams/backend/servers/'

{
  "id":0,
  "server":"172.17.0.3:80",
  "weight":1,
  "max_conns":0,
  "max_fails":1,
  "fail_timeout":"10s",
  "slow_start":"0s",
  "route":"",
  "backup":false,
  "down":false
}
```

The curl call in this example makes a request to NGINX Plus to add a new server to the backend upstream configuration. The HTTP method is a POST, a JSON object is passed as the body, and a JSON response is returned. The JSON response shows the server object configuration—note that a new id was generated, and other configuration settings were set to default values.

The NGINX Plus API is RESTful; therefore, there are parameters in the request URI.

The format of the URI is as follows:

```
/api/{version}/http/upstreams/{httpUpstreamName}/servers/
```

You can utilize the NGINX Plus API to list the servers in the upstream pool:

```
$ curl 'http://nginx.local/api/3/http/upstreams/backend/servers/'
[
  {
    "id":0,
    "server":"172.17.0.3:80",
    "weight":1,
    "max_conns":0,
    "max_fails":1,
    "fail_timeout":"10s",
    "slow_start":"0s",
    "route":"",
    "backup":false,
    "down":false
  }
]
```

The `curl` call in this example makes a request to NGINX Plus to list all of the servers in the upstream pool named `backend`. Currently, we have only the one server that we added in the previous `curl` call to the API. The request will return an upstream server object that contains all of the configurable options for a server.

Use the NGINX Plus API to drain connections from an upstream server, preparing it for a graceful removal from the upstream pool, as seen in the following code. You can find details about connection draining in Recipe 2.8.

```
$ curl -X PATCH -d '{"drain":true}' \
  'http://nginx.local/api/3/http/upstreams/backend/servers/0'
{
  "id":0,
  "server":"172.17.0.3:80",
  "weight":1,
  "max_conns":0,
  "max_fails":1,
  "fail_timeout":
  "10s","slow_start":
  "0s",
  "route":"",
  "backup":false,
  "down":false,
  "drain":true
}
```

In this `curl`, we specify that the request method is `PATCH`, pass a JSON body instructing it to drain connections for the server, and specify the server ID by appending it to the URI. We found the ID of the server by listing the servers in the upstream pool in the previous `curl` command.

NGINX Plus will begin to drain the connections. This process can take as long as the length of the sessions of the application. To check how many active connections are being served by the server you've begun to drain, use the following call and look for the `active` attribute of the server being drained:

```
$ curl 'http://nginx.local/api/3/http/upstreams/backend'
{
    "zone" : "http_backend",
    "keepalive" : 0,
    "peers" : [
        {
            "backup" : false,
            "id" : 0,
            "unavail" : 0,
            "name" : "172.17.0.3",
            "requests" : 0,
            "received" : 0,
            "state" : "draining",
            "server" : "172.17.0.3:80",
            "active" : 0,
            "weight" : 1,
            "fails" : 0,
            "sent" : 0,
            "responses" : {
                "4xx" : 0,
                "total" : 0,
                "3xx" : 0,
                "5xx" : 0,
                "2xx" : 0,
                "1xx" : 0
            },
            "health_checks" : {
                "checks" : 0,
                "unhealthy" : 0,
                "fails" : 0
            },
            "downtime" : 0
        }
    ],
    "zombies" : 0
}
```

After all connections have drained, utilize the NGINX Plus API to remove the server from the upstream pool entirely:

```
$ curl -X DELETE \
    'http://nginx.local/api/3/http/upstreams/backend/servers/0'
[]
```

The `curl` command makes a `DELETE` method request to the same URI used to update the server's state. The `DELETE` method instructs NGINX to remove the server. This

API call returns all of the servers and their IDs that are still left in the pool. Because we started with an empty pool, added only one server through the API, drained it, and then removed it, we now have an empty pool again.

Discussion

The NGINX Plus exclusive API enables dynamic application servers to add themselves to and remove themselves from the NGINX configuration on the fly. As servers come online, they can register themselves to the pool, and NGINX will start sending load to the newly added servers. When a server needs to be removed, the server can request NGINX Plus to drain its connections, and then remove itself from the upstream pool before it's shut down. This enables the infrastructure, through some automation, to scale in and out without human intervention.

Also See

NGINX Plus REST API documentation

5.2 Using the Key-Value Store with NGINX Plus

Problem

You need NGINX Plus to make dynamic traffic management decisions based on input from applications.

Solution

This section will use the example of a dynamic blocklist as a traffic management decision.

Set up the cluster-aware key-value store and API, and then add keys and values:

```
keyval_zone zone=blocklist:1M;
keyval $remote_addr $blocked zone=blocklist;

server {
    # ...
    location / {
        if ($blocked) {
            return 403 'Forbidden';
        }
        return 200 'OK';
    }
}
server {
    # ...
    # Directives limiting access to the API
    # See chapter 6
```

```
    location /api {
        api write=on;
    }
}
```

This NGINX Plus configuration uses the `keyval_zone` directive to build a key-value store shared memory zone named `blocklist` and sets a memory limit of 1 MB. The `keyval` directive then maps the value of the key, matching the first parameter `$remote_addr` to a new variable named `$blocked` from the zone. This new variable is then used to determine whether NGINX Plus should serve the request or return a 403 Forbidden code.

After starting the NGINX Plus server with this configuration, you can `curl` the local machine and expect to receive a 200 OK response:

```
$ curl 'http://127.0.0.1/'
OK
```

Now add the local machine's IP address to the key-value store with a value of 1:

```
$ curl -X POST -d '{"127.0.0.1":"1"}' \
  'http://127.0.0.1/api/3/http/keyvals/blocklist'
```

This `curl` command submits an HTTP POST request with a JSON object containing a key-value object to be submitted to the `blocklist` shared memory zone. The key-value store API URI is formatted as follows:

```
/api/{version}/http/keyvals/{httpKeyvalZoneName}
```

The local machine's IP address is now added to the key-value zone named `blocklist` with a value of 1. In the next request, NGINX Plus looks up the `$remote_addr` in the key-value zone, finds the entry, and maps the value to the variable `$blocked`. This variable is then evaluated in the `if` statement. When the variable has a value, the `if` evaluates to `True` and NGINX Plus returns the 403 Forbidden return code:

```
$ curl 'http://127.0.0.1/'
Forbidden
```

You can update or delete the key by making a `PATCH` method request:

```
$ curl -X PATCH -d '{"127.0.0.1":null}' \
  'http://127.0.0.1/api/3/http/keyvals/blocklist'
```

NGINX Plus deletes the key if the value is `null`, and requests will again return 200 OK.

Discussion

The key-value store, an NGINX Plus exclusive feature, enables applications to inject information into NGINX Plus. In the example provided, the `$remote_addr` variable is used to create a dynamic blocklist. You can populate the key-value store with any

key that NGINX Plus might have as a variable—a session cookie, for example—and provide NGINX Plus an external value. In NGINX Plus R16, the key-value store became cluster-aware, meaning that you have to provide your key-value update to only one NGINX Plus server, and all of them will receive the information.

In NGINX Plus R19, the key-value store enabled a `type` parameter, which enables indexing for specific types of keys. By default, the type is of value `string`, where `ip` and `prefix` are also options. The `string` type does not build an index and all key requests must be exact matches, whereas `prefix` will allow for partial key matches provided the prefix of the key is a match. An `ip` type enables the use of CIDR notation. In our example, if we had specified the `type=ip` as a parameter to our zone, we could have provided an entire CIDR range to block, such as `192.168.0.0/16` to block the entire RFC 1918 private range block, or `127.0.0.1/32` for localhost, which would have rendered the same effect as demonstrated in the example.

Also See

"Dynamic Bandwidth Limits Using the NGINX Plus Key-Value Store"

5.3 Using the NJS Module to Expose JavaScript Functionality Within NGINX

Problem

You need NGINX to perform custom logic on requests or responses.

Solution

Enable the use of JavaScript by installing the NJS module for NGINX. The following package installation steps assume that you've added the official NGINX repositories for your Linux distribution, as demonstrated in Chapter 1.

Debian/Ubuntu:

```
apt-get install nginx-module-njs
```

Debian/Ubuntu and NGINX Plus:

```
apt-get install nginx-plus-module-njs
```

RHEL/CentOS:

```
yum install nginx-module-njs
```

RHEL/CentOS and NGINX Plus:

```
yum install nginx-plus-module-njs
```

If you do not already have a directory for JavaScript files within the NGINX configuration, create one:

```
mkdir -p /etc/nginx/njs
```

Create a JavaScript file named */etc/nginx/njs/jwt.js* with the following contents:

```
function jwt(data) {
    var parts = data.split('.').slice(0,2)
        .map(v=>Buffer.from(v, 'base64url').toString())
        .map(JSON.parse);
    return { headers:parts[0], payload: parts[1] };
}
function jwt_payload_subject(r) {
    return jwt(r.headersIn.Authorization.slice(7)).payload.sub;
}
function jwt_payload_issuer(r) {
    return jwt(r.headersIn.Authorization.slice(7)).payload.iss;
}
export default {jwt_payload_subject, jwt_payload_issuer}
```

The provided JavaScript example defines a function that decodes JSON Web Tokens (JWTs). Further, two functions are defined that use the JWT decoder to return specific keys within the JWT. Those functions are exported to be made available to NGINX. These functions return common keys found in a JWT: the subject and issuer. The `.slice(7)` portion of the code accounts for removing the first seven characters of the `Authorization` header value. With the advent of JWT the type value is `Bearer`. Bearer is 6 characters, and we also need to remove the space delimiter, which is why we slice the first 7 characters. There are some authentication services that do not provide a type, such as AWS Cognito, because such services alter the slice count or remove it entirely so that the `jwt` function receives only the token value.

Within the core NGINX configuration, load the NJS module. Import and use the JavaScript in the `http` block:

```
load_module /etc/nginx/modules/ngx_http_js_module.so;

http {
    js_path "/etc/nginx/njs/";
    js_import main from jwt.js;
    js_set $jwt_payload_subject main.jwt_payload_subject;
    js_set $jwt_payload_issuer main.jwt_payload_issuer;
    ...
}
```

The provided NGINX configuration dynamically loads the NJS module and imports the JavaScript file that we previously defined. NGINX directives are used to set NGINX variables to the returned values of the JavaScript functions.

Use the variables to prove the JavaScript logic. Define a server that returns the variables set by JavaScript:

```
server {
    listen 80 default_server;
    listen [::]:80 default_server;
    server_name _;
    location / {
        return 200 "$jwt_payload_subject $jwt_payload_issuer";
    }
}
```

The provided configuration will produce a server that returns the subject and issuer values provided by the client through the `Authorization` header. These values are decoded by the defined JavaScript code.

To validate that the code works, you will make a request to the server with a given JWT. The following is a JSON format of the JWT, which you can use to verify that the code works:

```
{
  "iss": "nginx",
  "sub": "alice",
  "foo": 123,
  "bar": "qq",
  "zyx": false
}
```

Make a request to the server with a given JWT to verify that the JavaScript code runs and returns the correct values:

```
curl 'http://localhost/' -H \
"Authorization: Bearer eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzI1\
NiIsImV4cCI6MTU4NDcyMzA4NX0.eyJpc3MiOiJuZ2lueCIsInN1YiI6Im\
FsaWNlIiwiZm9vIjoxMjMsImJhciI6InFxIiwie\
nl4IjpmYWxzZX0.Kftl23Rvv9dIso1RuZ8uHaJ83BkKmMtTwch09rJtwgk"

alice nginx
```

Discussion

NGINX has provided a module that exposes standard JavaScript functionality during its processing of requests and responses. This module enables you to embed business logic into your proxy layer. JavaScript was chosen because of its volume of usage.

The NJS module provides the ability to inject logic during requests coming into NGINX, as well as responses coming from NGINX. You're able to validate and manipulate requests as they pass through the proxy, as demonstrated in this section by decoding the JSON Web Token. The NJS module is also able to manipulate responses from upstream services, by streaming the response data through the JavaScript logic. NJS additionally enables stream services to become application layer aware, as demonstrated in the Also See section.

Also See

5.4 Extending NGINX with a Common Programming Language

Problem

You need NGINX to perform some custom extension using a common programming language.

Solution

Before preparing to write a custom NGINX module from scratch in C, first evaluate if one of the other programming language modules will fit your use case. The C programming language is extremely powerful and performant. There are, however, many other languages available as modules that may enable the customization required. NGINX has introduced NGINX JavaScript (NJS), which exposes the power of Java-Script into the NGINX configuration by simply enabling a module. Lua and Perl modules are also available.

With these language modules, you either import a file including code or define a block of code directly within the configuration.

To use Lua, install the Lua module and the following NGINX configuration to define a Lua script inline:

```
load_module modules/ndk_http_module.so;
load_module modules/ngx_http_lua_module.so;

events {}

http {
    server {
        listen 8080;
        location / {
            default_type text/html;
            content_by_lua_block {
                ngx.say("hello, world")
            }
        }
    }
}
```

The Lua module provides its own NGINX API through an object defined by the module named `ngx`. Like the `request` object in NJS, the `ngx` object has attributes and methods to describe the request and manipulate the response.

With the Perl module installed, this example will use Perl to set an NGINX variable from the runtime environment:

```
load_module modules/ngx_http_perl_module.so;

events {}

http {
    perl_set $app_endpoint 'sub { return $ENV{"APP_DNS_ENDPOINT"}; }';
    server {
        listen 8080;
        location / {
                proxy_pass http://$app_endpoint
            }
        }
    }
}
```

The preceding example demonstrates that these language modules expose more functionality than just returning a response. The `perl_set` directive sets an NGINX variable to data returned from a Perl script. This limited example simply returns a system environment variable, which is used as the endpoint in which to proxy requests.

Discussion

The capabilities enabled by the extendability of NGINX are endless. NGINX is extendable with custom code through C modules, which can be compiled into NGINX when building from source, or dynamically loaded within the configuration. Existing modules that expose the functionality and syntax of JavaScript (njs), Lua, and Perl are already available. In many cases, unless distributing custom NGINX functionality to others, these pre-existing modules can suffice. Many scripts built for these modules already exist in the open source community.

This solution demonstrated basic usage of the Lua and Perl scripting languages available in NGINX and NGINX Plus. Whether looking to respond, set a variable, make a subrequest, or define a complex rewrite, these NGINX modules provide the capability.

Also See

NGINX Plus Lua module installation
NGINX Plus Perl module installation

NGINX Lua module documentation
NGINX Perl module documentation

5.5 Installing with Chef

Problem

You need to install and configure NGINX with Chef to manage NGINX configurations as code and conform with the rest of your Chef configurations.

Solution

Install the NGINX cookbook, maintained by the Sous Chefs, from the Chef Supermarket:

```
knife supermarket install nginx
```

This cookbook is resource-based, which means it provides Chef resources for you to use in your own cookbook. Create a cookbook for your NGINX use case. This cookbook will include the `nginx` cookbook installed from the Supermarket as a dependency. Once the dependency is in place you can use the resources provided. Create a recipe for installing NGINX:

```
nginx_install 'nginx' do
  source 'repo'
end
```

When `source` is set to `repo`, NGINX is installed from NGINX Inc.–maintained repositories, which provide the most up-to-date versions.

Use the `nginx_config` resource within a recipe to override core NGINX configurations:

```
nginx_config 'nginx' do
  default_site_enabled true
  keepalive_timeout 65
  worker_processes 'auto'
  action :create
  notifies :reload, 'nginx_service[nginx]', :delayed
end
```

Use the `nginx_site` resource within a recipe to configure an NGINX server block:

```
nginx_site 'test_site' do
  mode '0644'

  variables(
    'server' => {
      'listen' => [ '*:80' ],
      'server_name' => [ 'test.example.com' ],
```

```
        'access_log' => '/var/log/nginx/test_site.access.log',
        'locations' => {
          '/' => {
            'root' => '/var/www/nginx-default',
            'index' => 'index.html index.htm',
          },
        },
      }
    )

    action :create
    notifies :reload, 'nginx_service[nginx]', :delayed
  end
```

Discussion

Chef is a configuration management tool written in Ruby. It can be run in a client/server relationship or a solo configuration. Chef has a very large community called the Supermarket with many public cookbooks. Public cookbooks from the Supermarket can be installed and maintained via a command-line utility called Knife. Chef is extremely capable, and what we have demonstrated is just a small sample.

The public NGINX cookbook in the Supermarket is extremely flexible and provides the options to easily install NGINX from a package manager or from source, and the ability to compile and install many different modules as well as template out the basic configurations. In this section we installed NGINX from the NGINX Inc.– maintained repositories, and configured a server block to host an HTML file as a basic example. You can provide your own templates to the `nginx_site` resource to take further control of your NGINX configuration.

Also See

[Chef documentation](#)
[Chef Supermarket for NGINX](#)

5.6 Installing with Ansible

Problem

You need to install and configure NGINX with Ansible to manage NGINX configurations as code and conform with the rest of your Ansible configurations.

Solution

Install the Ansible NGINX collection from Ansible Galaxy:

```
ansible-galaxy collection install nginxinc.nginx_core
```

Create a playbook that uses the collection and `nginx` role to install NGINX:

```
---
- hosts: all
  collections:
    - nginxinc.nginx_core
  tasks:
    - name: Install NGINX
      include_role:
        name: nginx
```

To configure NGINX, add a task and use the `nginx_config` role and provide variable overrides to the default template to fit your use case:

```
    - name: Configure NGINX
      include_role:
        name: nginx_config
      vars:
        nginx_config_http_template_enable: true
        nginx_config_http_template:
          - template_file: http/default.conf.j2
            deployment_location: /etc/nginx/conf.d/default.conf
            config:
              servers:
                - core:
                    listen:
                      - port: 80
                    server_name: localhost
                  log:
                    access:
                      - path: /var/log/nginx/access.log
                        format: main
                  sub_filter:
                    sub_filters:
                      - string: server_hostname
                        replacement: $hostname
                    once: false
                  locations:
                    - location: /
                      core:
                        root: /usr/share/nginx/html
                        index: index.html

        nginx_config_html_demo_template_enable: true
        nginx_config_html_demo_template:
          - template_file: www/index.html.j2
            html_file_name: index.html
            html_file_location: /usr/share/nginx/html
            web_server_name: Ansible NGINX collection
```

Discussion

Ansible is a widely used and powerful configuration management tool written in Python. The configuration of tasks is in YAML, and you use the Jinja2 templating language for file templating. Ansible offers a server named Ansible Tower on a subscription model. However, it's commonly used from local machines or to build servers directly to the client or in a standalone model. Ansible will bulk SSH into servers and run the configuration. Much like other configuration management tools, there's a large community of public roles. Ansible calls this the Ansible Galaxy. You can find very sophisticated roles to utilize in your playbooks.

This solution used a collection of public roles, maintained by NGINX Inc., to install NGINX and produce a sample configuration. The configuration used in this example templates an NGINX demo HTML file and places it at */usr/share/nginx/html/index.html*. An NGINX configuration file is also templated to produce a server block listening on `localhost:80`, with a single location block configured to serve the demo file. The provided NGINX configuration templates are extremely comprehensive; however, the `nginx_config` role allows you to provide your own templates for complete control while taking advantage of the prebuilt and maintained Ansible configuration.

Also See

NGINX-provided Ansible collection
Ansible documentation

5.7 Automating Configurations with Consul Templating

Problem

You need to automate your NGINX configuration to respond to changes in your environment through use of Consul.

Solution

Use the `consul-template` daemon and a template file to template out the NGINX configuration file of your choice:

```
upstream backend { {{range service "app.backend"}}
    server {{.Address}};{{end}}
}
```

This example is a Consul template file that templates an upstream configuration block. The template will loop through nodes in Consul identified as `app.backend`. For

every node in Consul, the template will produce a server directive with that node's IP address.

The `consul-template` daemon is run via the command line and can be used to reload NGINX every time the configuration file is templated with a change:

```
# consul-template -consul-addr consul.example.internal -template \
  ./upstream.template:/etc/nginx/conf.d/upstream.conf:"nginx -s reload"
```

This command instructs the `consul-template` daemon to connect to a Consul cluster at `consul.example.internal` and to use a file named *upstream.template* in the current working directory to template the file and output the generated contents to */etc/nginx/conf.d/upstream.conf*, then to reload NGINX every time the templated file changes. The `-template` flag takes a string of the template file, the output location, and the command to run after the templating process takes place. These three variables are separated by colons. If the command being run has spaces, make sure to wrap it in double quotes. The `-consul` flag tells the daemon what Consul cluster to connect to.

Discussion

Consul is a powerful service discovery tool and configuration store. Consul stores information about nodes as well as key-value pairs in a directory-like structure and allows for RESTful API interaction. Consul also provides a DNS interface on each client, allowing for domain name lookups of nodes connected to the cluster. A separate project that utilizes Consul clusters is the `consul-template` daemon; this tool templates files in response to changes in Consul nodes, services, or key-value pairs. This makes Consul a very powerful choice for automating NGINX. With `consul-template` you can also instruct the daemon to run a command after a change to the template takes place. With this, you can reload the NGINX configuration and allow your NGINX configuration to come alive along with your environment. With Consul and `consul-template`, your NGINX configuration can be as dynamic as your environment. Infrastructure, configuration, and application information is centrally stored, and `consul-template` can subscribe and retemplate as necessary in an event-based manner. With this technology, NGINX can dynamically reconfigure in reaction to the addition and removal of servers, services, application versions, etc.

Also See

"Load Balancing with NGINX Plus' Service Discovery Integration"
"Load Balancing with NGINX and Consul Template"
Consul home page
"Service Configuration with Consul Template"
Consul Template GitHub

Authentication

6.0 Introduction

NGINX is able to authenticate clients. Authenticating client requests with NGINX offloads work and provides the ability to stop unauthenticated requests from reaching your application servers. Modules available for NGINX Open Source include basic authentication and authentication subrequests. The NGINX Plus exclusive module for verifying JSON Web Tokens (JWTs) enables integration with third-party authentication providers that use the authentication standard OpenID Connect.

6.1 HTTP Basic Authentication

Problem

You need to secure your application or content via HTTP basic authentication.

Solution

Generate a file in the following format, where the password is encrypted or hashed with one of the allowed formats:

```
# comment
name1:password1
name2:password2:comment
name3:password3
```

The username is the first field, the password the second field, and the delimiter is a colon. There is an optional third field, which you can use to comment on each user. NGINX can understand a few different formats for passwords, one of which is a password encrypted with the C function `crypt()`. This function is exposed to the

command line by the `openssl passwd` command. With `openssl` installed, you can create encrypted password strings by using the following command:

```
$ openssl passwd MyPassword1234
```

The output will be a string that NGINX can use in your password file.

Use the `auth_basic` and `auth_basic_user_file` directives within your NGINX configuration to enable basic authentication:

```
location / {
    auth_basic           "Private site";
    auth_basic_user_file conf.d/passwd;
}
```

You can use the `auth_basic` directives in the `http`, `server`, or `location` contexts. The `auth_basic` directive takes a string parameter, which is displayed on the basic authentication pop-up window when an unauthenticated user arrives. The `auth_basic_user_file` specifies a path to the user file.

To test your configuration, you can use `curl` with the `-u` or `--user` flag to build an `Authorization` header for the request:

```
$ curl --user myuser:MyPassword1234 https://localhost
```

Discussion

You can generate basic authentication passwords a few ways, and in a few different formats, with varying degrees of security. The `htpasswd` command from Apache can also generate passwords. Both the `openssl` and `htpasswd` commands can generate passwords with the `apr1` algorithm, which NGINX can also understand. The password can also be in the salted SHA-1 format that Lightweight Directory Access Protocol (LDAP) and Dovecot use. NGINX supports more formats and hashing algorithms; however, many of them are considered insecure because they can easily be defeated by brute-force attacks.

You can use basic authentication to protect the context of the entire NGINX host, specific virtual servers, or even just specific location blocks. Basic authentication won't replace user authentication for web applications, but it can help keep private information secure. Under the hood, basic authentication is done by the server returning a 401 unauthorized HTTP code with the response header `WWW-Authenticate`. This header will have a value of `Basic realm="`*`your string`*`"`. This response causes the browser to prompt for a username and password. The username and password are concatenated and delimited with a colon, then base64-encoded, and then sent in a request header named `Authorization`. The `Authorization` request header will specify a `Basic` and `user:password` encoded string. The server decodes the header and verifies against the provided `auth_basic_user_file`. Because the

username and password string is merely base64-encoded, it's recommended to use HTTPS with basic authentication.

6.2 Authentication Subrequests

Problem

You have a third-party authentication system for which you would like requests authenticated.

Solution

Use the `http_auth_request_module` to make a request to the authentication service to verify identity before serving the request:

```
location /private/ {
    auth_request     /auth;
    auth_request_set $auth_status $upstream_status;
}

location = /auth {
    internal;
    proxy_pass              http://auth-server;
    proxy_pass_request_body off;
    proxy_set_header        Content-Length "";
    proxy_set_header        X-Original-URI $request_uri;
}
```

The `auth_request` directive takes a URI parameter that must be a local internal location. The `auth_request_set` directive allows you to set variables from the authentication subrequest.

Discussion

The `http_auth_request_module` enables authentication on every request handled by the NGINX server. The module will use a subrequest to determine if the request is authorized to proceed. A subrequest is when NGINX passes the request to an alternate internal location and observes its response before routing the request to its destination. The `auth` location passes the original request, including the body and headers, to the authentication server. The HTTP status code of the subrequest is what determines whether or not access is granted. If the subrequest returns with an HTTP 200 status code, the authentication is successful and the request is fulfilled. If the subrequest returns HTTP 401 or 403, the same will be returned for the original request.

If your authentication service does not request the request body, you can drop the request body with the `proxy_pass_request_body` directive, as demonstrated. This

practice will reduce the request size and time. Because the response body is discarded, the `Content-Length` header must be set to an empty string. If your authentication service needs to know the URI being accessed by the request, you'll want to put that value in a custom header that your authentication service checks and verifies. If there are things you do want to keep from the subrequest to the authentication service, like response headers or other information, you can use the `auth_request_set` directive to make new variables out of response data.

6.3 Validating JWTs with NGINX Plus

Problem

You need to validate a JWT before the request is handled with NGINX Plus.

Solution

Use NGINX Plus's HTTP JWT authentication module to validate the token signature and embed JWT claims and headers as NGINX variables:

```
location /api/ {
    auth_jwt          "api";
    auth_jwt_key_file conf/keys.json;
}
```

This configuration enables validation of JWTs for this location. The `auth_jwt` directive is passed a string, which is used as the authentication realm. The `auth_jwt` configuration takes an optional token parameter of a variable that holds the JWT. By default, the `Authentication` header is used per the JWT standard. The `auth_jwt` directive can also be used to cancel the effects of required JWT authentication from inherited configurations. To turn `off` authentication, pass the keyword to the `auth_jwt` directive with nothing else. To cancel inherited authentication requirements, pass the `off` keyword to the `auth_jwt` directive with nothing else. The `auth_jwt_key_file` takes a single parameter. This parameter is the path to the key file in standard JSON Web Key (JWK) format.

Discussion

NGINX Plus is able to validate the JSON web-signature types of tokens, as opposed to the JSON web-encryption type, where the entire token is encrypted. NGINX Plus is able to validate signatures that are signed with the HS256, RS256, and ES256 algorithms. Having NGINX Plus validate the token can save the time and resources needed to make a subrequest to an authentication service. NGINX Plus deciphers the JWT header and payload, and captures the standard headers and claims into

embedded variables for your use. The `auth_jwt` directive can be used in the `http`, `server`, `location`, and `limit_except` contexts.

Also See

RFC Standard documentation of JSON Web Signature
RFC Standard documentation of JSON Web Algorithms
RFC Standard documentation of JSON Web Token
NGINX Plus JWT Authentication
"Authenticating API Clients with JWT and NGINX Plus"

6.4 Creating JSON Web Keys

Problem

You need a JSON Web Key (JWK) for NGINX Plus to use.

Solution

NGINX Plus utilizes the JWK format as specified in the RFC standard. This standard allows for an array of key objects within the JWK file.

The following is an example of what the key file may look like:

```json
{"keys":
  [
    {
        "kty":"oct",
        "kid":"0001",
        "k":"OctetSequenceKeyValue"
    },
    {
        "kty":"EC",
        "kid":"0002"
        "crv":"P-256",
        "x": "XCoordinateValue",
        "y": "YCoordinateValue",
        "d": "PrivateExponent",
        "use": "sig"
    },
    {
        "kty":"RSA",
        "kid":"0003"
        "n": "Modulus",
        "e": "Exponent",
        "d": "PrivateExponent"
    }
  ]
}
```

The JWK file shown demonstrates the three initial types of keys noted in the RFC standard. The format of these keys is also part of the RFC standard. The `kty` attribute is the key type. This file shows three key types: the Octet Sequence (`oct`), the EllipticCurve (`EC`), and the `RSA` type. The `kid` attribute is the key ID. Other attributes to these keys are specified in the standard for that type of key. Look to the RFC documentation of these standards for more information.

Discussion

There are numerous libraries available in many different languages to generate the JWK. It's recommended to create a key service that is the central JWK authority to create and rotate your JWKs at a regular interval. For enhanced security, it's recommended to make your JWKs as secure as your SSL/TLS certifications. Secure your key file with proper user and group permissions. Keeping them in memory on your host is best practice. You can do so by creating an in-memory filesystem like ramfs. Rotating keys on a regular interval is also important; you may opt to create a key service that creates public and private keys and offers them to the application and NGINX via an API.

Also See

RFC Standard documentation of JSON Web Key

6.5 Validate JSON Web Tokens with NGINX Plus

Problem

You want to validate JSON Web Tokens with NGINX Plus.

Solution

Use the JWT module that comes with NGINX Plus to secure a location or server, and instruct the `auth_jwt` directive to use `$cookie_auth_token` as the token to be validated:

```
location /private/ {
    auth_jwt "Google Oauth" token=$cookie_auth_token;
    auth_jwt_key_file /etc/nginx/google_certs.jwk;
}
```

This configuration directs NGINX Plus to secure the */private/* URI path with JWT validation. Google OAuth 2.0 OpenID Connect uses the cookie `auth_token` rather than the default bearer token. Thus, you must instruct NGINX to look for the token in this cookie rather than in the NGINX Plus default location. The

`auth_jwt_key_file` location is set to an arbitrary path, which is a step that we cover in Recipe 6.6.

Discussion

This configuration demonstrates how you can validate a Google OAuth 2.0 OpenID Connect JWT with NGINX Plus. The NGINX Plus JWT authentication module for HTTP is able to validate any JWT that adheres to the RFC for JSON Web Signature specification, instantly enabling any SSO authority that utilizes JWTs to be validated at the NGINX Plus layer. The OpenID 1.0 protocol is a layer on top of the OAuth 2.0 authentication protocol that adds identity, enabling the use of JWTs to prove the identity of the user sending the request. With the signature of the token, NGINX Plus can validate that the token has not been modified since it was signed. In this way, Google is using an asynchronous signing method and makes it possible to distribute public JWKs while keeping its private JWK secret.

Also See

"Authenticating API Clients with JWT and NGINX Plus"

6.6 Automatically Obtaining and Caching JSON Web Key Sets with NGINX Plus

Problem

You want NGINX Plus to automatically request the JSON Web Key Set (JWKS) from a provider and cache it.

Solution

Utilize a cache zone and the `auth_jwt_key_request` directive to automatically keep your key up to date:

```
proxy_cache_path /data/nginx/cache levels=1 keys_zone=foo:10m;

server {
    # ...

    location / {
        auth_jwt             "closed site";
        auth_jwt_key_request /jwks_uri;
    }

    location = /jwks_uri {
        internal;
        proxy_cache foo;
```

```
            proxy_pass  https://idp.example.com/keys;
        }
    }
```

In this example, the `auth_jwt_key_request` directive instructs NGINX Plus to retrieve the JWKS from an internal subrequest. The subrequest is directed to `/jwks_uri`, which will proxy the request to an identity provider. The request is cached for a default of 10 minutes to limit overhead.

Discussion

In NGINX Plus R17, the `auth_jwt_key_request` directive was introduced. This feature enables the NGINX Plus server to dynamically update its JWKs when a request is made. A subrequest method is used to fetch the JWKs, which means the location that the directive points to must be local to the NGINX Plus server. In the example, the subrequest location was locked down to ensure that only internal NGINX Plus requests would be served. A cache was also used to ensure the JWKs retrieval request is only made as often as necessary, and does not overload the identity provider. The `auth_jwt_key_request` directive is valid in the `http`, `server`, `location`, and `limit_except` contexts.

Also See

"Authenticating API Clients with JWT and NGINX Plus"
NGINX Plus "Faster JWT Validation with JSON Web Key Set Caching"

6.7 Authenticate Users via Existing OpenID Connect SSO with NGINX Plus

Problem

You want to integrate NGINX Plus with an OpenID Connect (OIDC) identity provider.

Solution

This solution consists of a number of configuration aspects and a bit of NGINX JavaScript code. The identity provider (IdP) must support OpenID Connect 1.0. NGINX Plus will act as a relaying party of your OIDC in an authorization code flow.

NGINX Inc. maintains a public GitHub repository containing configuration and code as a reference implementation of OIDC integration with NGINX Plus. The link to the repository in the Also See section has up-to-date instructions on how to set up the reference implementation with your own IdP.

Discussion

This solution simply linked to a reference implementation to ensure that you, the reader, have the most up-to-date solution. The reference provided configures NGINX Plus as a relaying party to an authorization code flow for OpenID Connect 1.0. When unauthenticated requests for protected resources are made to NGINX Plus in this configuration, NGINX Plus first redirects the request to the IdP. The IdP takes the client through its own login flow, and returns the client to NGINX Plus with an authentication code. NGINX Plus then communicates directly with the IdP to exchange the authentication code for a set of ID Tokens. These tokens are validated using JWTs, and stored in NGINX Plus's key-value store. By using the key-value store, the tokens are made available to all NGINX Plus nodes in a highly available (HA) configuration. During this process, NGINX Plus generates a session cookie for the client that is used as the key to look up the token in the key-value store. The client is then served a redirect with the cookie to the initial requested resource. Subsequent requests are validated by using the cookie to look up the ID Token in NGINX Plus's key-value store.

This capability enables integration with most major identity providers, including CA Single Sign On (formerly SiteMinder), ForgeRock OpenAM, Keycloak, Okta, OneLogin, and Ping Identity. OIDC as a standard is extremely relevant in authentica- tion—the aforementioned identity providers are only a subset of the integrations that are possible.

Also See

"Authenticating Users to Existing Applications with OpenID Connect and NGINX Plus"
OpenID Connect
NGINX OpenID Connect GitHub

Security Controls

7.0 Introduction

Security is done in layers, and there must be multiple layers to your security model for it to be truly hardened. In this chapter, we go through many different ways to secure your web applications with NGINX and NGINX Plus. You can use many of these security methods in conjunction with one another to help harden security. You might notice that this chapter does not touch upon the ModSecurity 3.0 NGINX module, which turns NGINX into a Web Application Firewall (WAF). To learn more about the WAF capabilities, download the *ModSecurity 3.0 and NGINX: Quick Start Guide*. Please note, the NGINX ModSecurity WAF for NGINX Plus is transitioning to End-of-Life (EoL) effective March 31, 2024. For more information please read this blog.

7.1 Access Based on IP Address

Problem

You need to control access based on the IP address of the client.

Solution

Use the HTTP or stream access module to control access to protected resources:

```
location /admin/ {
    deny  10.0.0.1;
    allow 10.0.0.0/20;
    allow 2001:0db8::/32;
    deny  all;
}
```

The given location block allows access from any IPv4 address in 10.0.0.0/20 except 10.0.0.1, allows access from IPv6 addresses in the `2001:0db8::/32` subnet, and returns a 403 for requests originating from any other address. The `allow` and `deny` directives are valid within the `http`, `server`, and `location` contexts, as well as in `stream` and `server` context for TCP/UDP. Rules are checked in sequence until a match is found for the remote address.

Discussion

Protecting valuable resources and services on the internet must be done in layers. NGINX functionality provides the ability to be one of those layers. The `deny` directive blocks access to a given context, while the `allow` directive can be used to allow subsets of the blocked access. You can use IP addresses, IPv4 or IPv6, classless inter-domain routing (CIDR) block ranges, the keyword `all`, and a Unix socket. Typically, when protecting a resource, one might allow a block of internal IP addresses and deny access from all.

7.2 Allowing Cross-Origin Resource Sharing

Problem

You're serving resources from another domain and need to allow cross-origin resource sharing (CORS) to enable browsers to utilize these resources.

Solution

Alter headers based on the `request` method to enable CORS:

```
map $request_method $cors_method {
  OPTIONS 11;
  GET  1;
  POST 1;
  default 0;
}
server {
  # ...
  location / {
    if ($cors_method ~ '1') {
        add_header 'Access-Control-Allow-Methods'
            'GET,POST,OPTIONS';
        add_header 'Access-Control-Allow-Origin'
            '*.example.com';
        add_header 'Access-Control-Allow-Headers'
                    'DNT,
                    Keep-Alive,
                    User-Agent,
                    X-Requested-With,
```

```
                   If-Modified-Since,
                   Cache-Control,
                   Content-Type';
    }
    if ($cors_method = '11') {
        add_header 'Access-Control-Max-Age' 1728000;
        add_header 'Content-Type' 'text/plain; charset=UTF-8';
        add_header 'Content-Length' 0;
        return 204;
    }
  }
}
```

There's a lot going on in this example, which has been condensed by using a `map` to group the `GET` and `POST` methods together. The `OPTIONS` request method returns a *preflight* request to the client about this server's CORS rules. `OPTIONS`, `GET`, and `POST` methods are allowed under CORS. Setting the `Access-Control-Allow-Origin` header allows for content being served from this server to also be used on pages of origins that match this header. The preflight request can be cached on the client for 1,728,000 seconds, or 20 days.

Discussion

Resources such as JavaScript make CORS when the resource they're requesting is of a domain other than its own. When a request is considered cross origin, the browser is required to obey CORS rules. The browser will not use the resource if it does not have headers that specifically allow its use. To allow our resources to be used by other subdomains, we have to set the CORS headers, which can be done with the `add_header` directive. If the request is a `GET`, `HEAD`, or `POST` with standard content type, and the request does not have special headers, the browser will make the request and only check for origin. Other request methods will cause the browser to make the preflight request to check the terms of the server it will obey for that resource. If you do not set these headers appropriately, the browser will give an error when trying to utilize that resource.

7.3 Client-Side Encryption

Problem

You need to encrypt traffic between your NGINX server and the client.

Solution

Utilize one of the SSL modules to encrypt traffic, such as `ngx_http_ssl_module` or `ngx_stream_ssl_module`:

```
http { # All directives used below are also valid in stream
    server {
        listen 8443 ssl;
        ssl_certificate /etc/nginx/ssl/example.crt;
        ssl_certificate_key /etc/nginx/ssl/example.key;
    }
}
```

This configuration sets up a server to listen on a port encrypted with SSL/TLS, 8443. The directive `ssl_certificate` defines the certificate and optional chain that is served to the client. The `ssl_certificate_key` directive defines the key used by NGINX to decrypt requests and encrypt responses. A number of SSL/TLS negotiation configurations are defaulted to secure presets for the NGINX version release date.

Discussion

Secure transport layers are the most common way of encrypting information in transit. As of this writing, the TLS protocol is preferred over the SSL protocol. That's because versions 1 through 3 of SSL are now considered insecure. Although the protocol name might be different, TLS still establishes a secure socket layer. NGINX enables your service to protect information between you and your clients, which in turn protects the client and your business. When using a CA-signed certificate, you need to concatenate the certificate with the certificate authority chain. When you concatenate your certificate and the chain, your certificate should be above the concatenated chain file. If your certificate authority has provided multiple files as intermediate certificates for the chain, there is an order in which they are layered. Refer to the certificate provider's documentation for the order.

Also See

Mozilla Security/Server Side TLS page
Mozilla SSL Configuration Generator
SSL Labs' SSL Server Test

7.4 Advanced Client-Side Encryption

Problem

You have advanced client-server encryption configuration needs.

Solution

The `http` and `stream` SSL modules for NGINX enable complete control of the accepted SSL/TLS handshake. Certificates and keys can be provided to NGINX by way of file path or variable value. NGINX presents the client with an accepted list of protocols, ciphers, and key types, per its configuration. The highest standard between the client and NGINX server is negotiated. NGINX can cache the result of client-server SSL/TLS negotiation for a period of time.

The following intentionally demonstrates many options at once to illustrate the available complexity of the client-server negotiation:

```
http { # All directives used below are also valid in stream
    server {
        listen 8443 ssl;
        # Set accepted protocol and cipher
        ssl_protocols TLSv1.2 TLSv1.3;
        ssl_ciphers HIGH:!aNULL:!MD5;

        # RSA certificate chain loaded from file
        ssl_certificate /etc/nginx/ssl/example.crt;
        # RSA encryption key loaded from file
        ssl_certificate_key /etc/nginx/ssl/example.pem;

        # Elliptic curve cert from variable value
        ssl_certificate $ecdsa_cert;
        # Elliptic curve key as file path variable
        ssl_certificate_key data:$ecdsa_key_path;

        # Client-Server negotiation caching
        ssl_session_cache shared:SSL:10m;
        ssl_session_timeout 10m;
    }
}
```

The server accepts the SSL protocol versions TLSv1.2 and TLSv1.3. The ciphers accepted are set to `HIGH`, which is a macro for the highest standard; explicit denies are demonstrated for `aNULL` and `MD5` by denotation of the `!`.

Two sets of certificate-key pairs are used. The values passed to the NGINX directives demonstrate different ways to provide NGINX certificate-key values. A variable is interpreted as a path to a file. When prefixed with `data:` the value of a variable is interpreted as a direct value. Multiple certificate-key formats may be provided to offer

reverse compatibility to the client. The strongest standard capable by the client and accepted by the server will be the result of the negotiation.

> If the SSL/TLS key is exposed as a direct value variable, it has the potential of being logged or exposed by the configuration. Ensure you have strict change and access controls if exposing the key value as a variable.

The SSL session cache and timeout allow NGINX worker processes to cache and store session parameters for a given amount of time. The NGINX worker processes share this cache between themselves as processes within a single instantiation, but not between machines. There are many other session cache options that can help with performance or security of all types of use cases. You can use session cache options in conjunction with one another. However, specifying one without the default will turn off the default built-in session cache.

Discussion

In this advanced example, NGINX provides the client with the SSL/TLS options of TLS version 1.2 or 1.3, highly regarded cipher algorithms, and the ability to use RSA or Elliptic Curve Cryptography (ECC) formatted keys. The strongest of the protocols, ciphers, and key formats the client is capable of is the result of the negotiation. The configuration instructs NGINX to cache the negotiation for a period of 10 minutes with the available memory allocation of 10 MB.

In testing, ECC certificates were found to be faster than the equivalent-strength RSA certificates. The key size is smaller, which results in the ability to serve more SSL/TLS connections, and with faster handshakes. NGINX allows you to configure multiple certificates and keys, and then serve the optimal certificate for the client browser. This allows you to take advantage of the newer technology but still serve older clients.

> NGINX is encrypting the traffic between itself and the client in this example. The connection to upstream servers may also be encrypted, however. The negotiation between NGINX and the upstream server is demonstrated in Recipe 7.5.

Also See

Mozilla Security/Server Side TLS page
Mozilla SSL Configuration Generator
SSL Labs' SSL Server Test

7.5 Upstream Encryption

Problem

You need to encrypt traffic between NGINX and the upstream service and set specific negotiation rules for compliance regulations, or the upstream service is outside of your secured network.

Solution

Use the SSL directives of the HTTP proxy module to specify SSL rules:

```
location / {
    proxy_pass https://upstream.example.com;
    proxy_ssl_verify on;
    proxy_ssl_verify_depth 2;
    proxy_ssl_protocols TLSv1.2;
}
```

These proxy directives set specific SSL rules for NGINX to obey. The configured directives ensure that NGINX verifies that the certificate and chain on the upstream service is valid up to two certificates deep. The `proxy_ssl_protocols` directive specifies that NGINX will only use TLS version 1.2. By default, NGINX does not verify upstream certificates and accepts all TLS versions.

Discussion

The configuration directives for the HTTP proxy module are vast, and if you need to encrypt upstream traffic, you should at least turn on verification. You can proxy over HTTPS simply by changing the protocol on the value passed to the `proxy_pass` directive. However, this does not validate the upstream certificate. Other directives, such as `proxy_ssl_certificate` and `proxy_ssl_certificate_key`, allow you to lock down upstream encryption for enhanced security. You can also specify `proxy_ssl_crl` or a certificate revocation list, which lists certificates that are no longer considered valid. These SSL proxy directives help harden your system's communication channels within your own network or across the public internet.

7.6 Securing a Location

Problem

You need to secure a location block using a secret.

Solution

Use the secure link module and the `secure_link_secret` directive to restrict access to resources to users who have a secure link:

```
location /resources {
    secure_link_secret mySecret;
    if ($secure_link = "") { return 403; }

    rewrite ^ /secured/$secure_link;
}

location /secured/ {
    internal;
    root /var/www;
}
```

This configuration creates an internal and public-facing location block. The public-facing location block `/resources` will return a 403 Forbidden unless the request URI includes an `md5` hash string that can be verified with the secret provided to the `secure_link_secret` directive. The `$secure_link` variable is an empty string unless the hash in the URI is verified.

Discussion

Securing resources with a secret is a great way to ensure your files are protected. The secret is used in conjunction with the URI. This string is then `md5` hashed, and the hex digest of that `md5` hash is used in the URI. The hash is placed into the link and evaluated by NGINX. NGINX knows the path to the file being requested because it's in the URI after the hash. NGINX also knows your secret as it's provided via the `secure_link_secret` directive. NGINX is able to quickly validate the `md5` hash and store the URI in the `$secure_link` variable. If the hash cannot be validated, the variable is set to an empty string. It's important to note that the argument passed to the `secure_link_secret` must be a static string; it cannot be a variable.

7.7 Generating a Secure Link with a Secret

Problem

You need to generate a secure link from your application using a secret.

Solution

The secure link module in NGINX accepts the hex digest of an `md5` hashed string, where the string is a concatenation of the URI path and the secret. Building on the last section, Recipe 7.6, we will create the secured link that will work with the

previous configuration example given that there's a file present at */var/www/secured/index.html*. To generate the hex digest of the `md5` hash, we can use the Unix `openssl` command:

```
$ echo -n 'index.htmlmySecret' | openssl md5 -hex
(stdin)= a53bee08a4bf0bbea978ddf736363a12
```

Here we show the URI that we're protecting, *index.html*, concatenated with our secret, `mySecret`. This string is passed to the `openssl` command to output an `md5` hex digest.

The following is an example of the same hash digest being constructed in Python using the `hashlib` library that is included in the Python Standard Library:

```python
import hashlib
hashlib.md5.(b'index.htmlmySecret').hexdigest()
'a53bee08a4bf0bbea978ddf736363a12'
```

Now that we have this hash digest, we can use it in a URL. Our example will be *www.example.com* making a request for the file */var/www/secured/index.html* through our */resources* location. Our full URL will be the following:

```
www.example.com/resources/a53bee08a4bf0bbea978ddf736363a12/\
index.html
```

Discussion

Generating the digest can be done in many ways, in many languages. Things to remember: the URI path goes before the secret, there are no carriage returns in the string, and use the hex digest of the `md5` hash.

7.8 Securing a Location with an Expire Date

Problem

You need to secure a location with a link that expires at some future time and is specific to a client.

Solution

Utilize the other directives included in the secure link module to set an expire time and use variables in your secure link:

```
location /resources {
    root /var/www;
    secure_link $arg_md5,$arg_expires;
    secure_link_md5 "$secure_link_expires$uri$remote_addrmySecret";
    if ($secure_link = "") { return 403; }
    if ($secure_link = "0") { return 410; }
}
```

The `secure_link` directive takes two parameters separated with a comma. The first parameter is the variable that holds the `md5` hash. This example uses an HTTP argument of `md5`. The second parameter is a variable that holds the time in which the link expires in Unix epoch time format. The `secure_link_md5` directive takes a single parameter that declares the format of the string that is used to construct the `md5` hash. Like the other configuration, if the hash does not validate, the `$secure_link` variable is set to an empty string. However, with this usage, if the hash matches but the time has expired, the `$secure_link` variable will be set to `0`.

Discussion

This usage of securing a link is more flexible and looks cleaner than the `secure_link_secret` shown in Recipe 7.6. With these directives, you can use any number of variables that are available to NGINX in the hashed string. Using user-specific variables in the hash string will strengthen your security, as users won't be able to trade links to secured resources. It's recommended to use a variable like `$remote_addr` or `$http_x_forwarded_for`, or a session cookie header generated by the application. The arguments to `secure_link` can come from any variable you prefer, and they can be named whatever best fits. The conditions are: Do you have access? Are you accessing it within the time frame? If you don't have access: Forbidden. If you have access but you're late: Gone. The HTTP 410, Gone, works great for expired links because the condition is to be considered permanent.

7.9 Generating an Expiring Link

Problem

You need to generate a link that expires.

Solution

Generate a timestamp for the expire time in the Unix epoch format. On a Unix system, you can test by using the date as demonstrated in the following:

```
$ date -d "2030-12-31 00:00" +%s --utc
1924905600
```

Next, you'll need to concatenate your hash string to match the string configured with the `secure_link_md5` directive. In this case, our string to be used will be `1924905600/resources/index.html127.0.0.1 mySecret`. The `md5` hash is a bit different than just a hex digest. It's an `md5` hash in binary format, base64-encoded, with plus signs (+) translated to hyphens (-), slashes (/) translated to underscores (_), and equal signs (=) removed. The following is an example on a Unix system:

```
$ echo -n '1924905600/resources/index.html127.0.0.1 mySecret' \
  | openssl md5 -binary \
  | openssl base64 \
  | tr +/ -_ \
  | tr -d =
sqysOw5kMvQBL3j9ODCyoQ
```

Now that we have our hash, we can use it as an argument along with the expire date:

```
/resources/index.html?md5=sqysOw5kMvQBL3j9ODCyoQ&expires=1924905600
```

The following is a more practical example in Python utilizing a relative time for the expiration, setting the link to expire one hour from generation. At the time of writing this example works with Python 2.7 and 3.x utilizing the Python Standard Library:

```python
from datetime import datetime, timedelta
from base64 import b64encode
import hashlib

# Set environment vars
resource = b'/resources/index.html'
remote_addr = b'127.0.0.1'
host = b'www.example.com'
mysecret = b'mySecret'

# Generate expire timestamp
now = datetime.utcnow()
expire_dt = now + timedelta(hours=1)
expire_epoch = str.encode(expire_dt.strftime('%s'))

# md5 hash the string
uncoded = expire_epoch + resource + remote_addr + mysecret
md5hashed = hashlib.md5(uncoded).digest()

# Base64 encode and transform the string
b64 = b64encode(md5hashed)
unpadded_b64url = b64.replace(b'+', b'-')\
    .replace(b'/', b'_')\
    .replace(b'=', b'')

# Format and generate the link
linkformat = "{}{}?md5={}?expires={}"
securelink = linkformat.format(
    host.decode(),
    resource.decode(),
    unpadded_b64url.decode(),
    expire_epoch.decode()
)
print(securelink)
```

Discussion

With this pattern, we're able to generate a secure link in a special format that can be used in URLs. The secret provides security through use of a variable that is never sent to the client. You're able to use as many other variables as you need to in order to secure the location. `md5` hashing and base64 encoding are common, lightweight, and available in nearly every language.

7.10 HTTPS Redirects

Problem

You need to redirect unencrypted requests to HTTPS.

Solution

Use a rewrite to send all HTTP traffic to HTTPS:

```
server {
    listen 80 default_server;
    listen [::]:80 default_server;
    server_name _;
    return 301 https://$host$request_uri;
}
```

This configuration listens on port `80` as the default server for both IPv4 and IPv6 and for any hostname. The `return` statement returns a 301 permanent redirect to the HTTPS server at the same host and request URI.

Discussion

It's important to always redirect to HTTPS where appropriate. You may find that you do not need to redirect all requests but only those with sensitive information being passed between client and server. In that case, you may want to put the `return` statement in particular locations only, such as */login*.

7.11 Redirecting to HTTPS Where SSL/TLS Is Terminated Before NGINX

Problem

You need to redirect to HTTPS, however, you've terminated SSL/TLS at a layer before NGINX.

Solution

Use the common `X-Forwarded-Proto` header to determine if you need to redirect:

```
server {
    listen 80 default_server;
    listen [::]:80 default_server;
    server_name _;
    if ($http_x_forwarded_proto = 'http') {
        return 301 https://$host$request_uri;
    }
}
```

This configuration is very much like HTTPS redirects. However, in this configuration we're only redirecting if the header `X-Forwarded-Proto` is equal to HTTP.

Discussion

It's a common use case that you may terminate SSL/TLS in a layer in front of NGINX. One reason you may do something like this is to save on compute costs. However, you need to make sure that every request is HTTPS, but the layer terminating SSL/TLS does not have the ability to redirect. It can, however, set proxy headers. This configuration works with layers such as the Amazon Web Services Elastic Load Balancer (AWS ELB), which will offload SSL/TLS at no additional cost. This is a handy trick to make sure that your HTTP traffic is secured.

7.12 HTTP Strict Transport Security

Problem

You need to instruct browsers to never send requests over HTTP.

Solution

Use the HTTP Strict Transport Security (HSTS) enhancement by setting the `Strict-Transport-Security` header:

```
add_header Strict-Transport-Security max-age=31536000;
```

This configuration sets the `Strict-Transport-Security` header to a max age of a year. This will instruct the browser to always do an internal redirect when HTTP requests are attempted to this domain, so that all requests will be made over HTTPS.

Discussion

For some applications, a single HTTP request trapped by a man-in-the-middle attack could be the end of the company. If a form post containing sensitive information

is sent over HTTP, the HTTPS redirect from NGINX won't save you; the damage is done. This opt-in security enhancement informs the browser to never make an HTTP request, and therefore the request is never sent unencrypted.

Also See

RFC Standard documentation of HTTP Strict Transport Security
OWASP HTTP Strict Transport Security Cheat Sheet

7.13 Satisfying Any Number of Security Methods

Problem

You need to provide multiple ways to pass security to a closed site.

Solution

Use the `satisfy` directive to instruct NGINX that you want to satisfy any or all of the security methods used:

```
location / {
    satisfy any;

    allow 192.168.1.0/24;
    deny  all;

    auth_basic           "closed site";
    auth_basic_user_file conf/htpasswd;
}
```

This configuration tells NGINX that the user requesting the `location /` needs to satisfy one of the security methods: either the request needs to originate from the *192.168.1.0/24* CIDR block or it must be able to supply a username and password that can be found in the *conf/htpasswd* file. The `satisfy` directive takes one of two options: `any` or `all`.

Discussion

The `satisfy` directive is a great way to offer multiple ways to authenticate to your web application. By specifying `any` to the `satisfy` directive, the user must meet one of the security challenges. By specifying `all` to the `satisfy` directive, the user must meet all of the security challenges. This directive can be used in conjunction with the `http_access_module` detailed in Recipe 7.1, the `http_auth_basic_module` detailed in Recipe 6.1, the `http_auth_request_module` detailed in Recipe 6.2, and the `http_auth_jwt_module` detailed in Recipe 6.3. Security is only truly secure if it's done

in multiple layers. The `satisfy` directive will help you achieve this for locations and servers that require deep security rules.

7.14 NGINX Plus Dynamic Application Layer DDoS Mitigation

Problem

You need a dynamic Distributed Denial of Service (DDoS) mitigation solution.

Solution

Use NGINX Plus to build a cluster-aware rate limit and automatic blocklist:

```
limit_req_zone    $remote_addr zone=per_ip:1M rate=100r/s sync;
                  # Cluster-aware rate limit
limit_req_status 429;

keyval_zone zone=sinbin:1M timeout=600 sync;
              # Cluster-aware "sin bin" with
              # 10-minute TTL
keyval $remote_addr $in_sinbin zone=sinbin;
              # Populate $in_sinbin with
              # matched client IP addresses

server {
    listen 80;
    location / {
        if ($in_sinbin) {
            set $limit_rate 50; # Restrict bandwidth of bad clients
        }

        limit_req zone=per_ip;
              # Apply the rate limit here
        error_page 429 = @send_to_sinbin;
              # Excessive clients are moved to
              # this location
        proxy_pass http://my_backend;
    }

    location @send_to_sinbin {
        rewrite ^ /api/3/http/keyvals/sinbin break;
              # Set the URI of the
              # "sin bin" key-val
        proxy_method POST;
        proxy_set_body '{"$remote_addr":"1"}';
        proxy_pass http://127.0.0.1:80;
    }

    location /api/ {
```

```
            api write=on;
            # directives to control access to the API
        }
    }
```

Discussion

This solution uses a synchronized rate limit by use of a synchronized key-value store to dynamically respond to DDoS attacks and mitigate their effects. The `sync` parameter provided to the `limit_req_zone` and `keyval_zone` directives synchronizes the shared memory zone with other machines in the active-active NGINX Plus cluster. This example identifies clients that send more than 100 requests per second, regardless of which NGINX Plus node receives the request. When a client exceeds the rate limit, its IP address is added to a "sin bin" key-value store by making a call to the NGINX Plus API. The sin bin is synchronized across the cluster. Further requests from clients in the sin bin are subject to a very low bandwidth limit, regardless of which NGINX Plus node receives them. Limiting bandwidth is preferable to rejecting requests outright because it does not clearly signal to the client that DDoS mitigation is in effect. After 10 minutes, the client is automatically removed from the sin bin.

7.15 Installing and Configuring NGINX Plus with the NGINX App Protect WAF Module

Problem

You need to install and configure the NGINX App Protect WAF module.

Solution

Follow the NGINX App Protect WAF administration guide for your platform. Make sure not to skip the portion about installing NGINX App Protect WAF signatures from the separate repository.

Ensure that the NGINX App Protect WAF module is dynamically loaded by NGINX Plus using the `load_module` directive in the main context, and enabled by using the `app_protect_*` directives:

```
user nginx;
worker_processes  auto;

load_module modules/ngx_http_app_protect_module.so;

# ... Other main context directives

http {
    app_protect_enable on;
```

```
    app_protect_policy_file "/etc/nginx/AppProtectTransparentPolicy.json";
    app_protect_security_log_enable on;
    app_protect_security_log "/etc/nginx/log-default.json"
      syslog:server=127.0.0.1:515;

    # ... Other http context directives
}
```

In this example, the `app_protect_enable` directive set to `on` enabled the module for the current context. This directive, and all of the following, are valid within the `http` context, as well as the `server` and `location` contexts with HTTP. The `app_protect_policy_file` directive points to an NGINX App Protect WAF policy file, which we will define next; if not defined, the default policy is used. Security logging is configured next, and requires a remote logging server. For the example, we've configured it to the local Syslog listener. The `app_protect_security_log` directive takes two parameters; the first is a JSON file that defines the logging settings, and the second is a log stream destination. The log settings file will be shown later in this section.

Build an NGINX App Protect WAF policy file, naming it */etc/nginx/AppProtectTransparentPolicy.json*:

```
{
    "policy": {
        "name": "transparent_policy",
        "template": { "name": "POLICY_TEMPLATE_NGINX_BASE" },
        "applicationLanguage": "utf-8",
        "enforcementMode": "transparent"
    }
}
```

This policy file configures the default NGINX App Protect WAF policy by use of a template, setting the policy name to `transparent_policy`, and setting the `enforcementMode` to `transparent`, which means NGINX Plus will log but not block. Transparent mode is great to test out new policies before putting them into effect.

Enable blocking by changing the `enforcementMode` to `blocking`. This policy file can be named */etc/nginx/AppProtectBlockingPolicy.json*. To switch between the files, update the `app_protect_policy_file` directive in your NGINX Plus configuration:

```
{
    "policy": {
        "name": "blocking_policy",
        "template": { "name": "POLICY_TEMPLATE_NGINX_BASE" },
        "applicationLanguage": "utf-8",
        "enforcementMode": "blocking"
    }
}
```

To enable some of the protection features of NGINX App Protect WAF, enable some violations:

```
{
    "policy": {
        "name": "blocking_policy",
        "template": { "name": "POLICY_TEMPLATE_NGINX_BASE" },
        "applicationLanguage": "utf-8",
        "enforcementMode": "blocking",
        "blocking-settings": {
            "violations": [
                {
                    "name": "VIOL_JSON_FORMAT",
                    "alarm": true,
                    "block": true
                },
                {
                    "name": "VIOL_PARAMETER_VALUE_METACHAR",
                    "alarm": true,
                    "block": false
                }
            ]
        }
    }
}
```

The preceding example demonstrates adding two violations to our policy. Take note that VIOL_PARAMETER_VALUE_METACHAR is not set to block, but only alarm; whereas VIOL_JSON_FORMAT is set to block and alarm. This functionality enables the overriding of the default enforcementMode when set to blocking. When enforcementMode is set to transparent, the default enforcement setting takes precedence.

Set up an NGINX Plus logging file, named */etc/nginx/log-default.json*:

```
{
    "filter":{
        "request_type":"all"
    },
    "content":{
        "format":"default",
        "max_request_size":"any",
        "max_message_size":"5k"
    }
}
```

This file was defined in the NGINX Plus configuration by the app_protect_secu rity_log directive and is necessary for NGINX App Protect WAF logging.

Discussion

This solution demonstrates the basis of configuring NGINX Plus with the NGINX App Protect WAF module. The NGINX App Protect WAF module enables an entire suite of Web Application Firewall (WAF) definitions. These definitions derive from the Advanced F5 Application Security functionality in F5. This comprehensive set of WAF attack signatures has been extensively field-tested and proven. By adding this to an NGINX Plus installation, it renders the best of F5 Application Security with the agility of the NGINX Platform.

Once the module is installed and enabled, most of the configuration is done in a policy file. The policy files in this section showed how to enable active blocking, passive monitoring, and transparent mode, as well as explained overrides to this functionality with violations. Violations are only one type of protection offered. Other protections include HTTP Compliance, Evasion Techniques, Attack Signatures, Server Technologies, Data Guard, and many more. To retrieve NGINX App Protect WAF logs, it's necessary to use the NGINX Plus logging format and send the logs to a remote listening service, a file, or `/dev/stderr`.

If you're using NGINX Controller ADC, you can enable NGINX App Protect WAF capabilities through the NGINX Controllers App Security component, and visualize the WAF metrics through the web interface.

Also See

NGINX App Protect WAF Administration Guides
NGINX App Protect WAF Configuration Guide
NGINX Controller App Security guide
NGINX App Protect DoS deployment guide

HTTP/2

8.0 Introduction

HTTP/2 is a major revision to the HTTP protocol. Much of the change to this revision is focused on the transport layer, such as enabling full request and response multiplexing over a single TCP connection. Efficiencies were gained through the use of compression on HTTP header fields, and support for request prioritization was added. Another large addition to the protocol was the ability for the server to push messages to the client. This chapter details the basic configuration for enabling HTTP/2 in NGINX as well as configuring Google's open source remote procedure call (gRPC) and HTTP/2 server push support.

8.1 Basic Configuration

Problem

You want to take advantage of HTTP/2.

Solution

Turn on HTTP/2 on your NGINX server:

```
server {
    listen 443 ssl http2 default_server;

    ssl_certificate    server.crt;
    ssl_certificate_key server.key;
    # ...
}
```

Discussion

To turn on HTTP/2, you simply need to add the `http2` parameter to the `listen` directive. The catch, however, is that although the protocol does not require the connection to be wrapped in SSL/TLS, some implementations of HTTP/2 clients only support HTTP/2 over an encrypted connection. Another caveat is that the HTTP/2 specification has blocked a number of TLS 1.2 cipher suites, and therefore will fail the handshake. The ciphers NGINX uses by default are not on the blocklist. The Application-Layer Protocol Negotiation of TLS allows the application layer to negotiate which protocol should be used over the secure connection to avoid additional round trips. To test that your setup is correct, you can install a plug-in for Chrome and Firefox browsers that indicates when a site is using HTTP/2, or on the command line with the `nghttp` utility.

Also See

[HTTP/2 RFC Cipher Suite Black List](#)
[Chrome HTTP/2 and SPDY indicator plug-in](#)
[Firefox HTTP Indicator add-on](#)

8.2 gRPC

Problem

You need to terminate, inspect, route, or load balance gRPC method calls.

Solution

Use NGINX to proxy gRPC connections:

```
server {
    listen 80 http2;

    location / {
        grpc_pass grpc://backend.local:50051;
    }
}
```

In this configuration, NGINX is listening on port `80` for unencrypted HTTP/2 traffic, and proxying that traffic to a machine named `backend.local` on port `50051`. The `grpc_pass` directive instructs NGINX to treat the communication as a gRPC call. The `grpc://` in front of our backend server location is not necessary; however, it does directly indicate that the backend communication is not encrypted.

To utilize TLS encryption between the client and NGINX, and terminate that encryption before passing the calls to the application server, turn on SSL and HTTP/2, as you did in the first section:

```
server {
    listen 443 ssl http2 default_server;

    ssl_certificate     server.crt;
    ssl_certificate_key server.key;
    location / {
        grpc_pass grpc://backend.local:50051;
    }
}
```

This configuration terminates TLS at NGINX and passes the gRPC communication to the application over unencrypted HTTP/2.

To configure NGINX to encrypt the gRPC communication to the application server, providing end-to-end encrypted traffic, simply modify the `grpc_pass` directive to specify `grpcs://` before the server information (note the addition of the `s` denoting secure communication):

```
grpc_pass grpcs://backend.local:50051;
```

You also can use NGINX to route calls to different backend services based on the gRPC URI, which includes the package, service, and method. To do so, utilize the `location` directive:

```
location /mypackage.service1 {
    grpc_pass grpc://$grpc_service1;
}
location /mypackage.service2 {
    grpc_pass grpc://$grpc_service2;
}
location / {
    root /usr/share/nginx/html;
    index index.html index.htm;
}
```

This configuration example uses the `location` directive to route incoming HTTP/2 traffic between two separate gRPC services, as well as a `location` to serve static content. Method calls for the `mypackage.service1` service are directed to the value of the variable `grpc_service1`, which may contain a hostname or IP and optional port. Calls for `mypackage.service2` are directed to the value of the variable `grpc_service2`. The `location /` catches any other HTTP request and serves static content. This demonstrates how NGINX is able to serve gRPC and non-gRPC under the same HTTP/2 endpoint and route accordingly.

Load balancing gRPC calls is also similar to non-gRPC HTTP traffic:

```
upstream grpcservers {
    server backend1.local:50051;
    server backend2.local:50051;
}
server {
    listen 443 ssl http2 default_server;

    ssl_certificate     server.crt;
    ssl_certificate_key server.key;
    location / {
        grpc_pass grpc://grpcservers;
    }
}
```

The `upstream` block works the exact same way for gRPC as it does for other HTTP traffic. The only difference is that the `upstream` is referenced by `grpc_pass`.

Discussion

NGINX is able to receive, proxy, load balance, route, and terminate encryption for gRPC calls. The gRPC module enables NGINX to set, alter, or drop gRPC call headers, set timeouts for requests, and set upstream SSL/TLS specifications. As gRPC communicates over the HTTP/2 protocol, you can configure NGINX to accept gRPC and non-gRPC web traffic on the same endpoint.

8.3 HTTP/2 Server Push

Problem

You need to preemptively push content to the client.

Solution

Use the HTTP/2 server push feature of NGINX:

```
server {
    listen 443 ssl http2 default_server;

    ssl_certificate     server.crt;
    ssl_certificate_key server.key;
    root /usr/share/nginx/html;

    location = /demo.html {
        http2_push /style.css;
        http2_push /image1.jpg;
    }
}
```

Discussion

To use HTTP/2 server push, your server must be configured for HTTP/2, as is done in Recipe 7.1. From there, you can instruct NGINX to push specific files preemptively with the `http2_push` directive. This directive takes one parameter, the full URI path of the file to push to the client.

NGINX can also automatically push resources to clients if proxied applications include an HTTP response header named `Link`. This header is able to instruct NGINX to preload the resources specified. To enable this feature, add `http2_push_preload on;` to the NGINX configuration.

Sophisticated Media Streaming

9.0 Introduction

This chapter covers streaming media with NGINX in MPEG-4 (MP4) or Flash Video (FLV) formats. NGINX is widely used to distribute and stream content to the masses. NGINX supports industry-standard formats and streaming technologies, which will be covered in this chapter. NGINX Plus enables the ability to fragment content on the fly with the HTTP Live Stream (HLS) module, as well as the ability to deliver HTTP Dynamic Streaming (HDS) of already fragmented media. NGINX natively allows for bandwidth limits, and NGINX Plus's advanced features offer bitrate limiting, enabling your content to be delivered in the most efficient manner while reserving the servers' resources to reach the most users.

9.1 Serving MP4 and FLV

Problem

You need to stream digital media, originating in MP4 or FLV.

Solution

Designate an HTTP location block to serve *.mp4* or *.flv* videos. NGINX will stream the media using progressive downloads or HTTP pseudostreaming and support seeking:

```
http {
    server {
        # ...

        location /videos/ {
```

```
            mp4;
        }
        location ~ \.flv$ {
            flv;
        }
    }
}
```

The example location block tells NGINX that files in the *videos* directory are in MP4 format type and can be streamed with progressive download support. The second location block instructs NGINX that any files ending in *.flv* are in FLV format and can be streamed with HTTP pseudostreaming support.

Discussion

Streaming video or audio files in NGINX is as simple as a single directive. Progressive download enables the client to initiate playback of the media before the file has finished downloading. NGINX provides support for seeking to an undownloaded portion of the media in both formats.

9.2 Streaming with HLS with NGINX Plus

Problem

You need to support HTTP Live Streaming (HLS) for H.264/AAC-encoded content packaged in MP4 files.

Solution

Utilize NGINX Plus's HLS module with real-time segmentation, packetization, and multiplexing, with control over fragmentation buffering and more, like forwarding HLS arguments:

```
location /hls/ {
    hls;  # Use the HLS handler to manage requests

    # Serve content from the following location
    alias /var/www/video;

    # HLS parameters
    hls_fragment            4s;
    hls_buffers         10 10m;
    hls_mp4_buffer_size     1m;
    hls_mp4_max_buffer_size 5m;
}
```

This location block directs NGINX to stream HLS media out of the */var/www/video* directory, fragmenting the media into 4-second segments. The number of HLS buffers is set to 10 with a size of 10 MB. The initial MP4 buffer size is set to 1 MB with a maximum of 5 MB.

Discussion

The HLS module available in NGINX Plus provides the ability to transmultiplex MP4 media files on the fly. There are many directives that give you control over how your media is fragmented and buffered. The location block must be configured to serve the media as an HLS stream with the HLS handler. The HLS fragmentation is set in number of seconds, instructing NGINX to fragment the media by time length. The amount of buffered data is set with the `hls_buffers` directive specifying the number of buffers and the size. The client is allowed to start playback of the media after a certain amount of buffering has accrued, specified by the `hls_mp4_buffer_size`. However, a larger buffer may be necessary because metadata about the video may exceed the initial buffer size. This amount is capped by the `hls_mp4_max_buffer_size`. These buffering variables allow NGINX to optimize the end-user experience. Choosing the right values for these directives requires knowing the target audience and your media. For instance, if the bulk of your media is large video files, and your target audience has high bandwidth, you may opt for a larger max buffer size and longer length fragmentation. This will allow for the metadata about the content to be downloaded initially without error and for your users to receive larger fragments.

9.3 Streaming with HDS with NGINX Plus

Problem

You need to support Adobe's HTTP Dynamic Streaming (HDS) files that have already been fragmented and separated from the metadata.

Solution

Use NGINX Plus's support for fragmented FLV files via the `F4F` module to offer Adobe Adaptive Streaming to your users:

```
location /video/ {
    alias /var/www/transformed_video;
    f4f;
    f4f_buffer_size 512k;
}
```

The example instructs NGINX Plus to serve previously fragmented media from a location on disk to the client using the NGINX Plus `F4F` module. The buffer size for the index file (*.f4x*) is set to 512 KB.

Discussion

The NGINX Plus `F4F` module enables NGINX to serve previously fragmented media to end users. The configuration of such is as simple as using the `f4f` handler inside of an HTTP location block. The `f4f_buffer_size` directive configures the buffer size for the index file of this type of media.

9.4 Bandwidth Limits with NGINX Plus

Problem

You need to limit bandwidth to downstream media streaming clients without affecting the viewing experience.

Solution

Utilize NGINX Plus's bitrate-limiting support for MP4 media files:

```
location /video/ {
    mp4;
    mp4_limit_rate_after 15s;
    mp4_limit_rate       1.2;
}
```

This configuration allows the downstream client to download for 15 seconds before applying a bitrate limit. After 15 seconds, the client is allowed to download media at a rate of 120% of the bitrate, which enables the client to always download faster than they play.

Discussion

NGINX Plus's bitrate limiting allows your streaming server to limit bandwidth dynamically, based on the media being served, allowing clients to download just as much as they need to ensure a seamless user experience. The MP4 handler described in Recipe 9.1 designates this location block to stream MP4 media formats. The rate-limiting directives, such as `mp4_limit_rate_after`, tell NGINX to only rate-limit traffic after a specified amount of time, in seconds. The other directive involved in MP4 rate limiting is `mp4_limit_rate`, which specifies the bitrate at which clients are allowed to download in relation to the bitrate of the media. A value of 1 provided to the `mp4_limit_rate` directive specifies that NGINX is to limit bandwidth (1-to-1) to the bitrate of the media. Providing a value of more than 1 to the `mp4_limit_rate` directive will allow users to download faster than they watch so they can buffer the media and watch seamlessly while they download.

Cloud Deployments

10.0 Introduction

The advent of cloud providers has changed the landscape of web application hosting. A process such as provisioning a new machine used to take anywhere from hours to months; now, you can create one with as little as a click or an API call. These cloud providers lease their virtual machines, called infrastructure as a service (IaaS), or managed software solutions such as databases, through a pay-per-usage model, which means you pay only for what you use. This enables engineers to build up entire environments for testing and tear them down when they're no longer needed. These cloud providers also enable applications to scale horizontally based on performance need at a moment's notice. This chapter covers basic NGINX and NGINX Plus deployments on a couple of the major cloud-provider platforms.

10.1 Auto-Provisioning on AWS

Problem

You need to automate the configuration of NGINX servers on Amazon Web Services (AWS) for machines to be able to automatically provision themselves.

Solution

Utilize EC2 `UserData` as well as a prebaked Amazon Machine Image (AMI). Create an Amazon Machine Image with NGINX and any supporting software packages installed. Utilize Amazon Elastic Compute Cloud (EC2) `UserData` to configure any environment-specific configurations at runtime.

Discussion

There are three patterns of thought when provisioning on AWS:

Provision at boot

Start from a common Linux image, then run configuration management or shell scripts at boot time to configure the server. This pattern is slow to start and can be prone to errors.

Fully baked AMIs

Fully configure the server, then burn an AMI to use. This pattern boots very fast and accurately. However, it's less flexible to the environment around it, and maintaining many images can be complex.

Partially baked AMIs

It's a mix of both worlds. Partially baked is where software requirements are installed and burned into an AMI, and environment configuration is done at boot time. This pattern is flexible compared to a fully baked pattern, and fast compared to a provision-at-boot solution.

Whether you choose to partially or fully bake your AMIs, you'll want to automate that process. To construct an AMI build pipeline, it's suggested to use a couple of tools:

Configuration management

Configuration management tools define the state of the server in code, such as what version of NGINX is to be run and what user it's to run as, what DNS resolver to use, and who to proxy upstream to. This configuration management code can be source controlled and versioned like a software project. Some popular configuration management tools are Chef and Ansible, which were described in Chapter 5.

Packer from HashiCorp

Packer is used to automate running your configuration management on virtually any virtualization or cloud platform and to burn a machine image if the run is successful. Packer basically builds a virtual machine (VM) on the platform of your choosing, then will SSH into the VM, run any provisioning you specify, and burn an image. You can utilize Packer to run the configuration management tool and reliably burn a machine image to your specification.

To provision environmental configurations at boot time, you can utilize the Amazon EC2 `UserData` to run commands the first time the instance is booted. If you're using the partially baked method, you can utilize this to configure environment-based items at boot time. Examples of environment-based configurations might be what server names to listen for, resolver to use, domain name to proxy to, or upstream server pool to start with. `UserData` is a base64-encoded string that is downloaded at the first boot and run. `UserData` can be as simple as an environment file accessed by other bootstrapping processes in your AMI, or it can be a script written in any

language that exists on the AMI. It's common for `UserData` to be a bash script that specifies variables, or downloads variables, to pass to configuration management. Configuration management ensures the system is configured correctly, templates configuration files based on environment variables, and reloads services. After `UserData` runs, your NGINX machine should be completely configured in a very reliable way.

10.2 Routing to NGINX Nodes Without an AWS ELB

Problem

You need to route traffic to multiple active NGINX nodes or create an active-passive failover set to achieve high availability without a load balancer in front of NGINX.

Solution

Use the Amazon Route 53 DNS service to route to multiple active NGINX nodes or configure health checks and failover between an active-passive set of NGINX nodes.

Discussion

DNS has balanced load between servers for a long time; moving to the cloud doesn't change that. The Route 53 service from Amazon provides a DNS service with many advanced features, all available through an API. All the typical DNS tricks are available, such as multiple IP addresses on a single A record and weighted A records. When running multiple active NGINX nodes, you'll want to use one of these A-record features to spread load across all nodes. The round-robin algorithm is used when multiple IP addresses are listed for a single A record. A weighted distribution can be used to distribute load unevenly by defining weights for each server IP address in an A record.

One of the more interesting features of Route 53 is its ability to health check. You can configure Route 53 to monitor the health of an endpoint by establishing a TCP connection or by making a request with HTTP or HTTPS. The health check is highly configurable with options for the IP, hostname, port, URI path, interval rates, monitoring, and geography. With these health checks, Route 53 can take an IP out of rotation if it begins to fail. You could also configure Route 53 to failover to a secondary record in case of a failure, which would achieve an active-passive, highly available setup.

Route 53 has a geological-based routing feature that will enable you to route your clients to the closest NGINX node to them, for the least latency. When routing by geography, your client is directed to the closest healthy physical location. When

running multiple sets of infrastructure in an active-active configuration, you can automatically failover to another geological location through the use of health checks.

When using Route 53 DNS to route your traffic to NGINX nodes in an Auto Scaling group, you'll want to automate the creation and removal of DNS records. To automate adding and removing NGINX machines to Route 53 as your NGINX nodes scale, you can use Amazon's Auto Scaling lifecycle hooks to trigger scripts within the NGINX box itself or scripts running independently on Amazon Lambda. These scripts would use the Amazon Command Line Interface (CLI) or software development kit (SDK) to interface with the Amazon Route 53 API to add or remove the NGINX machine IP and configured health check as it boots or before it is terminated.

Also See

"Global Server Load Balancing with Amazon Route 53 and NGINX Plus"

10.3 The NLB Sandwich

Problem

You need to autoscale your NGINX Open Source layer and distribute load evenly and easily between application servers.

Solution

Create a network load balancer (NLB). During creation of the NLB through the console, you are prompted to create a new target group. If you do not do this through the console, you will need to create this resource and attach it to a listener on the NLB. You create an Auto Scaling group with a launch configuration that provisions an EC2 instance with NGINX Open Source installed. The Auto Scaling group has a configuration to link to the target group, which automatically registers any instance in the Auto Scaling group to the target group configured on first boot. The target group is referenced by a listener on the NLB. Place your upstream applications behind another network load balancer and target group and then configure NGINX to proxy to the application NLB.

Discussion

This common pattern is called the NLB sandwich (see Figure 10-1), putting NGINX Open Source in an Auto Scaling group behind an NLB and the application Auto Scaling group behind another NLB. The reason for having NLBs between every layer is because the NLB works so well with Auto Scaling groups; they automatically

register new nodes and remove those being terminated as well as run health checks and pass traffic to only healthy nodes.
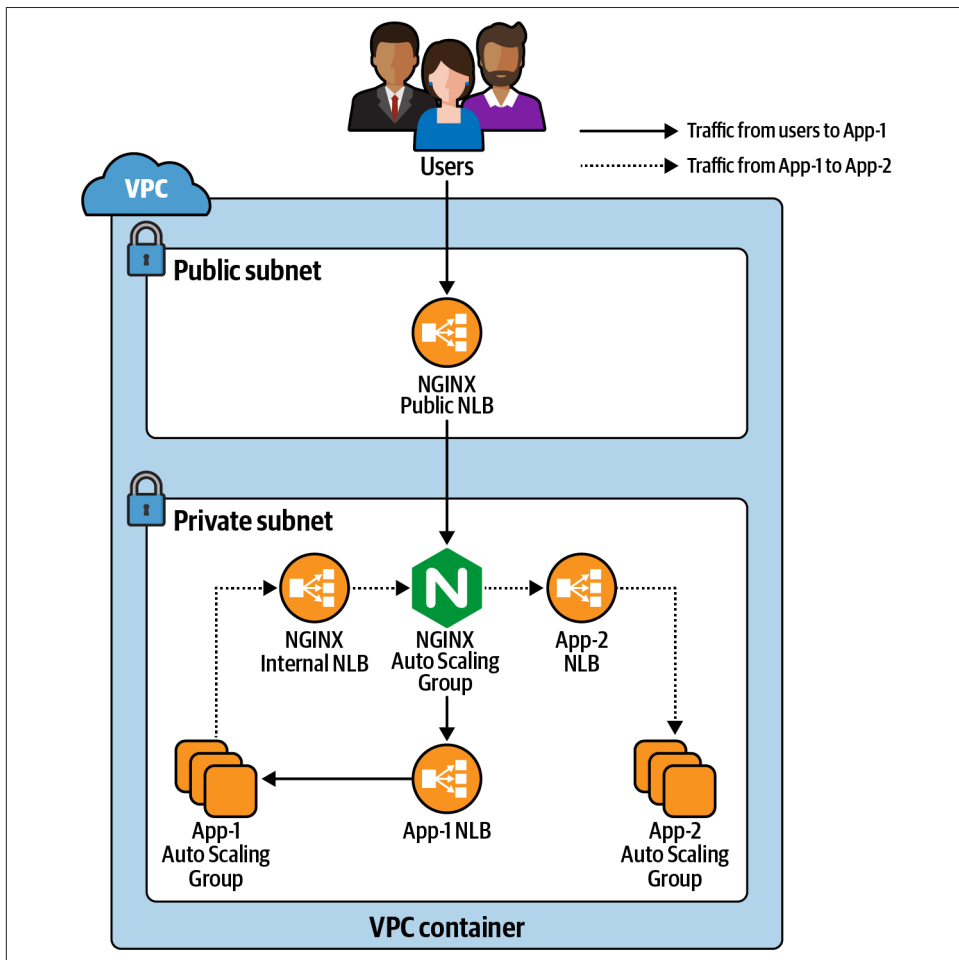
Figure 10-1. This image depicts NGINX in an NLB sandwich pattern with an internal NLB for internal applications to utilize. A user makes a request to App-1, and App-1 makes a request to App-2 through NGINX to fulfill the user's request.

It might be necessary to build a second, internal NLB for the NGINX Open Source layer because it allows services within your application to call out to other services through the NGINX Auto Scaling group without leaving the network and re-entering through the public NLB. This puts NGINX in the middle of all network traffic within your application, making it the heart of your application's traffic routing. This pattern used to be called the *elastic load balancer sandwich*; however, the NLB is preferred when working with NGINX because the NLB is a Layer 4 load balancer, whereas

ELBs and ALBs are Layer 7 load balancers. Layer 7 load balancers transform the request via the PROXY Protocol and are redundant with the use of NGINX. This pattern is needed only for NGINX Open Source because the feature set provided by the NLB is available in NGINX Plus.

You can use this pattern within other cloud providers as well. The concept is the same if you're using Azure load balancers and scale sets, or GCP load balancers and Auto Scaling groups. The core value of this pattern is to use cloud native services to provide automatic registration and load balancing of scaling application servers, and NGINX Open Source for proxy logic.

10.4 Deploying from the AWS Marketplace

Problem

You need to run NGINX Plus in AWS with ease, with a pay-as-you-go license.

Solution

Deploy through the AWS Marketplace. Visit the AWS Marketplace and search "NGINX Plus" (see Figure 10-2). Select the AMI that is based on the Linux distribution of your choice; review the details, terms, and pricing; then click the Continue link. On the next page you'll be able to accept the terms and deploy NGINX Plus with a single click, or accept the terms and use the AMI.

Discussion

The AWS Marketplace solution to deploying NGINX Plus provides ease of use and a pay-as-you-go license. Not only do you have nothing to install, but you also have a license without jumping through hoops like getting a purchase order for a year license. This solution enables you to try NGINX Plus without commitment. You can also use the NGINX Plus Marketplace AMI as overflow capacity. It's a common practice to purchase your expected workload worth of licenses and use the Marketplace AMI in an Auto Scaling group as overflow capacity. This strategy ensures you only pay for as much licensing as you use.
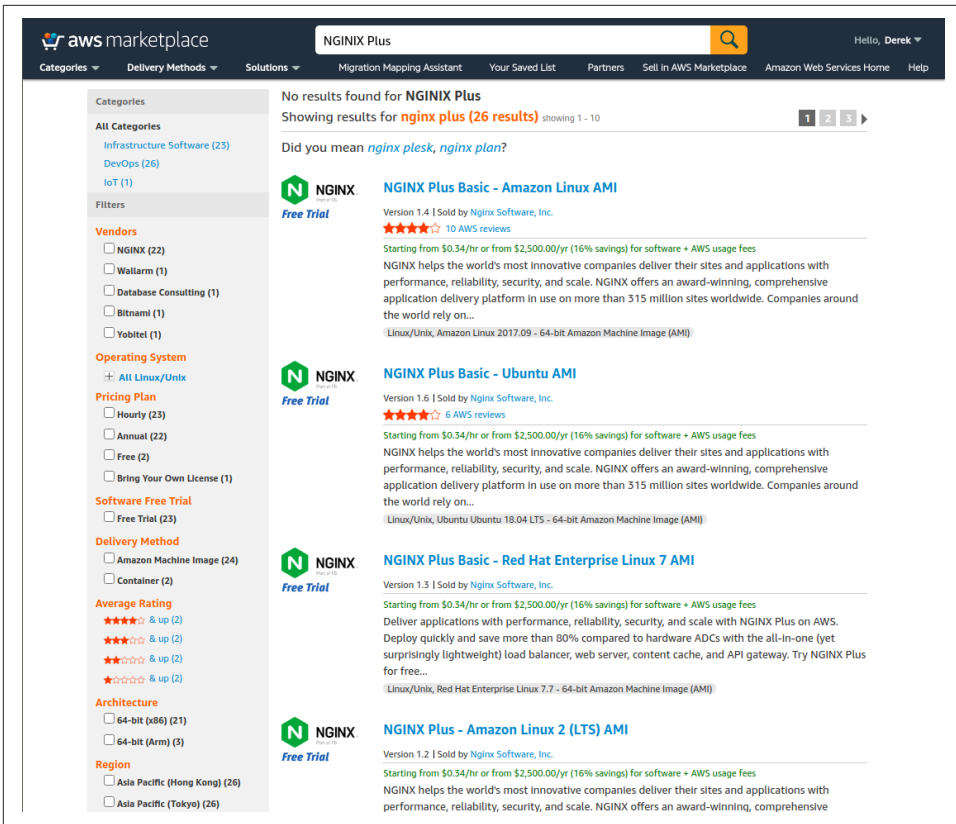
Figure 10-2. Searching for NGINX Plus on the AWS Marketplace

10.5 Creating an NGINX Virtual Machine Image on Azure

Problem

You need to create a virtual machine (VM) image of your own NGINX server configured as you see fit to quickly create more servers or to use in scale sets.

Solution

Create a VM from a Linux-based operating system of your choice. Once the VM is booted, log in and install NGINX or NGINX Plus in your preferred way, either from source or through the package management tool for the distribution you're running. Configure NGINX as desired and create a new VM image. To create a VM image, you must first generalize the VM. To generalize your VM, you need to remove the user that Azure provisioned, connect to it over SSH, and run the following command:

```
$ sudo waagent -deprovision+user -force
```

This command deprovisions the user that Azure provisioned when creating the VM. The `-force` option simply skips a confirmation step. After you've installed NGINX or NGINX Plus and removed the provisioned user, you can exit your session.

Connect your Azure CLI to your Azure account using the Azure login command, then ensure you're using the Azure Resource Manager mode. Now deallocate your VM:

```
$ azure vm deallocate -g <ResourceGroupName> \
                      -n <VirtualMachineName>
```

Once the VM is deallocated, you will be able to generalize it with the `azure vm generalize` command:

```
$ azure vm generalize -g <ResourceGroupName> \
                      -n <VirtualMachineName>
```

After your VM is generalized, you can create an image. The following command will create an image and also generate an Azure Resources Manager (ARM) template for you to use to boot this image:

```
$ azure vm capture <ResourceGroupName> <VirtualMachineName> \
    <ImageNamePrefix> -t <TemplateName>.json
```

The command line will produce output saying that your image has been created, that it's saving an ARM template to the location you specified, and that the request is complete. You can use this ARM template to create another VM from the newly created image. However, to use this template Azure has created, you must first create a new network interface:

```
$ azure network nic create <ResourceGroupName> \
    <NetworkInterfaceName> \
    <Region> \
    --subnet-name <SubnetName> \
    --subnet-vnet-name <VirtualNetworkName>
```

This command output will detail information about the newly created network interface. The first line of the output data will be the network interface ID, which you will need to utilize the ARM template created by Azure. Once you have the ID, you can create a deployment with the ARM template:

```
$ azure group deployment create <ResourceGroupName> \
    <DeploymentName> \
    -f <TemplateName>.json
```

You will be prompted for multiple input variables such as `vmName`, `adminUserName`, `adminPassword`, and `networkInterfaceId`. Enter a name for the VM and the admin username and password. Use the network interface ID harvested from the last command as the input for the `networkInterfaceId` prompt. These variables will be passed as parameters to the ARM template and used to create a new VM from the

custom NGINX or NGINX Plus image you've created. After entering the necessary parameters, Azure will begin to create a new VM from your custom image.

Discussion

Creating a custom image in Azure enables you to create copies of your preconfigured NGINX or NGINX Plus server at will. An Azure ARM template enables you to quickly and reliably deploy this same server time and time again as needed. With the VM image path that can be found in the template, you can create different sets of infrastructure such as VM scaling sets or other VMs with different configurations.

Also See

"How to install the Azure CLI"
"Sign in with Azure CLI"
"How to create a managed image of a virtual machine or VHD"

10.6 Load Balancing Over NGINX Scale Sets on Azure

Problem

You need to scale NGINX nodes behind an Azure load balancer to achieve high availability and dynamic resource usage.

Solution

Create an Azure load balancer that is either public-facing or internal. Deploy the NGINX VM image created in the prior section, or the NGINX Plus image from the Marketplace described in Recipe 10.7, into an Azure virtual machine scale set (VMSS). Once your load balancer and VMSS are deployed, configure a backend pool on the load balancer to the VMSS. Set up load-balancing rules for the ports and protocols you'd like to accept traffic on, and direct them to the backend pool.

Discussion

It's common to scale NGINX to achieve high availability or to handle peak loads without overprovisioning resources. In Azure you achieve this with VMSS. Using the Azure load balancer provides ease of management for adding and removing NGINX nodes to the pool of resources when scaling. With Azure load balancers, you're able to check the health of your backend pools and only pass traffic to healthy nodes. You can run internal Azure load balancers in front of NGINX where you want to enable access only over an internal network. You may use NGINX to proxy to an internal load balancer fronting an application inside of a VMSS, using the load balancer for the ease of registering and deregistering from the pool.

10.7 Deploying Through the Azure Marketplace

Problem

You need to run NGINX Plus in Azure with ease and a pay-as-you-go license.

Solution

Deploy an NGINX Plus VM image through the Azure Marketplace:

1. From the Azure dashboard, select the New icon, and use the search bar to search for "NGINX." Search results will appear.

2. From the list, select the NGINX Plus Virtual Machine Image published by NGINX Inc.

3. When prompted to choose your deployment model, select the Resource Manager option, and click the Create button.

4. You will then be prompted to fill out a form to specify the name of your VM, the disk type, the default username and password or SSH key-pair public key, which subscription to bill under, the resource group you'd like to use, and the location.

5. Once this form is filled out, you can click OK. Your form will be validated.

6. When prompted, select a VM size, and click the Select button.

7. On the next panel, you have the option to select optional configurations, which will be the default based on your resource group choice made previously. After altering these options and accepting them, click OK.

8. On the next screen, review the summary. You have the option of downloading this configuration as an ARM template so that you can create these resources again more quickly via a JSON template.

9. Once you've reviewed and downloaded your template, you can click OK to move to the purchasing screen. This screen will notify you of the costs you're about to incur from this VM usage. Click Purchase and your NGINX Plus box will begin to boot.

Discussion

Azure and NGINX have made it easy to create an NGINX Plus VM in Azure through just a few configuration forms. The Azure Marketplace is a great way to get NGINX Plus on demand with a pay-as-you-go license. With this model, you can try out the features of NGINX Plus or use it for on-demand overflow capacity of your already licensed NGINX Plus servers.

10.8 Deploying to Google Compute Engine

Problem

You need to create an NGINX server in Google Compute Engine to load balance or proxy for the rest of your resources in Google Compute or App Engine.

Solution

Start a new VM in Google Compute Engine. Select a name for your VM, zone, machine type, and boot disk. Configure identity and access management, firewall, and any advanced configuration you'd like. Create the VM.

Once the VM has been created, log in via SSH or through the Google Cloud Shell. Install NGINX or NGINX Plus through the package manager for the given OS type. Configure NGINX as you see fit and reload.

Alternatively, you can install and configure NGINX through the Google Compute Engine startup script, which is an advanced configuration option when creating a VM.

Discussion

Google Compute Engine offers highly configurable VMs at a moment's notice. Starting a VM takes little effort and enables a world of possibilities. Google Compute Engine offers networking and compute in a virtualized cloud environment. With a Google Compute instance, you have the full capabilities of an NGINX server wherever and whenever you need it.

10.9 Creating a Google Compute Image

Problem

You need to create a Google Compute Image to quickly instantiate a VM or create an instance template for an instance group.

Solution

Create a VM as described in Recipe 10.8. After installing and configuring NGINX on your VM instance, set the auto-delete state of the boot disk to `false`. To set the auto-delete state of the disk, edit the VM. On the Edit page under the disk configuration is a checkbox labeled "Delete boot disk when instance is deleted." Deselect this checkbox and save the VM configuration. Once the auto-delete state of the instance is set to `false`, delete the instance. When prompted, do not select the

checkbox that offers to delete the boot disk. By performing these tasks, you will be left with an unattached boot disk with NGINX installed.

After your instance is deleted and you have an unattached boot disk, you can create a Google Compute Image. From the Image section of the Google Compute Engine console, select Create Image. You will be prompted for an image name, family, description, encryption type, and the source. The source type you need to use is disk; for the source disk, select the unattached NGINX boot disk. Select Create, and Google Compute Cloud will create an image from your disk.

Discussion

You can utilize Google Cloud Images to create VMs with a boot disk identical to the server you've just created. The value in creating images is being able to ensure that every instance of this image is identical. When installing packages at boot time in a dynamic environment, unless using version locking with private repositories, you run the risk of package version and updates not being validated before being run in a production environment. With machine images, you can validate that every package running on this machine is exactly as you tested, strengthening the reliability of your service offering.

Also See

"Create, delete, and deprecate custom images"

10.10 Creating a Google App Engine Proxy

Problem

You need to create a proxy for Google App Engine to context switch between applications or serve HTTPS under a custom domain.

Solution

Utilize NGINX in Google Compute Cloud. Create a VM in Google Compute Engine, or create a virtual machine image with NGINX installed and create an instance template with this image as your boot disk. If you've created an instance template, follow up by creating an instance group that utilizes that template.

Configure NGINX to proxy to your Google App Engine endpoint. Make sure to proxy to HTTPS because Google App Engine is public, and you'll want to ensure you do not terminate HTTPS at your NGINX instance and allow information to travel between NGINX and Google App Engine unsecured. Because App Engine provides just a single DNS endpoint, you'll be using the `proxy_pass` directive rather than

upstream blocks in the open source version of NGINX because the open source version is not able to resolve DNS names of upstream servers. When proxying to Google App Engine, make sure to set the endpoint as a variable in NGINX, then use that variable in the `proxy_pass` directive to ensure NGINX does DNS resolution on every request. For NGINX to do any DNS resolution, you'll need to also utilize the `resolver` directive and point to your favorite DNS resolver. Google makes the IP address `8.8.8.8` available for public use. If you're using NGINX Plus, you'll be able to use the `resolve` flag on the `server` directive within the upstream block, keepalive connections, and other benefits of the upstream module when proxying to Google App Engine.

You may choose to store your NGINX configuration files in Google Storage, then use the startup script for your instance to pull down the configuration at boot time. This will allow you to change your configuration without having to burn a new image. However, it will add to the startup time of your NGINX server.

Discussion

You want to run NGINX in front of Google App Engine if you're using your own domain and want to make your application available via HTTPS. At this time, Google App Engine does not allow you to upload your own SSL certificates. Therefore, if you'd like to serve your app under a domain other than appspot.com with encryption, you'll need to create a proxy with NGINX to listen at your custom domain. NGINX will encrypt communication between itself and your clients, as well as between itself and Google App Engine.

Another reason you may want to run NGINX in front of Google App Engine is to host many App Engine apps under the same domain and use NGINX to do URI-based context switching. Microservices are a popular architecture, and it's common for a proxy like NGINX to conduct the traffic routing. Google App Engine makes it easy to deploy applications, and in conjunction with NGINX, you have a full-fledged application delivery platform.

Containers/Microservices

11.0 Introduction

Containers offer a layer of abstraction at the application layer, shifting the installation of packages and dependencies from the deploy to the build process. This is important because engineers are now shipping units of code that run and deploy in a uniform way regardless of the environment. Promoting containers as runnable units reduces the risk of dependency and configuration snafus between environments. Given this, there has been a large drive for organizations to deploy their applications on container platforms. When running applications on a container platform, it's common to containerize as much of the stack as possible, including your proxy or load balancer. NGINX and NGINX Plus containerize and ship with ease. They also include many features that make delivering containerized applications fluid. This chapter focuses on building NGINX and NGINX Plus container images, features that make working in a containerized environment easier, and deploying your image on Kubernetes and OpenShift.

When containerizing, it's often common to decompose services into smaller applications. When doing so, they're tied back together by an API gateway. The first section in this chapter provides a common case scenario of using NGINX as an API gateway to secure, validate, authenticate, and route requests to the appropriate service. The last section, while more of a security topic, discusses NGINX Service Mesh, an infrastructure pattern used in containerized environments to optimize and secure connections between such aforementioned services.

A couple of architecture considerations about running NGINX or NGINX Plus in a container should be called out. When containerizing a service, to make use of the Docker log driver, logs must be output to *dev/stdout* and error logs directed

to *dev/stderr*. By doing so, the logs are streamed to the Docker log driver, which is able to route them to consolidated logging servers natively.

Load-balancing methods are also of consideration when using NGINX Plus in a containerized environment. The `least_time` load-balancing method was designed with containerized networking overlays in mind. By favoring low response time, NGINX Plus will pass the incoming request to the upstream server with the fastest average response time. When all servers are adequately load balanced and performing equally, NGINX Plus can optimize by network latency, preferring servers in closest network proximity.

11.1 Using NGINX as an API Gateway

Problem

You need an API gateway to validate, authenticate, manipulate, and route incoming requests for your use case.

Solution

Use NGINX or NGINX Plus as an API gateway. An API gateway provides an entry point to one or more application programming interfaces (APIs). NGINX fits this role very well. This section will highlight some core concepts and reference other sections within this book for more detail on specifics. It's also important to note that NGINX has published an entire ebook on this topic: *Deploying NGINX as an API Gateway* by Liam Crilly.

Start by defining a server block for your API gateway within its own file. A name such as */etc/nginx/api_gateway.conf* will do.

```
server {
    listen 443 ssl;
    server_name api.company.com;
    # SSL Settings Chapter 7

    default_type application/json;
}
```

Add some basic error-handling responses to your server definition:

```
proxy_intercept_errors on;

error_page 400 = @400;
location @400 { return 400 '{"status":400,"message":"Bad request"}\n'; }

error_page 401 = @401;
location @401 { return 401 '{"status":401,"message":"Unauthorized"}\n'; }
```

```
error_page 403 = @403;
location @403 { return 403 '{"status":403,"message":"Forbidden"}\n'; }

error_page 404 = @404;
location @404 { return 404 '{"status":404,"message":"Resource not found"}\n'; }
```

The preceding section of NGINX configuration can be added directly to the server block in *ic/nginx/api_gateway.conf* or a separate file, and imported via an `include` directive. The `include` directive is covered in Recipe 18.1.

Use an `include` directive to import this server configuration into the main *nginx.conf* file within the `http` context:

```
include /etc/nginx/api_gateway.conf;
```

You now need to define your upstream service endpoints. Chapter 2 covers load balancing, which discusses the `upstream` block. As a reminder, `upstream` is valid within the `http` context, and not within `server` context. The following must be included or set outside of the `server` block:

```
upstream service_1 {
    server 10.0.0.12:80;
    server 10.0.0.13:80;
}
upstream service_2 {
    server 10.0.0.14:80;
    server 10.0.0.15:80;
}
```

Depending on the use case, you may want to declare your services inline, as an included file, or included per services. A case also exists where services should be defined as proxy location endpoints; in this case it's suggested to define the endpoint as a variable for use throughout. Chapter 5, "Programmability and Automation", discusses ways to automate adding and removing machines from `upstream` blocks.

Build an internally routable location within the `server` block for each service:

```
location = /_service_1 {
    internal;
    # Config common to service
    proxy_pass http://service_1/$request_uri;
}
location = /_service_2 {
    internal;
    # Config common to service
    proxy_pass http://service_2/$request_uri;
}
```

By defining internal routable locations for these services, configuration that is common to the service can be defined once, rather than repeatedly.

From here, we need to build up `location` blocks that define specific URI paths for a given service. These blocks will validate and route the request appropriately. An API gateway can be as simple as routing requests based on path, and as detailed as defining specific rules for every single accepted API URI. In the latter, you'll want to devise a file structure for organization and use NGINX includes to import your configuration files. This concept is discussed in Recipe 18.1.

Create a new directory for the API gateway:

```
mkdir /etc/nginx/api_conf.d/
```

Build a specification of a service use case by defining `location` blocks within a file at a path that makes sense for your configuration structure. Use the `rewrite` directive to direct the request to the prior configured `location` block that proxies the request to a service. The `rewrite` directive used in the following example instructs NGINX to reprocess the request with an altered URI. The example defines rules specific to an API resource, restricts HTTP methods, then uses the `rewrite` directive to send the request to the prior defined internal common proxy location for the service:

```
location /api/service_1/object {
    limit_except GET PUT { deny all; }
    rewrite ^ /_service_1 last;
}
location /api/service_1/object/[^/]*$ {
    limit_except GET POST { deny all; }
    rewrite ^ /_service_1 last;
}
```

Repeat this step for each service. Employ logical separation by means of file and directory structures to organize effectively for the use case. Use any and all information provided in this book to configure API `location` blocks to be as specific and restrictive as possible.

If separate files were used for the preceding `location` or `upstream` blocks, ensure they're included in your `server` context:

```
server {
    listen 443 ssl;
    server_name api.company.com;
    # SSL Settings Chapter 7

    default_type application/json;

    include api_conf.d/*.conf;
}
```

Enable authentication to protect private resources by using one of the many methods discussed in Chapter 6, or something as simple as preshared API keys as follows (note the `map` directive is only valid in the `http` context):

```
map $http_apikey $api_client_name {
    default "";

    "j7UqLLB+yRv2VTCXXDZ1M/N4" "client_one";
    "6B2kbyrrTiIN8S8JhSAxb63R" "client_two";
    "KcVgIDSY4Nm46m3tXVY3vbgA" "client_three";
}
```

Protect backend services from attack with NGINX by employing learnings from Chapter 2 to limit usage. In the `http` context, define one or many request limit shared memory zones:

```
limit_req_zone $http_apikey
        zone=limitbyapikey:10m rate=100r/s;
limit_req_status 429;
```

Protect a given context with rate limits and authentication:

```
location /api/service_2/object {
    limit_req zone=limitbyapikey;

    # Consider writing these if's to a file
    # and using an include were needed.
    if ($http_apikey = "") {
        return 401;
    }
    if ($api_client_name = "") {
        return 403;
    }

    limit_except GET PUT { deny all; }
    rewrite ^ /_service_2 last;
}
```

Test out some calls to your API gateway:

```
curl -H "apikey: 6B2kbyrrTiIN8S8JhSAxb63R" \
  https://api.company.com/api/service_2/object
```

Discussion

API gateways provide an entry point to an application programming interface (API). That sounds vague and basic, so let's dig in. Integration points happen at many different layers. Any two independent services that need to communicate (integrate) should hold an API version contract. Such version contracts define the compatibility of the services. An API gateway enforces such contracts—authenticating, authorizing, transforming, and routing requests between services.

This section demonstrated how NGINX can function as an API gateway by validating, authenticating, and directing incoming requests to specific services and limiting

their usage. This tactic is popular in microservice architectures, where a single API offering is split among different services.

Employ all of your learnings thus far to construct an NGINX server configuration to the exact specifications for your use case. By weaving together the core concepts demonstrated in this text, you have the ability to authenticate and authorize the use of URI paths, route or rewrite requests based on any factor, limit usage, and define what is and is not accepted as a valid request. There will never be a single solution to an API gateway, as each is intimately and infinitely definable to the use case it provides.

An API gateway provides an ultimate collaboration space between operations and application teams to form a true DevOps organization. Application development defines validity parameters of a given request. Delivery of such a request is typically managed by what is considered IT (networking, infrastructure, security, and middleware teams). An API gateway acts as an interface between those two layers. The construction of an API gateway requires input from all sides. Configuration of such should be kept in some sort of source control. Many modern-day source-control repositories have the concept of code owners. This concept allows you to require specific users' approval for certain files. In this way, teams can collaborate but verify changes specific to a given department.

Something to keep in mind when working with API gateways is the URI path. In the example configuration, the entire URI path is passed to the upstream servers. This means the `service_1` example needs to have handlers at the `/api/service_1/*` path. To perform path-based routing in this way, it's best that the application doesn't have conflicting routes with another application.

If conflicting routes do apply, there are a few things you can do. Edit the code to resolve the conflicts, or add a URI prefix configuration to one or both applications to move one of them to another context. In the case of off-the-shelf software that can't be edited, you can rewrite the requests URI upstream. However, if the application returns links in the body, you'll need to use regular expressions (regex) to rewrite the body of the request before providing it to the client—this should be avoided.

Also See

Deploying NGINX as an API Gateway ebook

11.2 Using DNS SRV Records with NGINX Plus

Problem

You'd like to use your existing DNS SRV record implementation as the source for upstream servers with NGINX Plus.

Solution

Specify the service directive with a value of `http` on an upstream server to instruct NGINX to utilize the SRV record as a load-balancing pool:

```
http {
    resolver 10.0.0.2 valid=30s;

    upstream backend {
        zone backends 64k;
        server api.example.internal service=http resolve;
    }
}
```

This feature is an NGINX Plus exclusive. The configuration instructs NGINX Plus to resolve DNS from a DNS server at `10.0.0.2` and set up an upstream server pool with a single `server` directive. This `server` directive specified with the `resolve` parameter is instructed to periodically re-resolve the domain name base on the DNS record TTL, or the `valid` override parameter of the `resolver` directive. The `service=http` parameter and value tells NGINX that this is an SRV record containing a list of IPs and ports, and to load balance over them as if they were configured with the `server` directive.

Discussion

Dynamic infrastructure is becoming ever more popular with the demand and adoption of cloud-based infrastructure. Auto Scaling environments scale horizontally, increasing and decreasing the number of servers in the pool to match the demand of the load. Scaling horizontally demands a load balancer that can add and remove resources from the pool. With an SRV record, you offload the responsibility of keeping the list of servers to DNS. This type of configuration is extremely enticing for containerized environments because you may have containers running applications on variable port numbers, possibly at the same IP address. It's important to note that UDP DNS record payload is limited to about 512 bytes.

11.3 Using the Official NGINX Image

Problem

You need to get up and running quickly with the NGINX image from Docker Hub.

Solution

Use the NGINX image from Docker Hub. This image contains a default configuration. You'll need to either mount a local configuration directory or create a Dockerfile and `ADD` in your configuration to the image build to alter the configuration. Here

we mount a volume where NGINX's default configuration serves static content to demonstrate its capabilities by using a single command:

```
$ docker run --name my-nginx -p 80:80 \
    -v /path/to/content:/usr/share/nginx/html:ro -d nginx
```

The `docker` command pulls the `nginx:latest` image from Docker Hub if it's not found locally. The command then runs this NGINX image as a Docker container, mapping `localhost:80` to port `80` of the NGINX container. It also mounts the local directory */path/to/content/* as a container volume at */usr/share/nginx/html/* as read only. The default NGINX configuration will serve this directory as static content. When specifying mapping from your local machine to a container, the local machine port or directory comes first, and the container port or directory comes second.

Discussion

NGINX has made an official Docker image available via Docker Hub. This official Docker image makes it easy to get up and going very quickly in Docker with your favorite application delivery platform, NGINX. In this section, we were able to get NGINX up and running in a container with a single command! The official NGINX Docker image mainline that we used in this example is built from the Debian Jessie Docker image. However, you can choose official images based on Alpine Linux. The Dockerfile and source for these official images are available on GitHub. You can extend the official image by building your own Dockerfile and specifying the official image in the `FROM` command. You can also mount an NGINX configuration directory as a Docker volume to override the NGINX configuration without modifying the official image.

Also See

Official NGINX Docker image, NGINX
Docker repo on GitHub

11.4 Creating an NGINX Dockerfile

Problem

You need to create an NGINX Dockerfile in order to create a Docker image.

Solution

Start with a `FROM` line of your favorite distribution's Docker image. Use the `RUN` command to install NGINX. Use the `ADD` command to add your NGINX configuration files. Use the `EXPOSE` command to instruct Docker to expose given ports, or do this manually when you run the image as a container. Use `CMD` to start NGINX when the

image is instantiated as a container. You'll need to run NGINX in the foreground. To do this, you'll need to start NGINX with `-g "daemon off;"` or add `daemon off;` to your configuration. This example will use the latter with `daemon off;` in the configuration file within the main context. You will also want to alter your NGINX configuration to log to *dev/stdout* for access logs and *dev/stderr* for error logs; doing so will put your logs into the hands of the Docker daemon, which will make them more easily available, based on the log driver you've chosen to use with Docker:

```
FROM centos:7

# Install epel repo to get nginx and install nginx
RUN yum -y install epel-release && \
    yum -y install nginx

# add local configuration files into the image
ADD /nginx-conf /etc/nginx

EXPOSE 80 443

CMD ["nginx"]
```

The directory structure looks as follows:

```
.
├── Dockerfile
└── nginx-conf
    ├── conf.d
    │   └── default.conf
    ├── fastcgi.conf
    ├── fastcgi_params
    ├── koi-utf
    ├── koi-win
    ├── mime.types
    ├── nginx.conf
    ├── scgi_params
    ├── uwsgi_params
    └── win-utf
```

I chose to host the entire NGINX configuration within this Docker directory for ease of access to all of the configurations with only one line in the Dockerfile to add all my NGINX configurations.

Discussion

You will find it useful to create your own Dockerfile when you require full control over the packages installed and updates. It's common to keep your own repository of images so that you know your base image is reliable and tested by your team before running it in production.

11.5 Building an NGINX Plus Docker Image

Problem

You need to build an NGINX Plus Docker image to run NGINX Plus in a container-ized environment.

Solution

Use this Dockerfile to build an NGINX Plus Docker image. You'll need to download your NGINX Plus repository certificates, named *nginx-repo.crt* and *nginx-repo.key*, respectively, and keep them in the directory with this Dockerfile. With that, this Dockerfile will do the rest of the work of installing NGINX Plus for your use and linking NGINX access and error logs to the Docker log collector.

```
FROM debian:stretch-slim

LABEL maintainer="NGINX <docker-maint@nginx.com>"

# Download certificate and key from the customer portal
# (https://cs.nginx.com) and copy to the build context

COPY nginx-repo.crt /etc/ssl/nginx/
COPY nginx-repo.key /etc/ssl/nginx/

# Install NGINX Plus
RUN set -x \
  && APT_PKG="Acquire::https::plus-pkgs.nginx.com::" \
  && REPO_URL="https://plus-pkgs.nginx.com/debian" \
  && apt-get update && apt-get upgrade -y \
  && apt-get install \
    --no-install-recommends --no-install-suggests\
    -y apt-transport-https ca-certificates gnupg1 \
  && \
  NGINX_GPGKEY=573BFD6B3D8FBC641079A6ABABF5BD827BD9BF62;\
  found=''; \
  for server in \
    ha.pool.sks-keyservers.net \
    hkp://keyserver.ubuntu.com:80 \
    hkp://p80.pool.sks-keyservers.net:80 \
    pgp.mit.edu \
  ; do \
   echo "Fetching GPG key $NGINX_GPGKEY from $server"; \
   apt-key adv --keyserver "$server" --keyserver-options \
     timeout=10 --recv-keys "$NGINX_GPGKEY" \
    && found=yes \
    && break;\
  done;\
  test -z "$found" && echo >&2 \
    "error: failed to fetch GPG key $NGINX_GPGKEY" && exit 1; \
```

```
        echo "${APT_PKG}Verify-Peer "true";"\
          >> /etc/apt/apt.conf.d/90nginx \
      && echo \
          "${APT_PKG}Verify-Host "true";">>\
          /etc/apt/apt.conf.d/90nginx \
      && echo "${APT_PKG}SslCert      \
          "/etc/ssl/nginx/nginx-repo.crt";" >> \
          /etc/apt/apt.conf.d/90nginx \
      && echo "${APT_PKG}SslKey       \
          "/etc/ssl/nginx/nginx-repo.key";" >> \
          /etc/apt/apt.conf.d/90nginx \
      && printf \
          "deb ${REPO_URL} stretch nginx-plus" \
          > /etc/apt/sources.list.d/nginx-plus.list \
      && apt-get update && apt-get install -y nginx-plus \
      && apt-get remove --purge --auto-remove -y gnupg1 \
      && rm -rf /var/lib/apt/lists/*

      # Forward request logs to Docker log collector
      RUN ln -sf /dev/stdout /var/log/nginx/access.log \
          && ln -sf /dev/stderr /var/log/nginx/error.log

      EXPOSE 80
      STOPSIGNAL SIGTERM

      CMD ["nginx", "-g", "daemon off;"]
```

To build this Dockerfile into a Docker image, run the following in the directory that
contains the Dockerfile and your NGINX Plus repository certificate and key:

```
$ docker build --no-cache -t nginxplus .
```

This `docker build` command uses the flag `--no-cache` to ensure that whenever
you build this, the NGINX Plus packages are pulled fresh from the NGINX Plus
repository for updates. If it's acceptable to use the same version on NGINX Plus as
the prior build, you can omit the `--no-cache` flag. In this example, the new Docker
image is tagged `nginxplus`.

Discussion

By creating your own Docker image for NGINX Plus, you can configure your
NGINX Plus container however you see fit and drop it into any Docker environment.
This opens up all of the power and advanced features of NGINX Plus to your
containerized environment. This Dockerfile does not use the Dockerfile property `ADD`
to add in your configuration; you will need to add in your configuration manually.

Also See

"Deploying NGINX and NGINX Plus with Docker"

11.6 Using Environment Variables in NGINX

Problem

You need to use environment variables inside your NGINX configuration in order to use the same container image for different environments.

Solution

Use the `ngx_http_perl_module` to set variables in NGINX from your environment:

```
daemon off;
env APP_DNS;
include /usr/share/nginx/modules/*.conf;
# ...
http {
  perl_set $upstream_app 'sub { return $ENV{"APP_DNS"}; }';
  server {
    # ...
    location / {
      proxy_pass https://$upstream_app;
    }
  }
}
```

To use `perl_set` you must have the `ngx_http_perl_module` installed; you can do so by loading the module dynamically or statically if building from source. NGINX by default wipes environment variables from its environment; you need to declare any variables you do not want removed with the `env` directive. The `perl_set` directive takes two parameters: the variable name you'd like to set and a Perl string that renders the result.

The following is a Dockerfile that loads the `ngx_http_perl_module` dynamically, installing this module from the package management utility. When installing modules from the package utility for CentOS, they're placed in the */usr/lib64/nginx/modules/* directory, and configuration files that dynamically load these modules are placed in the */usr/share/nginx/modules/* directory. This is why in the preceding configuration snippet we include all configuration files at that path:

```
FROM centos:7

# Install epel repo to get nginx and install nginx
RUN yum -y install epel-release && \
    yum -y install nginx nginx-mod-http-perl

# add local configuration files into the image
ADD /nginx-conf /etc/nginx
```

```
EXPOSE 80 443

CMD ["nginx"]
```

Discussion

A typical practice when using Docker is to utilize environment variables to change the way the container operates. You can use environment variables in your NGINX configuration so that your NGINX Dockerfile can be used in multiple, diverse environments.

11.7 Kubernetes Ingress Controller

Problem

You are deploying your application on Kubernetes and need an ingress controller.

Solution

Ensure that you have access to the ingress controller image. For NGINX, you can use the *nginx/nginx-ingress* image from Docker Hub. For NGINX Plus, you will need to build your own image and host it in your private Docker registry. You can find instructions on building and pushing your own NGINX Plus Kubernetes Ingress Controller in the NGINX documentation.

Visit the Kubernetes Ingress Controller Deployments folder in the kubernetes-ingress repository on GitHub. The commands that follow will be run from within this directory of a local copy of the repository.

Create a namespace and a service account for the ingress controller; both are named `nginx-ingress`:

```
$ kubectl apply -f common/ns-and-sa.yaml
```

Create a secret with a TLS certificate and key for the ingress controller:

```
$ kubectl apply -f common/default-server-secret.yaml
```

This certificate and key are self-signed and created by NGINX Inc. for testing and example purposes. It's recommended to use your own because this key is publicly available.

Optionally, you can create a config map for customizing NGINX configuration (the config map provided is blank; however, you can read more about customization of ConfigMaps and Annotations):

```
$ kubectl apply -f common/nginx-config.yaml
```

If Role-Based Access Control (RBAC) is enabled in your cluster, create a cluster role and bind it to the service account. You must be a cluster administrator to perform this step:

```
$ kubectl apply -f rbac/rbac.yaml
```

Now deploy the ingress controller. Two example deployments are made available in this repository: a Deployment and a DaemonSet. Use a Deployment if you plan to dynamically change the number of ingress controller replicas. Use a DaemonSet to deploy an ingress controller on every node or a subset of nodes.

If you plan to use the NGINX Plus Deployment manifests, you must alter the YAML file and specify your own registry and image.

For NGINX Deployment:

```
$ kubectl apply -f deployment/nginx-ingress.yaml
```

For NGINX Plus Deployment:

```
$ kubectl apply -f deployment/nginx-plus-ingress.yaml
```

For NGINX DaemonSet:

```
$ kubectl apply -f daemon-set/nginx-ingress.yaml
```

For NGINX Plus DaemonSet:

```
$ kubectl apply -f daemon-set/nginx-plus-ingress.yaml
```

Validate that the ingress controller is running:

```
$ kubectl get pods --namespace=nginx-ingress
```

If you created a DaemonSet, ports 80 and 443 of the ingress controller are mapped to the same ports on the node where the container is running. To access the ingress controller, use those ports and the IP address of any of the nodes on which the ingress controller is running. If you deployed a Deployment, continue with the next steps.

For the Deployment methods, there are two options for accessing the ingress controller pods. You can instruct Kubernetes to randomly assign a node port that maps to the ingress controller pod. This is a service with the type `NodePort`. The other option is to create a service with the type `LoadBalancer`. When creating a service of type `LoadBalancer`, Kubernetes builds a load balancer for the given cloud platform, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Compute.

To create a service of type `NodePort`, use the following:

```
$ kubectl create -f service/nodeport.yaml
```

To statically configure the port that is opened for the pod, alter the YAML and add the attribute `nodePort: {port}` to the configuration of each port being opened.

To create a service of type `LoadBalancer` for Google Cloud Compute or Azure, use this code:

```
$ kubectl create -f service/loadbalancer.yaml
```

To create a service of type `LoadBalancer` for Amazon Web Services:

```
$ kubectl create -f service/loadbalancer-aws-elb.yaml
```

On AWS, Kubernetes creates a classic ELB in TCP mode with the PROXY Protocol enabled. You must configure NGINX to use the PROXY Protocol. To do so, you can add the following to the config map mentioned previously in reference to the file *common/nginx-config.yaml*:

```
proxy-protocol: "True"
real-ip-header: "proxy_protocol"
set-real-ip-from: "0.0.0.0/0"
```

Then, update the config map:

```
$ kubectl apply -f common/nginx-config.yaml
```

You can now address the pod by its `NodePort` or by making a request to the load balancer created on its behalf.

Discussion

As of this writing, Kubernetes is the leading platform in container orchestration and management. The ingress controller is the edge pod that routes traffic to the rest of your application. NGINX fits this role perfectly and makes it simple to configure with its annotations. The NGINX Ingress project offers an NGINX Open Source ingress controller out of the box from a Docker Hub image, and NGINX Plus through a few steps to add your repository certificate and key. Enabling your Kubernetes cluster with an NGINX Ingress controller provides all the same features of NGINX but with the added features of Kubernetes networking and DNS to route traffic.

11.8 Prometheus Exporter Module

Problem

You are deploying NGINX into an environment using Prometheus monitoring and need NGINX statistics.

Solution

Use the NGINX Prometheus Exporter to harvest NGINX or NGINX Plus statistics and ship them to Prometheus.

The NGINX Prometheus Exporter Module is written in GoLang and distributed as a binary on GitHub and can be found as a prebuilt Docker Image on Docker Hub.

By default, the exporter will be started for NGINX and will only harvest the `stub_status` information. To run the exporter for NGINX Open Source, ensure stub status is enabled (if it's not, there is more information on how to do so in Recipe 13.1). Then use the following Docker command:

```
docker run -p 9113:9113 nginx/nginx-prometheus-exporter:0.8.0 \
  -nginx.scrape-uri http://{nginxEndpoint}:8080/stub_status
```

To use the exporter with NGINX Plus, a flag must be used to switch the exporter's context because much more data can be collected from the NGINX Plus API. You can learn how to turn on the NGINX Plus API in Recipe 13.2. Use the following Docker command to run the exporter for an NGINX Plus environment:

```
docker run -p 9113:9113 nginx/nginx-prometheus-exporter:0.8.0 \
  -nginx.plus -nginx.scrape-uri http://{nginxPlusEndpoint}:8080/api
```

Discussion

Prometheus is an extremely common metric monitoring solution that is very prevalent in the Kubernetes ecosystem. The NGINX Prometheus Exporter Module is a fairly simple component; however, it enables prebuilt integration between NGINX and common monitoring platforms. With NGINX, the stub status does not provide a vast amount of data, but important data to provide insight into the amount of work an NGINX node is handling. The NGINX Plus API enables many more statistics about the NGINX Plus server, all of which the exporter ships to Prometheus. With either case, the information gleaned is valuable monitoring data, and the work to ship this data to Prometheus is already done; you just need to wire it up and take advantage of the insight provided by NGINX statistics.

Also See

NGINX Prometheus Exporter GitHub
NGINX `stub_status` module documentation
NGINX Plus API module documentation
NGINX Plus Monitoring Dashboard

11.9 NGINX Service Mesh mTLS

Problem

You would like to enable mutual TLS (mTLS) authentication between services using NGINX Secure Service Mesh.

Solution

Use the `nginx-meshcli` tool to enable mTLS in permissive mode while validating:

```
nginx-meshctl deploy ... --mtls-mode permissive
```

Configure proper Public Key Infrastructure (PKI) by creating a YAML configuration file specifying a root certificate authority from disk. SPIRE will use this configuration for its upstream authorities:

```yaml
apiVersion: v1
upstreamAuthority: disk
config:
  cert_file_path: /path/to/rootCA.crt
  key_file_path: /path/to/rootCA.key
```

If you're deploying on AWS, use the AWS Secrets Manager plug-in to pull the certificate authority certificate and key from AWS Secrets Manager:

```yaml
apiVersion: "v1"
upstreamAuthority: "awssecret"
config:
  region: "us-east-1"
  cert_file_arn: "arn:aws:secretsmanager:us-east-1:123456789012:secret:
    /certificate-authority/test-certificate"
  key_file_arn: "arn:aws:secretsmanager:us-east-1:123456789012:secret:
    /certificate-authority/test-key"
```

Deploy the mTLS configuration to inform SPIRE of the upstream authorities using `nginx-meshcli`:

```
nginx-meshctl deploy ... --mtls-upstream-ca-conf /path/to/upstream_authority.yaml
```

Verify your deployment:

```
kubectl get pods -n nginx-mesh
NAME               READY   STATUS    RESTARTS   AGE
...
spire-agent-sv2tv  1/1     Running   0          2h
spire-server-0     2/2     Running   0          2h
...
```

Once you've validated your mTLS configuration, and that your applications are working as expected, it's recommended to enable the strict mTLS mode:

```
nginx-meshctl deploy ... --mtls-mode strict
```

Discussion

A service mesh is a dedicated infrastructure layer designed to handle communication between distributed applications. The NGINX Service Mesh utilizes a service-level sidecar pattern in Kubernetes. The NGINX sidecar handles interservice communication, monitoring, and security. NGINX Service Mesh and the `nginx-meshctl` tool

make deploying a service mesh for your environment extremely straightforward. Sidecar injection is used to automatically configure an NGINX sidecar for each service, allowing you to enable this functionality and continue focusing on your core service offering rather than the extensive configuration required to implement such an infrastructure layer from scratch.

This section focused specifically on enabling mTLS. mTLS provides authentication between client and server—the client verifies the server's identity and the server verifies the client's identity through a connection that is encrypted by both parties. In our service mesh, any given service could also be the client of another service, as the services may communicate between each other.

This section instructed you to enable NGINX Service Mesh in permissive mode, which will allow for communication to fall back to regular communication if misconfigured. At this stage, SPIRE, the service responsible for signing and distributing the certificates and keys, defaulted to generating a self-signed certificate. The configuration was then strengthened by providing a specific certificate authority (CA) certificate and key pair. This section showed two ways of providing the CA key pair: one directly from disk, then another through AWS Secrets Manager. If you happen to be using AWS Certificate Manager for a private CA, you could also instruct SPIRE to use AWS ACM to sign service certificates. Once configured, you validated your configuration, then enabled mTLS in strict mode, which requires all interservice communication to be authenticated through mTLS.

Also See

"The mTLS Architecture in NGINX Service Mesh"
"Secure Mesh Traffic using mTLS"
The Enterprise Path to Service Mesh Architectures (O'Reilly)

High-Availability Deployment Modes

12.0 Introduction

Fault-tolerant architecture separates systems into identical, independent stacks. Load balancers like NGINX are employed to distribute load, ensuring that what's provisioned is utilized. The core concepts of high availability are load balancing over multiple active nodes or an active-passive failover. Highly available applications have no single points of failure; every component must use one of these concepts, including the load balancers themselves. For us, that means NGINX. NGINX is designed to work in either configuration: multiple active or active-passive failover. This chapter details techniques on how to run multiple NGINX servers to ensure high availability in your load-balancing tier.

12.1 NGINX Plus HA Mode

Problem

You need a highly available (HA) load-balancing solution.

Solution

Use NGINX Plus's HA mode with keepalived by installing the `nginx-ha-keepalived` package from the NGINX Plus repository.

Discussion

The `nginx-ha-keepalived` package is based on keepalived and manages a virtual IP address exposed to the client. Another process is run on the NGINX server that ensures that NGINX Plus and the keepalived process are running. Keepalived

is a process that utilizes the Virtual Router Redundancy Protocol (VRRP), sending small messages, often referred to as heartbeats, to the backup server. If the backup server does not receive the heartbeat for three consecutive periods, the backup server initiates the failover, moving the virtual IP address to itself and becoming the primary. The failover capabilities of `nginx-ha-keepalived` can be configured to identify custom failure situations.

12.2 Load-Balancing Load Balancers with DNS

Problem

You need to distribute load between two or more NGINX servers.

Solution

Use DNS to round robin across NGINX servers by adding multiple IP addresses to a DNS A record.

Discussion

When running multiple load balancers, you can distribute load via DNS. The A record allows for multiple IP addresses to be listed under a single FQDN. DNS will automatically round robin across all the IPs listed. DNS also offers weighted round robin with weighted records, which works in the same way as weighted round robin in NGINX as described in Chapter 2. These techniques work great. However, a pitfall can be removing the record when an NGINX server encounters a failure. There are DNS providers—Amazon Route 53 for one, and Dyn DNS for another—that offer health checks and failover with their DNS offering, which alleviates these issues. If you are using DNS to load balance over NGINX, when an NGINX server is marked for removal, it's best to follow the same protocols that NGINX does when removing an upstream server. First, stop sending new connections to it by removing its IP from the DNS record, then allow connections to drain before stopping or shutting down the service.

12.3 Load Balancing on EC2

Problem

You're using NGINX on AWS, and the NGINX Plus HA does not support Amazon IPs.

Solution

Put NGINX behind an AWS NLB by configuring an Auto Scaling group of NGINX servers and linking the Auto Scaling group to a target group, and then attach the target group to the NLB. Alternatively, you can place NGINX servers into the target group manually by using the AWS console, command-line interface, or API.

Discussion

The HA solution from NGINX Plus based on keepalived will not work on AWS because it does not support the floating virtual IP address, since EC2 IP addresses work in a different way. This does not mean that NGINX can't be HA in the AWS cloud; in fact, the opposite is true. The AWS NLB is a product offering from Amazon that will natively load balance over multiple, physically separated data centers called *availability zones*, provide active health checks, and provide a DNS CNAME endpoint. A common solution for HA NGINX on AWS is to put an NGINX layer behind the NLB. NGINX servers can be automatically added to and removed from the target group as needed. The NLB is not a replacement for NGINX; there are many things NGINX offers that the NLB does not, such as multiple load-balancing methods, rate limiting, caching, and Layer 7 routing. The AWS ALB does perform Layer 7 load balancing based on the URI path and host header, but it does not by itself offer features that NGINX does, such as WAF caching, bandwidth limiting, HTTP/2 server push, and more. In the event that the NLB does not fit your need, there are many other options. One option is the DNS solution: Route 53 from AWS offers health checks and DNS failover.

12.4 NGINX Plus Configuration Synchronization

Problem

You're running an HA NGINX Plus tier and need to synchronize configuration across servers.

Solution

Use the NGINX Plus exclusive configuration synchronization feature. To configure this feature, follow these steps.

Install the `nginx-sync` package from the NGINX Plus package repository.

For RHEL or CentOS:

```
$ sudo yum install nginx-sync
```

For Ubuntu or Debian:

```
$ sudo apt-get install nginx-sync
```

Grant the primary machine SSH access as root to the peer machines.

Generate an SSH authentication key pair for root, and retrieve the public key:

```
$ sudo ssh-keygen -t rsa -b 2048
$ sudo cat /root/.ssh/id_rsa.pub
ssh-rsa AAAAB3Nz4rFgt...vgaD root@node1
```

Get the IP address of the primary node:

```
$ ip addr
1: lo:  mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
      valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
      valid_lft forever preferred_lft forever
2: eth0:  mtu 1500 qdisc pfifo_fast state UP group default qlen \
      1000
    link/ether 52:54:00:34:6c:35 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.2/24 brd 192.168.1.255 scope global eth0
      valid_lft forever preferred_lft forever
    inet6 fe80::5054:ff:fe34:6c35/64 scope link
      valid_lft forever preferred_lft forever
```

The `ip addr` command will dump information about interfaces on the machine. Disregard the loopback interface, which is normally the first. Look for the IP address following `inet` for the primary interface. In this example, the IP address is `192.168.1.2`.

Distribute the public key to the root user's *authorized_keys* file on each peer node, and specify to authorize only from the primary IP address:

```
$ sudo echo 'from="192.168.1.2" ssh-rsa AAAAB3Nz4rFgt...vgaD \
    root@node1' >> /root/.ssh/authorized_keys
```

Add the following line to */etc/ssh/sshd_config* and reload `sshd` on all nodes:

```
$ sudo echo 'PermitRootLogin without-password' >> \
    /etc/ssh/sshd_config
$ sudo service sshd reload
```

Verify that the root user on the primary node can `ssh` to each of the peer nodes without a password:

```
$ sudo ssh root@node2.example.com
```

Create the configuration file */etc/nginx-sync.conf* on the primary machine with the following configuration:

```
NODES="node2.example.com node3.example.com node4.example.com"
CONFPATHS="/etc/nginx/nginx.conf /etc/nginx/conf.d"
EXCLUDE="default.conf"
```

This example configuration demonstrates the three common configuration parameters for this feature: NODES, CONFIGPATHS, and EXCLUDE. The NODES parameter is set to a string of hostnames or IP addresses separated by spaces; these are the peer nodes to which the primary will push its configuration changes. The CONFIGPATHS parameter denotes which files or directories should be synchronized. Lastly, you can use the EXCLUDE parameter to exclude configuration files from synchronization. In our example, the primary pushes configuration changes of the main NGINX configuration file and includes the directory */etc/nginx/nginx.conf* and */etc/nginx/conf.d* to peer nodes named node2.example.com, node3.example.com, and node4.example.com. If the synchronization process finds a file named *default.conf*, it will not be pushed to the peers, because it's configured as an EXCLUDE.

There are advanced configuration parameters to configure the location of the NGINX binary, RSYNC binary, SSH binary, diff binary, lockfile location, and backup directory. There is also a parameter that utilizes sed to template the given files. For more information about the advanced parameters, see "Synchronizing NGINX Configuration in a Cluster".

Test your configuration:

```
$ nginx-sync.sh -h # display usage info
$ nginx-sync.sh -c node2.example.com # compare config to node2
$ nginx-sync.sh -C # compare primary config to all peers
$ nginx-sync.sh # sync the config & reload NGINX on peers
```

Discussion

This NGINX Plus exclusive feature enables you to manage multiple NGINX Plus servers in an HA configuration by updating only the primary node and synchronizing the configuration to all other peer nodes. By automating the synchronization of the configuration, you limit the risk of mistakes when transferring configurations. The *nginx-sync.sh* application provides some safeguards to prevent sending bad configurations to the peers. They include testing the configuration on the primary, creating backups of the configuration on the peers, and validating the configuration on the peer before reloading. Although it's preferable to synchronize your configuration by using configuration management tools or Docker, the NGINX Plus configuration synchronization feature is valuable if you have yet to make the big leap to managing environments in this way.

12.5 State Sharing with NGINX Plus and Zone Sync

Problem

You need NGINX Plus to synchronize its shared memory zones across a fleet of highly available servers.

Solution

Configure zone synchronization, then use the `sync` parameter when configuring an NGINX Plus shared memory zone:

```
stream {
    resolver 10.0.0.2 valid=20s;

    server {
        listen 9000;
        zone_sync;
        zone_sync_server nginx-cluster.example.com:9000 resolve;
        # ... Security measures
    }
}
http {
    upstream my_backend {
        zone my_backend 64k;
        server backends.example.com resolve;
        sticky learn zone=sessions:1m
                    create=$upstream_cookie_session
                    lookup=$cookie_session
                    sync;
    }

    server {
        listen 80;
        location / {
            proxy_pass http://my_backend;
        }
    }
}
```

Discussion

The `zone_sync` module is an NGINX Plus exclusive feature that enables NGINX Plus to truly cluster. As shown in the configuration, you must set up a `stream` server configured as the `zone_sync`. In the example, this is the server listening on port `9000`. NGINX Plus communicates with the rest of the servers defined by the `zone_sync_server` directive. You can set this directive to a domain name that resolves to multiple IP addresses for dynamic clusters, or statically define a series

of `zone_sync_server` directives to avoid single points of failure. You should restrict access to the zone sync server; there are specific SSL/TLS directives for this module for machine authentication. The benefit of configuring NGINX Plus into a cluster is that you can synchronize shared memory zones for rate limiting, sticky-learn sessions, and the key-value store. The example provided shows the `sync` parameter tacked on to the end of a `sticky learn` directive. In this example, a user is bound to an upstream server based on a cookie named `session`. Without the `zone_sync` module, if a user makes a request to a different NGINX Plus server, they could lose their session. With the `zone_sync` module, all of the NGINX Plus servers are aware of the session and which upstream server it's bound to.

Advanced Activity Monitoring

13.0 Introduction

To ensure that your application is running at optimal performance and precision, you need insight into the monitoring metrics about its activity. NGINX Plus offers an advanced monitoring dashboard and a JSON feed to provide in-depth monitoring about all requests that come through the heart of your application. The NGINX Plus activity monitoring provides insight into requests, upstream server pools, caching, health, and more. This chapter details the power and possibilities of the NGINX Plus dashboard, the NGINX Plus API, and the open source stub status module.

13.1 Enable NGINX Open Source Stub Status

Problem

You need to enable basic monitoring for NGINX.

Solution

Enable the `stub_status` module in a location block within an NGINX HTTP server:

```
location /stub_status {
    stub_status;
    allow 127.0.0.1;
    deny all;
    # Set IP restrictions as appropriate
}
```

Test your configuration by making a request for the status:

```
$ curl localhost/stub_status
Active connections: 1
server accepts handled requests
 1 1 1
Reading: 0 Writing: 1 Waiting: 0
```

Discussion

The stub_status module enables some basic monitoring of the NGINX OSS server. The information that is returned provides insight into the number of active connections, as well as the total connections accepted, connections handled, and requests served. The current number of connections being read, written, or in a waiting state is also shown. The information provided is global and is not specific to the parent server where the stub_status directive is defined. This means that you can host the status on a protected server. For security reasons we blocked all access to the monitoring feature except local traffic. This module provides active connection counts as embedded variables for use in logs and elsewhere. These variables are $connections_active, $connections_reading, $connections_writing, and $connections_waiting.

13.2 Enabling the NGINX Plus Monitoring Dashboard

Problem

You require in-depth metrics about the traffic flowing through your NGINX Plus server.

Solution

Utilize the real-time activity monitoring dashboard:

```
server {
    # ...
    location /api {
        api [write=on];
        # Directives limiting access to the API
        # See chapter 7
    }

    location = /dashboard.html {
        root   /usr/share/nginx/html;
    }
}
```

The NGINX Plus configuration serves the NGINX Plus status monitoring dashboard. This configuration sets up an HTTP server to serve the API and the status dashboard. The dashboard is served as static content out of the */usr/share/nginx/html* directory. The dashboard makes requests to the API at */api/* in order to retrieve and display the status in real time.

Discussion

NGINX Plus provides an advanced status monitoring dashboard. This status dashboard provides a detailed status of the NGINX system, such as number of active connections, uptime, upstream server pool information, and more. For a glimpse of the console, see Figure 13-1.

NGINX Controller ADC provides an app-centric view of monitoring a fleet of NGINX Plus servers across different locations. The Infrastructure tab provides information and metrics about servers, environments, and applications in a single interface. For a glimpse at the console, see Figure 13-2.

The landing page of the status dashboard provides an overview of the entire system. Clicking into the HTTP Zones tab lists details about all HTTP servers configured in the NGINX configuration, detailing the number of responses from 1XX to 5XX and an overall total as well as requests per second and the current traffic throughput. The HTTP Upstreams tab details upstream server status: if the server were in a failed state, how many requests it has served, and a total of how many responses have been served by status code, as well as other statistics such as how many health checks it has passed or failed. The TCP/UDP Zones tab details the amount of traffic flowing through the TCP or UDP streams and the number of connections. The TCP/UDP Upstreams tab shows information about how much each of the upstream servers in the TCP/UDP upstream pools is serving, as well as health check pass and fail details and response times. The Caches tab displays information about the amount of space utilized for cache; the amount of traffic served, written, and bypassed; as well as the hit ratio. The NGINX status dashboard is invaluable in monitoring the heart of your applications and traffic flow.
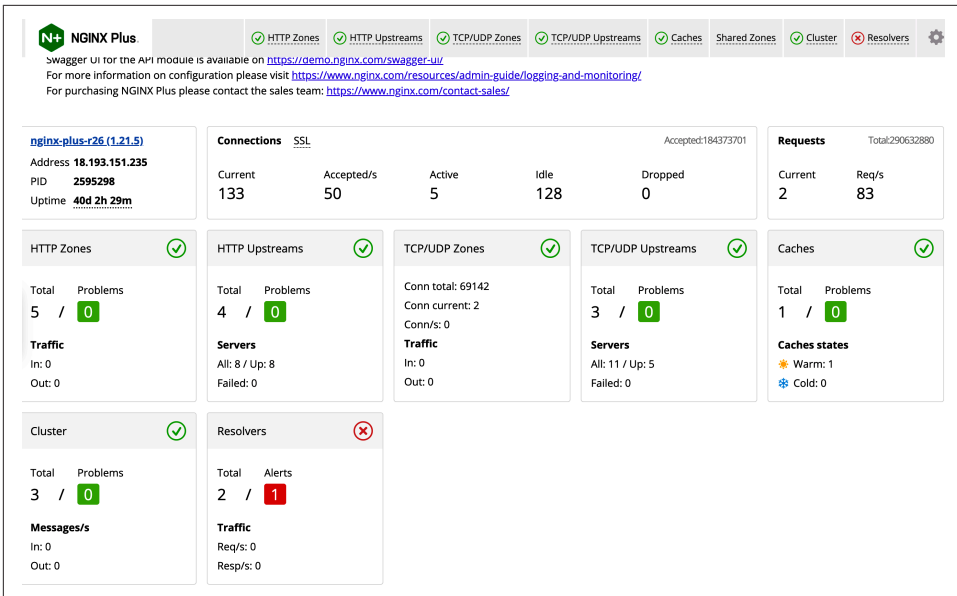
Figure 13-1. The NGINX Plus status dashboard
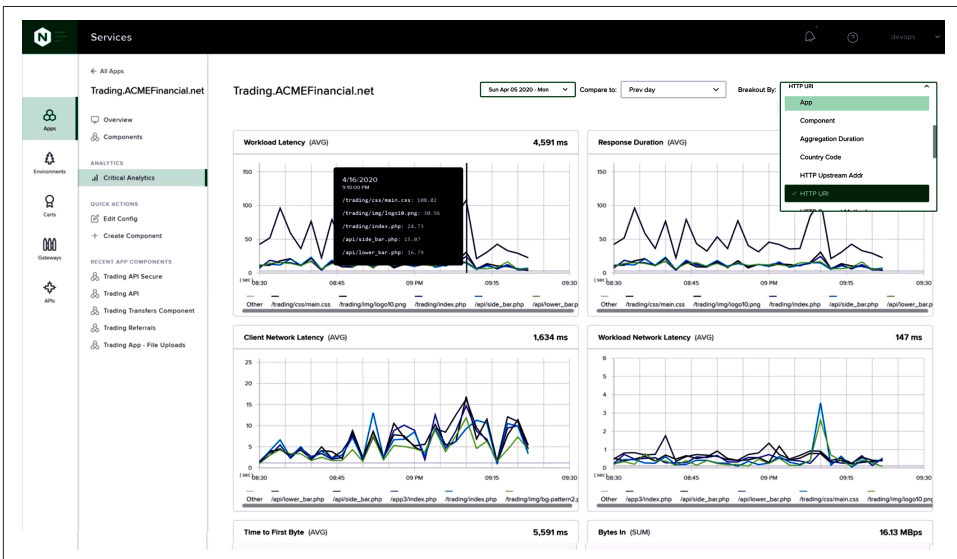
Figure 13-2. An NGINX Controller app-centric analytics dashboard

Also See

NGINX Plus Status Dashboard Demo
F5 NGINX Controller Product Page

13.3 Collecting Metrics Using the NGINX Plus API

Problem

You need API access to the detail metrics provided by the NGINX Plus status dashboard.

Solution

Utilize the RESTful API to collect metrics. The examples pipe the output through `json_pp` to make them easier to read:

```
$ curl "demo.nginx.com/api/3/" | json_pp
[
   "nginx",
   "processes",
   "connections",
   "ssl",
   "slabs",
   "http",
   "stream"
]
```

The `curl` call requests the top level of the API, which displays other portions of the API.

To get information about the NGINX Plus server, use the `/api/{version}/nginx` URI:

```
$ curl "demo.nginx.com/api/3/nginx" | json_pp
{
   "version" : "1.15.2",
   "ppid" : 79909,
   "build" : "nginx-plus-r16",
   "pid" : 77242,
   "address" : "206.251.255.64",
   "timestamp" : "2018-09-29T23:12:20.525Z",
   "load_timestamp" : "2018-09-29T10:00:00.404Z",
   "generation" : 2
}
```

To limit information returned by the API, use arguments:

```
$ curl "demo.nginx.com/api/3/nginx?fields=version,build" \
   | json_pp
{
   "build" : "nginx-plus-r16",
   "version" : "1.15.2"
}
```

You can request connection statistics from the `/api/{version}/connections` URI:

```
$ curl "demo.nginx.com/api/3/connections" | json_pp
{
   "active" : 3,
   "idle" : 34,
   "dropped" : 0,
   "accepted" : 33614951
}
```

You can collect request statistics from the `/api/{version}/http/requests` URI:

```
$ curl "demo.nginx.com/api/3/http/requests" | json_pp
{
   "total" : 52107833,
   "current" : 2
}
```

You can retrieve statistics about a particular server zone using the `/api/{version}/http/server_zones/{httpServerZoneName}` URI:

```
$ curl "demo.nginx.com/api/3/http/server_zones/hg.nginx.org" \
  | json_pp
{
   "responses" : {
      "1xx" : 0,
      "5xx" : 0,
      "3xx" : 938,
      "4xx" : 341,
      "total" : 25245,
      "2xx" : 23966
   },
   "requests" : 25252,
   "discarded" : 7,
   "received" : 5758103,
   "processing" : 0,
   "sent" : 359428196
}
```

The API can return any bit of data you can see on the dashboard. It has depth and follows a logical pattern. You can find links to resources at the end of this recipe.

Discussion

The NGINX Plus API can return statistics about many parts of the NGINX Plus server. You can gather information about the NGINX Plus server and its processes, connections, and slabs. You can also find information about `http` and `stream` servers running within NGINX, including servers, upstreams, upstream servers, and key-value stores, as well as information and statistics about HTTP cache zones. This provides you or third-party metric aggregators with an in-depth view of how your NGINX Plus server is performing.

With the NGINX Instance Manager API, you can query metrics from multiple NGINX Plus servers at once. The metrics provided by Controller offer a different, application-centric view of your metrics.

Also See

NGINX HTTP API module documentation
NGINX API REST UI
"Using the Metrics API" tutorial

Debugging and Troubleshooting with Access Logs, Error Logs, and Request Tracing

14.0 Introduction

Logging is the basis of understanding your application. With NGINX you have great control over logging information meaningful to you and your application. NGINX allows you to divide access logs into different files and formats for different contexts and to change the log level of error logging to get a deeper understanding of what's happening. The capability of streaming logs to a centralized server comes innately to NGINX through its Syslog logging capabilities. NGINX and NGINX Plus also enable tracing of requests as they move through a system. In this chapter, we discuss access and error logs, streaming over the Syslog protocol, and tracing requests end to end with request identifiers generated by NGINX and OpenTracing.

14.1 Configuring Access Logs

Problem

You need to configure access log formats to add embedded variables to your request logs.

Solution

Configure an access log format:

```
http {
    log_format  geoproxy
                '[$time_local] $remote_addr '
                '$realip_remote_addr $remote_user '
                '$proxy_protocol_server_addr $proxy_protocol_server_port '
                '$request_method $server_protocol '
                '$scheme $server_name $uri $status '
                '$request_time $body_bytes_sent '
                '$geoip_city_country_code3 $geoip_region '
                '"$geoip_city" $http_x_forwarded_for '
                '$upstream_status $upstream_response_time '
                '"$http_referer" "$http_user_agent"';
    # ...
}
```

This log format configuration is named `geoproxy` and uses a number of embedded variables to demonstrate the power of NGINX logging. This configuration shows the local time on the server when the request was made, the IP address that opened the connection, and the IP of the client, as NGINX understands it per `geoip_proxy` or `realip_header` instructions.

The variables prefixed with `$proxy_protocol_server_` provide information about the server from the PROXY Protocol header, when the `proxy_protocol` parameter is used on the `listen` directive of the server. `$remote_user` shows the username of the user, authenticated by basic authentication, followed by the request method and protocol, as well as the scheme, such as HTTP or HTTPS. The `server name` match is logged as well as the request URI and the return status code.

Statistics logged include the processing time in milliseconds and the size of the body sent to the client. Information about the country, region, and city are logged. The HTTP header `X-Forwarded-For` is included to show if the request is being forwarded by another proxy. The `upstream` module enables some embedded variables that we've used that show the status returned from the upstream server and how long the upstream request takes to return. Lastly, we've logged some information about where the client was referred from and what browser the client is using.

The `log_format` directive is only valid within the HTTP context. An optional `escape` parameter can specify what type of escaping is done on the string; `default`, `json`, and `none` are escape values. The `none` value disables escaping. For `default` escaping, characters ", \, and other characters with values less than 32 or above 126 are escaped as \xXX. If the variable value is not found, a hyphen (-) will be logged. For `json` escaping, all characters not allowed in JSON strings will be escaped: characters " and \ are escaped as \" and \\, characters with values less than 32 are escaped as \n, \r, \t, \b, \f, or \u00XX.

This log configuration renders a log entry that looks like the following:

```
[25/Nov/2016:16:20:42 +0000] 10.0.1.16 192.168.0.122 Derek
GET HTTP/1.1 http www.example.com / 200 0.001 370 USA MI
"Ann Arbor" - 200 0.001 "-" "curl/7.47.0"
```

To use this log format, use the `access_log` directive, providing a logfile path and the format name `geoproxy` as parameters:

```
server {
    access_log  /var/log/nginx/access.log  geoproxy;
    # ...
}
```

The `access_log` directive takes a logfile path and the format name as parameters. This directive is valid in many contexts and in each context can have a different log path and/or log format. Named parameters such as `buffer`, `flush`, and `gzip` configure how often logs are written to the logfile and if the file is gzipped or not. A parameter named `if` exists and takes a condition; if the condition evaluates to `0` or an empty string, the access will not be logged.

Discussion

The log module in NGINX allows you to configure log formats for many different scenarios to log to numerous logfiles as you see fit. You may find it useful to configure a different log format for each context, where you use different modules and employ those modules' embedded variables, or a single, catchall format that provides all the information you could ever want. It's also possible to log in JSON or XML, provided you construct the format string in that manner. These logs will aid you in understanding your traffic patterns, client usage, who your clients are, and where they're coming from. Access logs can also aid you in finding lag in responses and issues with upstream servers or particular URIs. Access logs can be used to parse and play back traffic patterns in test environments to mimic real user interaction. There are limitless possibilities for logs when troubleshooting, debugging, or analyzing your application or market.

14.2 Configuring Error Logs

Problem

You need to configure error logging to better understand issues with your NGINX server.

Solution

Use the `error_log` directive to define the log path and the log level:

```
error_log /var/log/nginx/error.log warn;
```

The `error_log` directive requires a path; however, the log level is optional, and defaults to `error`. This directive is valid in every context except for `if` statements. The log levels available are `debug`, `info`, `notice`, `warn`, `error`, `crit`, `alert`, or `emerg`. The order in which these log levels were introduced is also the order of severity from least to most. The `debug` log level is only available if NGINX is configured with the `--with-debug` flag.

Discussion

The error log is the first place to look when configuration files are not working correctly. The log is also a great place to find errors produced by application servers like FastCGI. You can use the error log to debug connections down to the worker, memory allocation, client IP, and server. The error log cannot be formatted. However, it follows a specific format of date, followed by the level, then the message.

14.3 Forwarding to Syslog

Problem

You need to forward your logs to a Syslog listener to aggregate logs to a centralized service.

Solution

Use the `error_log` and `access_log` directives to send your logs to a Syslog listener:

```
error_log syslog:server=10.0.1.42 debug;

access_log syslog:server=10.0.1.42,tag=nginx,severity=info geoproxy;
```

The `syslog` parameter for the `error_log` and `access_log` directives is followed by a colon and a number of options. These options include the required server flag that denotes the IP, DNS name, or Unix socket to connect to, as well as optional flags such as `facility`, `severity`, `tag`, and `nohostname`. The `server` option takes a port number, along with IP addresses or DNS names. However, it defaults to UDP 514. The `facility` option refers to the facility attribute of the log message, one of the 23 attributes defined in the RFC standard for Syslog; the default value is `local7`. The `tag` option tags the message with a value. This value defaults to `nginx`. `severity` defaults to `info` and denotes the severity of the message being sent. The `nohostname` flag disables the adding of the hostname field into the Syslog message header and does not take a value.

Discussion

Syslog is a standard protocol for sending log messages and collecting those logs on a single server or collection of servers. Sending logs to a centralized location helps in debugging when you've got multiple instances of the same service running on multiple hosts. This is called aggregating logs. Aggregating logs allows you to view logs together in one place without having to jump from server to server, and mentally stitch together logfiles by timestamp. A common log aggregation stack is Elasticsearch, Logstash, and Kibana, also known as the ELK Stack. NGINX makes streaming these logs to your Syslog listener easy with the `access_log` and `error_log` directives.

14.4 Request Tracing

Problem

You need to correlate NGINX logs with application logs to have an end-to-end understanding of a request.

Solution

Use the `request` identifying variable and pass it to your application to log as well:

```
log_format trace '$remote_addr - $remote_user [$time_local] '
                '"$request" $status $body_bytes_sent '
                '"$http_referer" "$http_user_agent" '
                '"$http_x_forwarded_for" $request_id';
upstream backend {
    server 10.0.0.42;
}
server {
    listen 80;
    # Add the header X-Request-ID to the response to the client
    add_header X-Request-ID $request_id;
    location / {
        proxy_pass http://backend;
        # Send the header X-Request-ID to the application
        proxy_set_header X-Request-ID $request_id;
        access_log /var/log/nginx/access_trace.log trace;
    }
}
```

In this example configuration, a `log_format` named `trace` is set up, and the variable `$request_id` is used in the log. This `$request_id` variable is also passed to the upstream application by use of the `proxy_set_header` directive to add the request ID to a header when making the upstream request. The request ID is also passed

back to the client through use of the `add_header` directive, setting the request ID in a response header.

Discussion

Made available in NGINX Plus R10 and NGINX version 1.11.0, the `$request_id` provides a randomly generated string of 32 hexadecimal characters that can be used to uniquely identify requests. By passing this identifier to the client as well as to the application, you can correlate your logs with the requests you make. From the frontend client, you will receive this unique string as a response header and can use it to search your logs for the entries that correspond. You will need to instruct your application to capture and log this header in its application logs to create a true end-to-end relationship between the logs. With this advancement, NGINX makes it possible to trace requests through your application stack.

14.5 OpenTracing for NGINX

Problem

You have a tracing server that supports OpenTracing and want to integrate NGINX or NGINX Plus.

Solution

Ensure you have an OpenTrace-compatible server available, and that the correct client is installed on the NGINX or NGINX Plus node.

A plug-in configuration file for the specific OpenTrace-compatible server will be needed. This solution will demonstrate Jaeger and Zipkin.

A Jaeger plug-in configuration example named */etc/jaeger/jaeger-config.json* is as follows:

```
{
  "service_name": "nginx",
  "sampler": {
    "type": "const",
    "param": 1
  },
  "reporter": {
    "localAgentHostPort": "Jaeger-server-IP-address:6831"
  }
}
```

A Zipkin plug-in configuration example named */etc/zipkin/zipkin-config.json* is as follows:

```
{
  "service_name": "nginx",
  "collector_host": "Zipkin-server-IP-address",
  "collector_port": 9411
}
```

If NGINX Plus is being used, install the OpenTracing module from the NGINX Plus repository by following the NGINX Plus OpenTracing Admin Guide.

If open source NGINX is being used, visit the NGINX OpenTracing Module Releases page to find a prebuilt dynamic module compatible with your system, or compile the module along with NGINX from source. Alternatively, in a Docker environment, an image named `opentracing/nginx-opentracing` is available on Docker Hub, and can be used to jump-start your testing.

When using a dynamically loaded module, which includes the NGINX Plus installation, ensure that you load it within your NGINX configuration by adding the following `load_module` directive, to tell NGINX where on the filesystem to find the module. As a reminder, the `load_module` directive is valid only in the main (top-level) context.

```
load_module modules/ngx_http_opentracing_module.so;
```

Once an OpenTrace-compatible server is listening, a client is installed on the NGINX node and plug-in configuration is in place, and the NGINX module is loaded, NGINX can be configured to start tracing requests. The following code provides examples of loading a tracer and configuring NGINX to tag requests. The directive to load a tracer plug-in is included, with a default location for the Jaeger and Zipkin plug-ins and the configuration files provided previously. Uncomment the appropriate vendor example for the use case.

```
# Load a vendor tracer
#opentracing_load_tracer /usr/local/libjaegertracing_plugin.so
#                         /etc/jaeger/jaeger-config.json;
#opentracing_load_tracer /usr/local/lib/libzipkin_opentracing_plugin.so
#                         /etc/zipkin/zipkin-config.json;

# Enable tracing for all requests
opentracing on;

# Set additional tags that capture the value of NGINX variables
opentracing_tag bytes_sent $bytes_sent;
opentracing_tag http_user_agent $http_user_agent;
opentracing_tag request_time $request_time;
opentracing_tag upstream_addr $upstream_addr;
opentracing_tag upstream_bytes_received $upstream_bytes_received;
opentracing_tag upstream_cache_status $upstream_cache_status;
opentracing_tag upstream_connect_time $upstream_connect_time;
opentracing_tag upstream_header_time $upstream_header_time;
opentracing_tag upstream_queue_time $upstream_queue_time;
opentracing_tag upstream_response_time $upstream_response_time;
```

```
server {
    listen 9001;

    location / {
        # The operation name used for OpenTracing Spans defaults
        # to the name of the 'location' block,
        # but uncomment this directive to customize it.
        #opentracing_operation_name $uri;

        # Propagate the active Span context upstream,
        # so that the trace can be continued by the backend.
        opentracing_propagate_context;

        # Example application location service
        proxy_pass http://10.0.0.2:8080;
    }
}
```

Discussion

An OpenTracing setup is by no means trivial, but it does provide enormous value in areas of distributed monitoring of performance and transactions. These tools enable teams to effectively provide root cause, and dependency analysis to pinpoint problem areas with data. It's natural for NGINX to serve as an API gateway, routing and authorizing requests between applications, and therefore has integral information to tracing requests through a complex system.

NGINX can tag a request with any variable available to itself, which enables the tracing system user to have a full view of how a request behaves. This example provided a limited sample of using OpenTracing for a proxied request. One can imagine the amount of data that can be harvested from NGINX as the opentracing_tag directive is valid in the http, server, and location contexts.

Also See

OpenTracing NGINX Module GitHub
NGINX Plus OpenTracing Dynamic Module Admin Guide
"Tracing a Proxy" plug-in guide from Datadog guide
"OpenTracing for NGINX and NGINX Plus" blog
"Enabling OpenTracing with NGINX Ingress Controller for Kubernetes" blog

Performance Tuning

15.0 Introduction

Tuning NGINX will make an artist of you. Performance tuning of any type of server or application is always dependent on a number of variable items, such as, but not limited to, the environment, use case, requirements, and physical components involved. It's common to practice bottleneck-driven tuning, meaning to test until you've hit a bottleneck, determine the bottleneck, tune for limitations, and repeat until you've reached your desired performance requirements. In this chapter, we suggest taking measurements when performance tuning by testing with automated tools and measuring results. This chapter also covers connection tuning for keeping connections open to clients as well as upstream servers, and serving more connections by tuning the operating system.

15.1 Automating Tests with Load Drivers

Problem

You need to automate your tests with a load driver to gain consistency and repeatability in your testing.

Solution

Use an HTTP load-testing tool such as Apache JMeter, Locust, Gatling, or whatever your team has standardized on. Create a configuration for your load-testing tool that runs a comprehensive test on your web application. Run your test against your service. Review the metrics collected from the run to establish a baseline. Slowly ramp up the emulated user concurrency to mimic typical production usage and identify

points of improvement. Tune NGINX and repeat this process until you achieve your desired results.

Discussion

Using an automated testing tool to define your test gives you a consistent test to build metrics from when tuning NGINX. You must be able to repeat your test and measure performance gains or losses to conduct science. Running a test before making any tweaks to the NGINX configuration to establish a baseline gives you a basis to work from so that you can measure if your configuration change has improved performance or not. Measuring for each change made will help you identify where your performance enhancements come from.

15.2 Keeping Connections Open to Clients

Problem

You need to increase the number of requests allowed to be made over a single connection from clients and increase the amount of time that idle connections are allowed to persist.

Solution

Use the `keepalive_requests` and `keepalive_timeout` directives to alter the number of requests that can be made over a single connection and change the amount of time idle connections can stay open:

```
http {
    keepalive_requests 320;
    keepalive_timeout 300s;
    # ...
}
```

The `keepalive_requests` directive defaults to 100, and the `keepalive_timeout` directive defaults to 75 seconds.

Discussion

Typically, the default number of requests over a single connection will fulfill client needs because browsers these days are allowed to open multiple connections to a single server per FQDN. The number of parallel open connections to a domain is still typically limited to a number less than 10, so in this regard, many requests over a single connection will happen. A trick for HTTP/1.1 commonly employed by content delivery networks is to create multiple domain names pointed to the content server and alternate which domain name is used within the code to enable the browser

to open more connections. You might find these connection optimizations helpful if your frontend application continually polls your backend application for updates, because an open connection that allows a larger number of requests and stays open longer will limit the number of connections that need to be made.

15.3 Keeping Connections Open Upstream

Problem

You need to keep connections open to upstream servers for reuse to enhance your performance.

Solution

Use the `keepalive` directive in the upstream context to keep connections open to upstream servers for reuse:

```
proxy_http_version 1.1;
proxy_set_header Connection "";

upstream backend {
  server 10.0.0.42;
  server 10.0.2.56;

  keepalive 32;
}
```

The `keepalive` directive in the `upstream` context activates a cache of connections that stay open for each NGINX worker. The directive denotes the maximum number of idle connections to keep open per worker. The proxy modules directives used above the upstream block are necessary for the `keepalive` directive to function properly for upstream server connections. The `proxy_http_version` directive instructs the proxy module to use HTTP version 1.1, which allows for multiple requests to be made in serial over a single connection while it's open. The `proxy_set_header` directive instructs the proxy module to strip the default header of `close`, allowing the connection to stay open.

Discussion

You want to keep connections open to upstream servers to save the amount of time it takes to initiate the connection, allowing the worker process to instead move directly to making a request over an idle connection. It's important to note that the number of open connections can exceed the number of connections specified in the `keepalive` directive because open connections and idle connections are not the same. The number of `keepalive` connections should be kept small enough to allow for

other incoming connections to your upstream server. This small NGINX tuning trick can save some cycles and enhance your performance.

15.4 Buffering Responses

Problem

You need to buffer responses between upstream servers and clients in memory to avoid writing responses to temporary files.

Solution

Tune proxy buffer settings to allow NGINX the memory to buffer response bodies:

```
server {
    proxy_buffering on;
    proxy_buffer_size 8k;
    proxy_buffers 8 32k;
    proxy_busy_buffer_size 64k;
    # ...
}
```

The `proxy_buffering` directive is either `on` or `off`; by default it's `on`. The `proxy_buffer_size` denotes the size of a buffer used for reading the first part of the response and response headers from the proxied server and defaults to either 4k or 8k, depending on the platform. The `proxy_buffers` directive takes two parameters: the number of buffers and the size of the buffers. By default, the `proxy_buffers` directive is set to a number of 8 buffers of size either 4k or 8k, depending on the platform. The `proxy_busy_buffer_size` directive limits the size of buffers that can be busy, sending a response to the client while the response is not fully read. The busy buffer size defaults to double the size of a proxy buffer or the buffer size. If proxy buffering is disabled, the request cannot be sent to the next upstream server in the event of a failure because NGINX has already started sending the request body.

Discussion

Proxy buffers can greatly enhance your proxy performance, depending on the typical size of your response bodies. Tuning these settings can have adverse effects and should be done by observing the average body size returned and performing thorough and repeated testing. Extremely large buffers, set when they're not necessary, can eat up the memory of your NGINX box. You can set these settings for specific locations that are known to return large response bodies for optimal performance.

Also See

15.5 Buffering Access Logs

Problem

You need to buffer logs to reduce the opportunity of blocks to the NGINX worker process when the system is under load.

Solution

Set the buffer size and flush time of your access logs:

```
http {
    access_log /var/log/nginx/access.log main buffer=32k
        flush=1m gzip=1;
}
```

The `buffer` parameter of the `access_log` directive denotes the size of a memory buffer that can be filled with log data before being written to disk. The `flush` parameter of the `access_log` directive sets the longest amount of time a log can remain in a buffer before being written to disk. When using `gzip`, the logs are compressed before being written to the log—values of level 1 (fastest, less compression) through 9 (slowest, best compression) are valid.

Discussion

Buffering log data into memory may be a small step toward optimization. However, for heavily requested sites and applications, this can make a meaningful adjustment to the usage of the disk and CPU. When using the `buffer` parameter to the `access_log` directive, logs will be written out to disk if the next log entry does not fit into the buffer. If using the `flush` parameter in conjunction with the `buffer` parameter, logs will be written to disk when the data in the buffer is older than the time specified. When tailing the log, and with buffering enabled, you may see delays up to the amount of time specified by the `flush` parameter.

15.6 OS Tuning

Problem

You need to tune your operating system to accept more connections to handle spike loads or highly trafficked sites.

Solution

Check the kernel setting for *net.core.somaxconn*, which is the maximum number of connections that can be queued by the kernel for NGINX to process. If you set this number over 512, you'll need to set the `backlog` parameter of the `listen` directive in your NGINX configuration to match. A sign that you should look into this kernel setting is if your kernel log explicitly says to do so. NGINX handles connections very quickly, and, for most use cases, you will not need to alter this setting.

Raising the number of open file descriptors is a more common need. In Linux, a file handle is opened for every connection; therefore, NGINX may open two if you're using it as a proxy or load balancer because of the open connection upstream. To serve a large number of connections, you may need to increase the file descriptor limit system-wide with the kernel option `sys.fs.file_max`, or, for the system user, NGINX is running as in the */etc/security/limits.conf* file. When doing so, you'll also want to bump the number of `worker_connections` and `worker_rlimit_nofile`. Both of these configurations are directives in the NGINX configuration.

Enable more ephemeral ports to support more connections. When NGINX acts as a reverse proxy or load balancer, every connection upstream opens a temporary port for return traffic. Depending on your system configuration, the server may not have the maximum number of ephemeral ports open. To check, review the setting for the kernel setting `net.ipv4.ip_local_port_range`. The setting is a lower- and upper-bound range of ports. It's typically OK to set this kernel setting from 1024 to 65535. 1024 is where the registered TCP ports stop, and 65535 is where dynamic or ephemeral ports stop. Keep in mind that your lower bound should be higher than the highest open listening service port.

Discussion

Tuning the operating system is one of the first places you look when you start tuning for a high number of connections. There are many optimizations you can make to your kernel for your particular use case. However, kernel tuning should not be done on a whim, and changes should be measured for their performance to ensure the changes are helping. As stated before, you'll know when it's time to start tuning your kernel from messages logged in the kernel log or when NGINX explicitly logs a message in its error log.

Introduction to NGINX Instance Manager

16.0 Introduction

NGINX Instance Manager is a service that works with an agent on your NGINX instances. The service can scan your networks for existing NGINX instances, as well as TLS certificates. By providing an inventory of instances and certificates, you can view your landscape of web servers and proxies from a single interface or API. The API enables you to scan, check, and update your data plane through automation that you can integrate into your existing tools and workflows.

The NGINX Instance Manager agent collects metrics from NGINX or NGINX Plus, and allows for direct configuration checks and updates from a centralized interface. You can use NGINX Instance Manager as a metric collection integration point, which reduces the complexity of adding NGINX stub status or NGINX Plus metric API endpoints to your consolidated metric-gathering platform.

NGINX Instance Manager requires a license; however, there is a free trial available so that you can explore its functionality and value for yourself.

16.1 Setup Overview

Problem

You would like to set up an NGINX Instance Manager.

Solution

Installing the NGINX Instance Manager server follows the same workflow as other licensed NGINX products. You must first obtain a license file, and either the package or a repository certificate key pair. Once you have the necessary files, you use the

package manager for your system to install the software. For the most up-to-date and accurate installation information, use the NGINX Instance Manager Installation Guide.

Discussion

Once you have NGINX Instance Manager installed, browse to its UI. At first you'll have an empty inventory (see Figure 16-1).
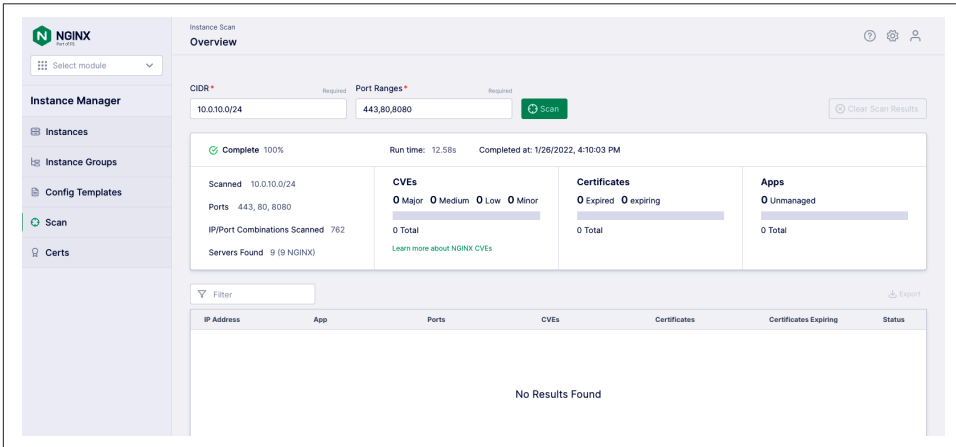
Figure 16-1. NGINX Instance Manager UI

To discover NGINX instances, navigate to the Scan tab (see Figure 16-2). You can use the scan functionality to discover NGINX and NGINX Plus servers on your network. The scan tool uses NMAP to probe for specific ports that you may run web servers on. If a host is listening on the ports you're scanning, NGINX Instance Manager attempts to determine if it's an NGINX server and, if so, what version. This is determined by making a request to the server and inspecting the Server header of the response.

After the scan is complete, you'll see a table populated with discovered instances. This table will detail information about each instance. When an NGINX instance is discovered, a lookup for any common vulnerabilities and exposures (CVEs) based on the NGINX version will be performed and shown. These results aid in keeping your NGINX servers secure by automating the discovery of necessary patches. A status column informs you which NGINX instances are managed by NGINX Instance Manager and which instances need the agent installed.
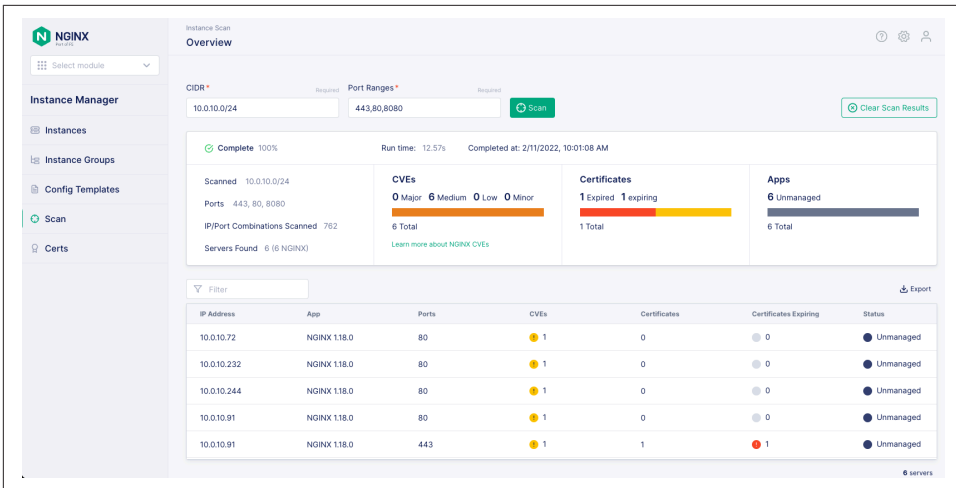
Figure 16-2. NGINX Instance Manager Scan

On the Certs page, you can scan for hosts configured for TLS. This tool works the same way as the host scan; however, it inspects the certificates in use and provides you information about the certificate, such as the common name and the expiration date. This is a helpful tool to discover where certificates live and where updates may be necessary.

Within Settings, you can view information about NGINX Instance Manager, your license, plus documentation for Instance Manager, the API, and how to configure metric telemetry. This is also where you can configure users and roles in order to control who has access and what they can do.

Before you can manage an instance, you must install the NGINX Instance Manager agent on your NGINX machines. To learn how to install the agent, continue to the next section.

Also See

NGINX Instance Manager Installation Guide
NGINX Instance Manager Admin Guide

16.2 Agent Installation

Problem

You would like to install the NGINX Instance Manager agent on your NGINX or NGINX Plus machines.

Solution

Installing the NGINX Instance Manager agent is done by downloading and executing the installation script from the NGINX Instance Manager server using the following command:

```
curl -k https://{NIM hostname or IP}/install/nginx-agent | sudo sh
```

For a basic installation, these are all the steps required. To set the displayed instance name, location, or instance group membership, you can update the */etc/nginx-agent/ dynamic-agent.conf* file and restart the nginx-agent service:

```
#
# /etc/nginx-agent/dynamic-agent.conf
# Dynamic configuration file for NGINX Agent.
# The purpose of this file is to track agent configuration
#  values that can be dynamically changed via the API and the agent
#  install script.
# You may edit this file, but API calls that modify the tags on
#  this system will overwrite the tag values in this file.
# The agent configuration values that API calls can modify are as
#  follows:
# - tags #
# The agent configuration values that the agent install script can
#  modify are as follows:
# - instance_name # - instance_group # - location
# Example Values
instance_name: nginx1.example.org
instance_group: internal
location: cloud1
```

Discussion

Once you have the NGINX Instance Manager agent installed and running on your NGINX or NGINX Plus instances, you will see them appear on the inventory page within NGINX Instance Manager. Information about the host will be detailed in a table, such as its hostname, type of NGINX installation, tags, and its connection status.

By selecting an instance, a more detailed view will appear on the righthand side, with the option to edit. By selecting Edit, you will be shown the current NGINX configuration of the instance. You can then use the editor to make changes or analyze the configuration. A green drop-down button in the top right of the editor allows you to switch between configurable files. Once a change is made you can save the configuration. To reload NGINX from the configuration you've saved, you must publish. The separation of the save and publish actions allows you to work on multiple files before NGINX attempts to reload the configuration.

16.3 Automating NGINX Discovery, Configuration, and Monitoring with the API

Problem

You would like to use the NGINX Instance Manager API to automate discovery, configuration, or monitoring.

Solution

Your NGINX Instance Manager server hosts an OpenAPI spec via Swagger UI for documentation of the installed versions API. You can find this documentation in the Docs Menu area.

You may find it useful to regularly scan your network for NGINX instances and their certificates. You can use this information to alert about CVEs that need to be patched, or certificates that are about to expire.

To start a scan for NGINX instances or certificates, send a `POST` to the NGINX Instance Manager endpoint `/api/platform/v1/servers/scan` with a body that details the CIDR range and ports you would like to scan. The instance scan also allows for a list of port ranges:

```
curl -X 'POST' \
  'https://nginx-manager.example.com/api/platform/v1/servers/scan' \
  -H 'accept: application/json' \
  -H 'Content-Type: application/json' \
  -d '{
  "cidr": "string",
  "hostDiscovery": "icmp",
  "portRanges": [
  "10.0.1.0/24"
  ]
}
```

You can use a `DELETE` request to the same API endpoint to cancel a scan. A `GET` request will return the status of the current scan.

Use a `GET` request to `/api/platform/v1/scan/servers` to return information about known NGINX instances or certificates after the scan is complete.

To list all managed NGINX instances, use a `GET` request to `/api/platform/v1/instances`. The response will provide you with instance IDs that can be used in subsequent requests to list configuration files. Those configuration files can be updated and published through the API. This functionality is useful in automating configuration updates to your NGINX instances. Reference the OpenAPI spec for up-to-date usage.

The NGINX Instance Manager also provides access to metric information of managed instances. External metric-gathering tools can request metric information directly from NGINX Instance Manager through the API. Limited metric information can be viewed by visiting the Instance Inventory page, selecting an instance, and clicking the View Metrics link.

Discussion

The operations landscape of today's web technology depends on automation and monitoring. NGINX Instance Manager's RESTful API provides a way for you to integrate NGINX management into your existing automation and DevOps workflows. Metric collection and analysis provides the basis for decision making about configuration changes to optimize your service delivery. You can use NGINX Instance Manager as an integration point for collecting metrics across all of your NGINX instances, which reduces the complexity of metric collection.

Also See

NGINX Instance Manager API Overview
"Using the Metrics API" tutorial

Introduction to NGINX Controller

17.0 Introduction

NGINX Controller is an application-centric control plane for your application environments. Controller provides an interface that allows you to view and configure an entire fleet of NGINX Plus servers, no matter their physical location. Controller allows teams to focus less on the raw NGINX Plus configuration, and more on the application they're using NGINX Plus to deliver.

In this chapter, you'll read an overview of the NGINX Controller setup, connect an NGINX Plus server instance, and learn about using the NGINX Controller API. NGINX Controller is an enterprise product that requires a license. You can request a free trial from the F5 NGINX Controller Product Page.

17.1 Setup Overview

Problem

You would like to set up an NGINX Controller environment.

Solution

Use the official NGINX Controller Installation Guide for an up-to-date installation process. The following are a few tips, observations, and callouts for items to look out for throughout the setup guide.

NGINX Controller 3.x installs as a Kubernetes stack. It is important to review all of the technical specifications before beginning. An external PostgreSQL database is required. The Controller installer is provided as a tarball. Once unpacked, an *install.sh* script will need to be run as a nonroot user.

Due to the way some OS images are distributed, there may be variance in the package repositories, which can cause some difficulty with the installation. Ubuntu 20.04 seems to be the most consistent in my testing, and is my recommendation for trials and exploration. Remember that NGINX Support is available to assist you in getting NGINX Controller up and running quickly.

There are a number of tools the installer needs installed before it's able to run correctly. Most of the tools are standard on many operating systems, but the `jq` tool is not. You will need to ensure all the required tools are installed on your system prior to running any of the installer scripts.

A *helper.sh* script is provided in the installation package that can aid with installation, or altering the base configuration once installed. For instance, the argument `supportpkg` will build a package of debug and log information for you to send to NGINX Support to enable them to have a quick overview of your situation. The `prereqs` argument will install required packages and set up Kubernetes. To view logs from NGINX Controller, you can use `./helper.sh logs`.

When the `installer` command is run, it will check the system for requirements and install any that it may additionally need. The installer will prompt for database information. Currently PostgreSQL is supported and must be on a remote server. The user information provided must be capable of creating databases. This information can be passed as command-line arguments to the installer.

A time series database volume will need to be supplied. You can use the local disk, an NSF volume, or an AWS EBS volume. If you choose to use an AWS EBS volume, the system will need appropriate AWS IAM permissions to attach the volume to the instance.

An end-user license agreement is presented and must be accepted to move forward. After reading, press `q` to exit the agreement, then `y` to accept.

An SMTP server is needed to invite users via email as well as for email notifications. In the event that an SMTP server is not yet available, these settings can be configured later by using the *helper.sh* script. Provide some generic values to these prompts, set the host to `localhost`, port to `25`, and decline authentication and TLS. NGINX Controller will not be able to send email until SMTP is configured.

The FQDN is used when generating agent configuration, and should be set to a domain that can be relied upon. The organization name prompt is a friendly name used for labeling—a team or company name will suffice. When providing values for the administrator user, note that the email is used for system login.

SSL/TLS certificate paths can be provided through command parameters of the installer, or as environment variables. If these are not found, the installer will prompt to generate self-signed certificates.

Once the installation completes, the installer will provide a link to Controller. Follow the link and log in with the administrator credentials.

Discussion

NGINX Controller provides a single control plane for management of your applications. The interface is sectioned into different views: Platform, Infrastructure, Services, and Analytics. By doing so, the view is clean and concise for the specific task at hand.

The Platform view is used for managing Controller and user access. The Infrastructure view provides details about the machines running the NGINX Controller agents. The next section will describe adding an NGINX Plus server to Controller by installing an agent.

In the Services view, the application-centric attributes of NGINX Controller come to light. Controller organizes your applications, environments, gateways, and APIs to enable you to reorganize and deploy rapidly.

Also See

NGINX Controller Installation Guide
NGINX Controller Tech Specs
"Installing and preparing PostgreSQL database for NGINX Controller"

17.2 Connecting NGINX Plus with Controller

Problem

You've installed Controller and need to connect an NGINX Plus instance with an agent.

Solution

If you have not yet installed NGINX Plus, use Recipe 1.3 to get an NGINX Plus node online.

The best way to find documentation for your Controller installation is to visit *https:// <Controller-FQDN>/docs/infrastructure/agent*. At this document location, you can find information about the technical specifications needed to run the NGINX Controller Agent, as well as how to install and manage it.

Installing the Controller Agent to an existing NGINX Plus server is straightforward. You will need to retrieve an installer script from the Controller API on port 8443, and run it with an API key. The Controller UI provides simple copy-and-paste

instructions for your environment. After the install is finished, you must start the Controller Agent by using the service manager for your system.

Once the Controller Agent service is running, you will see an instance running in the Controller Infrastructure view.

Discussion

In this section, you added an NGINX Plus server to the NGINX Controller as an instance. An inventory of the NGINX Plus systems is now shown in the Infrastructure view as well as with a list request to the API. When you have one or more running instances in NGINX Controller, you can monitor valuable server and NGINX Plus metrics with the Graphs tab within the Infrastructure view. In the Platform view, under the Agents tab, there's a setting to enable NGINX config analyzer. When turned on, the Infrastructure view enables an Analysis tab. The Analysis tab provides information relevant to the NGINX Plus installation and its current configuration.

Now that you have a fresh NGINX Plus node with the Controller Agent installed, you may want to take a bootable image of this machine, or build configuration management to support these installations, so that you can replicate the machine. With an instance configured, you're able to start setting up services, which consist of applications, their environments, and how they're served.

Also See

NGINX Controller Agent Installation Guide

17.3 Driving NGINX Controller with the API

Problem

You've learned how to configure NGINX Controller entities, and want to automate these processes with the API.

Solution

Ensure you have network connectivity to Controller on the API port, which is by default 8443.

NGINX Controller is 100% driven through its API. The interface simply uses that API to provide point-and-click access and dashboards. Use the API Overview in the documentation on your Controller installation by visiting *https://<Controller-FQDN>/docs/api/overview*. This will teach you the basics of the objects, permissions, and how to authenticate. From there, the API Reference can be found at *https://<Controller-FQDN>/docs/api/api-reference*.

One way to jump-start your automation through using the API is to view already configured entities in the NGINX Controller interface, edit the entity, and view the API Spec. This API Spec will show you the method, path, and payload needed to create that object. With some variable substitution, you're off to a good start on automating your Controller environment.

Discussion

For some engineers, the API will be the main interaction they have with NGINX Controller; for others, it will be the web interface. Either is valid and extremely powerful. The addition of showing the API call on the web interface lessens the frustration of digging through API reference and quickens your task automation. An Ansible collection for NGINX Controller exists to aid in Controller automation.

Also See

"Getting Started with the Ansible Collection for NGINX Controller"

17.4 Enable WAF Through Controller App Security

Problem

You're using NGINX Controller ADC and would like to enable Web Application Firewall (WAF) capabilities for your applications.

Solution

If you have not already, follow the NGINX App Protect WAF Installation Guide for your platform to install the App Protect module on your NGINX Plus node.

Navigate to the configuration of an existing App Component in NGINX Controller. Within the Security section, locate the WAF header or settings. Enable the WAF and save.

The WAF is now processing requests for the App through the default WAF policy. The default policy is set to alarm on all signatures, but will block for signatures that are considered highly accurate. Accuracy is determined by an algorithm that determines the probability of false positives. This means you can immediately start blocking harmful requests while gathering data about security events that are reported by the policy. Both flagged and blocked requests will show up in the NGINX Controller UI, appropriately labeled. NGINX Controller ADC will display WAF statistics and violation events for WAF violations triggered.

Ensure the application is handling some traffic. Test a request that would typically be blocked or flagged by a WAF. The following is an extremely basic SQL Injection request:

```
curl https://<appComponentEndpoint>/?query=9999999%20UNION%20SELECT%201%2C2
```

After a request is made that would be considered a security event, NGINX Controller ADC will report security analytics data. Locate these metrics in the Security Analytics section for this App and App Component. In Figure 17-1 you can see how NGINX Controller displays metric information about requests flagged by the WAF.
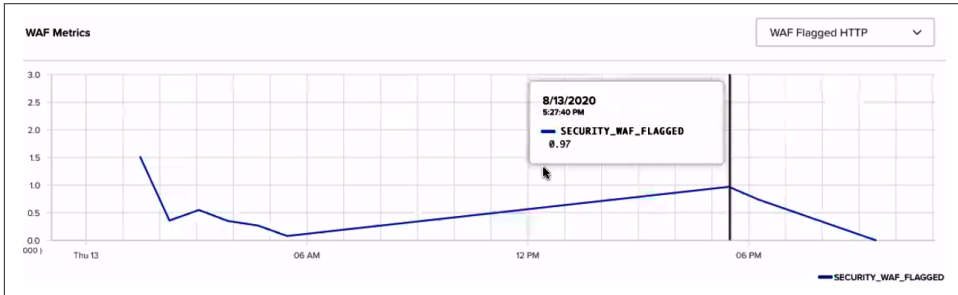
Figure 17-1. An NGINX App Security WAF has flagged some requests

Once some flagged requests start to show up, you can also view the event on the Security Events page. This is where you'll find detailed information about each request flagged or blocked by NGINX App Security WAF.

Before enabling more strict policies, you should verify that your normal valid application traffic is not getting flagged. If normal behavior is flagged, inspect the reasoning for the individual security event. If normal application traffic is being flagged, you may have vulnerabilities in your application that need to be addressed. In Figure 17-2 all traffic observed by the WAF is displayed to show a pattern between regular traffic and requests that violated WAF rules.
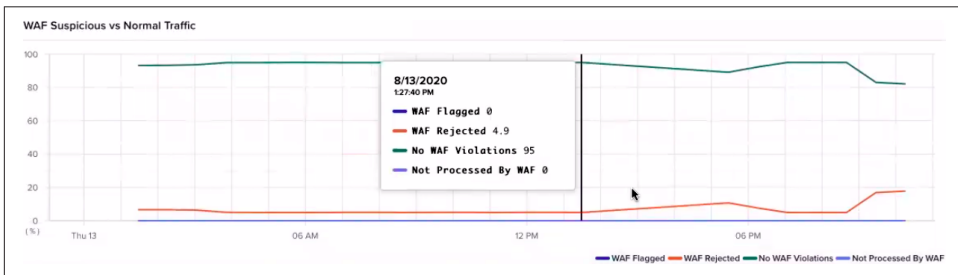
Figure 17-2. An example of statistics reported by Controller ADC with a WAF enabled

Discussion

Controller App Security provides a simplistic WAF experience to secure your applications. With the information provided by the monitoring, you can watch for trends over time for security attacks, investigate further by looking at the detail of the event, and decide what action must be taken.

Web Application Firewalls are extremely important in today's web security architecture. Applications are constantly being bombarded with attempts to compromise services with common vulnerabilities. By blocking these requests before they reach your application services, you're not only securing the web application, but reserving resources for legitimate client requests.

A WAF is not only for external clients; you should consider using a WAF on internal traffic as well, so that one compromised service does not affect another.

Practical Ops Tips and Conclusion

18.0 Introduction

This last chapter will cover practical operations tips and is the conclusion to this book. Throughout this book, we've discussed many concepts pertinent to operations engineers. However, I thought a few more ideas might be helpful to round things out. In this chapter, I'll cover making sure your configuration files are clean and concise, as well as debugging configuration files.

18.1 Using Includes for Clean Configs

Problem

You need to clean up bulky configuration files to keep your configurations logically grouped into modular configuration sets.

Solution

Use the `include` directive to reference configuration files, directories, or masks:

```
http {
    include config.d/compression.conf;
    include sites-enabled/*.conf
}
```

The `include` directive takes a single parameter of either a path to a file or a mask that matches many files. This directive is valid in any context.

Discussion

By using `include` statements you can keep your NGINX configuration clean and concise. You'll be able to logically group your configurations to avoid configuration files that go on for hundreds of lines. You can create modular configuration files that can be included in multiple places throughout your configuration to avoid duplication of configurations. Take the example *fastcgi_param* configuration file provided in most package management installs of NGINX. If you manage multiple FastCGI virtual servers on a single NGINX box, you can include this configuration file for any location or context where you require these parameters for FastCGI without having to duplicate this configuration. Another example is SSL configurations. If you're running multiple servers that require similar SSL configurations, you can simply write this configuration once and include it wherever needed. By logically grouping your configurations together, you can rest assured that your configurations are neat and organized. Changing a set of configuration files can be done by editing a single file rather than changing multiple sets of configuration blocks in multiple locations within a massive configuration file. Grouping your configurations into files and using `include` statements is good practice for your sanity and the sanity of your colleagues.

18.2 Debugging Configs

Problem

You're getting unexpected results from your NGINX server.

Solution

Debug your configuration, and remember these tips:

- NGINX processes requests looking for the most specific matched rule. This makes stepping through configurations by hand a bit harder, but it's the most efficient way for NGINX to work. There's more about how NGINX processes requests in the documentation link in the Also See section.

- You can turn on debug logging. For debug logging, you'll need to ensure that your NGINX package is configured with the `--with-debug` flag. Most of the common packages have it, but if you've built your own or are running a minimal package, you may want to at least double-check. Once you've ensured you have debug enabled, you can set the `error_log` directive's log level to `debug`: `error_log /var/log/nginx/error.log debug`.

- You can enable debugging for particular connections. The `debug_connection` directive is valid inside the `events` context and takes an IP or CIDR range as a parameter. The directive can be declared more than once to add multiple IP

addresses or CIDR ranges to be debugged. This may be helpful to debug an issue in production without degrading performance by debugging all connections.

- You can debug for only particular virtual servers. Because the `error_log` directive is valid in the main `http`, `mail`, `stream`, `server`, and `location` contexts, you can set the `debug` log level in only the contexts you need it.

- You can enable core dumps and obtain backtraces from them. Core dumps can be enabled through the operating system or through the NGINX configuration file. You can read more about this from the admin guide in the Also See section.

- You're able to log what's happening in rewrite statements with the `rewrite_log` directive on: `rewrite_log on`.

Discussion

The NGINX platform is vast, and the configuration enables you to do many amazing things. However, with the power to do amazing things, there's also the power to shoot yourself in the foot. When debugging, make sure you know how to trace your request through your configuration; if you have problems, add the debug log level to help. The debug log is quite verbose but very helpful in finding out what NGINX is doing with your request and where in your configuration you've gone wrong.

Also See

"How nginx processes a request"
Debugging NGINX Admin Guide
NGINX `rewrite_log` module documentation

Conclusion

NGINX Cookbook has focused on high-performance load balancing, security, and deploying and maintaining NGINX and NGINX Plus servers. The book has demonstrated some of the most powerful features of the NGINX application delivery platform. NGINX Inc. continues to develop amazing features and stay ahead of the curve.

We've included many short recipes that enable you to better understand some of the directives and modules that make NGINX the heart of the modern web. The NGINX server is not just a web server, nor just a reverse proxy, but an entire application delivery platform, fully capable of authentication and coming alive with the environments that it's employed in.

Index

Symbols

NODES parameter, configuration synchronization, 135

note for working in containerized environment, 113

NTP (network time protocol) servers, 13

O

OIDC (OpenID Connect) identity provider, 66

open file descriptors, raising number of, 160

openssl command, 76

openssl passwd command, 60

OpenTracing for NGINX, 152-154

opentracing_tag directive, 154

OpenVPN service, 13

original client IP address, finding, 30

OS tuning, 159

P

Packer, HashiCorp, 100

passive health checks, 10, 20

passwords, authentication, 59-61

performance tuning, 155-160

    automating tests with load drivers, 155

    buffering access logs, 159

    buffering responses, 158

    caching, 40

    keeping client connections open, 156

    keeping upstream connections open, 157

    OS tuning, 159

Perl module, 53

perl_set directive, 53, 124

PKI (Public Key Infrastructure), 129

PostgreSQL database, NGINX Controller, 167

programmability

    Chef installing and configuring, 54-55

    Consul templating, automating configurations with, 57-58

    dynamic environment configuration for NGINX Plus, 43-47

    extending NGINX with common programming language, 52-53

    installing with Ansible, 55-57

    key-value store setup, 47-49

    NJS module for exposing JavaScript functionality, 49-51

Prometheus Exporter Module, 127

proxies

    and finding original client IP address, 30

    Google App Engine, 110

HTTP proxy module SSL rules, 75

PROXY Protocol, 127

PROXY Protocol header, 148

proxy_buffering directive, 158

proxy_busy_buffer_size directive, 158

proxy_cache directive, 36

proxy_cache_bypass directive, 38

proxy_cache_key directive, 37, 41

proxy_cache_lock directive, 36

proxy_cache_path directive, 35

proxy_cache_purge directive, 39

proxy_http_version directive, 157

proxy_pass directive, 75, 110

proxy_pass_request_body directive, 61

proxy_protocol parameter, 148

proxy_response directive, 14

proxy_set_header directive, 151, 157

proxy_ssl_certificate directive, 75

proxy_ssl_certificate_key directive, 75

proxy_ssl_crl directive, 75

proxy_timeout directive, 14

ps command, 4

Public Key Infrastructure (see PKI)

purging cache, 39

Python, 57, 77, 79

R

random directive, 16, 23

random load balancing, 16

rate-limiting module, 32-34, 117

RBAC (Role-Based Access Control), 126

RedHat (see RHEL)

reload method, 7

reloading configuration, 7

request tracing, 151-154

request_id variable, 151

resolve parameter, server directive, 119

resolver directive, 110, 119

reuseport parameter, 13, 14

rewrite directive, NGINX as API gateway, 116

rewrite_log directive, 177

RHEL (RedHat Enterprise Linux), 2, 27, 49, 133

Role-Based Access Control (see RBAC)

root directive, 7

round robin load balancing, 15, 132

Route 53 DNS service, 101-102, 132

RSA formatted key, 74

S

satisfy directive, 82
scale sets, virtual machine, 107
secrets, 75, 76
secure link directive, 78
secure link module, 76
secure_link_md5 directive, 78
secure_link_secret directive, 76
security controls, 69-87
    authentication (see authentication)
    client-side encryption, 72-74
    CORS, allowing, 70-71
    dynamic DDoS mitigation, 83-84
    expire date to secure a location, 77
    expiring link, generating, 78-80
    HSTS, 81
    HTTPS redirects, 80
    IP address limits, 31-34, 69
    NGINX as API gateway, 116
    NGINX Plus App Protect Module, 84-87
    rate-limiting module, 32
    secrets, 75, 76
    security methods, satisfying multiple, 82
    upstream encryption, 75
segmenting files to increase efficiency, 40
send directive, match block, 22
server block
    serving static files over HTTP, 7
    setting upstream service outside of, 115
server directive, 11, 119
server drain directive, 20
server slow_start directive, 23
servers, adding and removing with NGINX
    Plus, 19, 43-47
server_name directive, 7
session state, 9, 16, 18, 20
shared-memory zones, 117, 136
sidecar pattern, NGINX Service Mesh, 129
slice directive, 40
slicing, cache, 40
slow start, load balancing, 23
SMTP server, NGINX Controller installation,
    168
SPIRE, 130
split_clients module, 25
SSL modules, 72, 73
SSL/TLS
    client-side encryption, 72-74
    HTTP/2 consideration, 90, 91

redirecting to HTTPS, 81
upstream encryption, 75
ssl_certificate directive, 72
ssl_certificate_key directive, 72
state sharing with NGINX Plus and zone sync,
    136
stateful versus stateless applications, 9
static content, serving, 6
sticky cookie directive, 16
sticky learn directive, 17
sticky route directive, 18
stream health checks, 22
stream module, 11-13, 14
stream SSL module, 73
streaming media, 95-98
    HDS, 97
    MP4 and FLV, serving, 95-97
Strict-Transport-Security header, 81
stub status information, Prometheus Exporter
    Module, 128
stub status module, 139
stub_status directive, 140
subrequests, authentication, 61
Swagger UI, 165
sync parameter, DDoS attack mitigation, 84
sys.fs.file_max kernel option, 160
Syslog listener, forwarding to, 150
syslog parameter, 150

T

TCP load balancing, 11-13
TCP/UDP health checks, 21-23, 141
TLS (see SSL/TLS)
traffic management, 25-34
    A/B testing, 25
    access restrictions based on country, 29
    GeoIP module and database, 27-28
    limiting bandwidth per client, 34, 98
    limiting connections, 31
    limiting rate of requests by predefined key,
        32-34
troubleshooting (see debugging)
type parameter, key-value store, 49

U

Ubuntu, 1, 49, 134
UDP load balancing, 13, 21-23, 141
upstream block, 92, 115
upstream connections, keeping open, 157

About the Author

Derek DeJonghe has a passion for technology. His background and experience in web development, system administration, and networking give him a well-rounded understanding of modern web architecture. Derek leads a team of site reliability and cloud solution engineers and produces self-healing, autoscaling infrastructure for numerous applications. While designing, building, and maintaining highly available applications for clients, he consults for larger organizations as they embark on their journey to the cloud. Derek and his team are on the forefront of a technology tidal wave and are engineering cloud best practices every day. With a proven track record for resilient cloud architecture, Derek pioneers cloud deployments for security and maintainability that are in the best interest of his clients.

Colophon

The animal on the cover of *NGINX Cookbook* is the Eurasian lynx (*Lynx lynx*), the largest of the lynx species, found in a broad geographical range from Western Europe to Central Asia.

This wild cat has striking vertical tufts of dark fur atop its ears, with rough long hair on its face. Its fur is yellow-gray to gray-brown, with white coloring on the underbelly. This lynx is spattered with dark spots, with northern-dwelling variants tending to be grayer and less spotted than their southern counterparts.

Unlike other lynx species, the Eurasian lynx preys on larger ungulates—hooved animals—such as wild deer, moose, and even domesticated sheep. Adults require two to five pounds of meat each day and will feed on a single source of food for up to a week.

The Eurasian lynx came close to extinction in the mid-twentieth century but subsequent conservation efforts have brought the cat's conservation status to Least Concern. Many of the animals on O'Reilly covers are endangered; all of them are important to the world.

The cover illustration is by Karen Montgomery, based on a black and white engraving from Shaw's *Zoology*. The cover fonts are Gilroy Semibold and Guardian Sans. The text font is Adobe Minion Pro; the heading font is Adobe Myriad Condensed; and the code font is Dalton Maag's Ubuntu Mono.