


O'REILLY®

Compliments of
NGIN+

Complete NGINX Cookbook



Derek DeJonghe

see more please visit: <https://homeofpdf.com>

flawless application delivery



Load
Balancer




Content
Cache



Web
Server



Security
Controls



Monitoring &
Management

FREE TRIAL

LEARN MORE

NGINX+

Complete NGINX Cookbook

Advanced Recipes for Operations

Derek DeJonghe

Beijing • Boston • Farnham • Sebastopol • Tokyo

O'REILLY®

NGINX Cookbook

by Derek DeJonghe

Copyright © 2017 O'Reilly Media Inc. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://oreilly.com/safari>). For more information, contact our corporate/institutional sales department: 800-998-9938 or corporate@oreilly.com.

Editor: Virginia Wilson

Acquisitions Editor: Brian Anderson

Production Editor: Shiny Kalapurakkal

Copyeditor: Amanda Kersey

Proofreader: Sonia Saruba

Interior Designer: David Futato

Cover Designer: Karen Montgomery

Illustrator: Rebecca Demarest

March 2017: First Edition

Revision History for the First Edition

2017-05-26: First Release

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. NGINX Cookbook, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

While the publisher and the author have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the author disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

978-1-491-96895-6

[LSI]

Table of Contents

Part I. Part I: Load Balancing and HTTP Caching

1. High-Performance Load Balancing.....	1
1.0 Introduction	1
1.1 HTTP Load Balancing	2
1.2 TCP Load Balancing	3
1.3 Load-Balancing Methods	4
1.4 Connection Limiting	6
2. Intelligent Session Persistence.....	9
2.0 Introduction	9
2.1 Sticky Cookie	10
2.2 Sticky Learn	11
2.3 Sticky Routing	12
2.4 Connection Draining	13
3. Application-Aware Health Checks.....	15
3.0 Introduction	15
3.1 What to Check	15
3.2 Slow Start	16
3.3 TCP Health Checks	17
3.4 HTTP Health Checks	18
4. High-Availability Deployment Modes.....	21
4.0 Introduction	21
4.1 NGINX HA Mode	21

4.2 Load-Balancing Load Balancers with DNS	22
4.3 Load Balancing on EC2	23
5. Massively Scalable Content Caching.....	25
5.0 Introduction	25
5.1 Caching Zones	25
5.2 Caching Hash Keys	27
5.3 Cache Bypass	28
5.4 Cache Performance	29
5.5 Purging	30
6. Sophisticated Media Streaming.....	31
6.0 Introduction	31
6.1 Serving MP4 and FLV	31
6.2 Streaming with HLS	32
6.3 Streaming with HDS	34
6.4 Bandwidth Limits	34
7. Advanced Activity Monitoring.....	37
7.0 Introduction	37
7.1 NGINX Traffic Monitoring	37
7.2 The JSON Feed	39
8. DevOps On-the-Fly Reconfiguration.....	41
8.0 Introduction	41
8.1 The NGINX API	41
8.2 Seamless Reload	43
8.3 SRV Records	44
9. UDP Load Balancing.....	47
9.0 Introduction	47
9.1 Stream Context	47
9.2 Load-Balancing Algorithms	49
9.3 Health Checks	49
10. Cloud-Agnostic Architecture.....	51
10.0 Introduction	51
10.1 The Anywhere Load Balancer	51
10.2 The Importance of Versatility	52

Part II. Part II: Security and Access

11. Controlling Access.....	57
11.0 Introduction.....	57
11.1 Access Based on IP Address.....	57
11.2 Allowing Cross-Origin Resource Sharing.....	58
12. Limiting Use.....	61
12.0 Introduction.....	61
12.1 Limiting Connections.....	61
12.2 Limiting Rate.....	63
12.3 Limiting Bandwidth.....	64
13. Encrypting.....	67
13.0 Introduction.....	67
13.1 Client-Side Encryption.....	67
13.2 Upstream Encryption.....	69
14. HTTP Basic Authentication.....	71
14.0 Introduction.....	71
14.1 Creating a User File.....	71
14.2 Using Basic Authentication.....	72
15. HTTP Authentication Subrequests.....	75
15.0 Introduction.....	75
15.1 Authentication Subrequests.....	75
16. Secure Links.....	77
16.0 Introduction.....	77
16.1 Securing a Location.....	77
16.2 Generating a Secure Link with a Secret.....	78
16.3 Securing a Location with an Expire Date.....	80
16.4 Generating an Expiring Link.....	81
17. API Authentication Using JWT.....	83
17.0 Introduction.....	83
17.1 Validating JWTs.....	83
17.2 Creating JSON Web Keys.....	84
18. OpenId Connect Single Sign-On.....	87
18.0 Introduction.....	87

18.1 Authenticate Users via Existing OpenId Connect	
Single Sign-On (SSO)	87
18.2 Obtaining JSON Web Key from Google	89
19. ModSecurity Web Application Firewall.....	91
19.0 Introduction	91
19.1 Installing ModSecurity for NGINX Plus	91
19.2 Configuring ModSecurity in NGINX Plus	92
19.3 Installing ModSecurity from Source for a Web	
Application Firewall	93
20. Practical Security Tips.....	97
20.0 Introduction	97
20.1 HTTPS Redirects	97
20.2 Redirecting to HTTPS Where SSL/TLS Is Terminated	
Before NGINX	98
20.3 HTTP Strict Transport Security	99
20.4 Satisfying Any Number of Security Methods	100

Part III. Part III: Deployment and Operations

21. Deploying on AWS.....	103
21.0 Introduction	103
21.1 Auto-Provisioning on AWS	103
21.2 Routing to NGINX Nodes Without an ELB	105
21.3 The ELB Sandwich	106
21.4 Deploying from the Marketplace	108
22. Deploying on Azure.....	111
22.0 Introduction	111
22.1 Creating an NGINX Virtual Machine Image	111
22.2 Load Balancing Over NGINX Scale Sets	113
22.3 Deploying Through the Marketplace	114
23. Deploying on Google Cloud Compute.....	117
23.0 Introduction	117
23.1 Deploying to Google Compute Engine	117
23.2 Creating a Google Compute Image	118
23.3 Creating a Google App Engine Proxy	119

24. Deploying on Docker.....	123
24.0 Introduction	123
24.1 Running Quickly with the NGINX Image	123
24.2 Creating an NGINX Dockerfile	124
24.3 Building an NGINX Plus Image	126
24.4 Using Environment Variables in NGINX	128
25. Using Puppet/Chef/Ansible/SaltStack.....	131
25.0 Introduction	131
25.1 Installing with Puppet	131
25.2 Installing with Chef	133
25.3 Installing with Ansible	135
25.4 Installing with SaltStack	136
26. Automation.....	139
26.0 Introduction	139
26.1 Automating with NGINX Plus	139
26.2 Automating Configurations with Consul Templating	140
27. A/B Testing with split_clients.....	143
27.0 Introduction	143
27.1 A/B Testing	143
28. Locating Users by IP Address Using the GeoIP Module.....	145
28.0 Introduction	145
28.1 Using the GeoIP Module and Database	146
28.2 Restricting Access Based on Country	147
28.3 Finding the Original Client	148
29. Debugging and Troubleshooting with Access Logs, Error Logs, and Request Tracing.....	151
29.0 Introduction	151
29.1 Configuring Access Logs	151
29.2 Configuring Error Logs	153
29.3 Forwarding to Syslog	154
29.4 Request Tracing	155
30. Performance Tuning.....	157
30.0 Introduction	157
30.1 Automating Tests with Load Drivers	157
30.2 Keeping Connections Open to Clients	158

30.3 Keeping Connections Open Upstream	159
30.4 Buffering Responses	160
30.5 Buffering Access Logs	161
30.6 OS Tuning	162
31. Practical Ops Tips and Conclusion.	165
31.0 Introduction	165
31.1 Using Includes for Clean Configs	165
31.2 Debugging Configs	166
31.3 Conclusion	168

Part I: Load Balancing and HTTP Caching

This is Part I of III of *NGINX Cookbook*. This book is about NGINX the web server, reverse proxy, load balancer, and HTTP cache. Part I will focus mostly on the load-balancing aspect and the advanced features around load balancing, as well as some information around HTTP caching. This book will touch on NGINX Plus, the licensed version of NGINX that provides many advanced features, such as a real-time monitoring dashboard and JSON feed, the ability to add servers to a pool of application servers with an API call, and active health checks with an expected response. The following chapters have been written for an audience that has some understanding of NGINX, modern web architectures such as n-tier or microservice designs, and common web protocols such as TCP, UDP, and HTTP. I wrote this book because I believe in NGINX as the strongest web server, proxy, and load balancer we have. I also believe in NGINX's vision as a company. When I heard Owen Garrett, head of products at NGINX, Inc. explain that the core of the NGINX system would continue to be developed and open source, I knew NGINX, Inc. was good for all of us, leading the World Wide Web with one of the most powerful software technologies to serve a vast number of use cases.

Throughout this book, there will be references to both the free and open source NGINX software, as well as the commercial product

from NGINX, Inc., NGINX Plus. Features and directives that are only available as part of the paid subscription to NGINX Plus will be denoted as such. Most readers in this audience will be users and advocates for the free and open source solution; this book's focus is on just that, free and open source NGINX at its core. However, this first part provides an opportunity to view some of the advanced features available in the paid solution, NGINX Plus.

High-Performance Load Balancing

1.0 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 load increases, another copy of the system can be brought online. The architecture technique is called *horizontal scaling*. Software-based infrastructure is increasing in popularity because of its flexibility, opening up a vast world of possibility. Whether the use case is as small as a set of two for high availability or as large as thousands world wide, 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 UDP load balancing, the last of which is discussed in [Chapter 9](#).

This chapter discusses load-balancing configurations for HTTP and TCP in NGINX. In this chapter, you will learn about the NGINX load-balancing algorithms, such as round robin, least connection, least time, IP hash, and generic hash. They will aid you in distributing load in ways more useful to your application. When balancing load, you also want to control the amount of load being served to the application server, which is covered in [Recipe 1.4](#).

1.1 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 {  
    location / {  
        proxy_pass http://backend;  
    }  
}
```

This configuration balances load across two HTTP servers on port 80. The `weight` parameter instructs NGINX to pass twice as many connections 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, either a list of Unix sockets, IP addresses, and DNS 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 an 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 reso-

lution control, and the ability to slowly ramp up connections to a server after it starts.

1.2 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 as a backup that will be passed traffic if the primaries are down.

Discussion

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's to 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.

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, max num-

ber of connections, DNS resolvers, and 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 offered in NGINX Plus can be found throughout Part I of this book. Features available in NGINX Plus, such as connection limiting, can be found later in this chapter. Health checks for all load balancing will be covered in [Chapter 2](#). Dynamic reconfiguration for upstream pools, a feature available in NGINX Plus, is covered in [Chapter 8](#).

1.3 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, or IP hash:

```
upstream backend {  
    least_conn;  
    server backend.example.com;  
    server backend1.example.com;  
}
```

This sets the load-balancing algorithm for the backend upstream pool to be least connections. All load-balancing algorithms, with the exception of generic hash, will be standalone directives like the preceding example. Generic hash takes a single parameter, which can be a concatenation of variables, to build the hash from.

Discussion

Not all requests or packets carry an equal weight. Given this, round robin, or even the weighted round robin used in examples prior, will not fit the need of all applications or traffic flow. NGINX provides a number of load-balancing algorithms that can be used to fit particular use cases. These load-balancing algorithms or methods can not only be chosen, but also configured. The following load-balancing methods are available for upstream HTTP, TCP, and UDP pools:

Round robin

The default load-balancing method, which distributes requests in order of the list of servers in the upstream pool. Weight can be taken into consideration for a weighted round robin, which could be used 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 statistical probability of a weighted average. Round robin is the default load-balancing algorithm and is used if no other algorithm is specified.

Least connections

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

Least time

Available only in NGINX Plus, 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 out there and fits the need 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. 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 amongst 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 determining what 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`.

IP hash

Only supported for HTTP, is the last of the bunch. 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 get 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`.

1.4 Connection Limiting

Problem

You have too much load overwhelming your upstream servers.

Solution

Use NGINX Plus's `max_conns` parameter to limit connections to upstream servers:

```
upstream backend {
    zone backends 64k;
    queue 750 timeout=30s;

    server webserver1.example.com max_conns=25;
    server webserver2.example.com max_conns=15;
}
```

The connection-limiting feature is currently only available in NGINX Plus. This NGINX Plus configuration sets an integer on each upstream server that specifies the max number of connections to be handled at any given time. If the max number of connections has been reached on each server, the request can be placed into the queue for further processing, provided the optional queue directive is specified. The optional queue directive sets the maximum number of requests that can be simultaneously in the queue. A *shared memory zone* is created by use of the zone directive. The shared memory zone allows NGINX Plus worker processes to share information

about how many connections are handled by each server and how many requests are queued.

Discussion

In dealing with distribution of load, one concern is overload. Overloading a server will cause it to queue connections in a listen queue. If the load balancer has no regard for the upstream server, it can load the server's listen queue beyond repair. The ideal approach is for the load balancer to be aware of the connection limitations of the server and queue the connections itself so that it can send the connection to the next available server with understanding of load as a whole. Depending on the upstream server to process its own queue will lead to poor user experience as connections start to timeout. NGINX Plus provides a solution by allowing connections to queue at the load balancer and by making informed decisions on where it sends the next request or session.

The `max_conns` parameter on the `server` directive within the `upstream` block provides NGINX Plus with a limit of how many connections each upstream server can handle. This parameter is configurable in order to match the capacity of a given server. When the number of current connections to a server meets the value of the `max_conns` parameter specified, NGINX Plus will stop sending new requests or sessions to that server until those connections are released.

Optionally, in NGINX Plus, if all upstream servers are at their `max_conns` limit, NGINX Plus can start to queue new connections until resources are freed to handle those connections. Specifying a queue is optional. When queuing, we must take into consideration a reasonable queue length. Much like in everyday life, users and applications would much rather be asked to come back after a short period of time than wait in a long line and still not be served. The `queue` directive in an `upstream` block specifies the max length of the queue. The `timeout` parameter of the `queue` directive specifies how long any given request should wait in queue before giving up, which defaults to 60 seconds.

Intelligent Session Persistence

2.0 Introduction

While HTTP may be a stateless protocol, if the context it's to convey were stateless, the internet would be a much less interesting place. 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 may be stored locally for a number of reasons; for example, in applications where 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 portion of the problem 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.

NGINX Plus's `sticky` directive alleviates difficulties of server affinity at the traffic controller, allowing the application to focus on its core. NGINX tracks session persistence in three ways: by creating and tracking its own cookie, detecting when applications prescribe cookies, or routing based on runtime variables.

2.1 Sticky Cookie

Problem

You need to bind a downstream client to an upstream server.

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. The cookie in this example is named `affinity`, is set for `example.com`, persists an hour, cannot be consumed client-side, can only be sent over HTTPS, and is valid for all paths.

Discussion

Using the `cookie` parameter on the `sticky` directive will create a cookie on first request containing 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 expire time, domain, path, and whether the cookie can be consumed client-side or if it can be passed over unsecure protocols.

2.2 Sticky Learn

Problem

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

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=$upstream_cookie_cookie_name  
        lookup=$cookie_cookie_name  
        zone=client_sessions:2m;  
}
```

The example instructs NGINX to look for and track sessions by looking for a cookie named `COOKIE_NAME` 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 megabytes 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 told to look for cookies in the response from the upstream server with specification of the `create` parameter, and searches for prior registered server affinity by the `lookup` parameter. The value of these parameters are variables exposed by the HTTP module.

2.3 Sticky Routing

Problem

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

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;  
}
```

The 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 non-zero 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. While 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 you may want to direct traffic to a particular server with a bit more granular control. 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 another variable 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.

2.4 Connection Draining

Problem

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

Solution

Use the `drain` parameter through the NGINX Plus API, described in more detail in [Chapter 8](#), to instruct NGINX to stop sending new connections that are not already tracked:

```
$ curl 'http://localhost/upstream_conf\
?upstream=backend&id=1&drain=1'
```

Discussion

When session state is stored locally to a server, connections and persistent sessions must be drained before it's removed from the pool. Draining connections is the process of letting sessions to that server expire natively before removing the server from the upstream pool. Draining can be configured for a particular server by adding the

drain parameter to the server directive. When the drain parameter is set, NGINX Plus will stop sending new sessions to this server but will allow current sessions to continue being served for the length of their session.

Application-Aware Health Checks

3.0 Introduction

For a number of reasons, applications 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 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 the open source version; as well as active, available only in NGINX Plus. Active health checks on a regular interval will make a connection or request to the upstream server and have the ability to 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 may want to use passive health checks to reduce the load of your upstream servers, and you may want to use active health checks to determine failure of an upstream server before a client is served a failure.

3.1 What to Check

Problem

You need to check your application for health but don't know what to check.

Solution

Use a simple but direct indication of the application health. For example, a handler that simply returns an HTTP 200 response tells the load balancer that the application process is running.

Discussion

It's important to check the core of the service you're load balancing for. A single comprehensive health check that ensures all of the systems are available can be problematic. Health checks should check that the application directly behind the load balancer is available over the network and that the application itself is running. With application-aware health checks, you want to pick an endpoint that simply ensures that the processes on that machine are running. It may be tempting to make sure that the database connection strings are correct or that the application can contact its resources. However, this can cause a cascading effect if any particular service fails.

3.2 Slow Start

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.

3.3 TCP Health Checks

Problem

You need to check your upstream TCP server for health and remove unhealthy servers from the pool.

Solution

Use the `health_check` directive in the server block for an active health check:

```
stream {
    server {
        listen      3306;
        proxy_pass  read_backend;
        health_check interval=10 passes=2 fails=3;
    }
}
```

The example monitors the upstream servers actively. The upstream server will be considered unhealthy if it fails to respond to three or more TCP connections initiated by NGINX. NGINX performs the check every 10 seconds. The server will only be considered healthy after passing two health checks.

Discussion

TCP health can be verified by NGINX Plus either passively or actively. Passive health monitoring is done by noting the communication between the client and the upstream server. If the upstream server is timing out or rejecting connections, a passive health check will deem that server unhealthy. Active health checks will initiate their own configurable checks to determine health. Active health

checks not only test a connection to the upstream server, but can expect a given response.

3.4 HTTP Health Checks

Problem

You need to actively check your upstream HTTP servers for health.

Solution

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 request to the URI `/` every two seconds. The upstream servers must pass five consecutive health checks to be considered healthy and will be 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.

Discussion

HTTP health checks in NGINX Plus 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. Active health check monitoring can be configured for how often upstream servers are checked, the URI to check, how many times it 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. The `match` parameter points to a `match` block that defines the acceptance criteria for the response. The `match` block has three directives: `status`, `header`, and `body`. All three of these directives have comparison flags as well.

High-Availability Deployment Modes

4.0 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 will detail techniques on how to run multiple NGINX servers to ensure high availability in your load-balancing tier.

4.1 NGINX HA Mode

Problem

You need a highly available 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 Plus repository includes a package called `nginx-ha-keepalived`. This package, based on `keepalived`, 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 master. The failover capabilities of `nginx-ha-keepalived` can be configured to identify custom failure situations.

4.2 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, fully qualified domain name. 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 described in [Chapter 1](#). These techniques work great. However, a pitfall can be removing the

record when an NGINX server encounters a failure. There are DNS providers—Amazon Route53 for one, and Dyn DNS for another—that offer health checks and failover with their DNS offering, which alleviates these issues. If 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.

4.3 Load Balancing on EC2

Problem

You're using NGINX in AWS, and the NGINX Plus HA does not support Amazon IPs.

Solution

Put NGINX behind an elastic load balancer by configuring an Auto Scaling group of NGINX servers and linking the Auto Scaling group to the elastic load balancer. Alternatively, you can place NGINX servers into the elastic load balancer manually through the Amazon Web Services console, command-line interface, or API.

Discussion

The HA solution from NGINX Plus based on keepalived will not work on Amazon Web Services because it does not support the floating virtual IP address, as EC2 IP addresses work in a different way. This does not mean that NGINX can't be HA in the AWS cloud; in fact, it's the opposite. The AWS elastic load balancer 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 ELB. NGINX servers can be automatically added to and removed from the ELB pool as needed. The ELB is not a replacement for NGINX; there are many things NGINX offers that the ELB does not, such as multiple load-balancing methods, context switching, and UDP load balancing. In the event that the ELB will not fit

your need, there are many other options. One option is the DNS solution, Route53. The DNS product from AWS offers health checks and DNS failover. Amazon also has a white paper about high-availability NGINX Plus, with use of Corosync and Pacemaker, that will cluster the NGINX servers and use an elastic IP to float between boxes for automatic failover.¹

¹ Amazon also has a white paper about NGINX Plus failover on AWS: <http://bit.ly/2aWAqW8>.

Massively Scalable Content Caching

5.0 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. Scaling and distributing 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.

5.1 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 megabytes 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 gigabytes. 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 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 use of the `max_size` parameter. Other parameters are in relation to the cache loading process, which loads the cache keys into the shared memory zone from the files cached on disk.

5.2 Caching Hash Keys

Problem

You need to control how your content is cached and looked up.

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` is `"$scheme$proxy_host$request_uri"`. This default will fit most use cases. The variables used include the scheme, HTTP or HTTPS, the `proxy_host`, where the request is being sent, and the request URI. All together, 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

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 host-name 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. For 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

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: <http://nginx.org/en/docs/varindex.html>.

`proxy_cache_key` can be set in the context of HTTP, server, and location blocks, providing flexible control on how requests are cached.

5.3 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 is by setting a variable within location blocks that you do not want cached to equal 1:

```
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.

Discussion

There are many scenarios that demand that the request is not cached. For this, NGINX exposes a `proxy_cache_bypass` directive that when the value is nonempty or nonzero, the request will be sent to an upstream server rather than be pulled from cache. Interesting techniques and solutions for cache bypass are derived from the need of the client and application. These can be as simple as a request variable or as intricate as a number of map blocks.

For many reasons, you may want to bypass the cache. One important reason is troubleshooting and 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 cache when a particular cookie, header, or request argument is set. You can also turn off cache completely for a given context such as a location block by setting `proxy_cache off;`.

5.4 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 JavaScript 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 to do with many variables, disk speed being high on the list. There are many things within the NGINX configuration 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.

5.5 Purging

Problem

You need to invalidate an object from the cache.

Solution

Use NGINX Plus's purge feature, the `proxy_cache_purge` directive, and a nonempty or zero value variable:

```
map $request_method $purge_method {
    PURGE 1;
    default 0;
}
server {
    ...
    location / {
        ...
        proxy_cache_purge $purge_method;
    }
}
```

Discussion

A common concept for 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 this 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 non-empty 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 a simple authentication to ensure that not 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.

Sophisticated Media Streaming

6.0 Introduction

This chapter covers streaming media with NGINX in MPEG-4 or Flash Video 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 module, as well as the ability to deliver HTTP Dynamic Streaming of already fragmented media. NGINX natively allows for bandwidth limits, and NGINX Plus's advanced feature offers bitrate limiting, enabling your content to be delivered in the most efficient manner while reserving the servers' resources to reach the most users.

6.1 Serving MP4 and FLV

Problem

You need to stream digital media, originating in MPEG-4 (MP4) or Flash Video (FLV).

Solution

Designate an HTTP location block as *.mp4* or *.flv*. 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 of MP4 format type and can be streamed with progressive download support. The second location block instructs NGINX that any files ending in *.flv* are of Flash Video 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 supports seeking to an undownloaded portion of the media in both formats.

6.2 Streaming with HLS

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;
}

```

The location block demonstrated directs NGINX to stream HLS media out of the `/var/www/video` directory, fragmenting the media into four-second segments. The number of HLS buffers is set to 10 with a size of 10 megabytes. The initial MP4 buffer size is set to one megabyte with a maximum of five megabytes.

Discussion

The HLS module available in NGINX Plus provides the ability to transmux 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 as 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 your users to receive larger fragments.

6.3 Streaming with HDS

Problem

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

Solution

Use NGINX Plus's support for fragmented FLV files (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 kilobytes.

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.

6.4 Bandwidth Limits

Problem

You need to limit bandwidth to downstream media streaming clients without impacting 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 a previous section 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 one 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.

Advanced Activity Monitoring

7.0 Introduction

To ensure 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 will detail the power and possibilities of the NGINX Plus dashboard and JSON feed.

7.1 NGINX Traffic Monitoring

Problem

You require in-depth metrics about the traffic flowing through your system.

Solution

Utilize NGINX Plus's real-time activity monitoring dashboard:

```

server {
    listen 8080;
    root /usr/share/nginx/html;

    # Redirect requests for / to /status.html
    location = / {
        return 301 /status.html;
    }

    location = /status.html { }

    # Everything beginning with /status
    # (except for /status.html) is
    # processed by the status handler
    location /status {
        status;
    }
}

```

The NGINX Plus configuration serves the NGINX Plus status monitoring dashboard. This configuration sets up an HTTP server to listen on port 8080, serve content out of the `/usr/share/nginx/html` directory, and redirect `/` requests to `/status.html`. All other `/status` requests will be served by the `/status` location that serves the NGINX Plus status API.

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 7-1](#).

The landing page of the status dashboard provides an overview of the entire system. Clicking into the Server 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 Upstream tab details upstream server status, as in if it's 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 stats 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 Upstream 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.

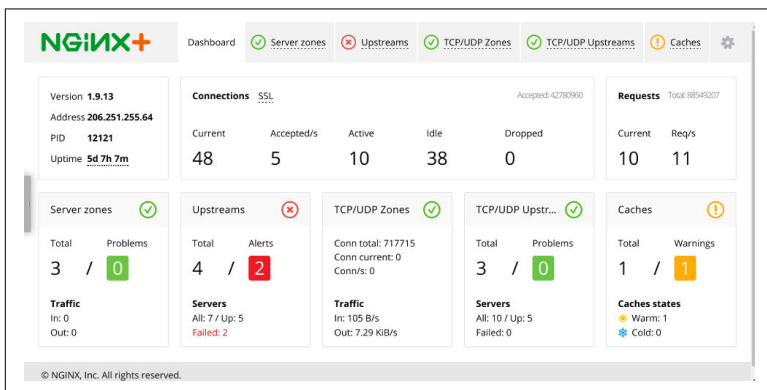


Figure 7-1. The NGINX Plus status dashboard

Also See

[NGINX Plus Status Dashboard Demo](#)

7.2 The JSON Feed

Problem

You need API access to the detail metrics provided by the NGINX Plus status dashboard.

Solution

Utilize the JSON feed provided by NGINX Plus's status API:

```
$ curl "demo.nginx.com/status/upstreams\
/demo-backend/peers/0/responses"
{
  "1xx":0,
  "2xx":199237,
  "3xx":7404,
  "4xx":104415,
  "5xx":19574,
```

```
"total":330630  
}
```

The `curl` call requests a JSON feed from the NGINX Plus status API for information about an upstream HTTP server pool, and in particular about the first server in the pool's responses.

Discussion

The NGINX Plus status API is vast, and requesting just the status will return a JSON object with all the information that can be found on the status dashboard in whole. The JSON feed API allows you to drill down to particular information you may want to monitor or use in custom logic to make application or infrastructure decisions. The API is intuitive and RESTful, and you're able to make requests for objects within the overall status JSON feed to limit the data returned. This JSON feed enables you to feed the monitoring data into any other number of systems you may be utilizing for monitoring, such as Graphite, Datadog, and Splunk.

DevOps On-the-Fly Reconfiguration

8.0 Introduction

The term DevOps has been tossed and spun around more than your favorite pizza crust. To the people actually doing the work, the term has nearly lost meaning; the origin of this term comes from a culture of developers and operations folk working together in an Agile workflow to enhance quality and productivity and share responsibility. If you ask a recruiter, it's a job title; ask someone in marketing, it's a hit-generating Swiss army knife. In this context, we mean DevOps to be developing software and tools to solve operational tasks in the ever-evolving dynamic technology landscape. In this chapter, we'll discuss the NGINX Plus API that allows you to dynamically reconfigure the NGINX Plus load balancer, as well as other tools and patterns to allow your load balancer to evolve with the rest of your environment, such as the seamless reload and NGINX Plus's ability to utilize DNS SRV records.

8.1 The NGINX API

Problem

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

Solution

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

```
location /upstream_conf {
    upstream_conf;
    allow 10.0.0.0/8; # permit access from private network
    deny all;        # deny access from everywhere else
}

...
upstream backend {
    zone backend 64k;
    state /var/lib/nginx/state/backend.state;
    ...
}
```

The NGINX Plus configuration enables the upstream configuration API and only allows access from a private network. The configuration of the `upstream` block defines a shared memory zone named `backend` of 64 kilobytes. The `state` directive tells NGINX to persist these changes through a restart by saving them to the filesystem.

Utilize the API to add servers when they come online:

```
$ curl 'http://nginx.local/upstream_conf?\
add=&upstream=backend&server=10.0.0.42:8080'
```

The `curl` call demonstrated makes a request to NGINX Plus and requests a new server be added to the backend upstream configuration.

Utilize the NGINX Plus API to list the servers in the upstream pool:

```
$ curl 'http://nginx.local/upstream_conf?upstream=backend'\
server 10.0.0.42:8080; # id=0
```

The `curl` call demonstrated makes a request to NGINX Plus to list all of the servers in the upstream pool named `backend`. Currently we only have the one server that we added in the previous `curl` call to the API. The list request will show the IP address, port, and ID of each server in the pool.

Use the NGINX Plus API to drain connections from an upstream server, preparing it for a graceful removal from the upstream pool. Details about connection draining can be found in [Chapter 2, Recipe 2.4](#):

```
$ curl 'http://nginx.local/upstream_conf?\  
upstream=backend&id=0&drain=1'  
server 10.0.0.42:8080; # id=0 draining
```

In this `curl`, we specify arguments for the upstream pool, backend, the ID of the server we wish to drain, 0, and set the `drain` argument to equal 1. 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 in on how many active connections are being served by the server you've begun to drain, you can use the NGINX Plus JSON feed that was detailed in [Chapter 7, Recipe 7.2](#).

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

```
$ curl 'http://nginx.local/upstream_conf?\  
upstream=backend&id=0&remove=1'
```

The `curl` command passes arguments to the NGINX Plus API to remove server 0 from the upstream pool named backend. This API call will return all of the servers and their IDs that are still left in the pool. As 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

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

8.2 Seamless Reload

Problem

You need to reload your configuration without dropping packets.

Solution

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

```
service nginx reload
```

The command-line example reloads the NGINX system using the NGINX init script generally located in the `/etc/init.d/` directory.

Discussion

Reloading the NGINX configuration without stopping the server provides the ability to change configuration 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 the needs of the environment.

8.3 SRV Records

Problem

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

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;

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


The configuration instructs NGINX 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. 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. Autoscaling 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.

UDP Load Balancing

9.0 Introduction

User Datagram Protocol (UDP) is used in many contexts, such as DNS, NTP, and Voice over IP. NGINX can load balance over upstream servers with all the load-balancing algorithms provided to the other protocols. In this chapter, we'll cover the UDP load balancing in NGINX.

9.1 Stream Context

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 NTP servers using the UDP protocol. Specifying UDP load balancing is as simple as using the `udp` parameter on the `listen` directive.

Discussion

One might ask, “Why do you need a load balancer when you can have multiple hosts in a DNS A or SRV record?” The answer is that not only are there alternative balancing algorithms we can balance with, 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, and Voice over IP. UDP load balancing may be less common to some but just as useful in the world of scale.

UDP load balancing will be found in the `stream` module, just like TCP, and configured 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 may apply where they would not in TCP, such as the `proxy_response` directive that tells NGINX how many expected responses may be sent from the upstream server, by default being unlimited until the `proxy_timeout` limit is reached.

9.2 Load-Balancing Algorithms

Problem

You need to distribute load of a UDP service with control over the destination or for best performance.

Solution

Utilize the different load-balancing algorithms, like IP hash or least conn, described in [Chapter 1](#):

```
upstream dns {  
    least_conn;  
    server ns1.example.com:53;  
    server ns2.example.com:53;  
}
```

The configuration load balances over two DNS name servers and directs the request to the name server with the least number of current connections.

Discussion

All of the load-balancing algorithms that were described in [Recipe 9.2](#) are available in UDP load balancing as well. These algorithms, such as least connections, least time, generic hash, or IP hash, are useful tools to provide the best experience to the consumer of the service or application.

9.3 Health Checks

Problem

You need to check the health of upstream UDP servers.

Solution

Use NGINX health checks with UDP load balancing to ensure only healthy upstream servers are sent datagrams:

```
upstream ntp {  
    server ntp1.example.com:123 max_fails=3 fail_timeout=3s;  
    server ntp2.example.com:123 max_fails=3 fail_timeout=3s;  
}
```

This configuration passively monitors the upstream health, setting the `max_fails` directive to 3, and `fail_timeout` to 3 seconds.

Discussion

Health checking is important on all types of load balancing not only from a user experience standpoint but also for business continuity. NGINX can actively and passively monitor upstream UDP servers to ensure they're healthy and performing. Passive monitoring watches for failed or timed-out connections as they pass through NGINX. Active health checks send a packet to the specified port, and can optionally expect a response.

Cloud-Agnostic Architecture

10.0 Introduction

One thing many companies request when moving to the cloud is to be cloud agnostic. Being cloud agnostic in their architectures enables them to pick up and move to another cloud or instantiate the application in a location that one cloud provider may have that another does not. Cloud-agnostic architecture also reduces risk of vendor lock-in and enables an insurance fallback for your application. It's very common for disaster-recovery plans to use an entirely separate cloud, as failure can sometimes be systematic and affect a cloud as a whole. For cloud-agnostic architecture, all of your technology choices must be able to be run in all of those environments. In this chapter, we'll talk about why NGINX is the right technology choice when architecting a solution that will fit in any cloud.

10.1 The Anywhere Load Balancer

Problem

You need a load-balancer solution that can be deployed in any data center, cloud environment, or even local hosts.

Solution

Load balance with NGINX. NGINX is software that can be deployed anywhere. NGINX runs on Unix; and on multiple flavors of Linux such as Cent OS and Debian, BSD variants, Solaris, macOS, Windows, and others. NGINX can be built from source on Unix and Linux derivatives as well as installed through package managers such as yum, aptitude, and zypper. On Windows, it can be installed by downloading a ZIP archive and running the `.exe` file.

Discussion

The fact that NGINX is a software load balancer rather than strictly hardware allows it to be deployed on almost any infrastructure.¹ Cross-cloud environments and hybrid cloud architectures are on the rise, applications are distributed between different clouds for high availability, and vendor-agnostic architecture limits risk of production outages and reduces network latency between the end user and the application. In these scenarios, the application being hosted typically doesn't change and neither should your load-balancing solution. NGINX can be run in all of these environments with all of the power of its configuration.²

10.2 The Importance of Versatility

Problem

You need versatility in your architecture and the ability to build in an iterative manner.

Solution

Use NGINX as your load balancer or traffic router. NGINX provides versatility on the platform it runs on or its configuration. If you're architecting a solution, and you're not sure where it's going to live or

¹ NGINX provides a page to download its software: <http://nginx.org/en/download.html>.

² Linux packages and repositories can be found at http://nginx.org/en/linux_packages.html.

need the flexibility to be able to move it to another provider, NGINX will fit this need. If you're working in an iterative workflow, and new services or configurations are continually changing during the development cycle, NGINX is a prime resource, as its configuration can change; and with a reload of the service, the new configuration is online without concern of stopping the service. An example might be planning to build out a data center, and then for cost and flexibility, switching gears into a cloud environment. Another example might be refactoring an existing monolithic application and slowly decoupling the application into microservices, deploying service by service as the smaller applications become ready for production.

Discussion

Agile workflows have changed how development work is done. The idea of an Agile workflow is an iterative approach where it's OK if requirements or scope change. Infrastructure architecture can also follow an Agile workflow: you may start out aiming to go into a particular cloud provider and then have to switch to another partway through the project, or want to deploy to multiple cloud providers. NGINX being able to run anywhere makes it an extremely versatile tool. The importance of versatility is that with the inevitable onset of cloud, things are always changing. In the ever-evolving landscape of software, NGINX is able to efficiently serve your application needs as it grows with your features and user base.

Part II: Security and Access

This is Part II of III of *NGINX Cookbook*. This part will focus on security aspects and features of NGINX and NGINX Plus, the licensed version of the NGINX server. Throughout this part, you will learn the basics about controlling access and limiting abuse and misuse of your web assets and applications. Security concepts such as encryption of your web traffic as well as basic HTTP authentication will be explained as applicable to the NGINX server. More advanced topics are covered as well, such as setting up NGINX to verify authentication via third-party systems as well as through JSON Web Token Signature validation and integrating with Single sign-on providers. This part covers some amazing features of NGINX and NGINX Plus such as securing links for time-limited access and security as well as enabling Web Application Firewall capabilities of NGINX Plus with the ModSecurity module. Some of the plug-and-play modules in this part are only available through the paid NGINX Plus subscription, however this does not mean that the core open source NGINX server is not capable of these securities.

Controlling Access

11.0 Introduction

Controlling access to your web applications or subsets of your web applications is important business. Access control takes many forms in NGINX, such as denying it at the network level, allowing it based on authentication mechanisms, or HTTP responses instructing browsers how to act. In this chapter we will discuss access control based on network attributes, authentication, and how to specify *Cross-Origin Resource Sharing* (CORS) rules.

11.1 Access Based on IP Address

Problem

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

Solution

Use the HTTP 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. 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 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, 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.

11.2 Allowing Cross-Origin Resource Sharing

Problem

You're serving resources from another domain and need to allow CORS to enable browsers to utilize these resources.

Solution

Alter headers based on the request method to enable CORS:

```
map $request_method $cors_method {
    OPTIONS 1;
    GET 1;
    POST 1;
    default 0;
}
server {
    ...
    location / {
        if ($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 information called 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 cross-origin resource requests when the resource they're requesting is of a domain other than its own origin. 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 to which 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.

Limiting Use

12.0 Introduction

Limiting use or abuse of your system can be important for throttling heavy users or stopping attacks. NGINX has multiple modules built in to help control the use of your applications. This chapter focuses on limiting use and abuse, the number of connections, the rate at which requests are served, and the amount of bandwidth used. It's important to differentiate between connections and requests: connections (TCP connections) are the transport layer on which requests are made and therefore are not the same thing. A browser may open multiple connections to a server to make multiple requests. However, in HTTP/1 and HTTP/1.1, requests can only be made one at a time on a single connection; whereas in HTTP/2, multiple requests can be made in parallel over a single TCP connection. This chapter will help you restrict usage of your service and mitigate abuse.

12.1 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 megabytes. 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 of 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 on which to limit by. 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.

12.2 Limiting Rate

Problem

You need to limit the rate of requests by 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=1r/s;
    limit_req_status 429;
    ...
    server {
        ...
        limit_req zone=limitbyaddr burst=10 nodelay;
        ...
    }
}
```

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 megabytes. The zone sets the rate with a keyword argument. The `limit_req` directive takes two optional keyword arguments: `zone` and `burst`. `zone` is required to instruct the directive on which shared memory request limit zone to use. When the request rate for a given zone is exceeded, requests are delayed until their maximum burst size is reached, denoted by the `burst` keyword argument. The `burst` keyword argument defaults to zero. `limit_req` also takes a third optional parameter, `nodelay`. This parameter enables the client to use its burst without delay before being limited. `limit_req_status` sets the status returned to the client to a particular HTTP status code; the default is 503. `limit_req_status` and `limit_req` are valid in the context of HTTP, server, and location. `limit_req_zone` is only valid in the HTTP context.

Discussion

The rate-limiting module is very powerful in 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 disable the plans of malicious users that might try to deny service to your application or to waste resources by setting a sane limit on all requests. The configuration of the rate-limit module is much like the preceding connection-limiting module described in [Recipe 12.1](#), and much of the same concerns apply. The rate at which requests are limited can be done in requests per second or requests per minute. When the rate limit is hit, the incident is logged. There's a directive not in the example: `limit_req_log_level`, which defaults to error, but can be set to info, notice, or warn.

12.3 Limiting Bandwidth

Problem

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

Solution

Utilize NGINX's `limit_rate` and `limit_rate_after` directives to limit the rate 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 megabytes to a rate of 1 megabyte 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 con-

texts as `limit_rate_after`; however, it can alternatively be set by setting 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.

Encrypting

13.0 Introduction

The internet can be a scary place, but it doesn't have to be. Encryption for information in transit has become easier and more attainable in that signed certificates have become less costly with the advent of Let's Encrypt and Amazon Web Services. Both offer free certificates with limited usage. With free signed certificates, there's little standing in the way of protecting sensitive information. While not all certificates are created equal, any protection is better than none. In this chapter, we discuss how to secure information between NGINX and the client, as well as NGINX and upstream services.

13.1 Client-Side Encryption

Problem

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

Solution

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

```

http { # All directives used below are also valid in stream
    server {
        listen 8443 ssl;
        ssl_protocols      TLSv1.2;
        ssl_ciphers         HIGH:!aNULL:!MD5;
        ssl_certificate      /usr/local/nginx/conf/cert.pem;
        ssl_certificate_key  /usr/local/nginx/conf/cert.key;
        ssl_session_cache    shared:SSL:10m;
        ssl_session_timeout  10m;
    }
}

```

This configuration sets up a server to listen on a port encrypted with SSL, 8443. The server accepts the SSL protocol version TLSv1.2. The SSL certificate and key locations are disclosed to the server for use. The server is instructed to use the highest strength offered by the client while restricting a few that are insecure. The SSL session cache and timeout allow for workers to cache and store session parameters for a given amount of time. There are many other session cache options that can help with performance or security of all types of use cases. Session cache options can be used in conjunction. However, specifying one without the default will turn off that default, built-in session cache.

Discussion

Secure transport layers are the most common way of encrypting information in transit. At the time of writing, the *Transport Layer Security* protocol (TLS) is the default over the *Secure Socket Layer* (SSL) protocol. That's because versions 1 through 3 of SSL are now considered insecure. While the protocol name may 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 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 chain in the file. If your certificate authority has provided many files in the chain, it is also able to provide the order in which they are layered. The SSL session cache enhances performance by not having to negotiate for SSL/TLS versions and ciphers.

Also See

[Mozilla Server Side TLS Page](#)

[Mozilla SSL Configuration Generator](#)

[Test your SSL Configuration with SSL Labs SSL Test](#)

13.2 Upstream Encryption

Problem

You need to encrypt traffic between NGINX and the upstream service and set specific negotiation rules for compliance regulations or if the upstream 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 available, 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.

HTTP Basic Authentication

14.0 Introduction

Basic authentication is a simple way to protect private content. This method of authentication can be used to easily hide development sites or keep privileged content hidden. Basic authentication is pretty unsophisticated, not extremely secure, and, therefore, should be used with other layers to prevent abuse. It's recommended to set up a rate limit on locations or servers that require basic authentication to hinder the rate of brute force attacks. It's also recommended to utilize HTTPS, as described in [Chapter 13](#), whenever possible, as the username and password are passed as a base64-encoded string to the server in a header on every authenticated request. The implications of basic authentication over an unsecured protocol such as HTTP means that the username and password can be captured by any machine the request passes through.

14.1 Creating a User File

Problem

You need an HTTP basic authentication user file to store usernames and passwords.

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. An optional third field can be used for comment on each user. NGINX can understand a few different formats for passwords, one of which is if the password is 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 with the following command:

```
$ openssl passwd MyPassword1234
```

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

Discussion

Basic authentication passwords can be generated a few ways and in a few different formats to 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 LDAP and Dovecot use. NGINX supports more formats and hashing algorithms, however, many of them are considered insecure because they can be easily brute-forced.

14.2 Using Basic Authentication

Problem

You need basic authentication to protect an NGINX location or server.

Solution

Use the `auth_basic` and `auth_basic_user_file` directives to enable basic authentication:

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

The `auth_basic` directives can be used 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, which was just described in [Recipe 14.1](#).

Discussion

Basic authentication can be used 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 a response header `WWW-Authenticate`. This header will have a value of `Basic realm="your string"`. This response will cause the browser to prompt for a username and password. The username and password are concatenated and delimited with a colon, then base64 encoded, and 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 `auth_basic_user_file` provided. Because the username password string is merely base64 encoded, it's recommended to use HTTPS with basic authentication.

HTTP Authentication Subrequests

15.0 Introduction

With many different approaches to authentication, NGINX makes it easy to validate against a wide range of authentication systems by enabling a subrequest mid-flight to validate identity. The HTTP authentication request module is meant to enable authentication systems like LDAP or custom authentication microservices. The authentication mechanism proxies the request to the authentication service before the request is fulfilled. During this proxy you have the power of NGINX to manipulate the request as the authentication service requires. Therefore, it is extremely flexible.

15.1 Authentication Subrequests

Problem

You have a third-party authentication system to 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 makes a subrequest before serving the original to determine if the request has access to the resource it's requesting. The entire original request is proxied to this subrequest location. The authentication location acts as a typical proxy to the subrequest and sends the original request, including the original request body and headers. 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.

Secure Links

16.0 Introduction

Secure links are a way to keep static assets protected with the `md5` hashing algorithm. With this module, you can also put a limit on the length of time for which the link is accepted. Using secure links enables your NGINX application server to serve static content securely while taking this responsibility off of the application server. This module is included in the free and open source NGINX. However, it is not built into the standard NGINX package but instead the `nginx-extras` package. Alternatively, it can be enabled with the `--with-http_secure_link_module` configuration parameter when building NGINX from source.

16.1 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 as 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.

16.2 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 16.1](#), we will create the secured link that will work with the previous configuration example given 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:

```
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 for `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.

16.3 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$remote_addr
mySecret";
    if ($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 16.1](#). 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 around what the `$secure_link` variable is set to returns known HTTP codes for Forbidden and Gone. The HTTP 410, Gone, works great for expired links as the condition is to be considered permanent.

16.4 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 "2020-12-31 00:00" +%s --utc
1609372800
```

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 `1293771600/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 '1609372800/resources/index.html127.0.0.1 mySecret' \
| openssl md5 -binary \
| openssl base64 \
| tr +/ -_ \
| tr -d =
TG6ck30pAttQ1d7jW3J0cw
```

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

```
/resources/index.html?md5=TG6ck30pAttQ1d7jW3J0cw&expires=1609372800'
```

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:

```

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 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.

API Authentication Using JWT

17.0 Introduction

JSON Web Tokens (JWTs) are quickly becoming a widely used and preferred authentication method. These authentication tokens have the ability to store some information about the user as well as information about the user's authorization into the token itself. These tokens can also be validated asymmetrically, which means load balancers and proxies are able to validate the token with a public key and do not need the private key that the token was signed with, thus enhancing security and flexibility. An advantage of offloading authentication verification to your NGINX Plus layer is that you're saving cycles on your authentication service, as well as speeding up your transactions. The JWT authentication module described in this chapter is available only with an NGINX Plus subscription.

17.1 Validating JWTs

Problem

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

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` 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 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 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 of making 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.

Also See

- [RFC standard documentation of JSON Web Signature](#)
- [RFC standard documentation of JSON Web Algorithms](#)
- [RFC standard documentation of JSON Web Token](#)
- [NGINX embedded variables](#)
- [Detailed NGINX blog](#)

17.2 Creating JSON Web Keys

Problem

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

Solution

NGINX Plus utilizes the JSON Web Key (JWK) format as specified in the RFC standard. The 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:

```
{ "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 to 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 JSON Web Key. 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 standardization documentation of JSON Web Key

OpenId Connect Single Sign-On

18.0 Introduction

Single sign-on (SSO) authentication providers are a great way to reduce authentication requests to your application and provide your users with seamless integration into an application they already log in to on a regular basis. As more authentication providers bring themselves to market, your application can be ready to integrate by using NGINX Plus to validate the signature of their JSON Web Tokens. In this chapter we'll explore using the NGINX Plus JWT authentication module for HTTP in conjunction with an existing OpenId Connect OAuth 2.0 provider from Google. As in [Chapter 17](#), this chapter describes the JWT authentication module, which is only available with an NGINX Plus subscription.

18.1 Authenticate Users via Existing OpenId Connect Single Sign-On (SSO)

Problem

You want to offload OpenId Connect authentication validation to NGINX Plus.

Solution

Use NGINX Plus's JWT module to secure a location or server and tell 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 tells 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, we must tell NGINX to look for the token in this cookie rather than the NGINX Plus Default location. The `auth_jwt_key_file` location is set to an arbitrary path, a step that we will cover in [Recipe 18.2](#).

Discussion

This configuration demonstrates how you can validate a Google OAuth 2.0 OpenId Connect JSON Web Token with NGINX Plus. The NGINX Plus JWT authentication module for HTTP is able to validate any JSON Web Token that adheres to the RFC for JSON Web Signature specification, instantly enabling any single sign-on authority that utilizes JSON Web Tokens 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 JSON Web Tokens 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

[Detailed NGINX Blog on OpenId Connect](#)
[OpenId Connect](#)

18.2 Obtaining JSON Web Key from Google

Problem

You need to obtain the JSON Web Key from Google to use when validating OpenId Connect tokens with NGINX Plus.

Solution

Utilize Cron to request a fresh set of keys every hour to ensure your keys are always up-to-date:

```
0 * * * * root wget https://www.googleapis.com/oauth2/v3/ \
certs-0 /etc/nginx/google_certs.jwk
```

This code snippet is a line from a crontab file. Unix-like systems have many options for where crontab files can live. Every user will have a user-specific crontab, and there's also a number of files and directories in the */etc/* directory.

Discussion

Cron is a common way to run a scheduled task on a Unix-like system. JSON Web Keys should be rotated on a regular interval to ensure the security of the key, and in turn, the security of your system. To ensure that you always have the most up-to-date key from Google, you'll want to check for new JWKs at a regular interval. This cron solution is one way of doing so.

Also See

[Cron](#)

ModSecurity Web Application Firewall

19.0 Introduction

ModSecurity is an open source web application firewall (WAF) that was first built for Apache web servers. It was made available to NGINX as a module in 2012 and added as an optional feature to NGINX Plus in 2016. This chapter will detail installing ModSecurity 3.0 with NGINX Plus through dynamic modules. It will also cover compiling and installing the ModSecurity 2.9 module and NGINX from source. ModSecurity 3.0 with NGINX Plus is far superior to ModSecurity 2.x in terms of security and performance. When running ModSecurity 2.9 configured from open source, it's still wrapped in Apache and, therefore, requires much more overhead than 3.0, which was designed for NGINX natively. The plug-and-play ModSecurity 3.0 module for NGINX is only available with an NGINX Plus subscription.

19.1 Installing ModSecurity for NGINX Plus

Problem

You need to install the ModSecurity module for NGINX Plus.

Solution

Install the module from the NGINX Plus repository. The package name is `nginx-plus-module-modsecurity`. On an Ubuntu-based system, you can install NGINX Plus and the ModSecurity module through the advanced packaging tool, also known as `apt-get`:

```
$ apt-get update
$ apt-get install nginx-plus
$ apt-get install nginx-plus-module-modsecurity
```

Discussion

Installing NGINX Plus and the ModSecurity module is as easy as pulling it from the NGINX Plus repository. Your package management tool, such as `apt-get` or `yum`, will install NGINX Plus as well as the module and place the module in the *modules* directory within the default NGINX Plus configuration directory */etc/nginx/*.

19.2 Configuring ModSecurity in NGINX Plus

Problem

You need to configure NGINX Plus to use the ModSecurity module.

Solution

Enable the dynamic module in your NGINX Plus configuration, and use the `modsecurity_rules_file` directive to point to a ModSecurity rule file:

```
load_module modules/ngx_http_modsecurity.so;
```

The `load_module` directive is applicable in the main context, which means that this directive is to be used before opening the HTTP or Stream blocks.

Turn on ModSecurity and use a particular rule set:

```
modsecurity on;
location / {
    proxy_pass http://backend;
    modsecurity_rules_file rule-set-file;
}
```

The `modsecurity` directive turns on the module for the given context when passed the `on` parameter. The `modsecurity_rules_file`

directive instructs NGINX Plus to use a particular ModSecurity rule set.

Discussion

The rules for ModSecurity can prevent common exploits of web servers and applications. ModSecurity is known to be able to prevent application-layer attacks such as HTTP violations, SQL injection, cross-site scripting, distributed-denial-of-service, and remote and local file-inclusion attacks. With ModSecurity, you're able to subscribe to real-time blacklists of malicious user IPs to help block issues before your services are affected. The ModSecurity module also enables detailed logging to help identify new patterns and anomalies.

Also See

[OWASP ModSecurity Core Rule Set](#)

[TrustWave ModSecurity Paid Rule Set](#)

19.3 Installing ModSecurity from Source for a Web Application Firewall

Problem

You need to run a web application firewall with NGINX using ModSecurity and a set of ModSecurity rules on a CentOS or RHEL-based system.

Solution

Compile ModSecurity and NGINX from source and configure NGINX to use the ModSecurity module.

First update security and install prerequisites:

```
$ yum --security update -y && \  
yum -y install automake \  
autoconf \  
curl \  
curl-devel \  
gcc \  
gcc-c++ \  
httpd-devel \  
libxml2 \
```

```
libxml2-devel \
make \
openssl \
openssl-devel \
perl \
wget
```

Next, download and install PERL 5 regular expression pattern matching:

```
$ cd /opt && \
  wget http://ftp.exim.org/pub/pcre/pcre-8.39.tar.gz && \
  tar -zxf pcre-8.39.tar.gz && \
  cd pcre-8.39 && \
  ./configure && \
  make && \
  make install
```

Download and install zlib from source:

```
$ cd /opt && \
  wget http://zlib.net/zlib-1.2.8.tar.gz && \
  tar -zxf zlib-1.2.8.tar.gz && \
  cd zlib-1.2.8 && \
  ./configure && \
  make && \
  make install
```

Download and install ModSecurity from source:

```
$ cd /opt && \
  wget \
  https://www.modsecurity.org/tarball/2.9.1/modsecurity-2.9.1.\
  tar.gz&& \
  tar -zxf modsecurity-2.9.1.tar.gz && \
  cd modsecurity-2.9.1 && \
  ./configure --enable-standalone-module && \
  make
```

Download and install NGINX from source and include any modules you may need with the configure script. Our focus here is the ModSecurity module:

```
$ cd /opt && \
  wget http://nginx.org/download/nginx-1.11.4.tar.gz && \
  tar zxf nginx-1.11.4.tar.gz && \
  cd nginx-1.11.4 && \
  ./configure \
    --sbin-path=/usr/local/nginx/nginx \
    --conf-path=/etc/nginx/nginx.conf \
    --pid-path=/usr/local/nginx/nginx.pid \
    --with-pcre=../pcre-8.39 \
    --with-zlib=../zlib-1.2.8 \
```

```

--with-http_ssl_module \
--with-stream \
--with-http_ssl_module \
--with-http_secure_link_module \
--add-module=../modsecurity-2.9.1/nginx/modsecurity \
&& \
make && \
make install && \
ln -s /usr/local/nginx/nginx /usr/sbin/nginx

```

This will yield NGINX compiled from source with the ModSecurity version 2.9.1 module installed. From here we are able to use the `ModSecurityEnabled` and `ModSecurityConfig` directives in our configurations:

```

server {
    listen 80 default_server;
    listen [::]:80 default_server;
    server_name _;
    location / {
        ModSecurityEnabled on;
        ModSecurityConfig modsecurity.conf;
    }
}

```

This configuration for an NGINX server turns on ModSecurity for the `location /` and uses a ModSecurity configuration file located at the base of the NGINX configuration.

Discussion

This section compiles NGINX from source with the ModSecurity for NGINX. It's advised when compiling NGINX from source to always check that you're using the latest stable packages available. With the preceding example, you can use the open source version of NGINX along with ModSecurity to build your own open source web application firewall.

Also See

[ModSecurity Source](#)

[Updated and maintained ModSecurity Rules from SpiderLabs](#)

Practical Security Tips

20.0 Introduction

Security is done in layers, and much like an onion, there must be multiple layers to your security model for it to be truly hardened. In Part II of this book, we've gone through many different ways to secure your web applications with NGINX and NGINX Plus. Many of these security methods can be used in conjunction to help harden security. The following are a few more practical security tips to ensure your users are using HTTPS and to tell NGINX to satisfy one or more security methods.

20.1 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`.

20.2 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 standard `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, which will offload SSL/TLS at no additional cost. This is a handy trick to make sure that your HTTP traffic is secured.

20.3 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, therefore the request is never sent unencrypted.

Also See

[RFC-6797 HTTP Strict Transport Security](#)

[OWASP HSTS Cheat Sheet](#)

20.4 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 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 [Chapter 11](#), the `http_auth_basic_module` detailed in [Chapter 14](#), the `http_auth_request_module` detailed in [Chapter 15](#), and the `http_auth_jwt_module` detailed in [Chapter 17](#). 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.

Part III: Deployment and Operations

This is the third and final part of the *NGINX Cookbook*. This part will focus on deployment and operations of NGINX and NGINX Plus, the licensed version of the server. Throughout this part, you will learn about deploying NGINX to Amazon Web Services, Microsoft Azure, and Google Cloud Compute, as well as working with NGINX in Docker containers. This part will dig into using configuration management to provision NGINX servers with tools such as Puppet, Chef, Ansible, and SaltStack. It will also get into automating with NGINX Plus through the NGINX Plus API for on-the-fly reconfiguration and using Consul for service discovery and configuration templating. We'll use an NGINX module to conduct A/B testing and acceptance during deployments. Other topics covered are using NGINX's GeoIP module to discover the geographical origin of our clients, including it in our logs, and using it in our logic. You'll learn how to format access logs and set log levels of error logging for debugging. Through a deep look at performance, this part will provide you with practical tips for optimizing your NGINX configuration to serve more requests faster. It will help you install, monitor, and maintain the NGINX application delivery platform.

Deploying on AWS

21.0 Introduction

Amazon Web Services (AWS), in many opinions, has led the cloud infrastructure landscape since the arrival of S3 and EC2 in 2006. AWS provides a plethora of *infrastructure-as-a-service* (IaaS) and *platform-as-a-service* (PaaS) solutions. Infrastructure as a service, such as Amazon EC2 or Elastic Cloud Compute, is a service providing virtual machines in as little as a click or API call. This chapter will cover deploying NGINX into an Amazon Web Service environment, as well as some common patterns.

21.1 Auto-Provisioning on AWS

Problem

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

Solution

Utilize EC2 UserData as well as a pre-baked Amazon Machine Image. Create an Amazon Machine Image with NGINX and any supporting software packages installed. Utilize Amazon EC2 UserData to configure any environment-specific configurations at runtime.

Discussion

There are three patterns of thought when provisioning on Amazon Web Services:

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 Amazon Machine Images (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 Ansible, Chef, Puppet, and SaltStack, which will be described in [Chapter 25](#).

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 on the platform of your choosing, SSH's into the virtual machine, runs any provisioning you specify, and burns an image. You can utilize Packer to run the con-

figuration 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. The `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.

21.2 Routing to NGINX Nodes Without an 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 Route53 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 Route53 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 Route53 is its ability to health check. You can configure Route53 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, Route53 can take an IP out of rotation if it begins to fail. You could also configure Route53 to failover to a secondary record in case of a failure, which would achieve an active-passive, highly available setup.

Route53 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 Route53 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 Route53 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 CLI or SDK to interface with the Amazon Route53 API to add or remove the NGINX machine IP and configured health check as it boots or before it is terminated.

21.3 The ELB Sandwich

Problem

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

Solution

Create an *elastic load balancer* (ELB) or two. Create an Auto Scaling group with a launch configuration that provisions an EC2 instance with NGINX installed. The Auto Scaling group has a configuration to link to the elastic load balancer, which will automatically register any instance in the Auto Scaling group to the load balancers configured on first boot. Place your upstream applications behind another elastic load balancer and configure NGINX to proxy to that ELB.

Discussion

This common pattern is called the ELB sandwich (see [Figure 21-1](#)), putting NGINX in an Auto Scaling group behind an ELB and the application Auto Scaling group behind another ELB. The reason for having ELBs between every layer is because the ELB works so well with Auto Scaling groups; they automatically register new nodes and remove ones being terminated, as well as run health checks and only pass traffic to healthy nodes. The reason behind building a second ELB for NGINX is because it allows services within your application to call out to other services through the NGINX Auto Scaling group without leaving the network and reentering through the public ELB. This puts NGINX in the middle of all network traffic within your application, making it the heart of your application's traffic routing.

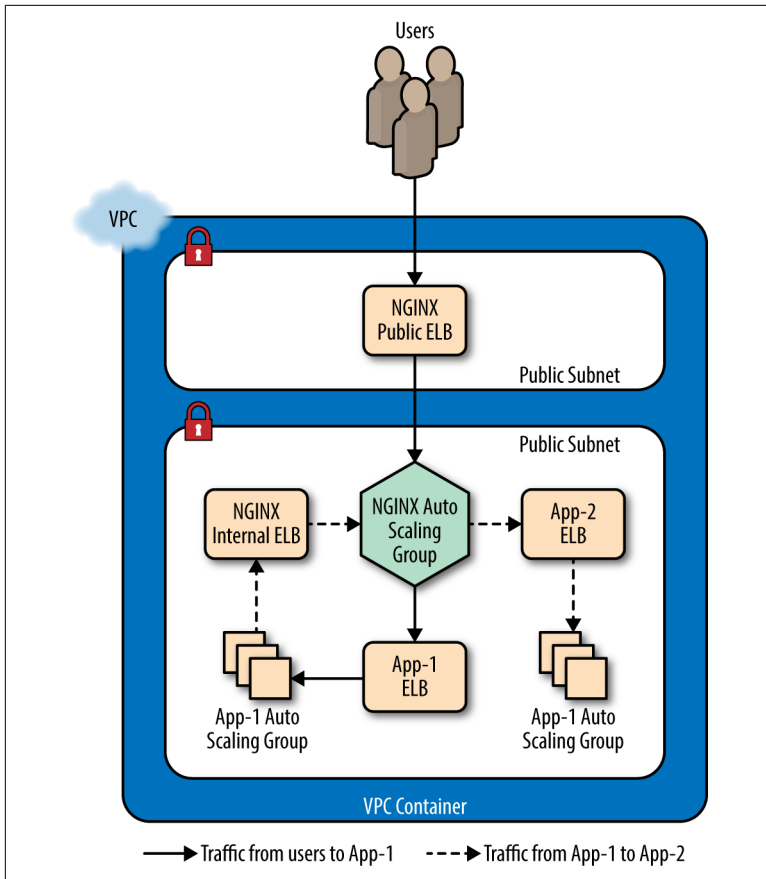


Figure 21-1. This image depicts NGINX in an ELB sandwich pattern with an internal ELB for internal applications to utilize. A user makes a request to App-1, and App-1 makes a request to App-2 through NGINX to fulfill the user's request.

21.4 Deploying from the 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 21-2](#)). Select the Amazon Machine Image (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.

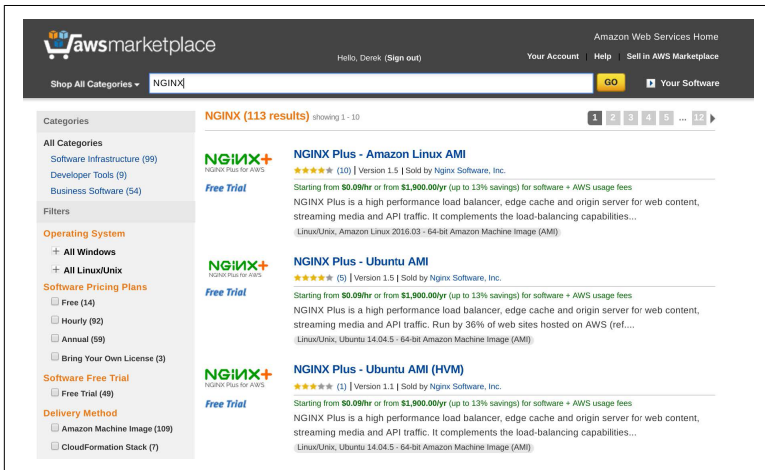


Figure 21-2. The AWS Marketplace after searching for NGINX

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 easily 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.

Deploying on Azure

22.0 Introduction

Azure is a powerful cloud platform offering from Microsoft. Azure enables cross-platform virtual machine hosting inside of virtual cloud networks. NGINX is an amazing application delivery platform for any OS or application type and works seamlessly in Microsoft Azure. NGINX has provided a pay-per-usage NGINX Plus Market-place offering, which this chapter will explain how to use, making it easy to get up and running quickly with on-demand licensing in Microsoft Azure.

22.1 Creating an NGINX Virtual Machine Image

Problem

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

Solution

Create a virtual machine from a base 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 virtual machine image. To create a virtual machine image, you must first generalize

the VM. To generalize your virtual machine, 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 virtual machine. 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 virtual machine:

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

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

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

After your virtual machine 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 virtual machine 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 tem-

plate 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 of your choosing for the virtual machine name, 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 virtual machine from the custom NGINX or NGINX Plus image you've created. After entering the necessary parameters, Azure will begin to create a new virtual machine 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. Azure creating an ARM template enables you to quickly and reliably deploy this same server time and time again as needed. With the virtual machine image path that can be found in the template, you can use this image to create different sets of infrastructure such as virtual machine scaling sets or other VMs with different configurations.

Also See

[Installing Azure cross-platform CLI](#)

[Azure cross-platform CLI login](#)

[Capturing Linux virtual machine images](#)

22.2 Load Balancing Over NGINX Scale Sets

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 virtual machine image created in the prior section or the NGINX Plus image from the Marketplace described in [Recipe 22.3](#) 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 virtual machine scaling sets. 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.

22.3 Deploying Through the Marketplace

Problem

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

Solution

Deploy an NGINX Plus virtual machine 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 virtual machine, 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 virtual machine 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 virtual machine 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 virtual machine 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.

Deploying on Google Cloud Compute

23.0 Introduction

Google Cloud Compute is an advanced cloud platform that enables its customers to build diverse, high-performing web applications at will on hardware they provide and manage. Google Cloud Compute offers virtual networking and machines, a tried-and-true platform-as-a-service (PaaS) offering, as well as many other managed service offerings such as Bigtable, BigQuery, and SQL. In this chapter, we will discuss how to deploy NGINX servers to Google Cloud Compute, how to create virtual machine images, and how and why you might want to use NGINX to serve your Google App Engine applications.

23.1 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 virtual machine in Google Compute Engine. Select a name for your virtual machine, zone, machine type, and boot disk.

Configure identity and access management, firewall, and any advanced configuration you'd like. Create the virtual machine.

Once the virtual machine 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 virtual machine.

Discussion

Google Compute Engine offers highly configurable virtual machines at a moment's notice. Starting a virtual machine 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.

23.2 Creating a Google Compute Image

Problem

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

Solution

Create a virtual machine as described in the previous section. After installing and configuring NGINX on your virtual machine instance, set the auto-delete state of the boot disk to `false`. To set the auto-delete state of the disk, edit the virtual machine. 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 virtual machine 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; and 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 virtual machines 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 depreciate private images

23.3 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 virtual machine 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. 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 would 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 common 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.

Deploying on Docker

24.0 Introduction

Docker is an open source project that automates the deployment of Linux applications inside software containers. Docker provides an additional layer of abstraction and automation of operating-system-level virtualization on Linux. Containerized environments have made a huge break into the production world, and I'm excited about it. Docker and other container platforms enable fast, reliable, cross-platform application deployments. In this chapter we'll discuss the official NGINX Docker image, creating your own Dockerfile to run NGINX, and using environment variables within NGINX, a common Docker practice.

24.1 Running Quickly with the 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, so you'll need to either mount a local configuration directory or create a Dockerfile and ADD in your configuration to the image build. We'll mount a volume and get NGINX running in a Docker container locally in two commands:

```
$ docker pull nginx:latest
$ docker run -it -p 80:80 -v $PWD/nginx-conf:/etc/nginx \
    nginx:latest
```

The first docker command pulls the `nginx:latest` image from Docker Hub. The second docker command runs this NGINX image as a Docker container in the foreground, mapping `localhost:80` to port `80` of the NGINX container. It also mounts the local directory `nginx-conf` as a container volume at `/etc/nginx`. `nginx-conf` is a local directory that contains the necessary files for NGINX configuration. 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 only two commands! The official NGINX Docker image mainline that we used in this example is built off of the Debian Jessie Docker image. However, you can choose official images built off of Alpine Linux. The Dockerfile and source for these official images are available on GitHub.

Also See

Official NGINX Docker image, NGINX
Docker repo on GitHub

24.2 Creating an NGINX Dockerfile

Problem

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

Solution

Start FROM 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 available to you more easily based on the log driver you've chosen to use with Docker:

Dockerfile:

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 choose 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.

24.3 Building an NGINX Plus Image

Problem

You need to build an NGINX Plus Docker image to run NGINX Plus in a containerized environment.

Solution

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

Ubuntu:

```
FROM ubuntu:14.04

MAINTAINER NGINX Docker Maintainers "docker-maint@nginx.com"

# Set the debconf frontend to Noninteractive
RUN echo 'debconf debconf/frontend select Noninteractive' \
    | debconf-set-selections

RUN apt-get update && apt-get install -y -q wget \
    apt-transport-https lsb-release ca-certificates

# Download certificate and key from the customer portal
# (https://cs.nginx.com) and copy to the build context
ADD nginx-repo.crt /etc/ssl/nginx/
ADD nginx-repo.key /etc/ssl/nginx/

# Get other files required for installation
RUN wget -q -O - http://nginx.org/keys/nginx_signing.key \
    | apt-key add -
RUN wget -q -O /etc/apt/apt.conf.d/90nginx \
    https://cs.nginx.com/static/files/90nginx
```

```

RUN printf "deb https://plus-pkgs.nginx.com/ubuntu \
`lsb_release -cs` nginx-plus\n" \
>/etc/apt/sources.list.d/nginx-plus.list

# Install NGINX Plus
RUN apt-get update && apt-get install -y nginx-plus

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

EXPOSE 80 443

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

```

CentOS 7:

```

FROM centos:centos7

MAINTAINER NGINX Docker Maintainers "docker-maint@nginx.com"

RUN yum install -y ca-certificates

# Download certificate and key from the customer portal
# (https://cs.nginx.com) and copy to the build context
ADD nginx-repo.crt /etc/ssl/nginx/
ADD nginx-repo.key /etc/ssl/nginx/

# Get other files required for installation
RUN wget -q -O /etc/yum.repos.d/nginx-plus-7.repo \
https://cs.nginx.com/static/files/nginx-plus-7.repo

# Install NGINX Plus
RUN yum install -y nginx-plus

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

EXPOSE 80 443

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

```

To build these Dockerfiles into Docker images, 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. These Dockerfiles do not use the Dockerfile property `ADD` to add in configuration; you will need to add in your configuration manually.

Also See

[NGINX blog on Docker images](#)

24.4 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.

Using Puppet/Chef/ Ansible/SaltStack

25.0 Introduction

Configuration management tools have been an invaluable utility in the age of the cloud. Engineers of large-scale web applications are no longer configuring servers by hand but rather by code, using one of the many configuration management tools available. Configuration management tools enable engineers to write configurations and code one time to produce many servers with the same configuration in a repeatable, testable, and modular fashion. In this chapter we'll discuss 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.

25.1 Installing with Puppet

Problem

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

Solution

Create a module that installs NGINX, manages the files you need, and ensures that NGINX is running:

```
class nginx {
  package {"nginx": ensure => 'installed',}
  service {"nginx":
    ensure => 'true',
    hasrestart => 'true',
    restart => '/etc/init.d/nginx reload',
  }
  file { "nginx.conf":
    path    => '/etc/nginx/nginx.conf',
    require => Package['nginx'],
    notify  => Service['nginx'],
    content => template('nginx/templates/nginx.conf.erb'),
    user=>'root',
    group=>'root',
    mode='0644';
  }
}
```

This module uses the package management utility to ensure the NGINX package is installed. It also ensures NGINX is running and enabled at boot time. The configuration informs Puppet that the service does have a restart command with the `hasrestart` directive, and we override the `restart` command with an NGINX reload. It will manage and template the *nginx.conf* file with the Embedded Ruby (ERB) templating language. The templating of the file will happen after the NGINX package is installed due to the `require` directive. However, it will notify the NGINX service to reload because of the `notify` directive. The templated configuration file is not included. However, it can be simple to install a default NGINX configuration file, or very complex if using Embedded Ruby (ERB) or Embedded Puppet (EPP) templating language loops and variable substitution.

Discussion

Puppet is a configuration management tool based in the Ruby programming language. Modules are built in a domain-specific language and called via a manifest file that defines the configuration for a given server. Puppet can be run in a master-slave or masterless configuration. With Puppet, the manifest is run on the master and then sent to the slave. This is important because it ensures that the

slave is only delivered the configuration meant for it and no extra configurations meant for other servers. There are a lot of extremely advanced public modules available for Puppet. Starting from these modules will help you get a jump-start on your configuration. A public NGINX module from voxpupuli on GitHub will template out NGINX configurations for you.

Also See

[Puppet documentation](#)

[Puppet package documentation](#)

[Puppet service documentation](#)

[Puppet file documentation](#)

[Puppet templating documentation](#)

[Voxpupuli NGINX module](#)

25.2 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

Create a cookbook with a recipe to install NGINX and configure configuration files through templating, and ensure NGINX reloaded after the configuration is put in place. The following is an example recipe:

```
package 'nginx' do
  action :install
end

service 'nginx' do
  supports :status => true, :restart => true, :reload => true
  action [ :start, :enable ]
end

template 'nginx.conf' do
  path   "/etc/nginx.conf"
  source "nginx.conf.erb"
  owner  'root'
  group  'root'
```

```
mode '0644'  
  notifies :reload, 'service[nginx]', :delayed  
end
```

The `package` block installs NGINX. The `service` block ensures that NGINX is started and enabled at boot, then declares to the rest of Chef what the `nginx` service will support as far as actions. The `template` block templates an ERB file and places it at `/etc/nginx.conf` with an owner and group of `root`. The `template` block also sets the mode to 644 and notifies the `nginx` service to `reload`, but waits until the end of the Chef run declared by the `:delayed` statement. The templated configuration file is not included. However, it can be as simple as a default NGINX configuration file or very complex with Embedded Ruby (ERB) templating language loops and variable substitution.

Discussion

Chef is a configuration management tool based in Ruby. Chef can be run in a master-slave, or solo configuration, now known as Chef Zero. Chef has a very large community and many public cookbooks, called the Supermarket. Public cookbooks from the Supermarket can be installed and maintained via a command-line utility called Berkshelf. Chef is extremely capable, and what we have demonstrated is just a small sample. The public NGINX cookbook for NGINX 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.

Also See

- [Chef documentation](#)
- [Chef package](#)
- [Chef service](#)
- [Chef template](#)
- [Chef Supermarket for NGINX](#)

25.3 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

Create an Ansible playbook to install NGINX and manage the *nginx.conf* file. The following is an example task file for the playbook to install NGINX. Ensure it's running and template the configuration file:

```
- name: NGINX | Installing NGINX
  package: name=nginx state=present

- name: NGINX | Starting NGINX
  service:
    name: nginx
    state: started
    enabled: yes

- name: Copy nginx configuration in place.
  template:
    src: nginx.conf.j2
    dest: "/etc/nginx/nginx.conf"
    owner: root
    group: root
    mode: 0644
  notify:
    - reload nginx
```

The `package` block installs NGINX. The `service` block ensures that NGINX is started and enabled at boot. The `template` block templates a *Jinja2* file and places the result at */etc/nginx.conf* with an owner and group of root. The `template` block also sets the mode to 644 and notifies the `nginx` service to reload. The templated configuration file is not included. However, it can be as simple as a default NGINX configuration file or very complex with Jinja2 templating language loops and variable substitution.

Discussion

Ansible is a widely used and powerful configuration management tool based in Python. The configuration of tasks is in YAML, and you use the Jinja2 templating language for file templating. Ansible offers a master named Ansible Tower on a subscription model. However, it's commonly used from local machines or build servers directly to the client or in a masterless model. Ansible bulk SSH's into its servers and runs the configurations. 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.

Also See

[Ansible documentation](#)

[Ansible packages](#)

[Ansible service](#)

[Ansible template](#)

[Ansible Galaxy](#)

25.4 Installing with SaltStack

Problem

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

Solution

Install NGINX through the package management module and manage the configuration files you desire. The following is an example state file (*sls*) that will install the `nginx` package and ensure the service is running, enabled at boot, and reloads if a change is made to the configuration file:

```
nginx:
  pkg:
    - installed
  service:
    - name: nginx
    - running
    - enable: True
```

```

- reload: True
- watch:
  - file: /etc/nginx/nginx.conf

/etc/nginx/nginx.conf:
  file:
    - managed
    - source: salt://path/to/nginx.conf
    - user: root
    - group: root
    - template: jinja
    - mode: 644
    - require:
      - pkg: nginx

```

This is a basic example of installing NGINX via a package management utility and managing the *nginx.conf* file. The NGINX package is installed and the service is running and enabled at boot. With SaltStack you can declare a file managed by Salt as seen in the example and templated by many different templating languages. The templated configuration file is not included. However, it can be as simple as a default NGINX configuration file or very complex with the Jinja2 templating language loops and variable substitution. This configuration also specifies that NGINX must be installed prior to managing the file because of the `require` statement. After the file is in place, NGINX is reloaded because of the `watch` directive on the service and reloads as opposed to restarts because the `reload` directive is set to `True`.

Discussion

SaltStack is a powerful configuration management tool that defines server states in YAML. Modules for SaltStack can be written in Python. Salt exposes the Jinja2 templating language for states as well as for files. However, for files there are many other options, such as Mako, Python itself, and others. Salt works in a master-slave configuration as well as a masterless configuration. Slaves are called minions. The master-slave transport communication, however, differs from others and sets SaltStack apart. With Salt you're able to choose ZeroMQ, TCP, or Reliable Asynchronous Event Transport (RAET) for transmissions to the Salt agent; or you can not use an agent, and the master can SSH instead. Because the transport layer is by default asynchronous, SaltStack is built to be able to deliver its message to a large number of minions with low load to the master server.

Also See

SaltStack

Installed packages

Managed files

Templating with Jinja

Automation

26.0 Introduction

There are many ways to automate NGINX and NGINX Plus configuration files, such as rerunning your configuration management tools or cron jobs that retemplate configuration files. As dynamic environments increase in popularity and necessity, the need for configuration automation becomes more relevant. In [Chapter 25](#), we made sure that NGINX was reloaded after the configuration file was templated. In this chapter, we'll discuss further on-the-fly reconfiguration of NGINX configuration files with the NGINX Plus API and Consul Template.

26.1 Automating with NGINX Plus

Problem

You need to reconfigure NGINX Plus on the fly to load balance for a dynamic environment.

Solution

Use the NGINX Plus API to reconfigure NGINX Plus upstream pools:

```
$ curl 'http://nginx.local/upstream_conf?\  
add=&upstream=backend&server=10.0.0.42:8080'
```

The `curl` call demonstrated makes a request to NGINX Plus and requests a new server be added to the backend upstream configuration.

Discussion

Covered in great detail in [Chapter 8](#) of Part I, NGINX Plus offers an API to reconfigure NGINX Plus on the fly. The NGINX Plus API enables adding and removing servers from upstream pools as well as draining connections. You can use this API to automate your NGINX Plus configuration as application servers spawn and release in the environment.

26.2 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 of a Consul Template file that templates an upstream configuration block. This template will loop through nodes in Consul identifying 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 consul.example.internal -template \
  template:/etc/nginx/conf.d/upstream.conf:"nginx -s reload"
```

The command demonstrated instructs the `consul-template` daemon to connect to a Consul cluster at `consul.example.internal`

and to use a file named *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 a colon. If the command being run has spaces, make sure to wrap it in double quotes. The `-consul` flag instructs the daemon to 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, we can reload the NGINX configuration and allow your NGINX configuration to come alive along with your environment. With Consul you're able to set up health checks on each client to check the health of the intended service. With this failure detection, you're able to template your NGINX configuration accordingly to only send traffic to healthy hosts.

Also See

[Consul home page](#)

[Introduction to Consul Template](#)

[Consul template GitHub](#)

A/B Testing with `split_clients`

27.0 Introduction

NGINX has a module named `split_clients` that allows you to programmatically divide up your users based on a variable key. NGINX splits users by using a lightweight hashing algorithm to hash a given string. Then it mathematically divides them by percentages, mapping predefined values to a variable you can use to change the response of your server. This chapter covers the `split_clients` module.

27.1 A/B Testing

Problem

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

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 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.

Discussion

This type of A/B testing is useful when testing different types of marketing and frontend 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. 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 case of an error. Whatever the reason for splitting clients between two different application sets, NGINX makes this simple through use of this `split_clients` module.

Also See

[split_client documentation](#)

Locating Users by IP Address Using the GeoIP Module

28.0 Introduction

Tracking, analyzing, and utilizing the location of your clients in your application or your metrics can take your understanding of them to the next level. There are many implementations for the location of your clients, and NGINX makes locating them easy through use of the GeoIP module and a couple directives. This module makes it easy to log location, control access, or make decisions based on client locations. It also enables the geography of the client to be looked up initially upon ingestion of the request and passed along to any of the upstream applications so they don't have to do the lookup. This NGINX module is not installed by default and will need to be statically compiled from source, dynamically imported, or included in the NGINX package by installing `nginx-full` or `nginx-extras` in Ubuntu, which are both built with this module. On RHEL derivatives such as CentOS, you can install the `nginx-mod-http-geoip` package and dynamically import the module with the `load_module` directive. This chapter will cover importing the GeoIP dynamic module, installing the GeoIP database, the embedded variables available in this module, controlling access, and working with proxies.

28.1 Using the GeoIP Module and Database

Problem

You need to install the GeoIP database and enable its embedded variables within NGINX to log and tell your application the location of your clients.

Solution

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, we 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/nginx_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

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`, `$geoip_city_continent_code`, `$geoip_latitude`, `$geoip_longitude`, and `$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 update](#)

28.2 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/nginx_http_geoip_module.so";

http {
    map $geoip_country_code $country_access {
        "US"    0;
        "RU"    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 or Russia, 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 or Russia:

```
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 US or Russia.

Discussion

This is a short but simple example of how to only allow access from a couple countries. This example can be expounded upon 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.

28.3 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/nginx_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 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.

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 the AWS ELB, Google's load balancer, or Azure's load balancer in front of NGINX.

Debugging and Troubleshooting with Access Logs, Error Logs, and Request Tracing

29.0 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. In this chapter, we'll discuss access and error logs, streaming over the Syslog protocol, and tracing requests end to end with request identifiers generated by NGINX.

29.1 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 $remote_user '
        '$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. `$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.

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.

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 necessary information you could ever want. It's also possible to structure format to log in JSON or XML. 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's limitless possibility to logs when troubleshooting, debugging, or analyzing your application or market.

29.2 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. 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.

29.3 Forwarding to Syslog

Problem

You need to forward your logs to a Syslog listener to aggregate logs to a centralized service.

Solution

Use the `access_log` and `error_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 of the log message defined as one of the 23 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 adding 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.

29.4 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_header X-Request-ID $request_id; # Return to client
    location / {
        proxy_pass http://backend;
        proxy_set_header X-Request-ID $request_id; #Pass to app
        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 front-end 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.

Performance Tuning

30.0 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 limitation, and repeat until you've reached your desired performance requirements. In this chapter we'll suggest taking measurements when performance tuning by testing with automated tools and measuring results. This chapter will also cover connection tuning for keeping connections open to clients as well as upstream servers, and serving more connections by tuning the operating system.

30.1 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 off of 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.

30.2 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 the amount of time 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 the 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 fully qualified domain name. The number of parallel open connections to a domain is still limited typically to a number less than 10, so in this regard, many requests over a single connection will happen. A trick 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, as 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.

30.3 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 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 would want to keep connections open to upstream servers to save the amount of time it takes to initiate the connection, and the worker process can 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 as 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.

30.4 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 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.

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 thoroughly and repeatedly 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.

30.5 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;
}
```

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.

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 buffering logs in this way, when tailing the log, you may see delays up to the amount of time specified by the `flush` parameter.

30.6 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; and 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. 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 systems 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.

Practical Ops Tips and Conclusion

31.0 Introduction

This last chapter will cover practical operations tips and is the conclusion to this book. Throughout these three parts, we've discussed many ideas and concepts pertinent to operations engineers. However, I thought a few more 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.

31.1 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.

31.2 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 section [“Also See” on page 168](#).

- 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, 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 section [“Also See” on page 168](#).
- 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 your own foot. When debugging, make sure you know how to trace your request through your configuration; and 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 requests](#)

[Debugging admin guide](#)

[Rewrite log](#)

31.3 Conclusion

This book's three parts have focused on high-performance load balancing, security, and deploying and maintaining NGINX and NGINX Plus servers. This book has demonstrated some of the most powerful features of the NGINX application delivery platform. NGINX continues to develop amazing features and stay ahead of the curve.

This book has demonstrated 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 sever 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. May you now know that.

About the Author

Derek DeJonghe has had a lifelong 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 engineers and produces self-healing, auto-scaling infrastructure for numerous applications. He specializes in Linux cloud environments. 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 helps RightBrain Networks be one of the strongest cloud consulting agencies and managed service providers in partnership with AWS today.