amazon web services

Architecting for High Availability

Attila Narin
AWS Solutions Architecture

QCon London, March 2013

Session Feedback ID 1927

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

LET’S BUILD A SYSTEM

LET'S BUILD A

HIGHLY AVAILABLE

SYSTEM

AWS BUILDING BLOCKS

Inherently Highly Available and Fault Tolerant Services

- Amazon S3
- Amazon DynamoDB
- Amazon CloudFront
- Amazon Route53
- Elastic Load Balancing
- Amazon SQS
- Amazon SNS
- Amazon SES
- Amazon SWF
- …

Highly Available with the right architecture

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon VPC

AWS BUILDING BLOCKS

Inherently Highly Available and Fault Tolerant Services

- Amazon S3
- Amazon DynamoDB
- Amazon CloudFront
- Amazon Route53
- Elastic Load Balancing
- Amazon SQS
- Amazon SNS
- Amazon SES
- Amazon SWF
- …

Highly Available with the right architecture

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon VPC


Web Server
Amazon
EC2

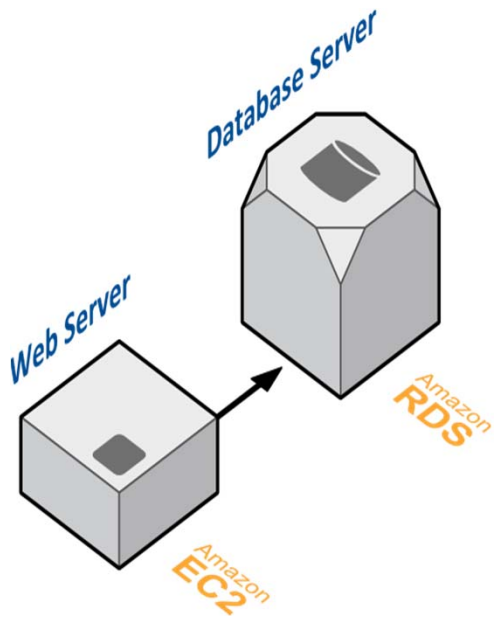
Database Server
Web Server
Amazon
RDS
Amazon
EC2

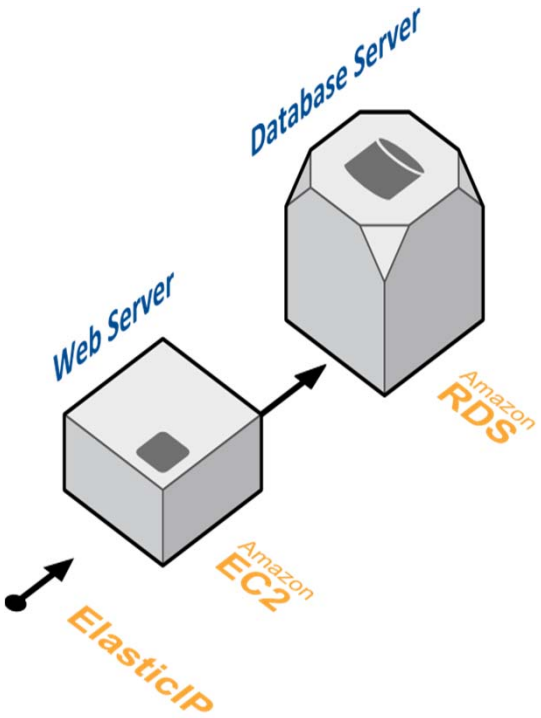
Database Server
Web Server
Amazon
RDS
Amazon
EC2
ElasticIP

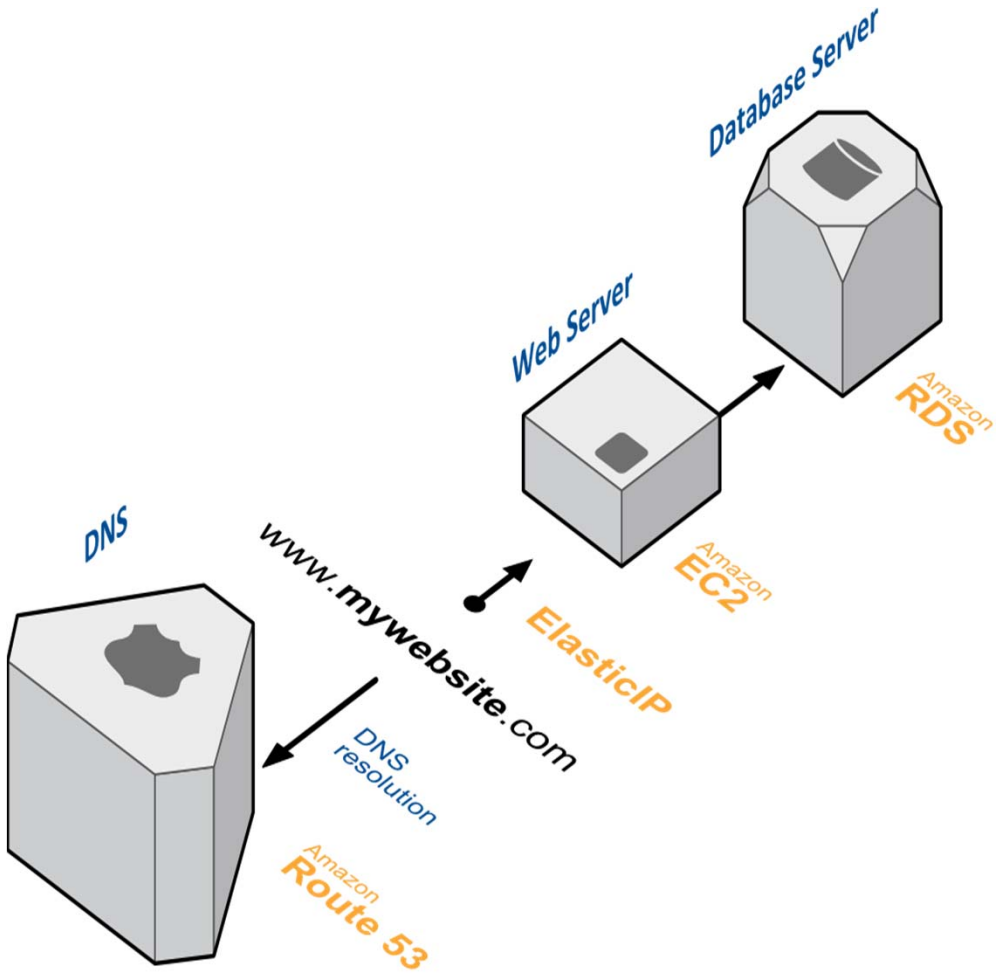
Database Server
Amazon
RDS
Web Server
Amazon
EC2
ElasticIP
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

#1

DESIGN FOR FAILURE

●○○○○

« Everything fails **all the time** »

Werner Vogels
CTO of Amazon

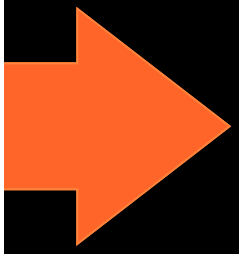

AVOID SINGLE POINTS OF FAILURE

AVOID SINGLE POINTS OF FAILURE

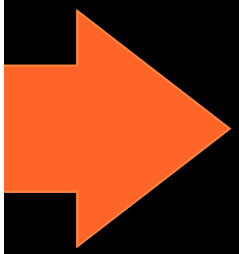

ASSUME EVERYTHING FAILS, AND WORK BACKWARDS

YOUR GOAL

Applications should
continue to function

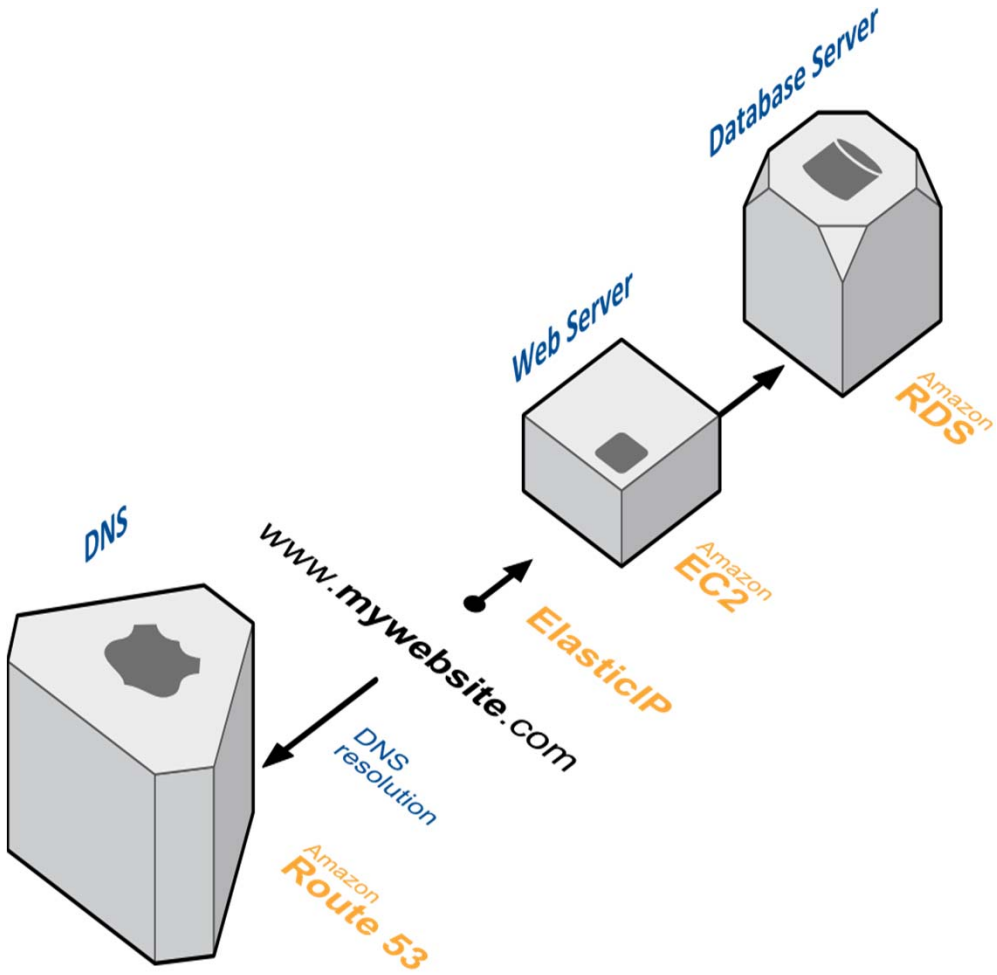
Database Server
Amazon
RDS
Web Server
Amazon
EC2
ElasticIP
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

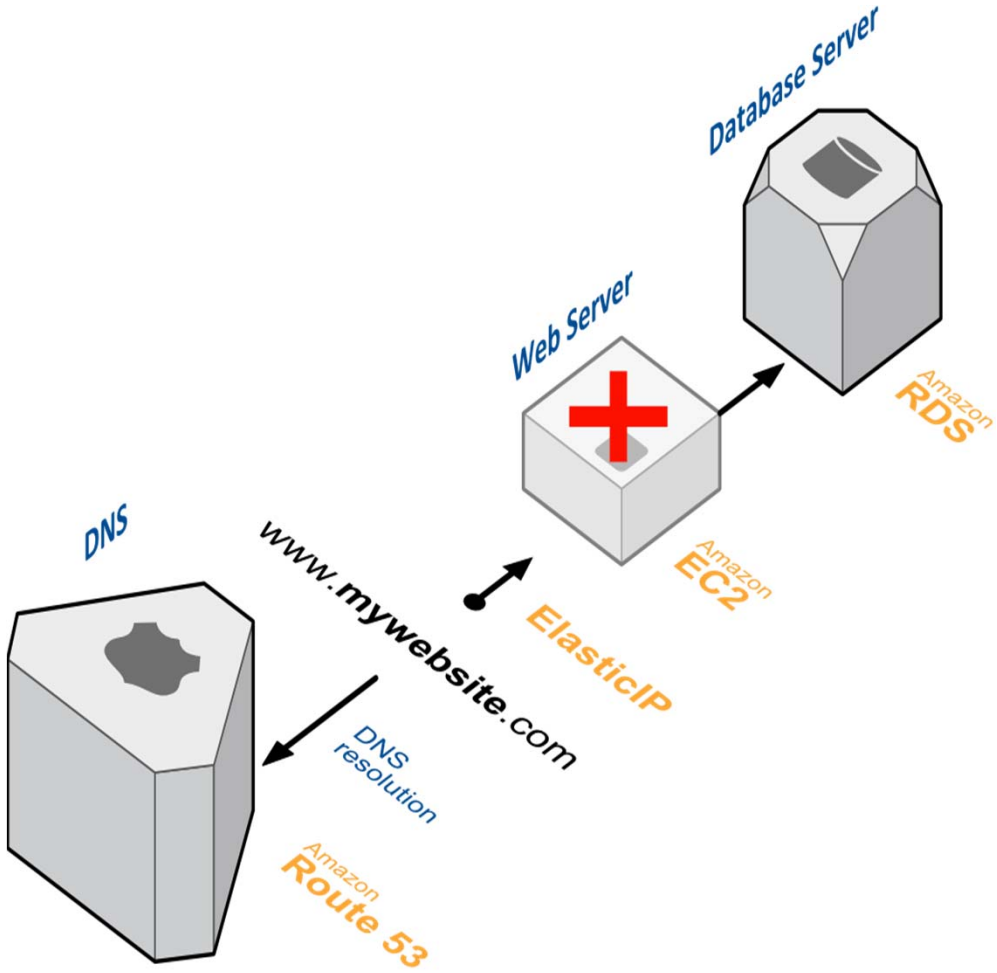
Database Server
Amazon
RDS
Web Server
Amazon
EC2
ElasticIP
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

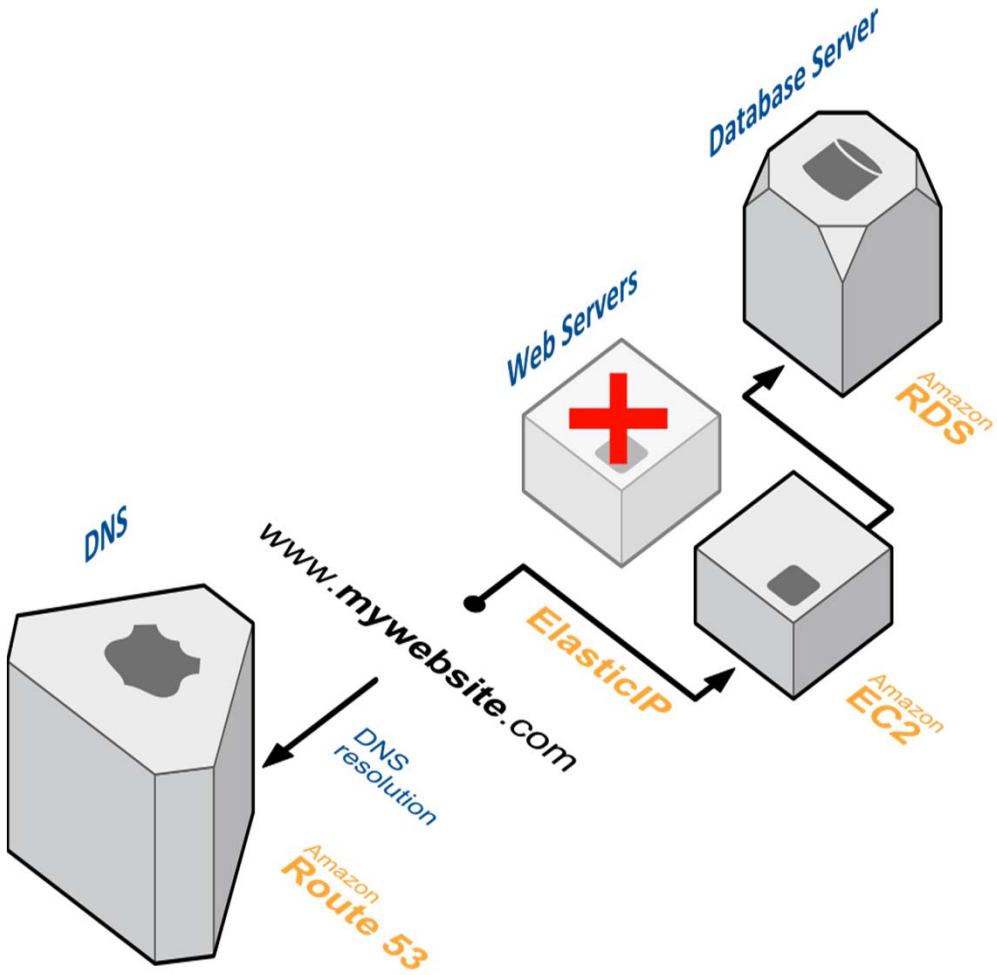
Database Server
Web Servers
Amazon
RDS
DNS
www.mywebsite.com
ElasticIP
Amazon
EC2
DNS
resolution
Amazon
Route 53

AMAZON EBS

ELASTIC BLOCK STORE

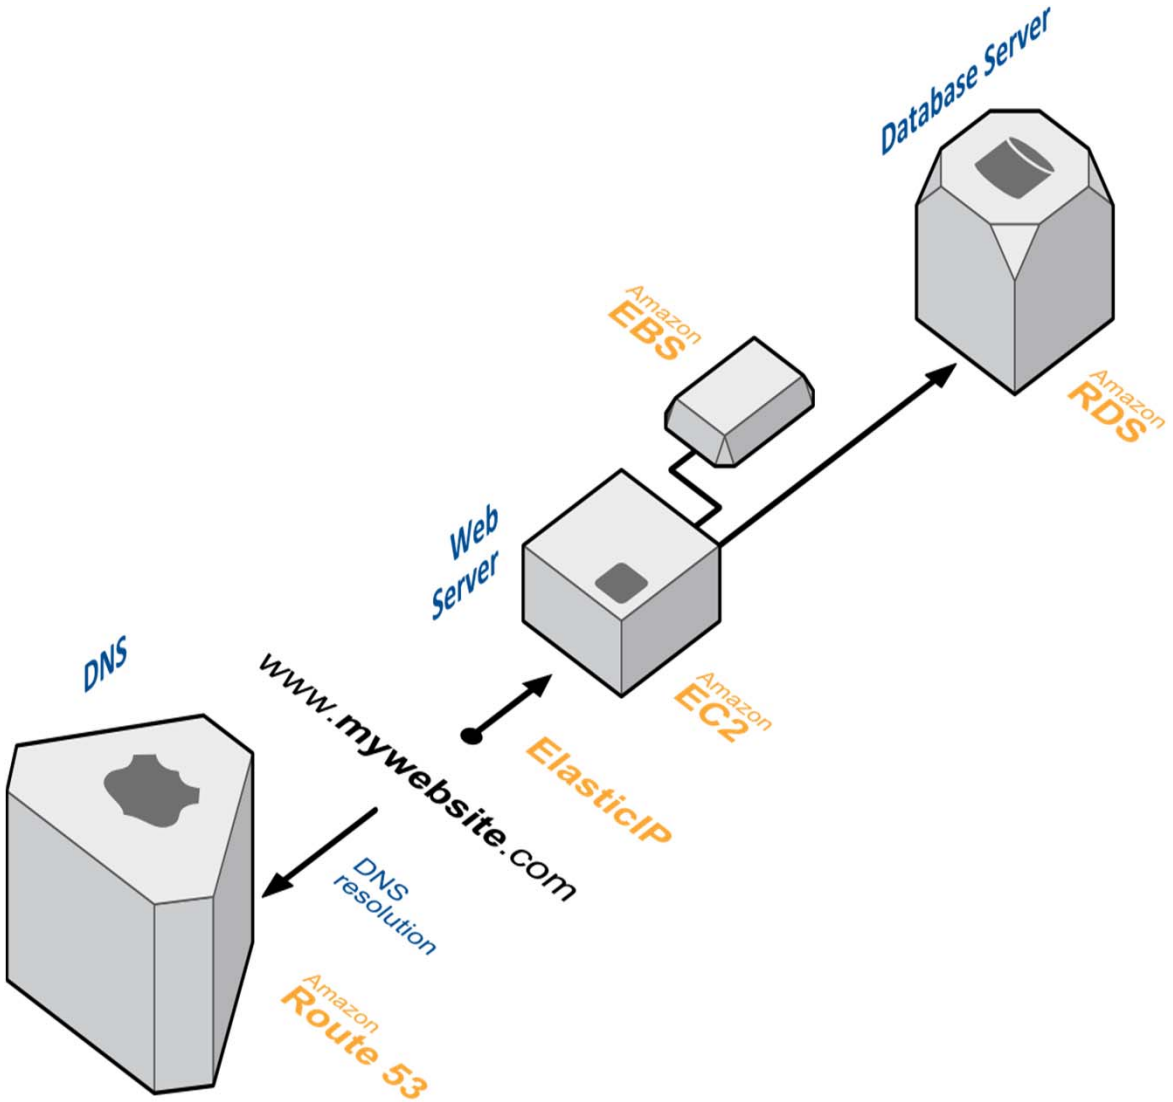
Database Server
Amazon
EBS
Amazon
RDS
Web
Server
DNS
www.mywebsite.com
Amazon
EC2
ElasticIP
DNS
resolution
Amazon
Route 53

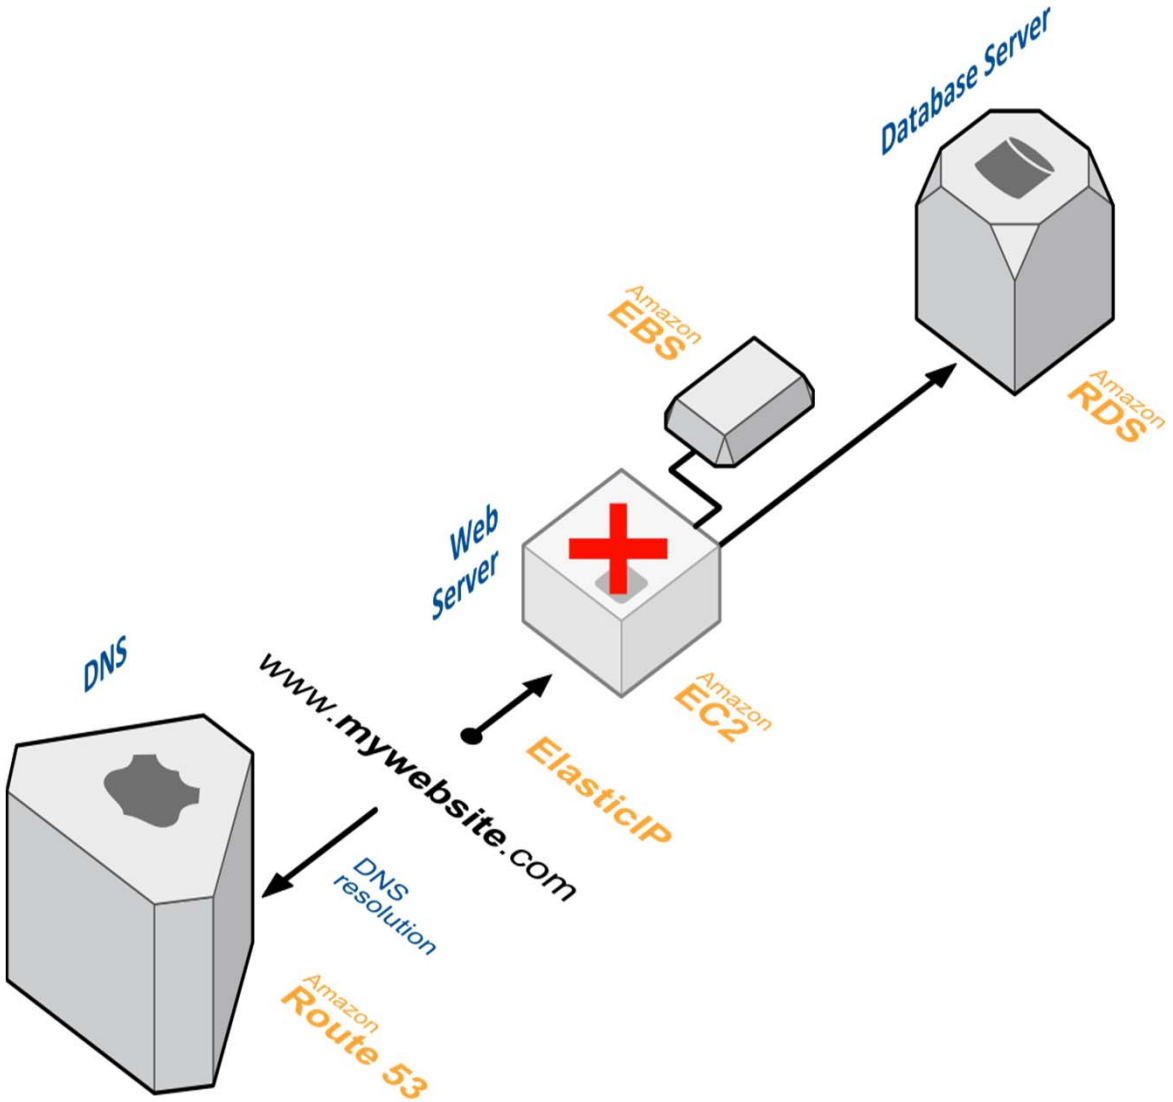
Database Server
Amazon
EBS
Amazon
RDS
Web
Server
Amazon
EC2
DNS
www.mywebsite.com
ElasticIP
DNS
resolution
Amazon
Route 53

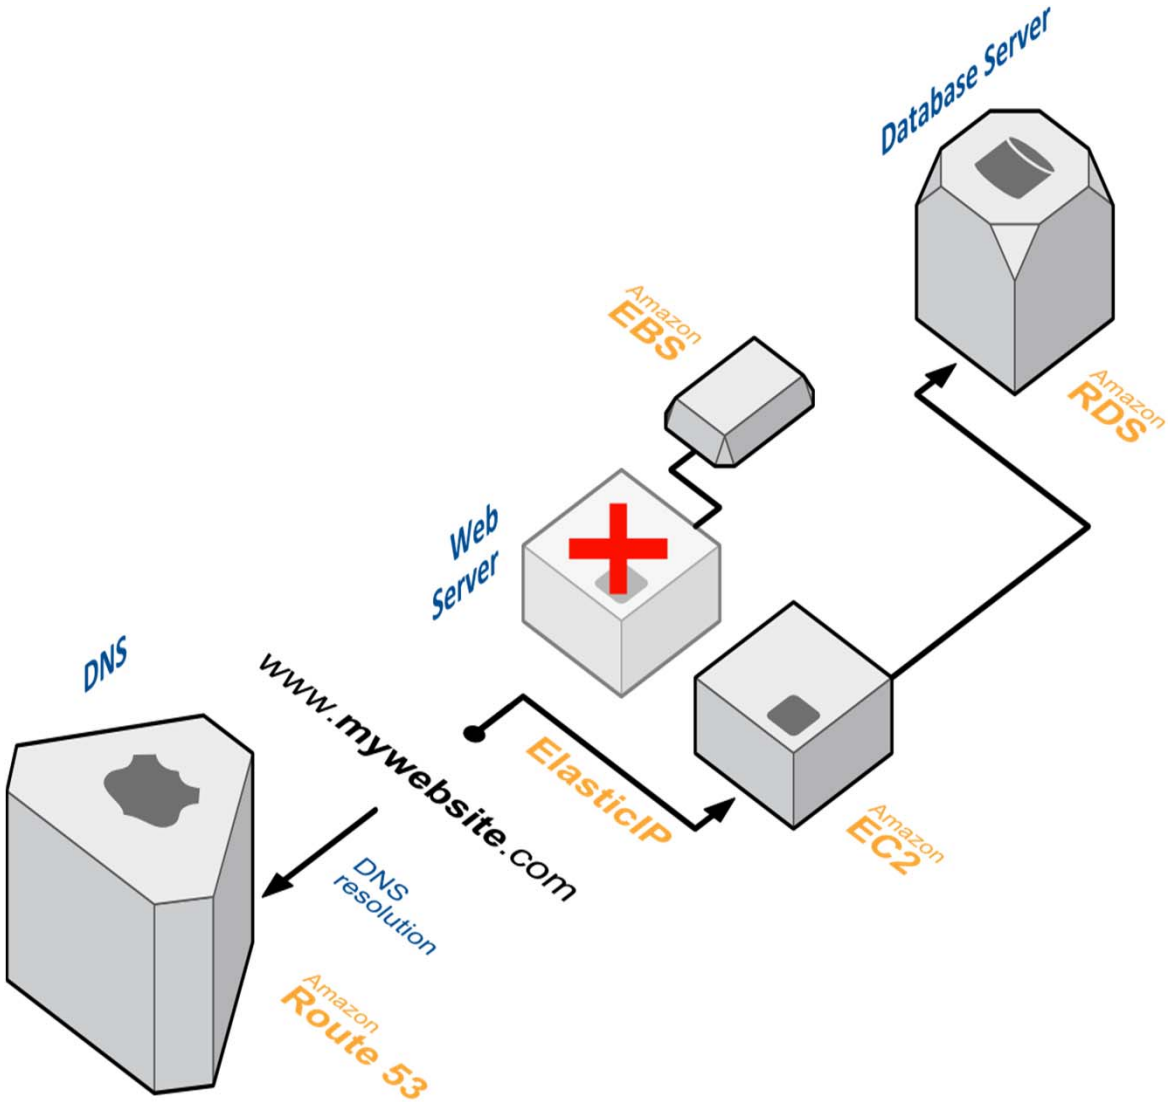
Database Server
Amazon
EBS
Amazon
RDS
Web
Server
DNS
www.mywebsite.com
ElasticIP
Amazon
EC2
DNS
resolution
Amazon
Route 53

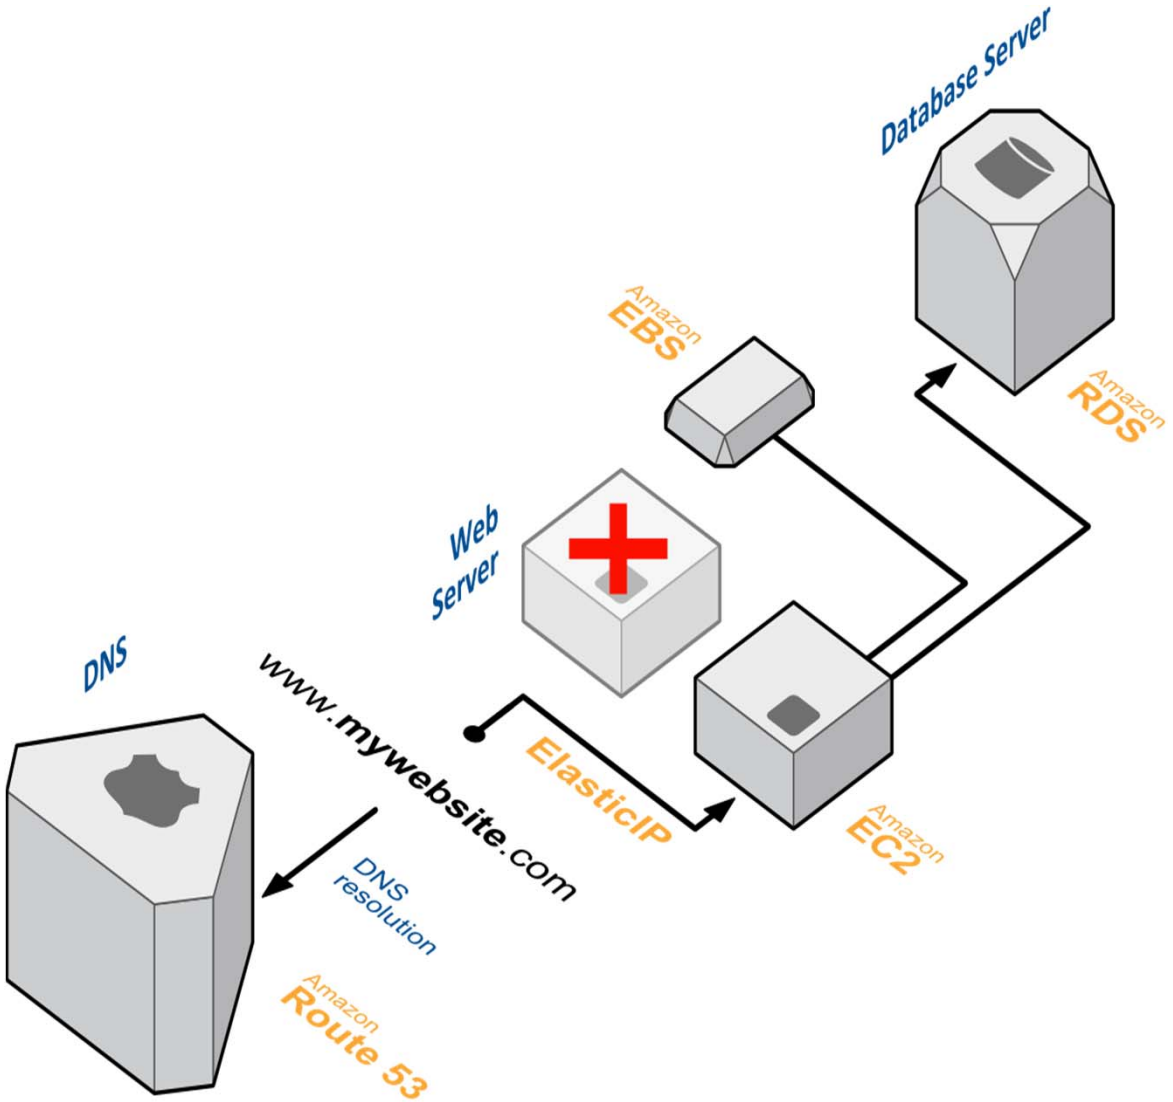
Database Server
Amazon
EBS
Amazon
RDS
Web
Server
DNS
www.mywebsite.com
ElasticIP
Amazon
EC2
DNS
resolution
Amazon
Route 53

AMAZON ELB

ELASTIC LOAD BALANCING

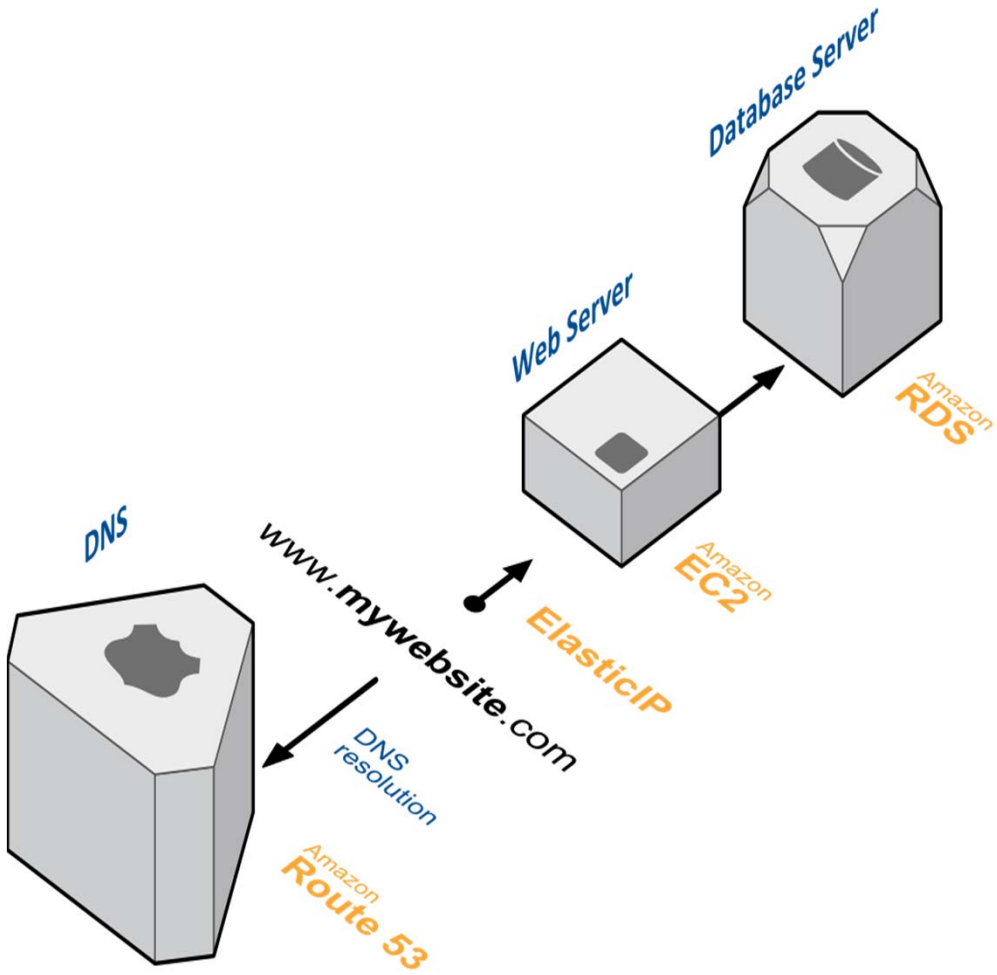
Database Server
Amazon
RDS
Web Server
Amazon
EC2
ElasticIP
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

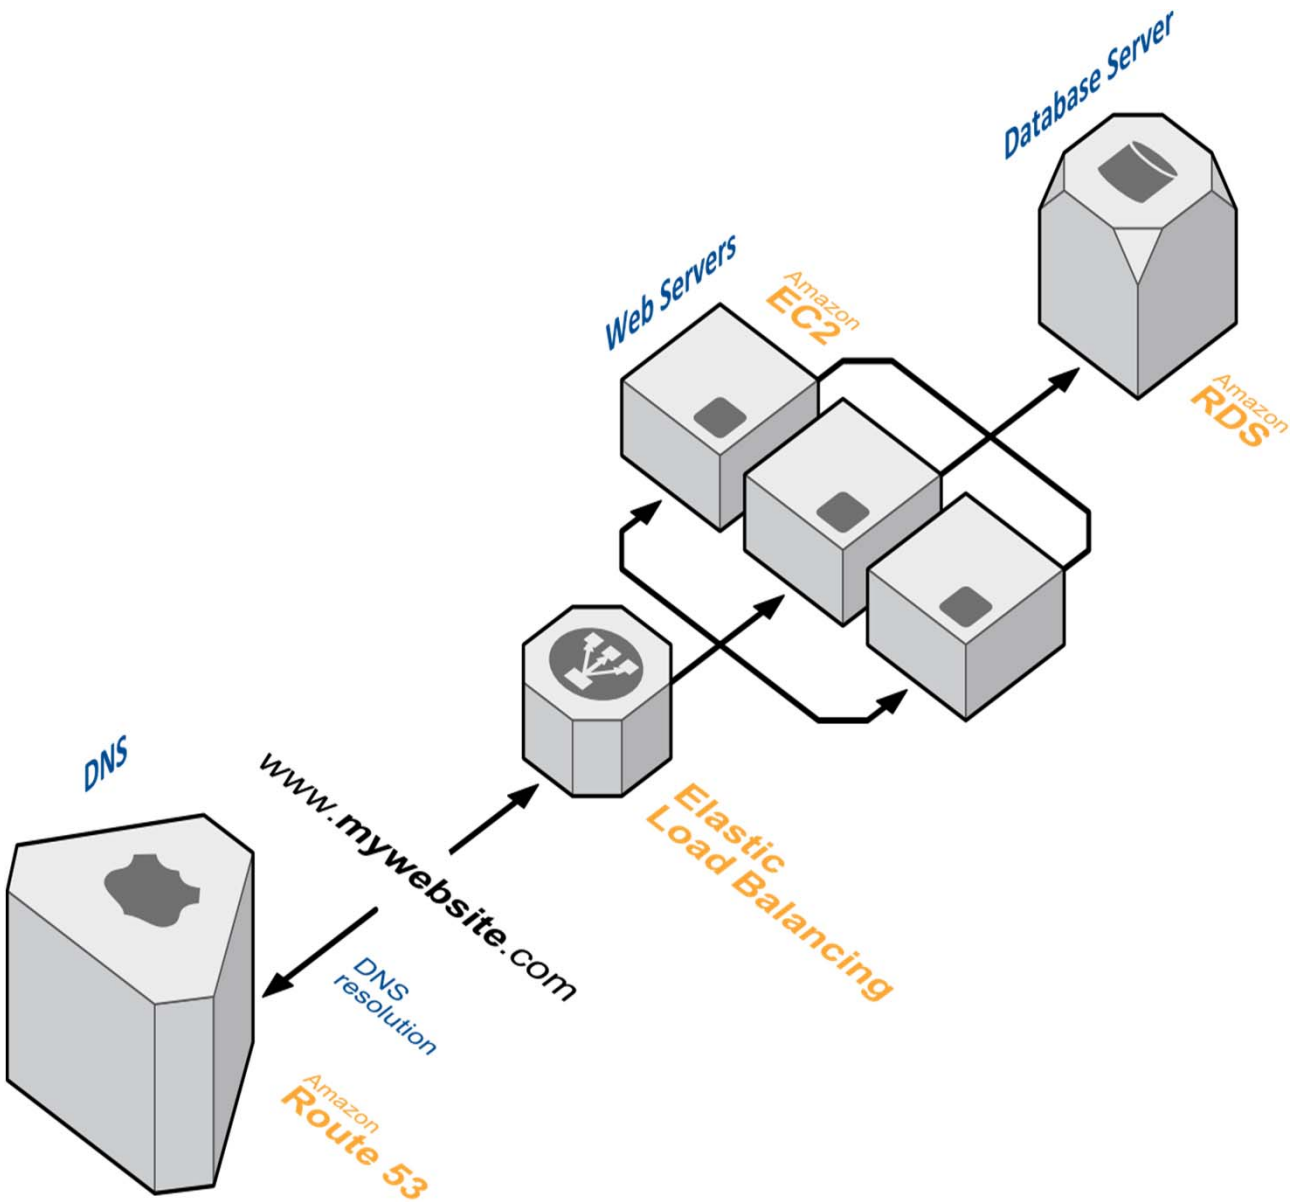
Database Server
Web Servers
Amazon
EC2
Amazon
RDS
DNS
www.mywebsite.com
Elastic
Load Balancing
DNS
resolution
Amazon
Route 53

HEALTH CHECKS

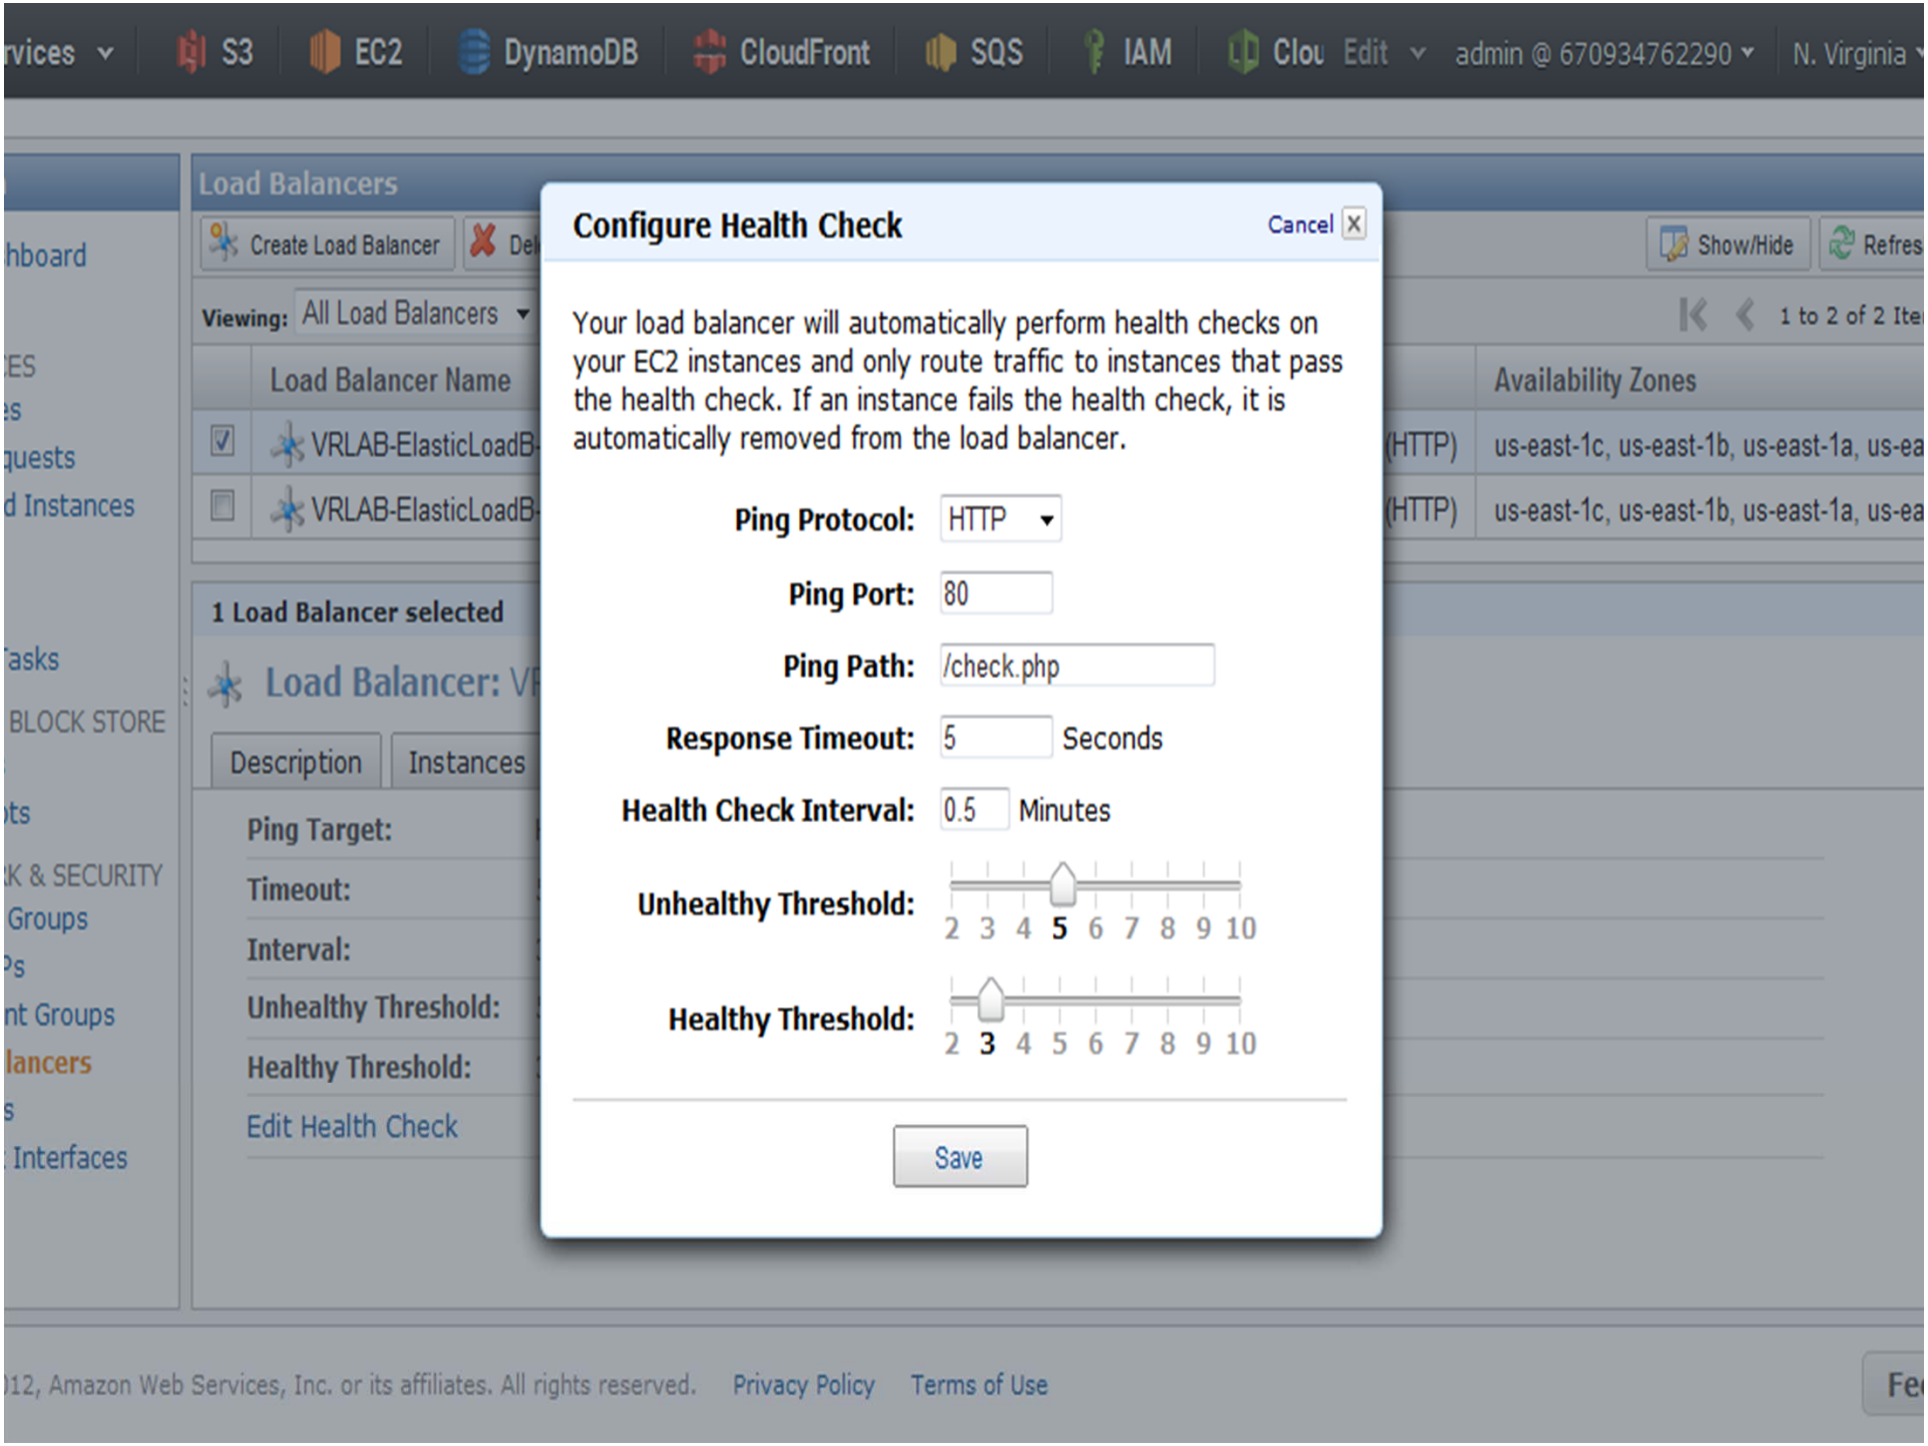
S3
EC2
DynamoDB
CloudFront
SQS
IAM
admin @ 670934762290
N. Virginia
Load Balancers
Create Load Balancer
Show/Hide
Viewing: All Load Balancers
Load Balancer Name
Availability Zones
VRLAB-ElasticLoadB
(HTTP)
us-east-1c, us-east-1b, us-east-1a, us-ea
1 Load Balancer selected
Load Balancer:
Description
Instances
Ping Target:
Timeout:
Interval:
Unhealthy Threshold:
Healthy Threshold:
Edit Health Check
Configure Health Check
Cancel
Your load balancer will automatically perform health checks on your EC2 instances and only route traffic to instances that pass the health check. If an instance fails the health check, it is automatically removed from the load balancer.
Ping Protocol: HTTP
Ping Port: 80
Ping Path: /check.php
Response Timeout: 5 Seconds
Health Check Interval: 0.5 Minutes
Unhealthy Threshold: 2 3 4 5 6 7 8 9 10
Healthy Threshold: 2 3 4 5 6 7 8 9 10
Save
Privacy Policy
Terms of Use

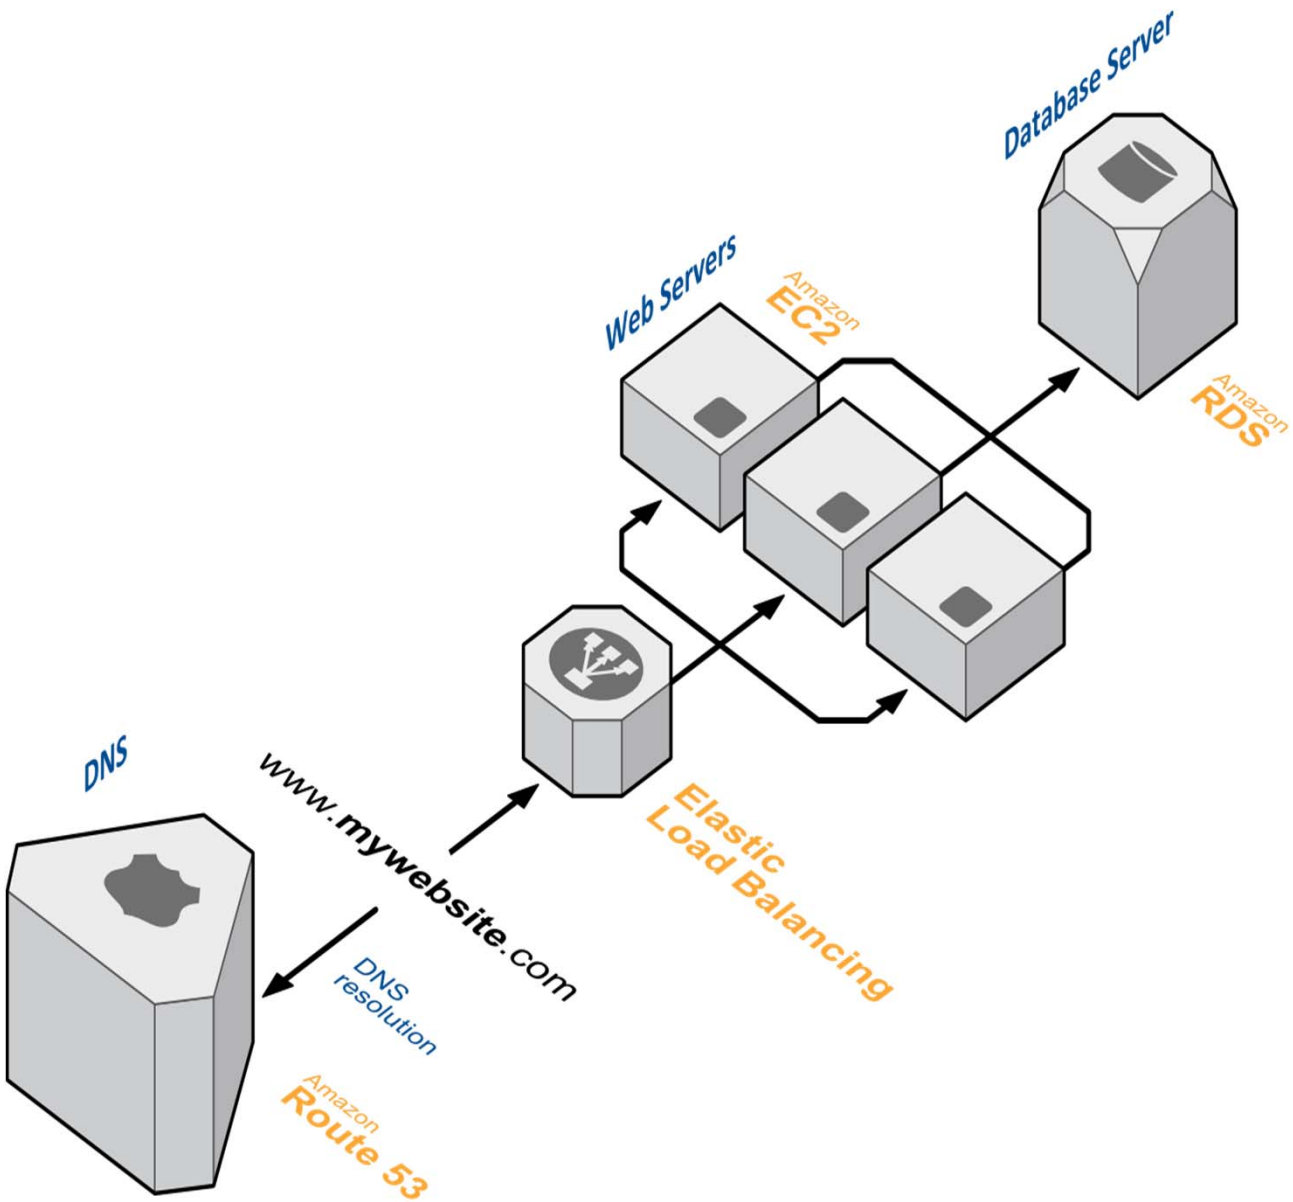
Database Server
Web Servers
Amazon
EC2
Amazon
RDS
DNS
www.mywebsite.com
Elastic
Load Balancing
DNS
resolution
Amazon
Route 53

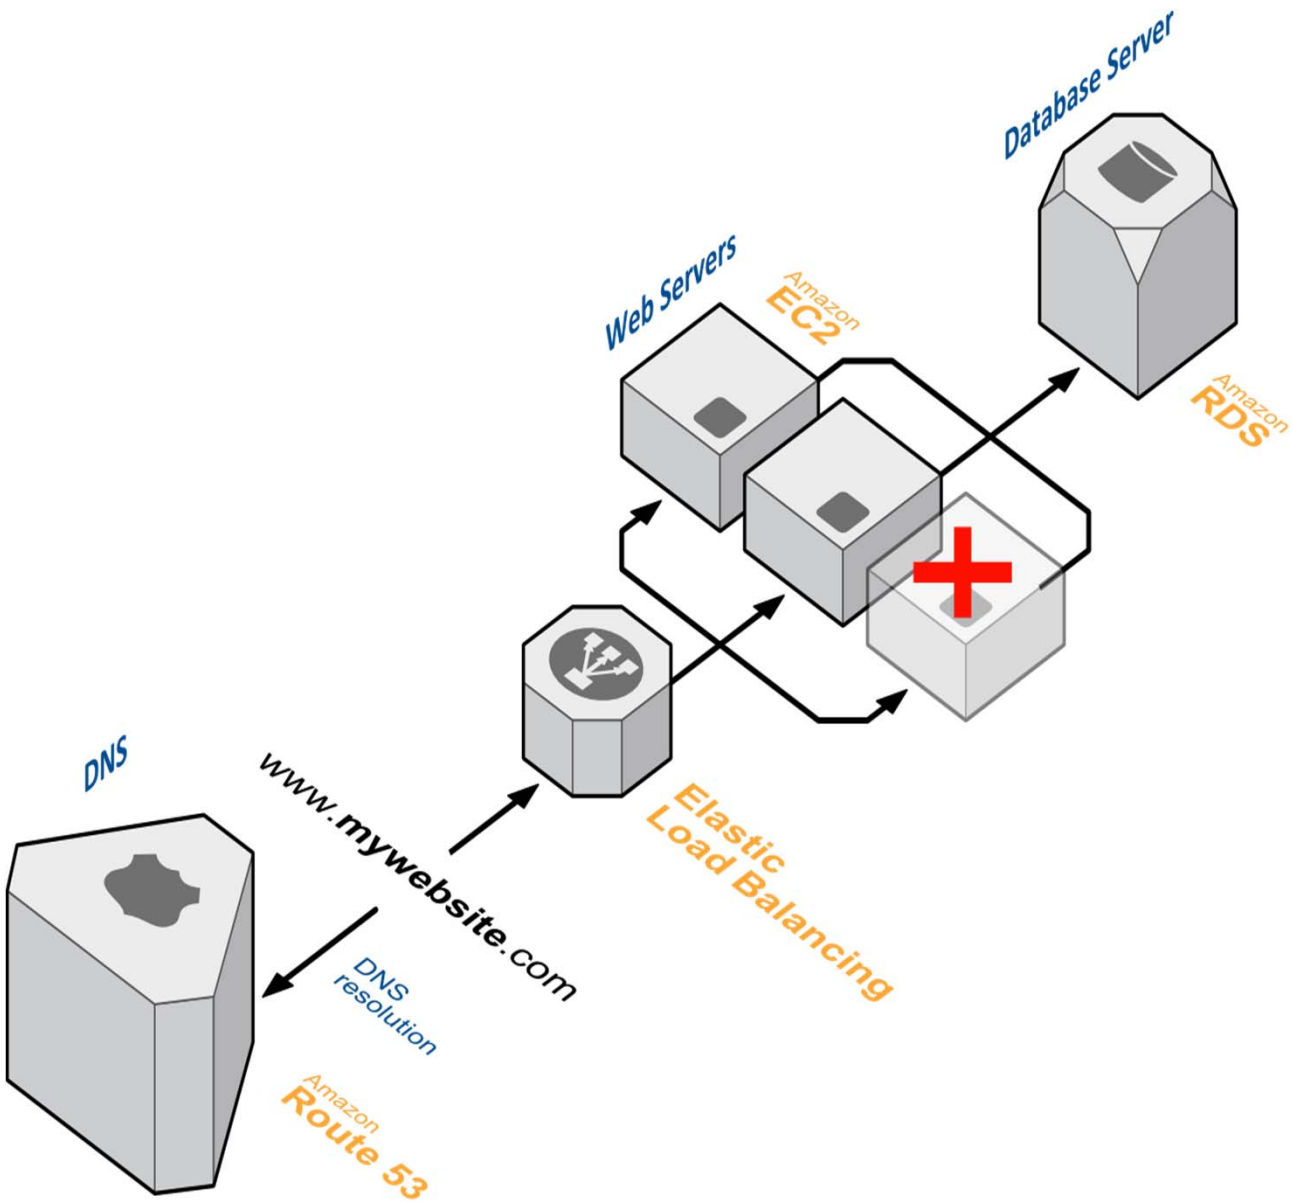
Database Server
Amazon
RDS
Web Servers
Amazon
EC2
Elastic
Load Balancing
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

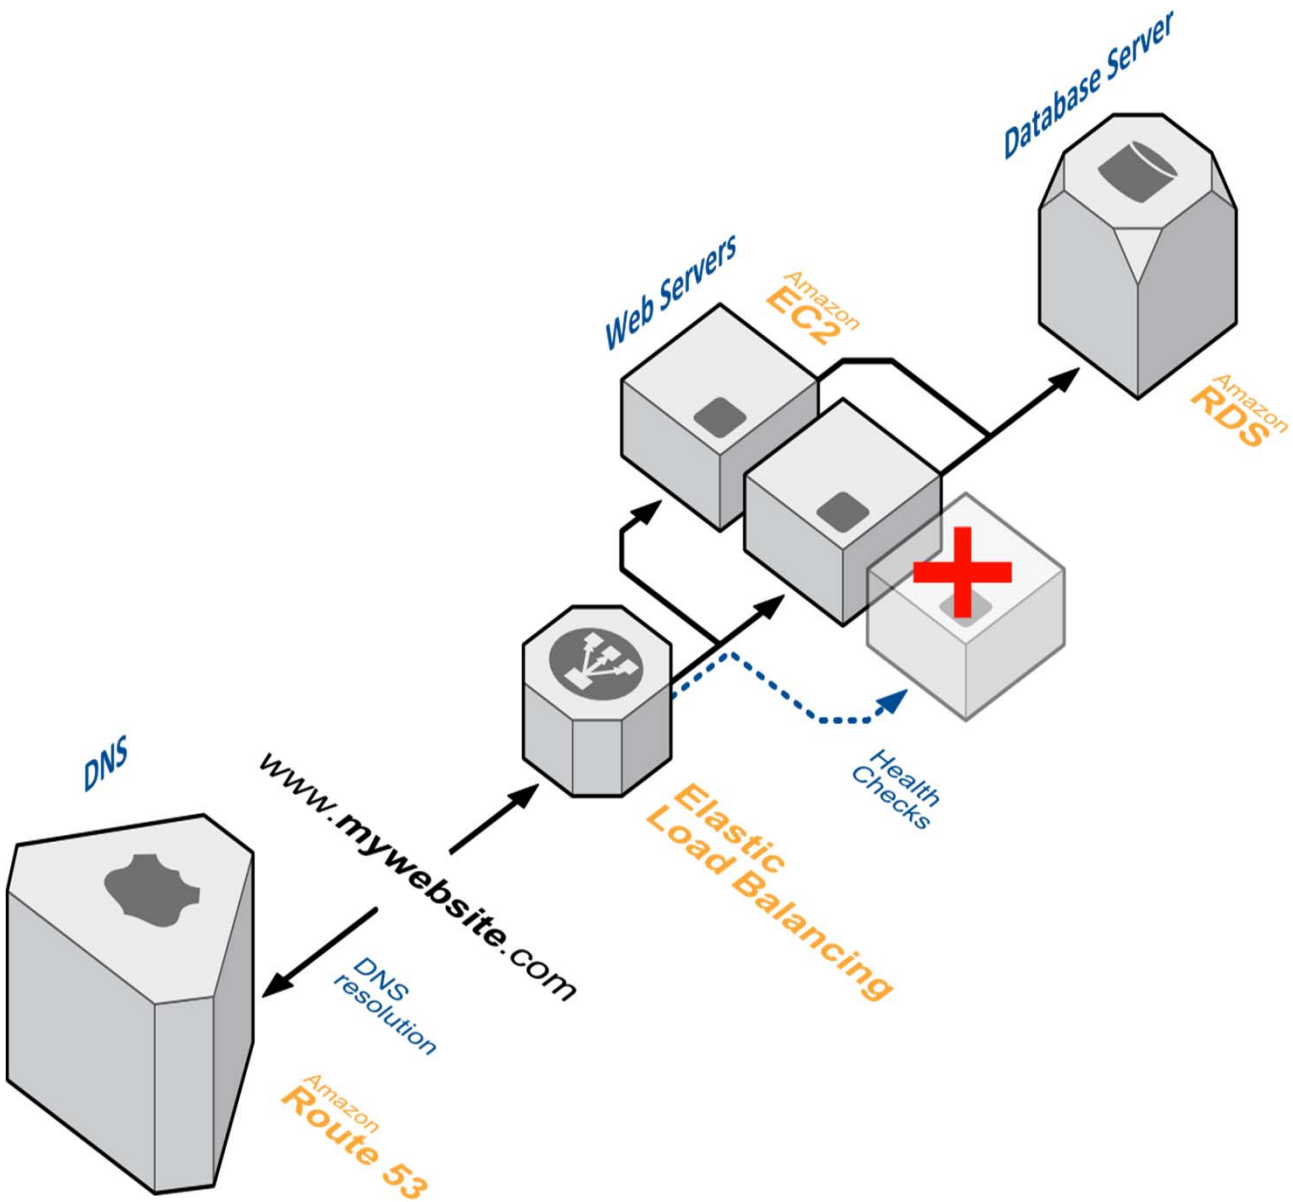
Database Server
Amazon
RDS
Web Servers
Amazon
EC2
Health
Checks
Elastic
Load Balancing
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

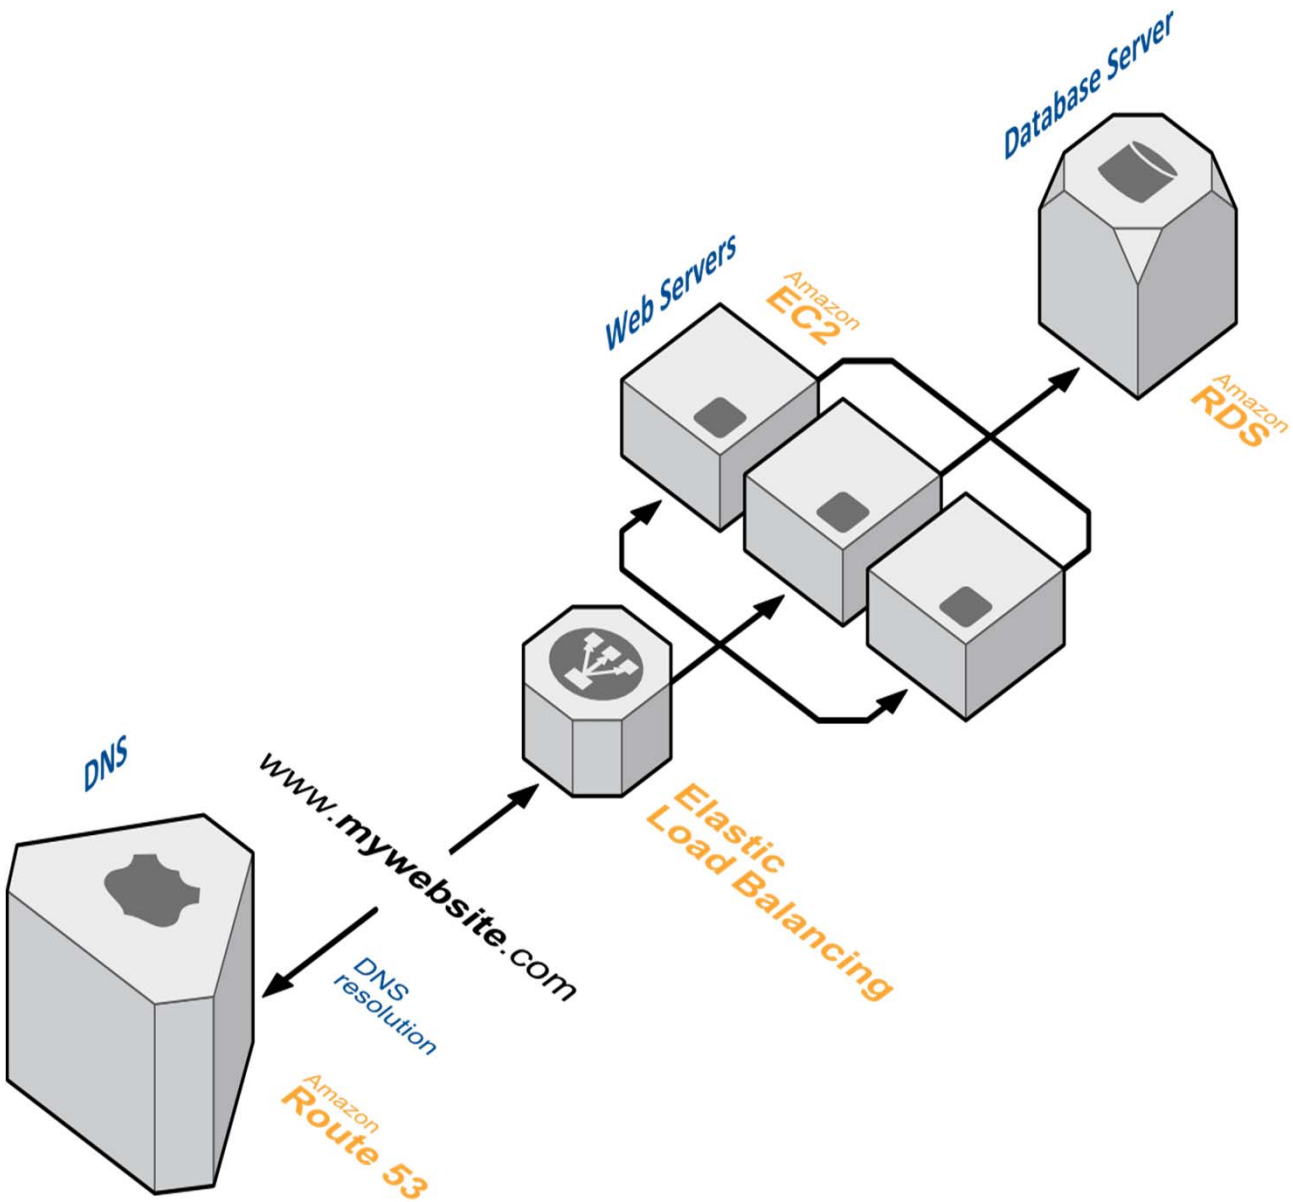
Database Server
Web Servers
Amazon
EC2
Amazon
RDS
DNS
www.mywebsite.com
Elastic
Load Balancing
DNS
resolution
Amazon
Route 53

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

#2

MULTIPLE AVAILABILITY ZONES

AVAILABILITY ZONES

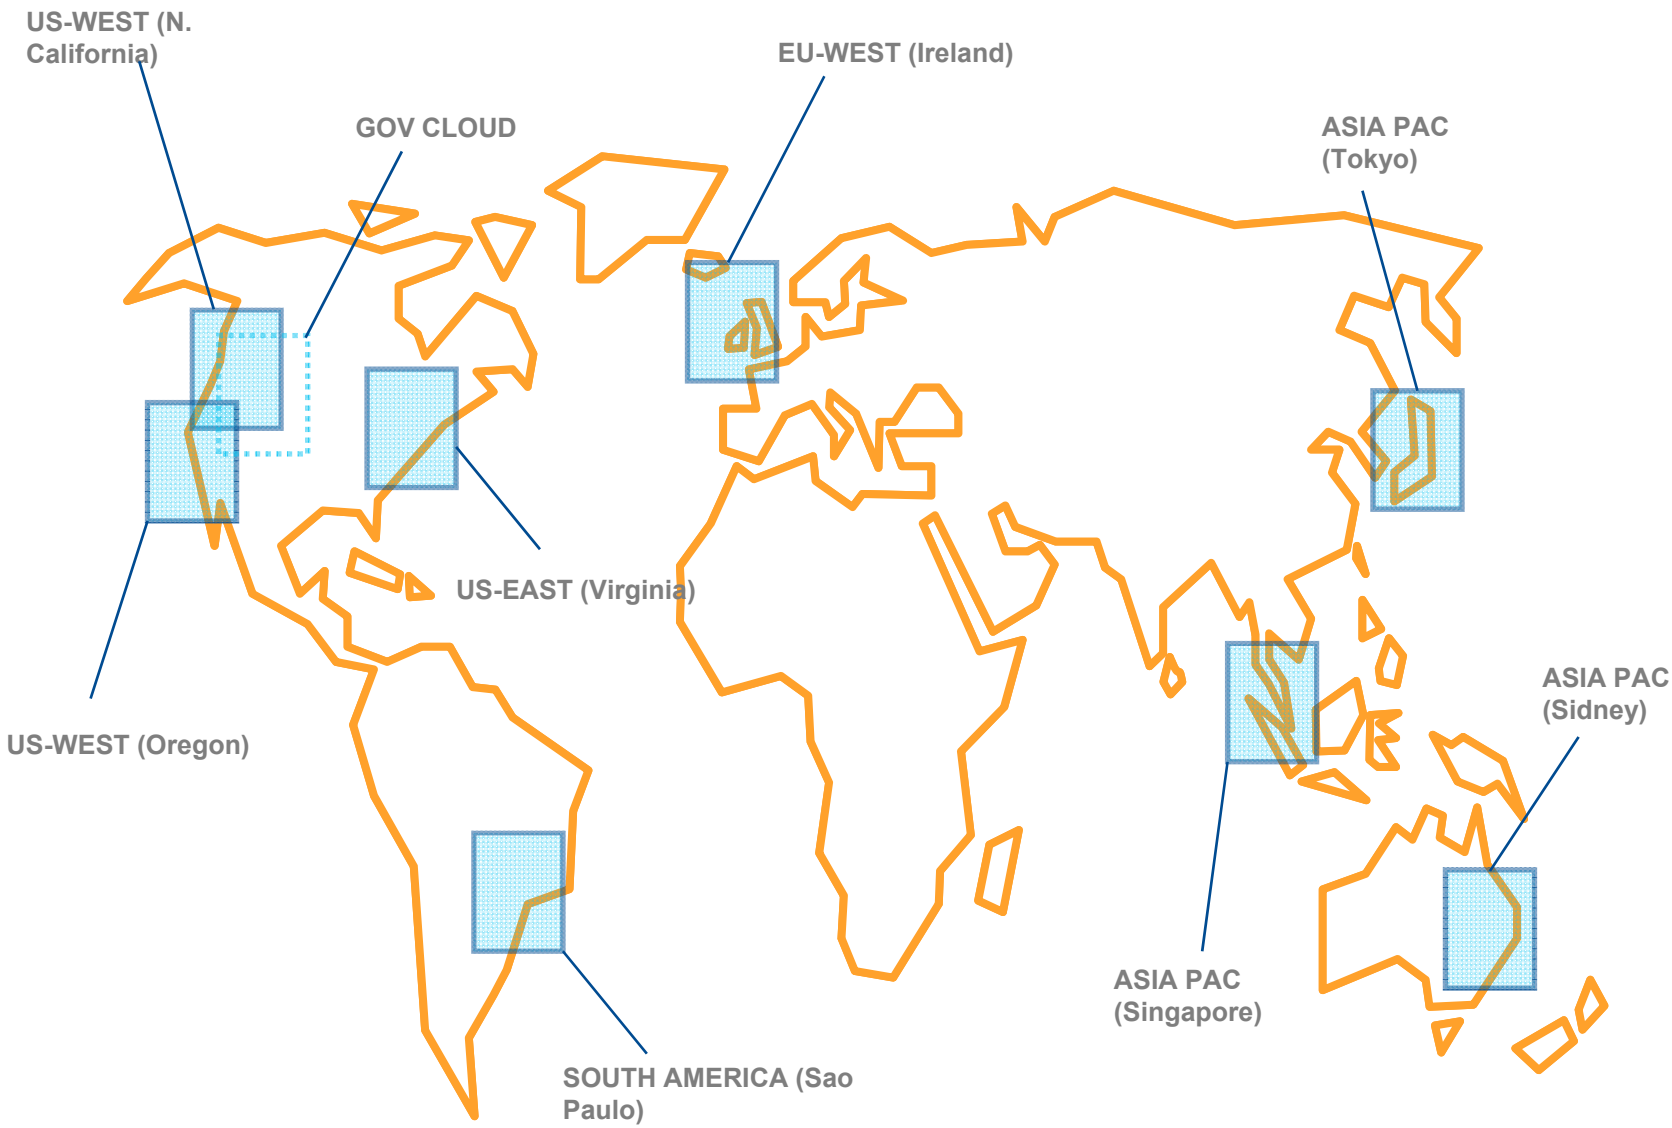
US-WEST (N. California)
GOV CLOUD
EU-WEST (Ireland)
ASIA PAC (Tokyo)
US-EAST (Virginia)
US-WEST (Oregon)
ASIA PAC (Sidney)
ASIA PAC (Singapore)
SOUTH AMERICA (Sao Paulo)

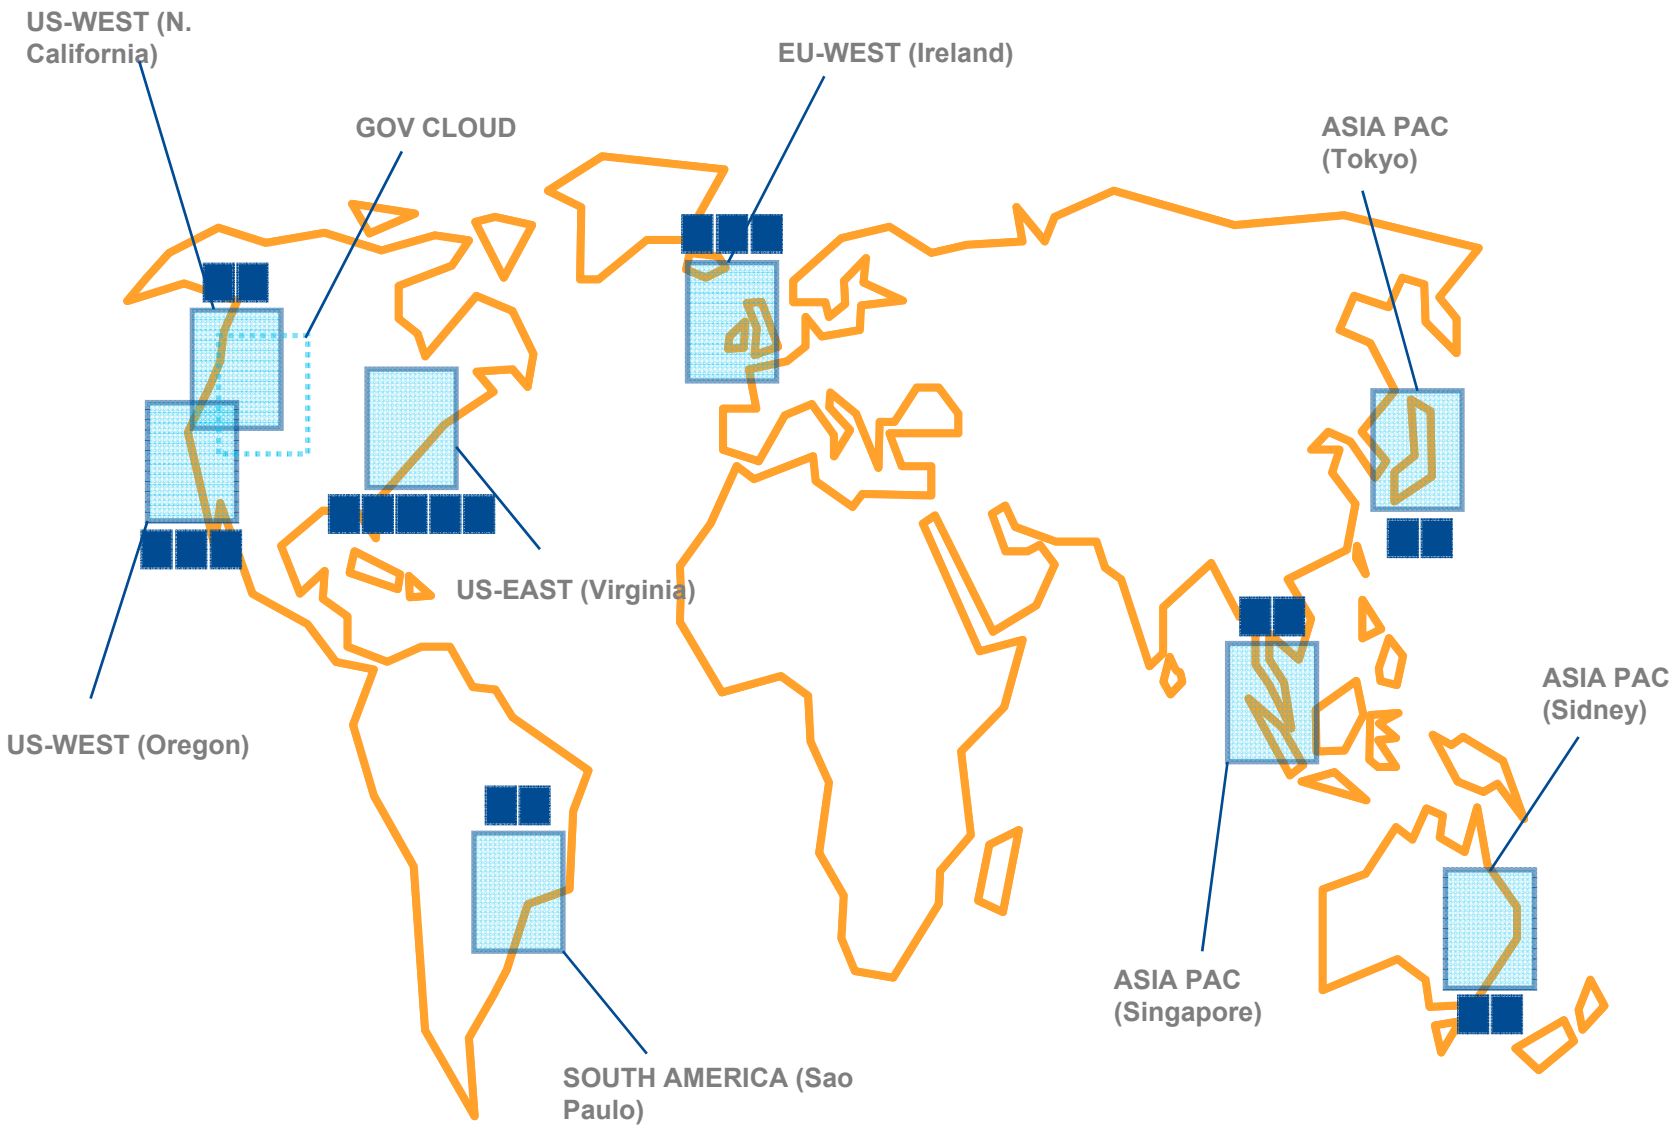
US-WEST (N. California)
GOV CLOUD
EU-WEST (Ireland)
ASIA PAC (Tokyo)
US-EAST (Virginia)
US-WEST (Oregon)
ASIA PAC (Sidney)
ASIA PAC (Singapore)
SOUTH AMERICA (Sao Paulo)

AMAZON RDS
MULTI-AZ

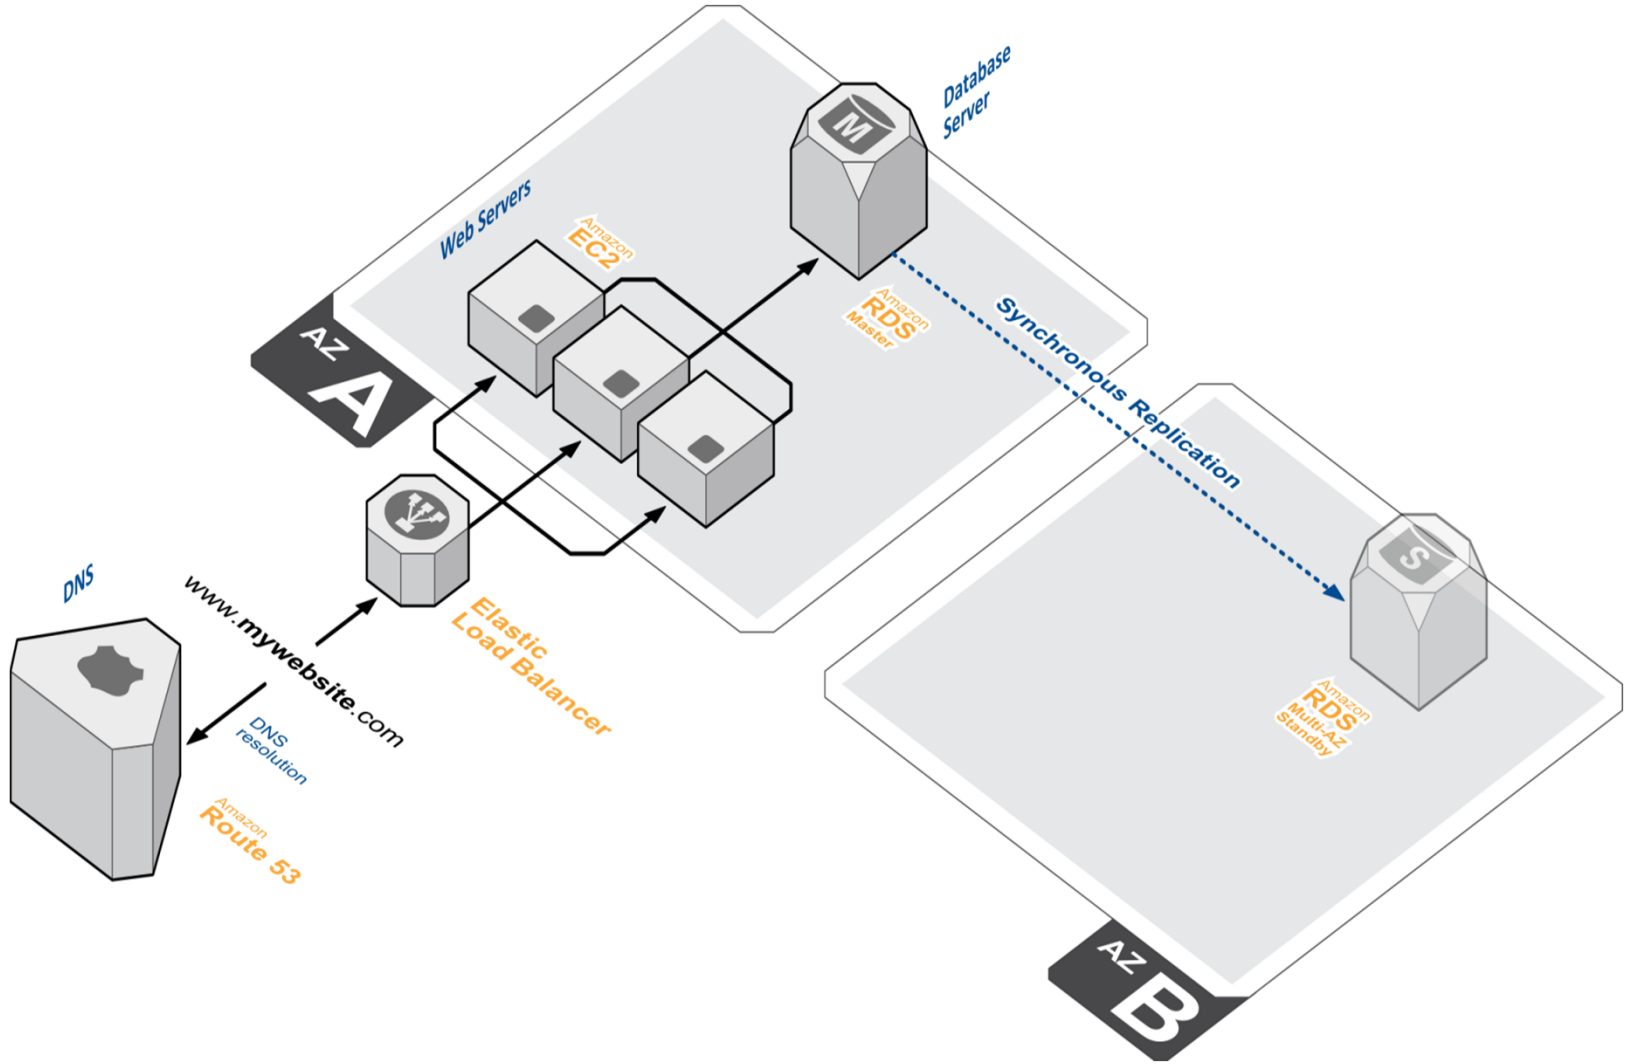
Web Servers
Amazon EC2
Database Server
Amazon RDS Master
Synchronous Replication
AZ A
AZ B
Amazon RDS Multi-AZ Standby
Elastic Load Balancer
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

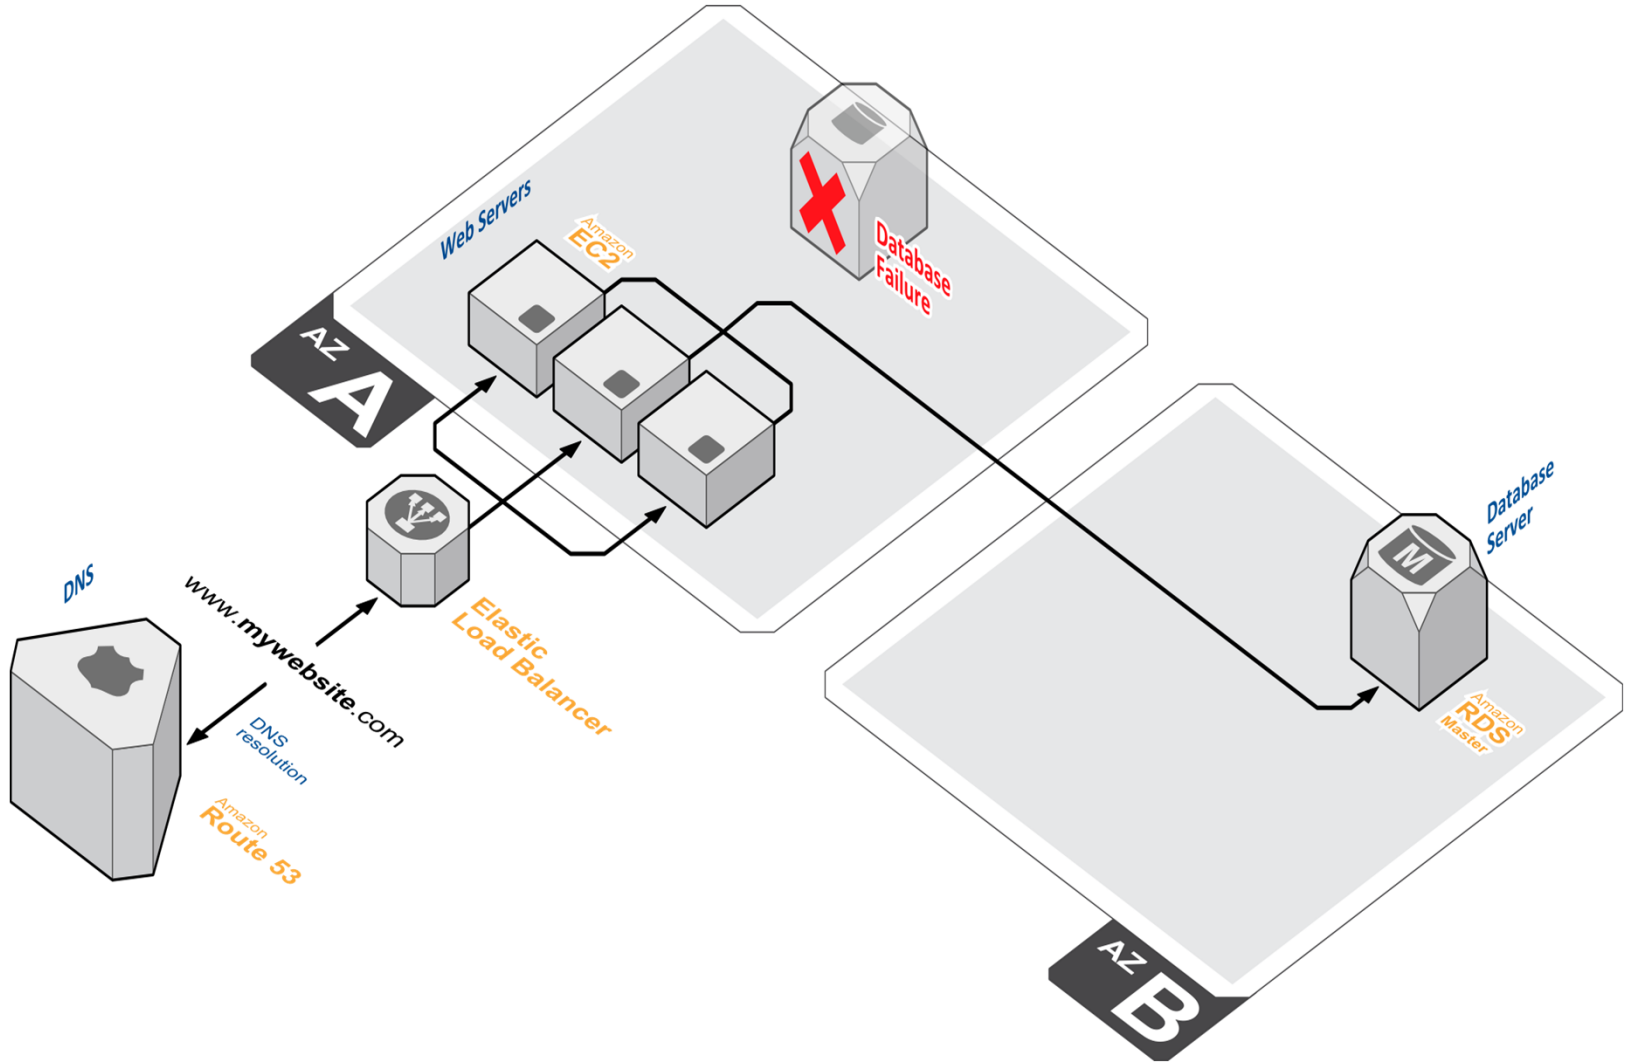
Web Servers
Amazon EC2
Database Failure
AZ A
Database Server
Amazon RDS Master
AZ B
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancer

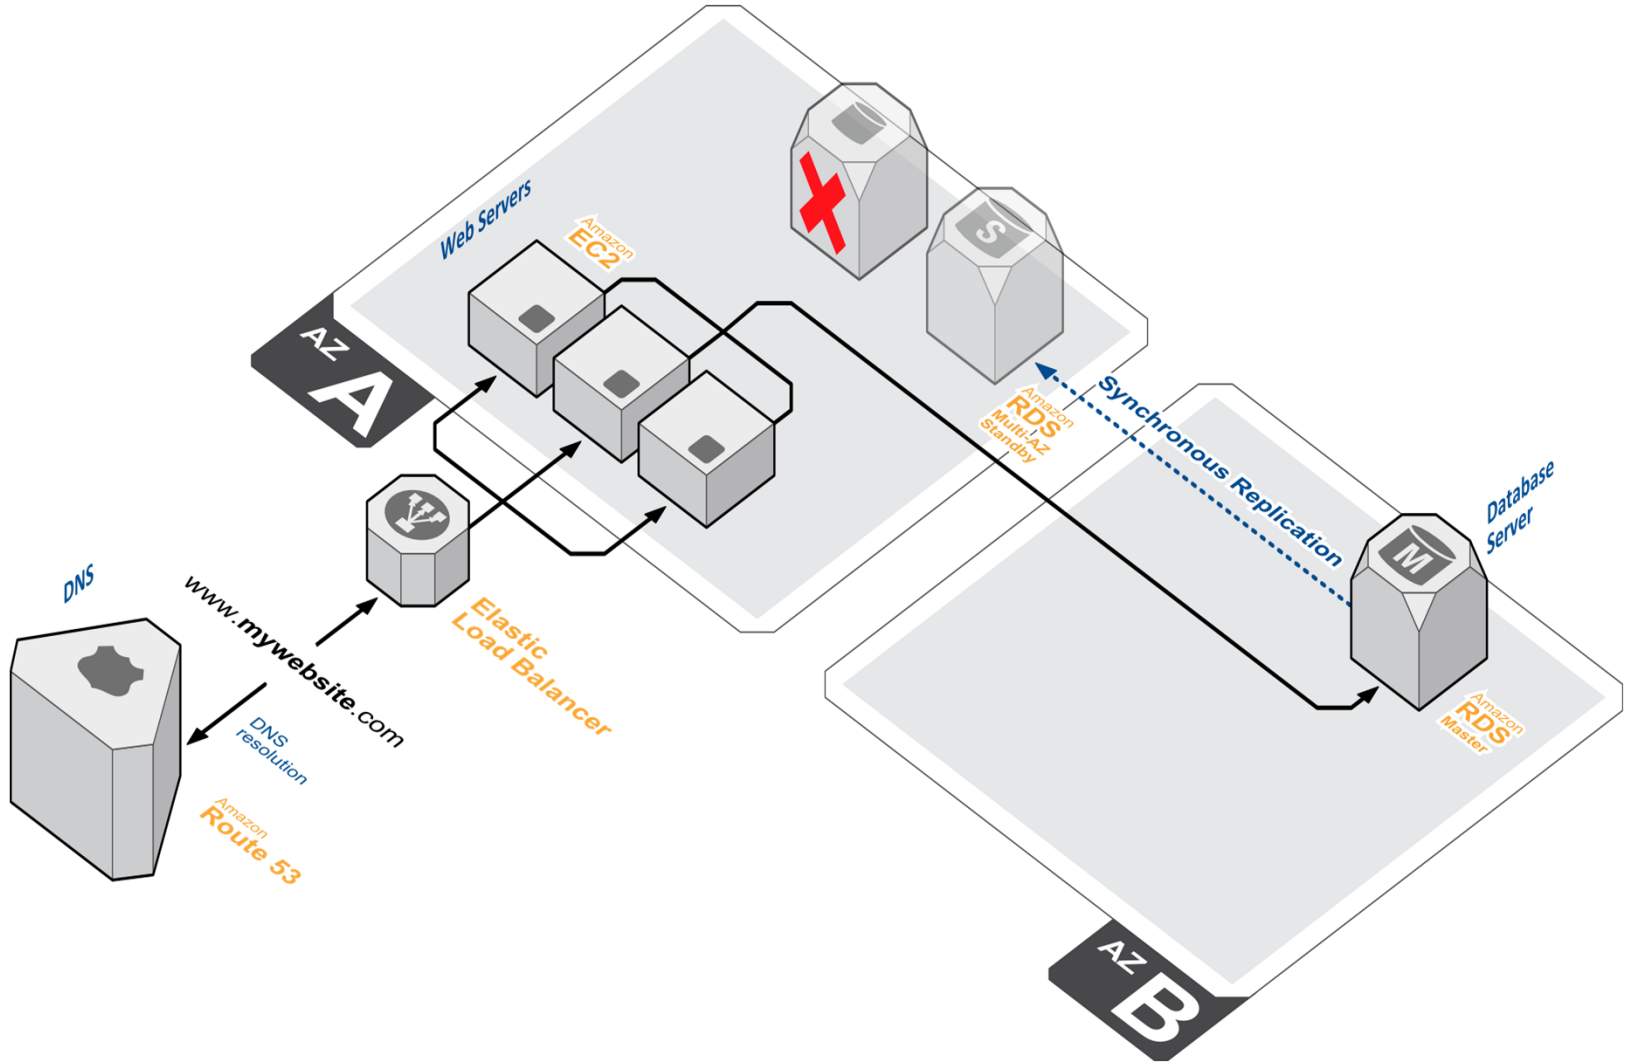
Web Servers
Amazon EC2
AZ A
Amazon RDS Multi-AZ Standby
Synchronous Replication
Database Server
Amazon RDS Master
AZ B
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancer

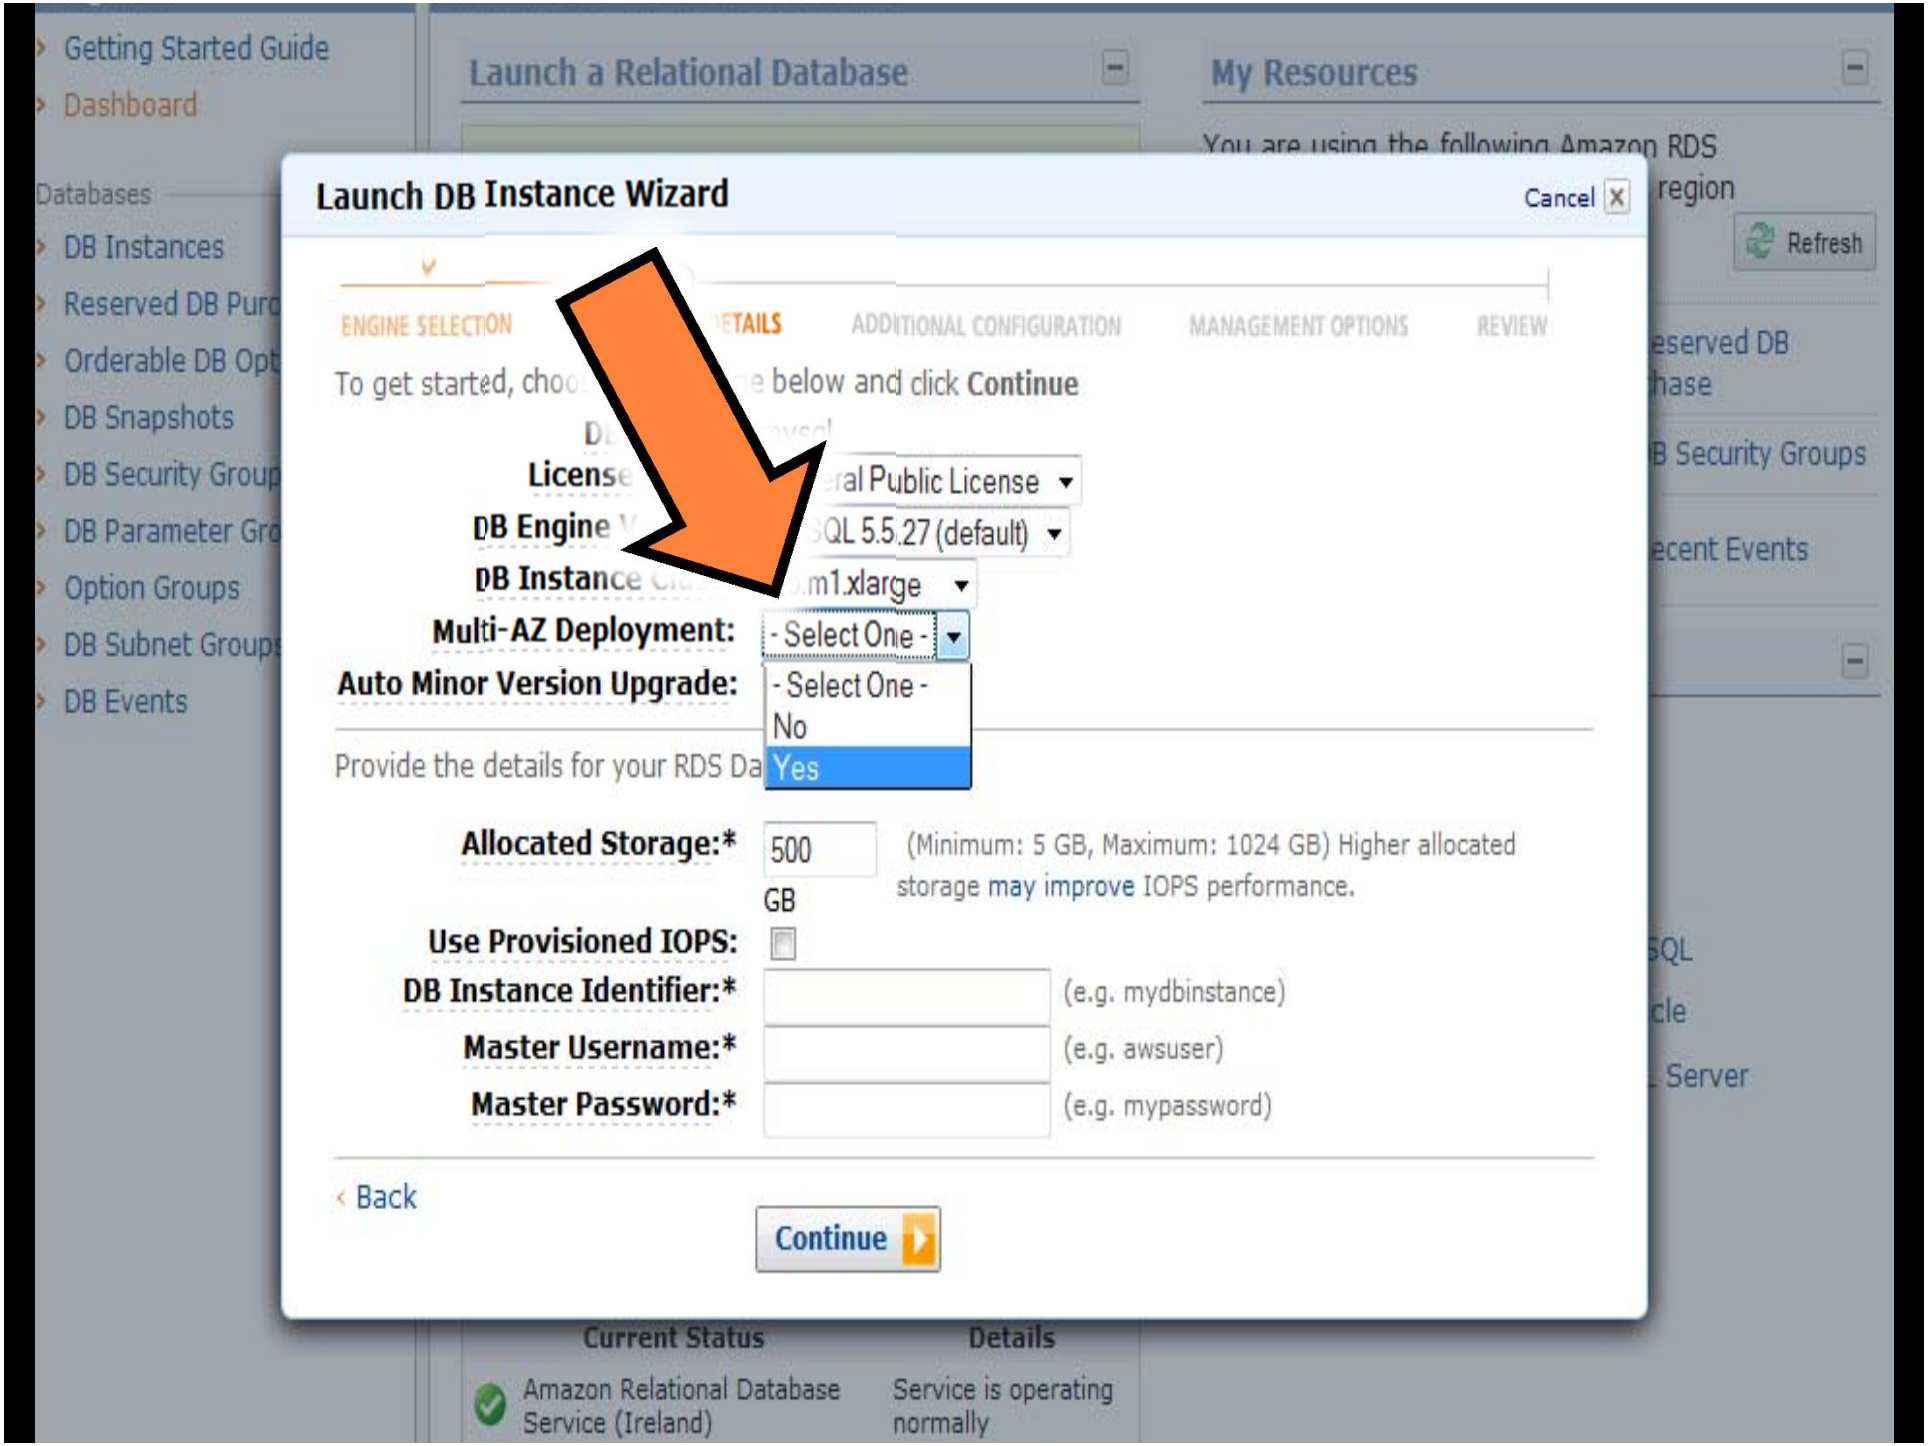
Launch DB Instance Wizard
Cancel
ENGINE SELECTION
ADDITIONAL CONFIGURATION
MANAGEMENT OPTIONS
REVIEW
To get started, choo
below and click Continue
Licens
ral Public License
DB Engine V
QL 5.5.27 (default)
DB Instance
m1.xlarge
Multi-AZ Deployment:
- Select One -
Auto Minor Version Upgrade:
- Select One -
No
Yes
Provide the details for your RDS Da
Allocated Storage:*
500
GB
(Minimum: 5 GB, Maximum: 1024 GB) Higher allocated storage may improve IOPS performance.
Use Provisioned IOPS:
DB Instance Identifier:*
(e.g. mydbinstance)
Master Username:*
(e.g. awsuser)
Master Password:*
(e.g. mypassword)
Back
Continue
Launch a Relational Database
My Resources
You are using the following Amazon RDS
region
Refresh
Getting Started Guide
Dashboard
Databases
DB Instances
Reserved DB Purc
Orderable DB Opt
DB Snapshots
DB Security Group
DB Parameter Gro
Option Groups
DB Subnet Groups
DB Events
eserved DB
hase
B Security Groups
ecent Events
SQL
cle
Server
Current Status
Details
Amazon Relational Database Service (Ireland)
Service is operating normally

AMAZON ELB AND MULTIPLE AZs

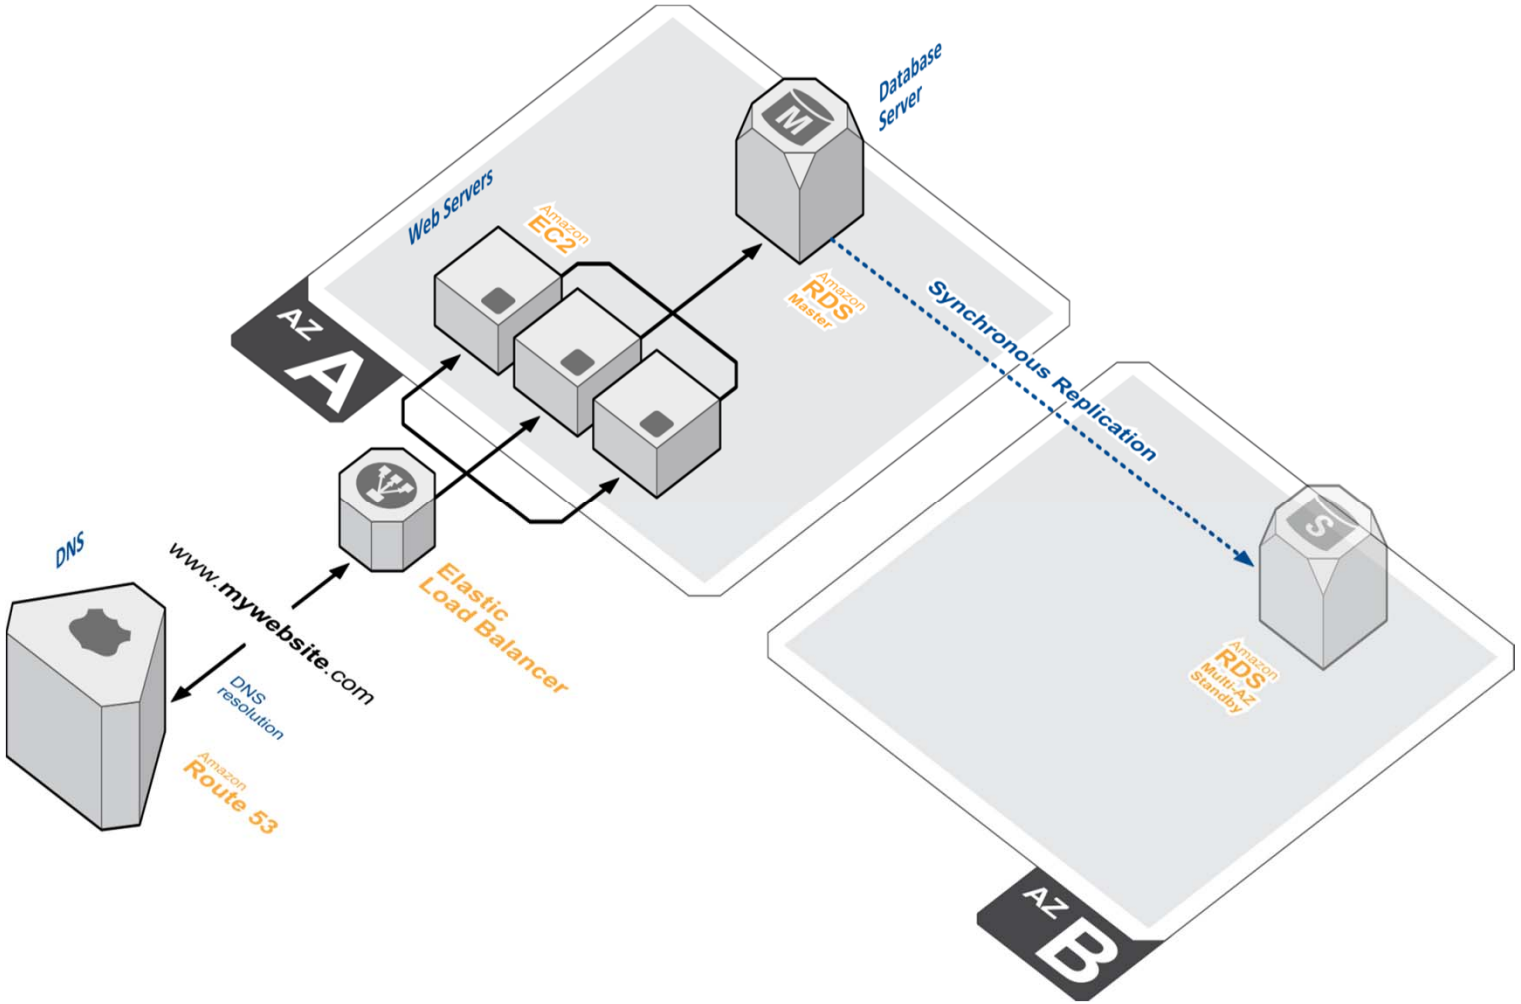
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
AZ B
Amazon RDS Multi-AZ Standby
Elastic Load Balancer
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

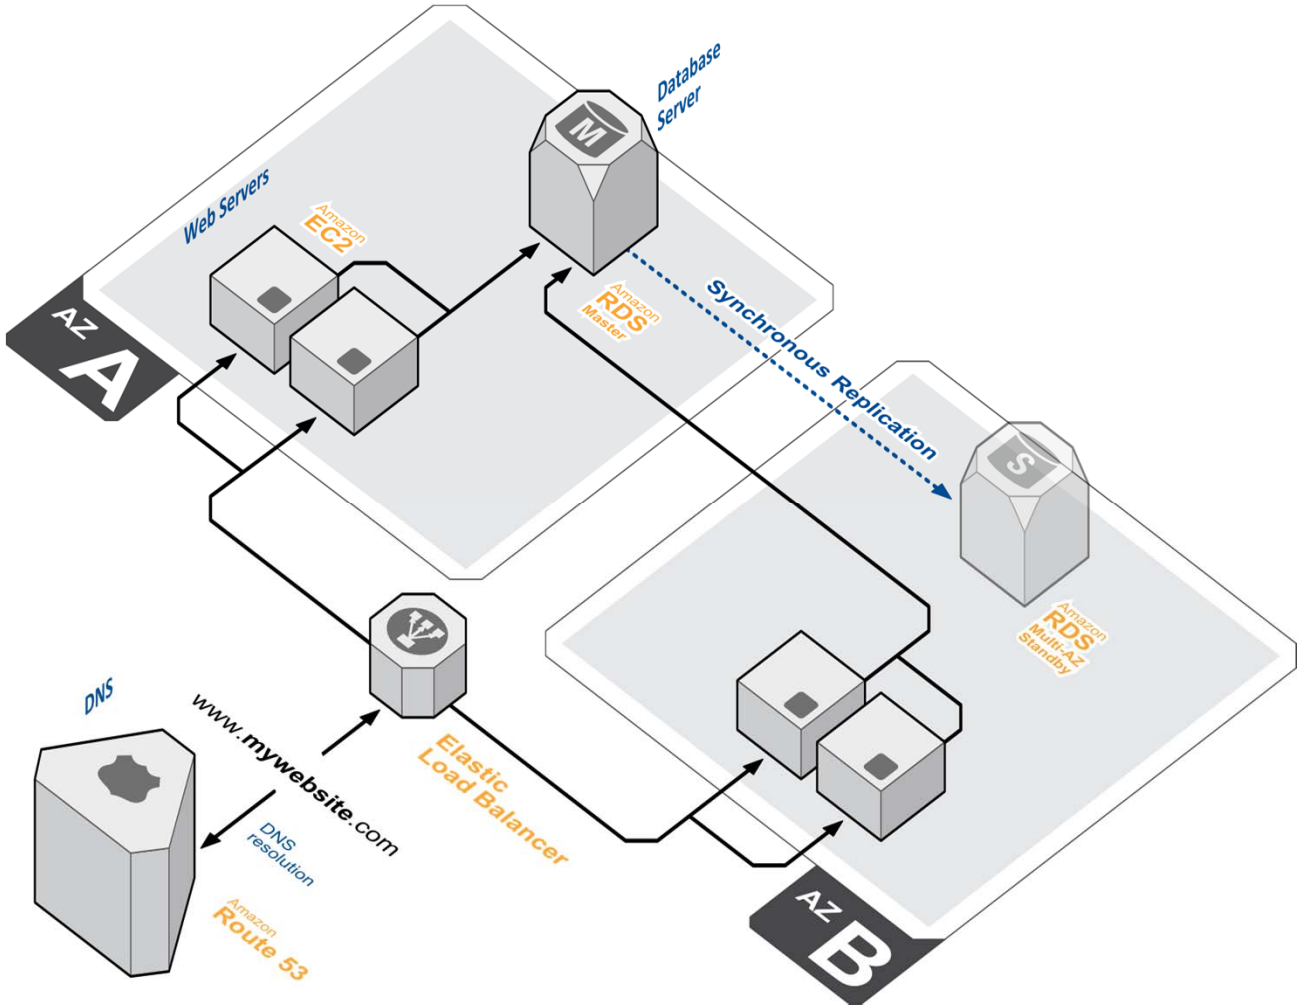
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
AZ B
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

#3

SCALING

●●●○○

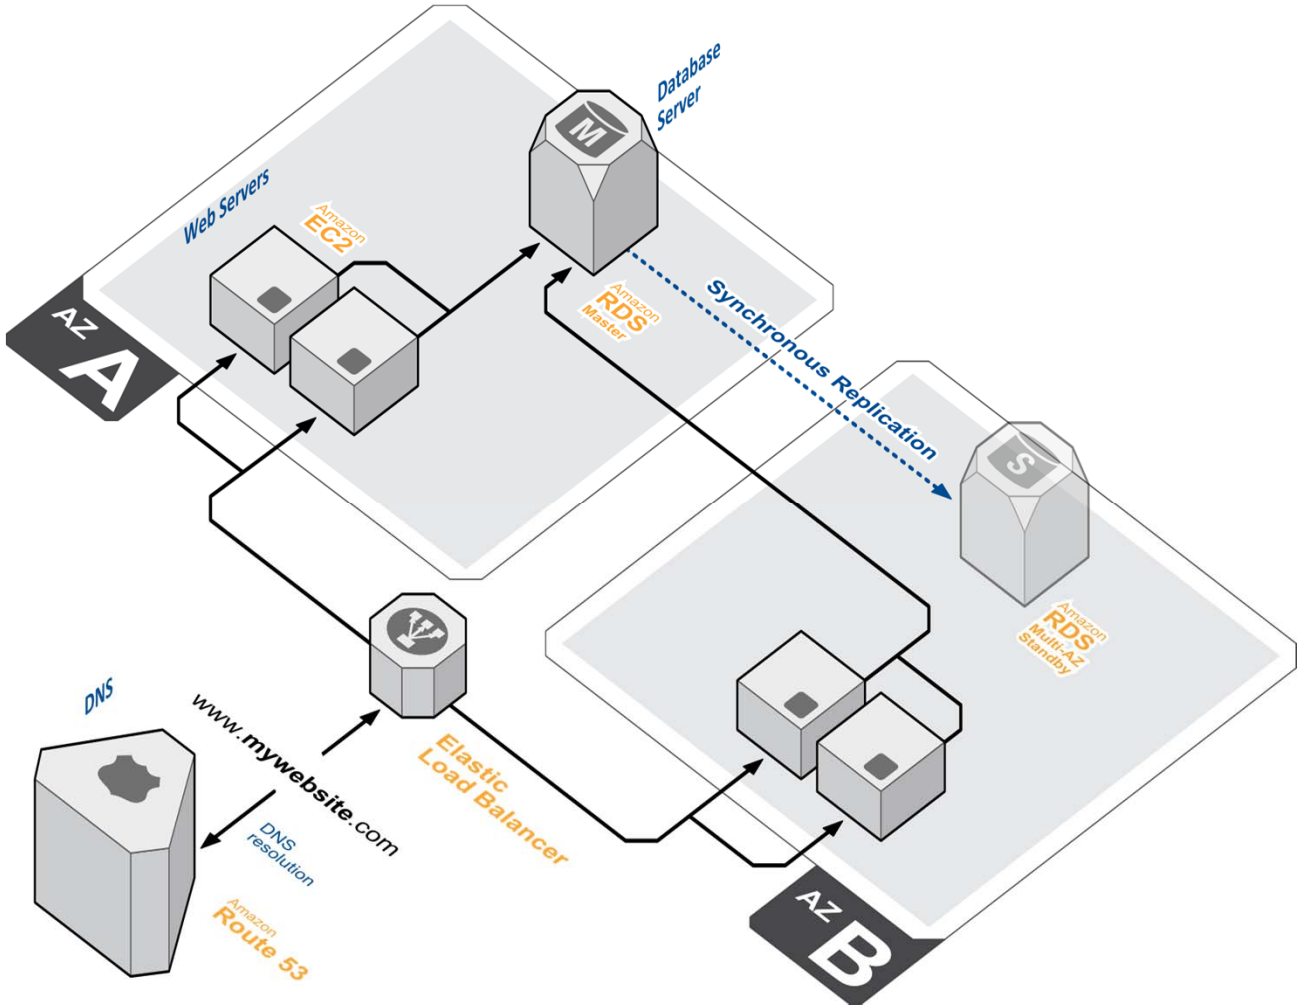
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
AZ B
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53

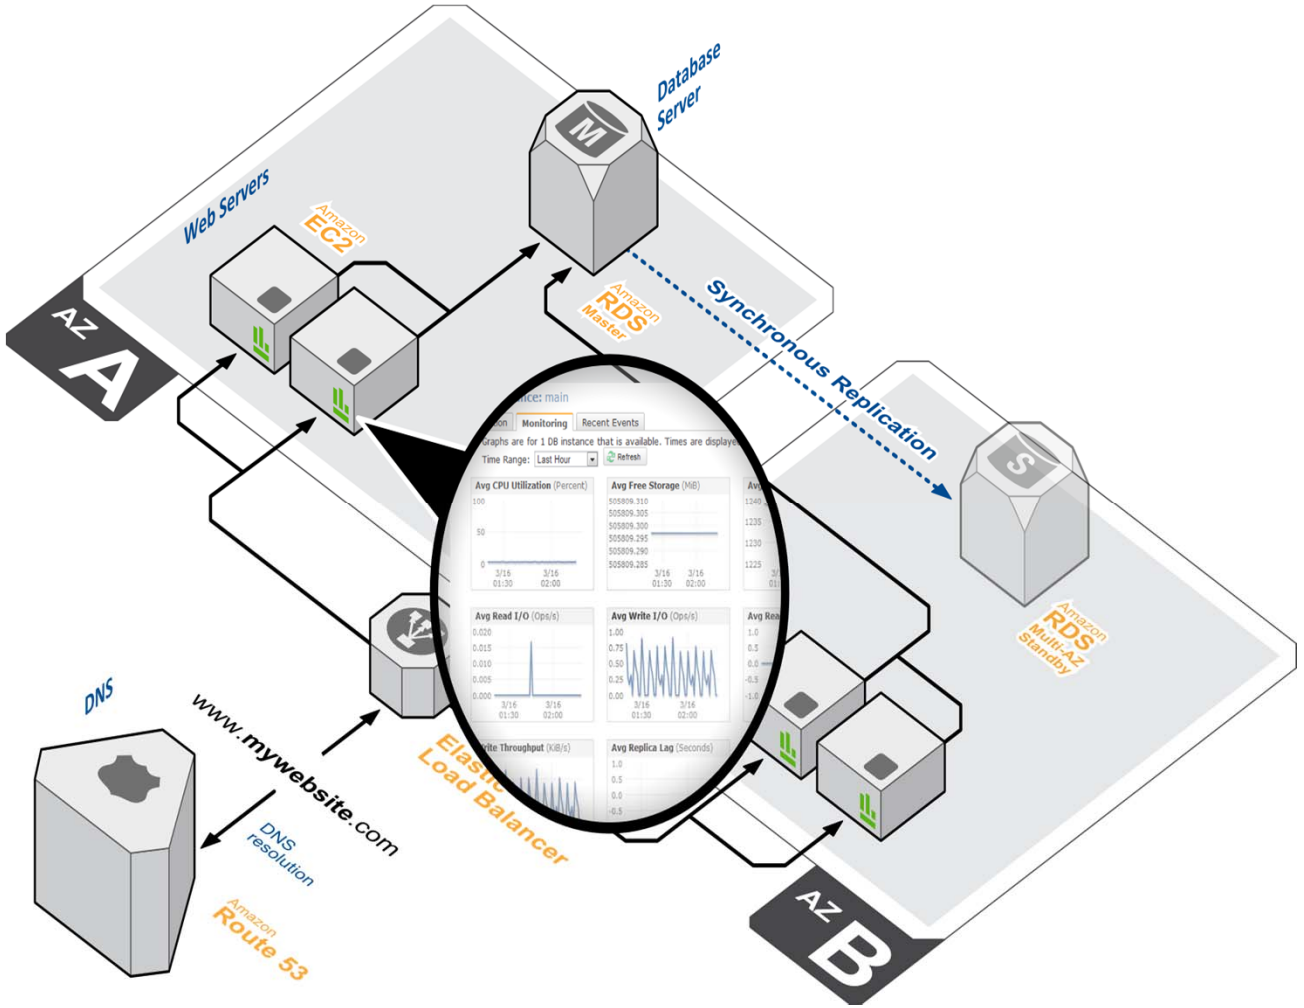
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
AZ B
Amazon RDS Multi-AZ Standby
Monitoring
Recent Events
Time Range: Last Hour
Refresh
Avg CPU Utilization (Percent)
Avg Free Storage (MiB)
Avg Read I/O (Ops/s)
Avg Write I/O (Ops/s)
Avg Replica Lag (Seconds)
Load Balancer
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

AUTO SCALING

SCALE UP/DOWN EC2 CAPACITY

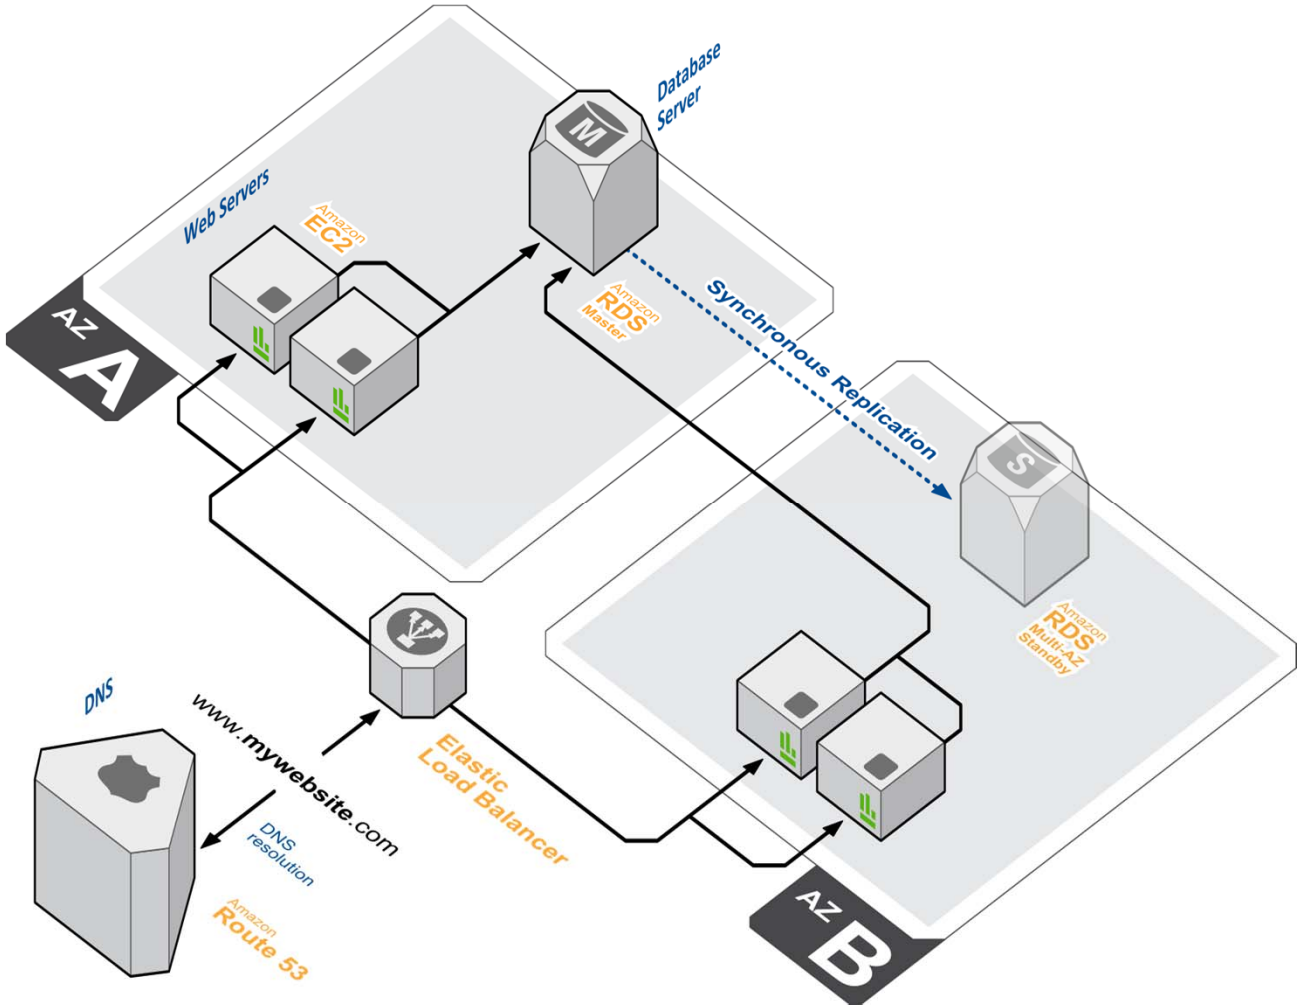
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53
AZ B

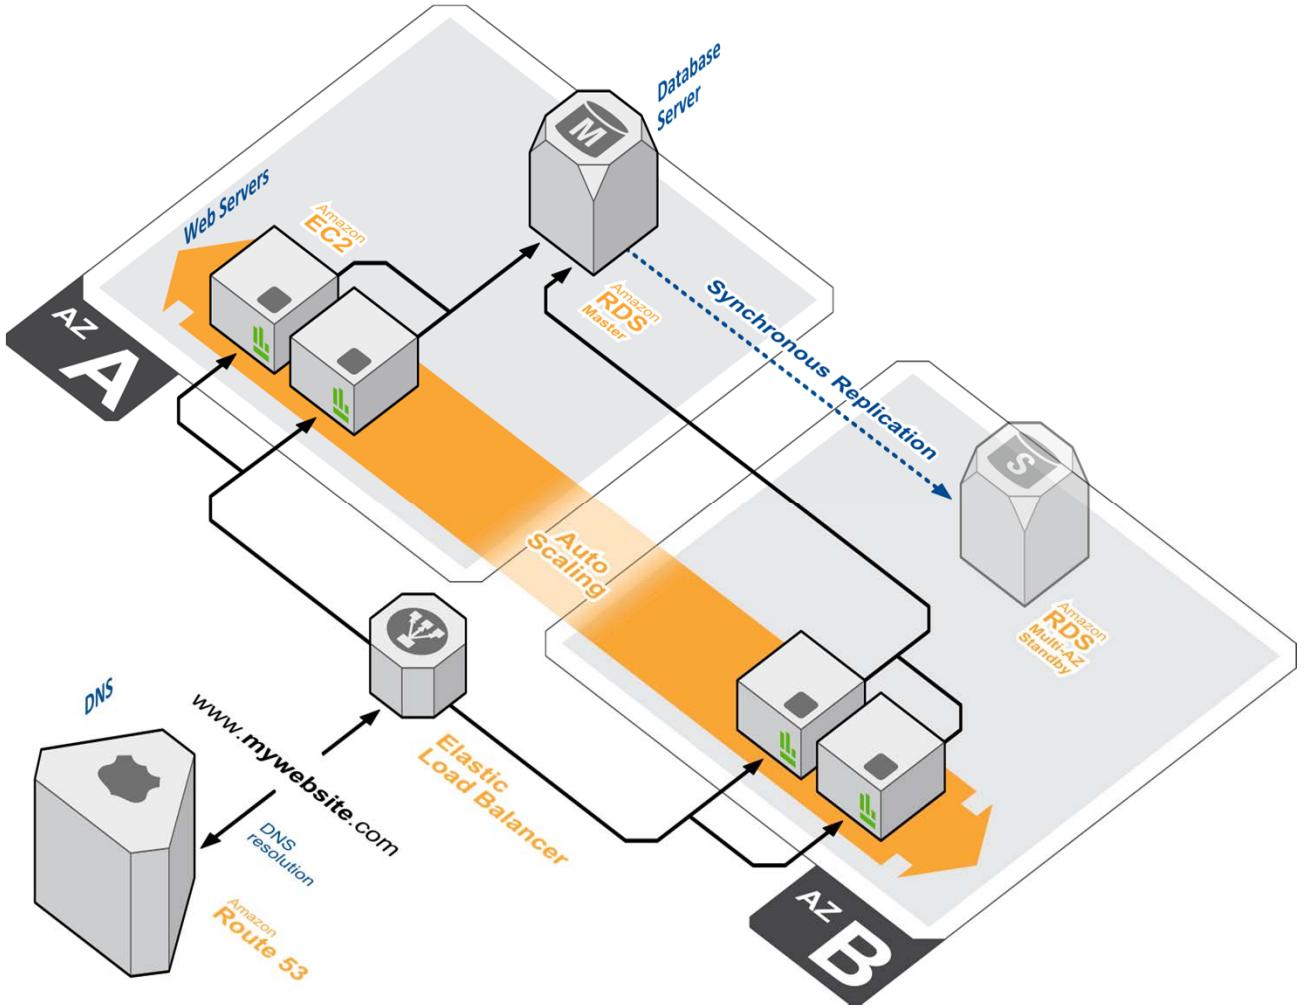
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
DNS
www.mywebsite.com
DNS resolution
Elastic Load Balancer
Amazon Route 53

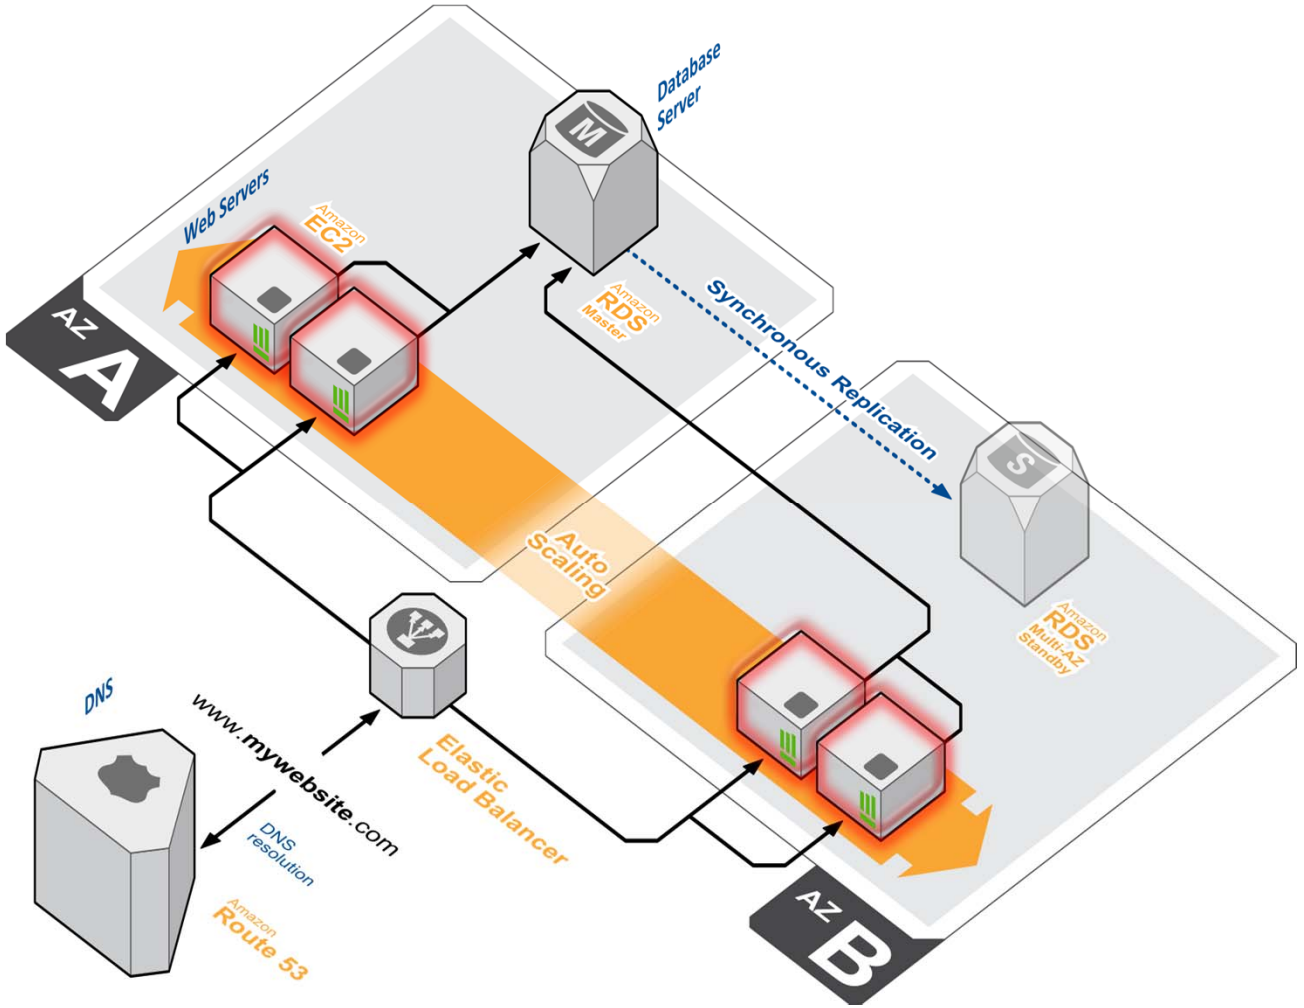
Database Server
M
Web Servers
Amazon EC2
AZ A
Amazon RDS Master
Synchronous Replication
S
Auto Scaling
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53
AZ B

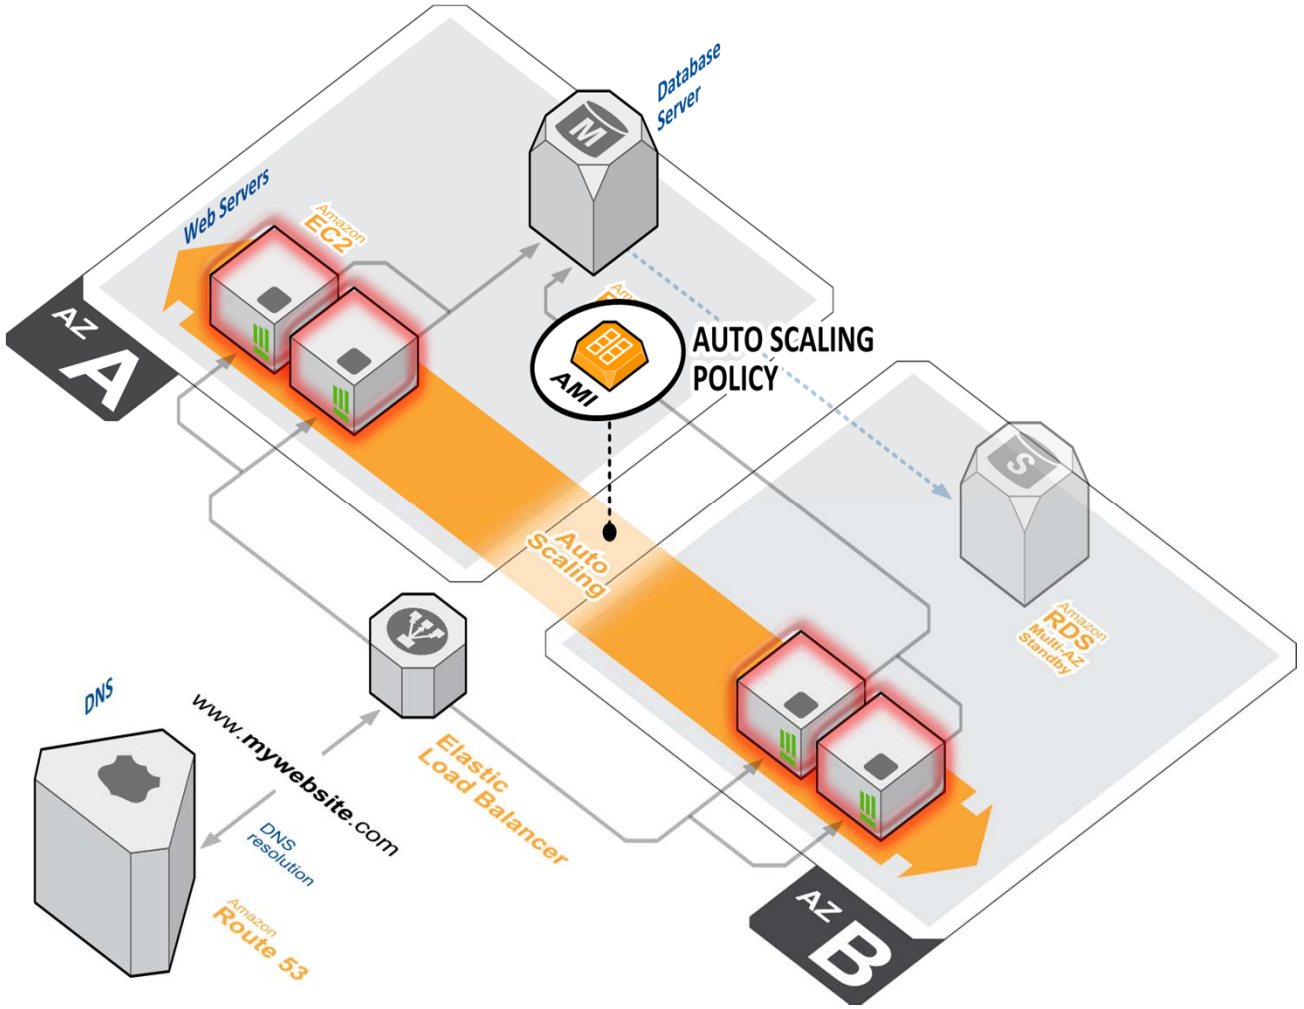
Database Server
Web Servers
Amazon EC2
AZ A
AUTO SCALING POLICY
AMI
Auto Scaling
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancer
AZ B

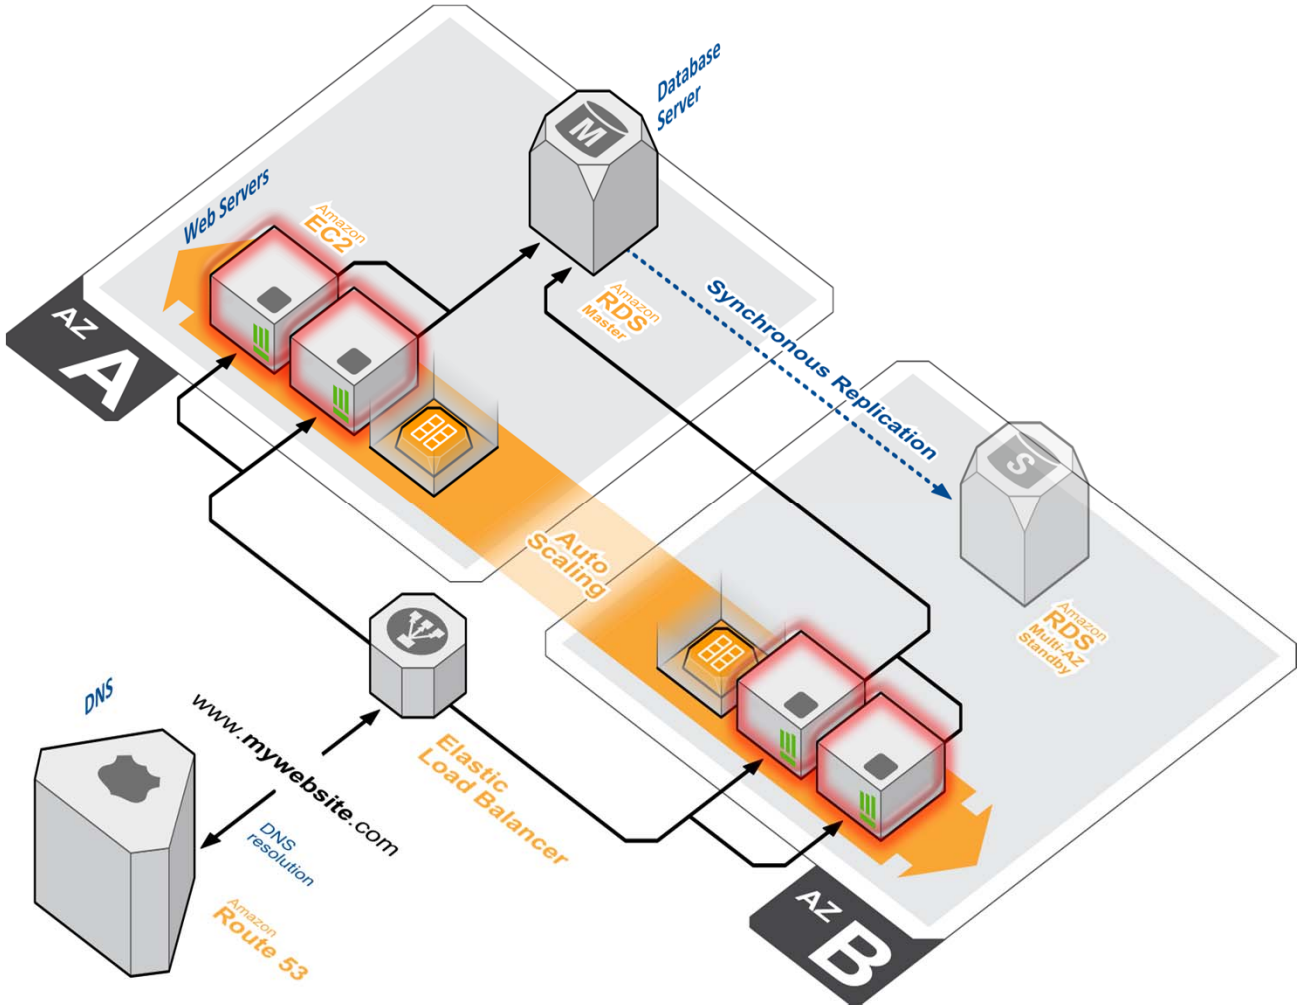
Database Server
M
Web Servers
Amazon EC2
AZ A
Amazon RDS Master
Synchronous Replication
S
Auto Scaling
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53
AZ B

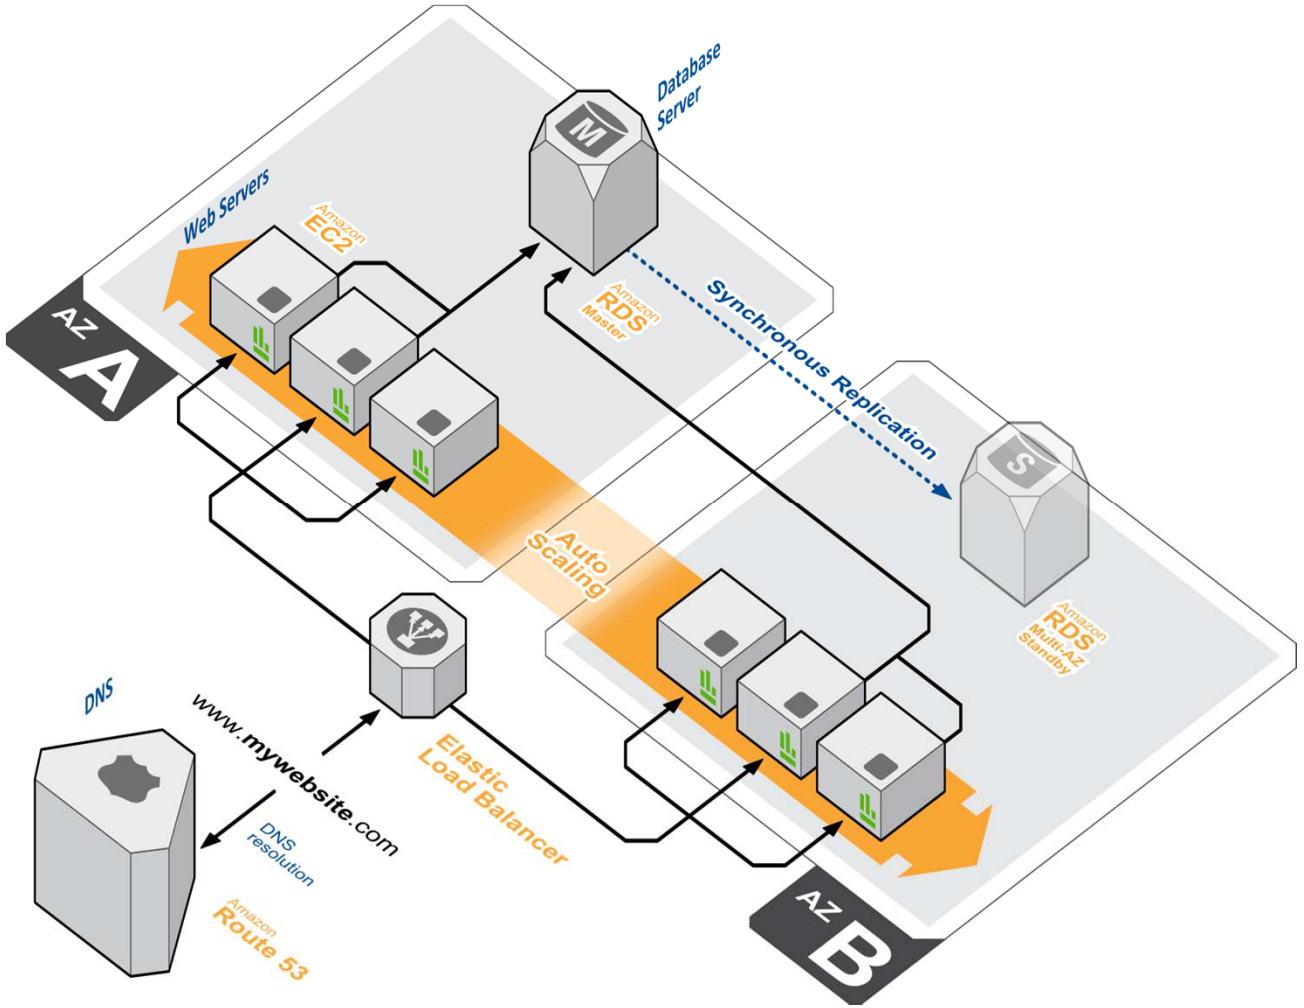
Database Server
M
Web Servers
Amazon EC2
AZ A
Amazon RDS Master
Synchronous Replication
S
Auto Scaling
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53
AZ B

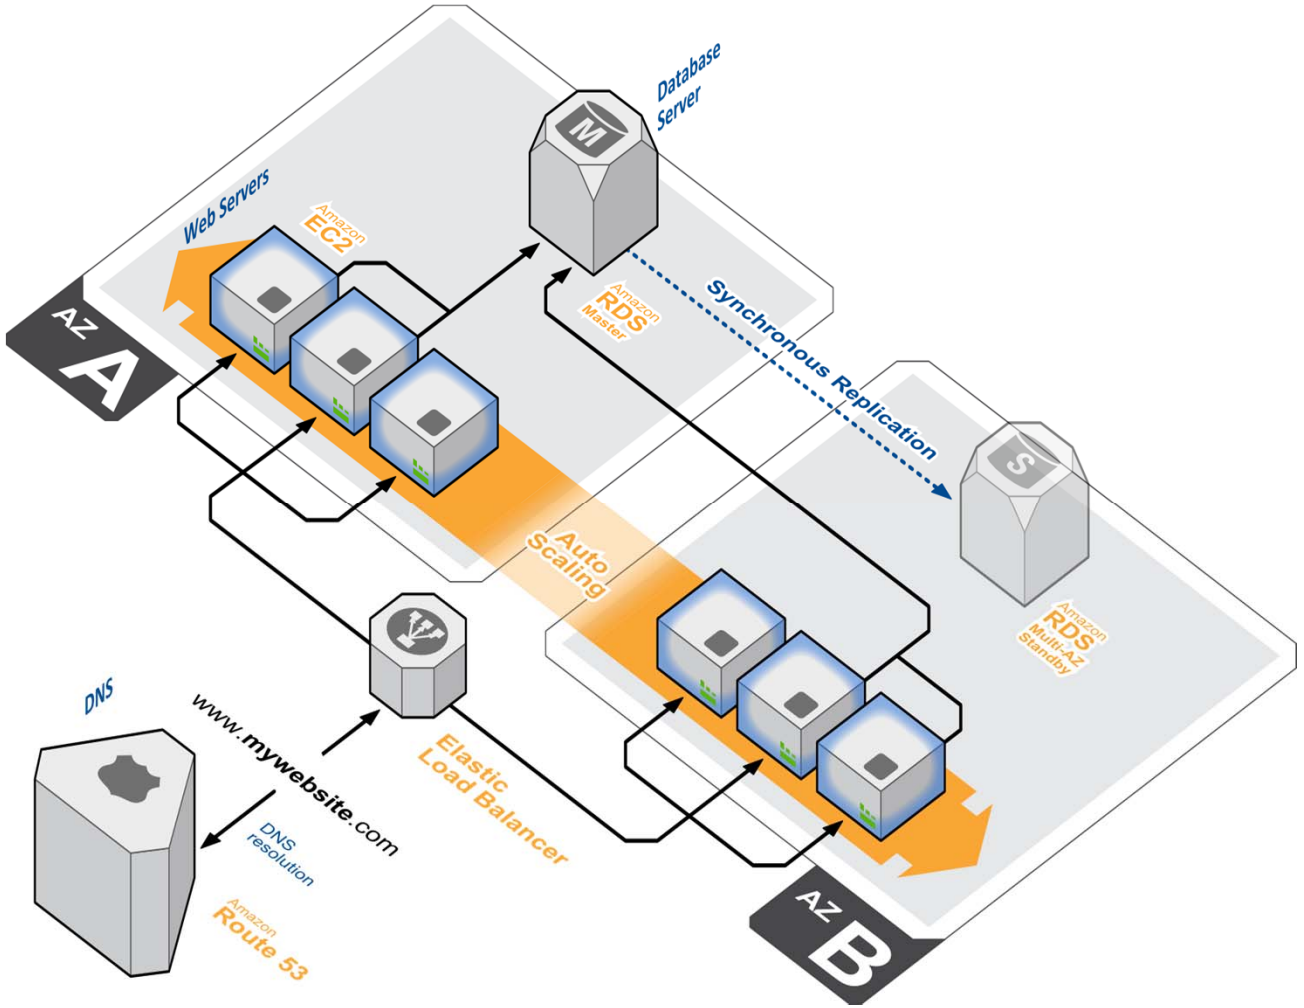
Database Server
Web Servers
Amazon EC2
AZ A
Amazon RDS Master
Synchronous Replication
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53

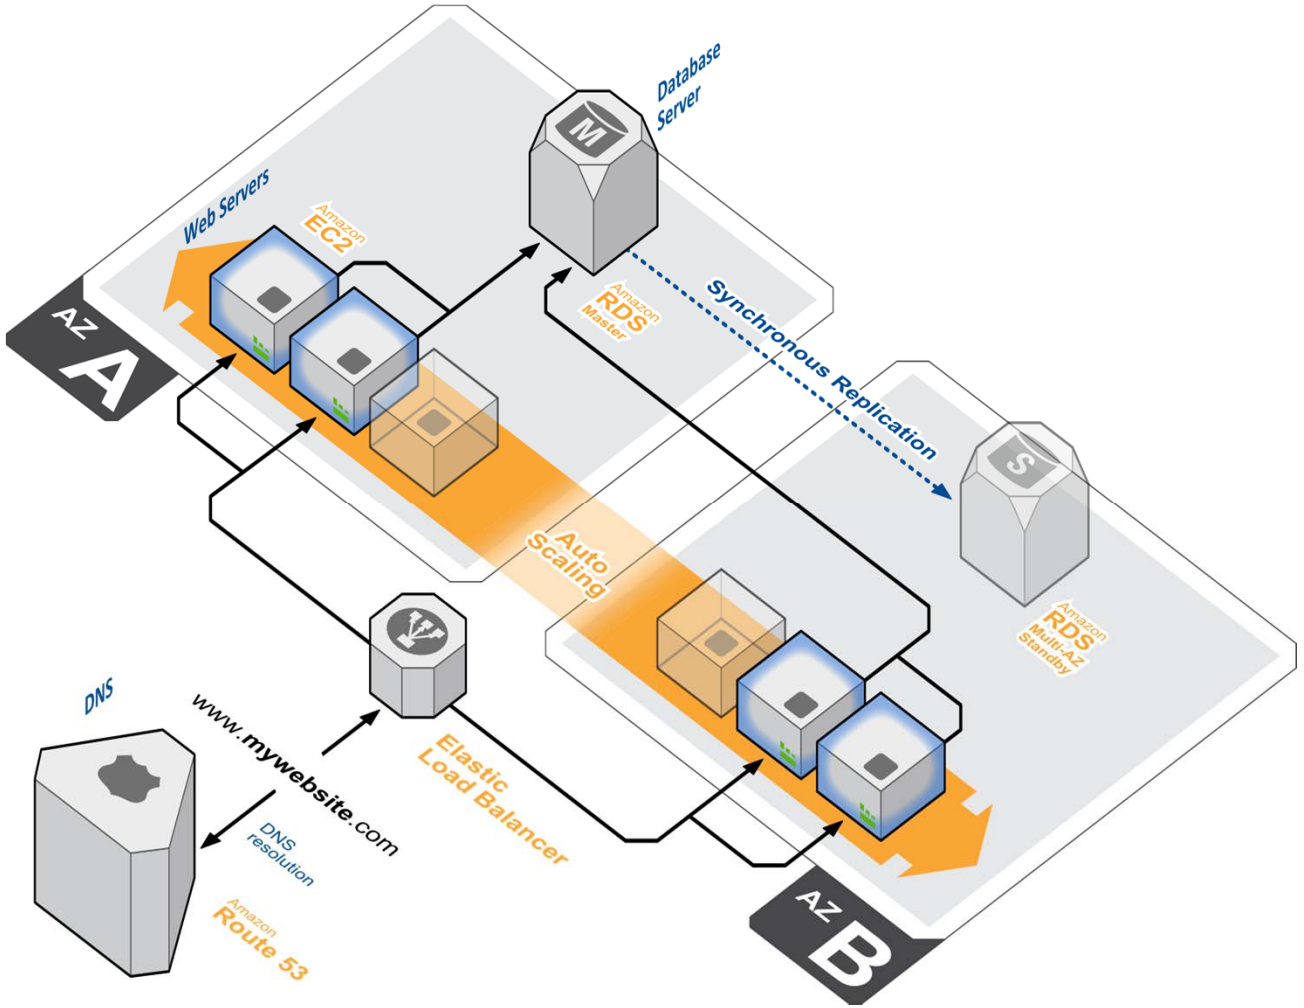
Database Server
Web Servers
Amazon EC2
AZ A
Amazon RDS Master
Synchronous Replication
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
Elastic Load Balancer
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

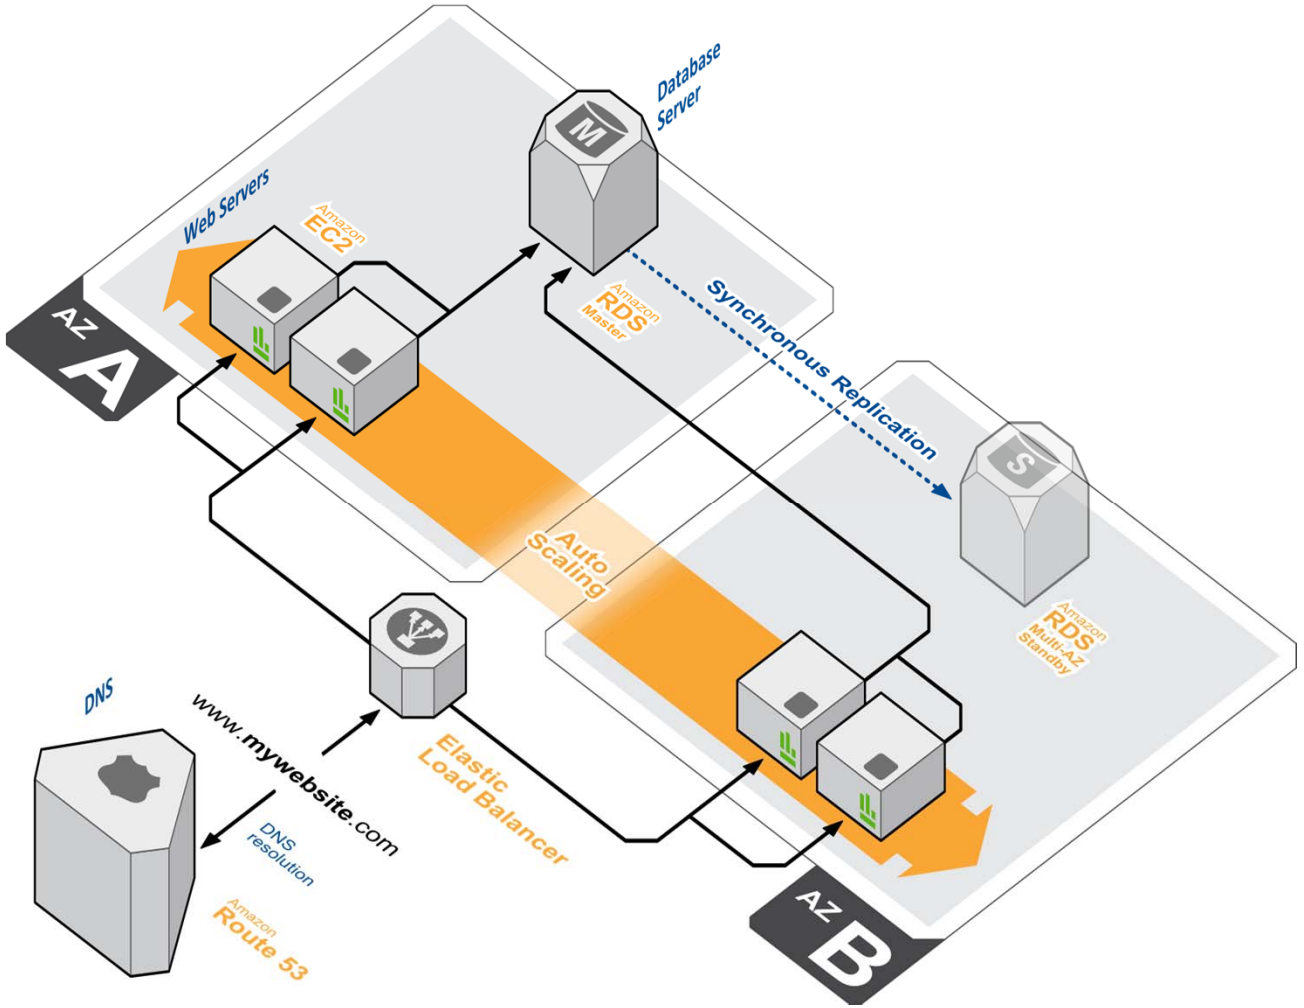
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
DNS
www.mywebsite.com
DNS resolution
Elastic Load Balancer
Amazon Route 53

UNEXPECTED
SPIKES

EXPECTED
SPIKES

DATA TIER

AWS BUILDING BLOCKS

Inherently Highly Available and Fault Tolerant Services

- Amazon S3
- Amazon DynamoDB
- Amazon CloudFront
- Amazon Route53
- Elastic Load Balancing
- Amazon SQS
- Amazon SNS
- Amazon SES
- Amazon SWF
- …

Highly Available with the right architecture

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon VPC

AWS BUILDING BLOCKS

Inherently Highly Available and Fault Tolerant Services

- Amazon S3
- Amazon DynamoDB
- Amazon CloudFront
- Amazon Route53
- Elastic Load Balancing
- Amazon SQS
- Amazon SNS
- Amazon SES
- Amazon SWF
- …

Highly Available with the right architecture

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon VPC

AMAZON S3

SIMPLE STORAGE SERVICE

99.999999%

DURABILITY

SCALABLE
& AVAILABLE
NO CAPACITY PLANNING REQUIRED

ZERO ADMINISTRATION

AWS BUILDING BLOCKS

Inherently Highly Available and Fault Tolerant Services

- Amazon S3
- Amazon DynamoDB
- Amazon CloudFront
- Amazon Route53
- Elastic Load Balancing
- Amazon SQS
- Amazon SNS
- Amazon SES
- Amazon SWF
- …

Highly Available with the right architecture

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon VPC

AMAZON DYNAMODB

HIGH-PERFORMANCE, FULLY MANAGED NoSQL DATABASE SERVICE

LOW LATENCY

AVERAGE READS < **5**MS,
WRITES < **10**MS

PREDICTABLE PERFORMANCE

PROVISIONES THROUGHPUT

SEAMLESS SCALABILITY

LIVE REPARTITIONING

ZERO ADMINISTRATION

AWS BUILDING BLOCKS

Inherently Highly Available and Fault Tolerant Services

- Amazon S3
- Amazon DynamoDB
- Amazon CloudFront
- Amazon Route53
- Elastic Load Balancing
- Amazon SQS
- Amazon SNS
- Amazon SES
- Amazon SWF
- …

Highly Available with the right architecture

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon VPC

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

#4

SELF-HEALING

●●●●○

HEALTH CHECKS

+

AUTO SCALING

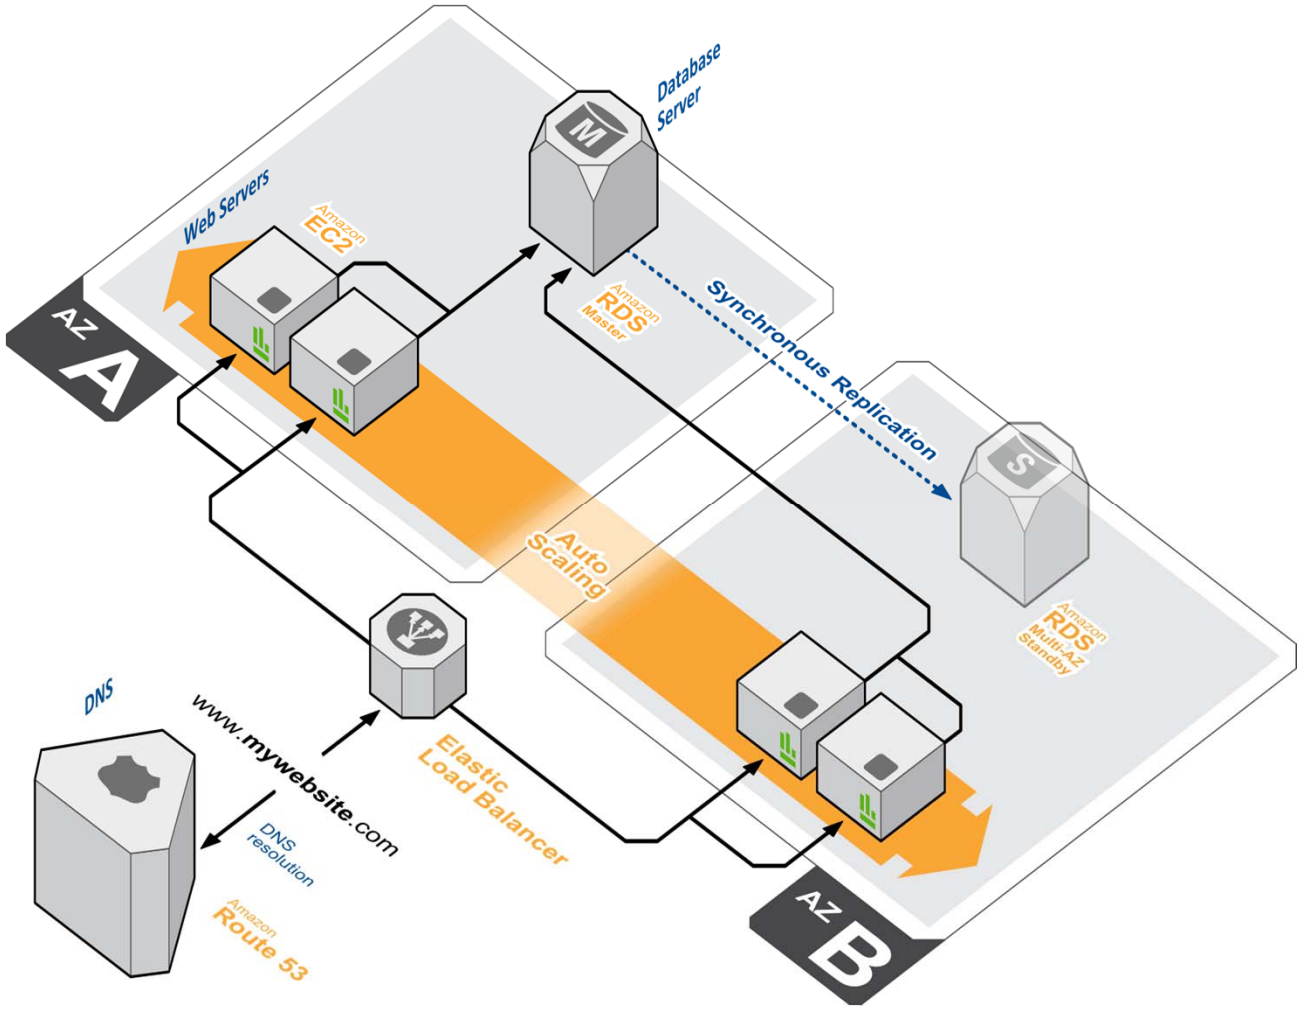
Database Server
M
Web Servers
Amazon EC2
AZ A
Amazon RDS Master
Synchronous Replication
S
Auto Scaling
Amazon RDS Multi-AZ Standby
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53
AZ B

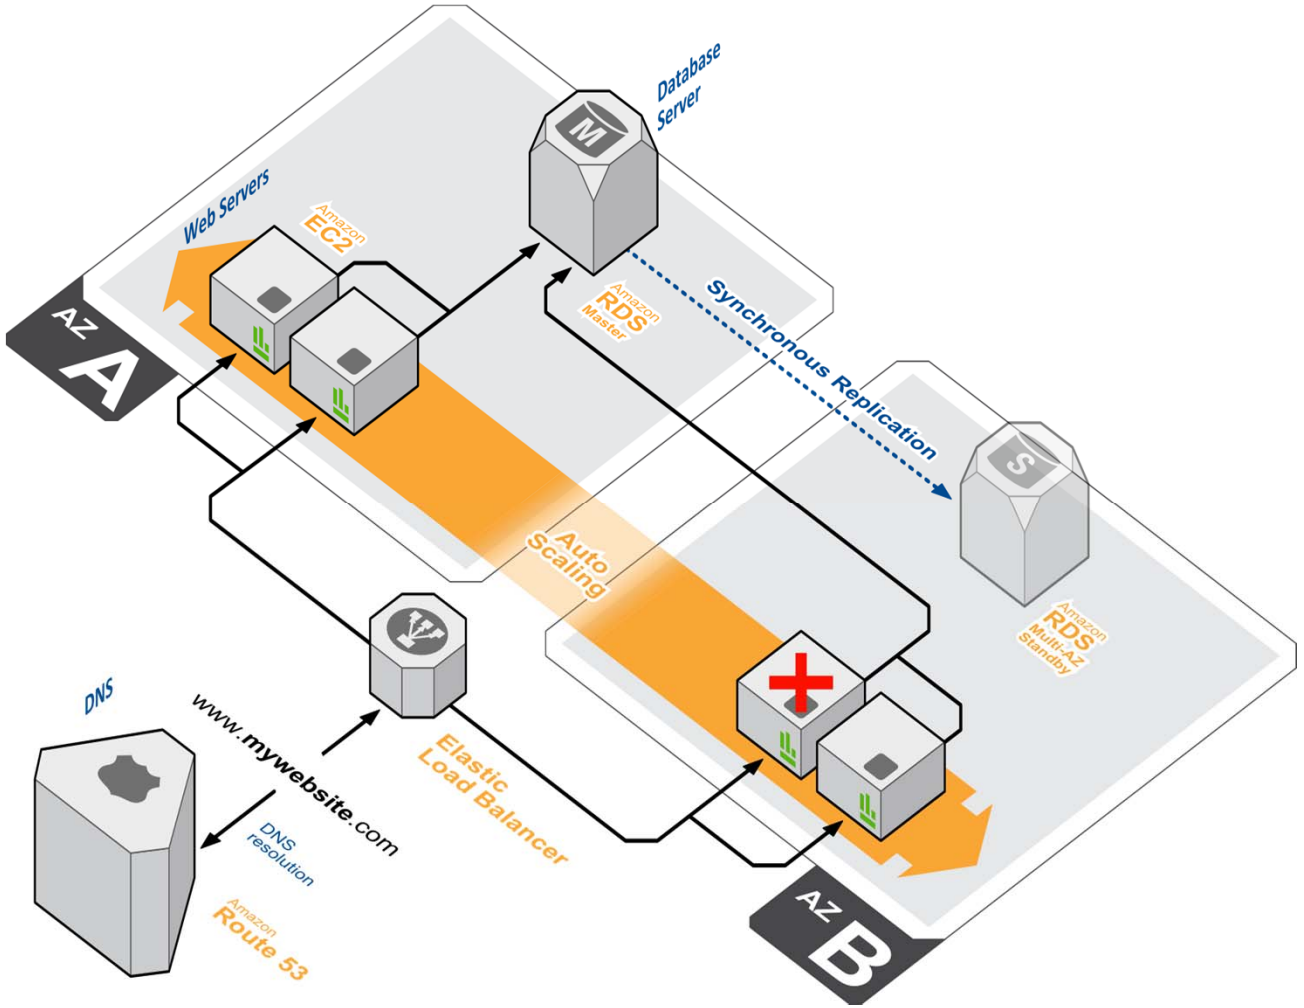
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53

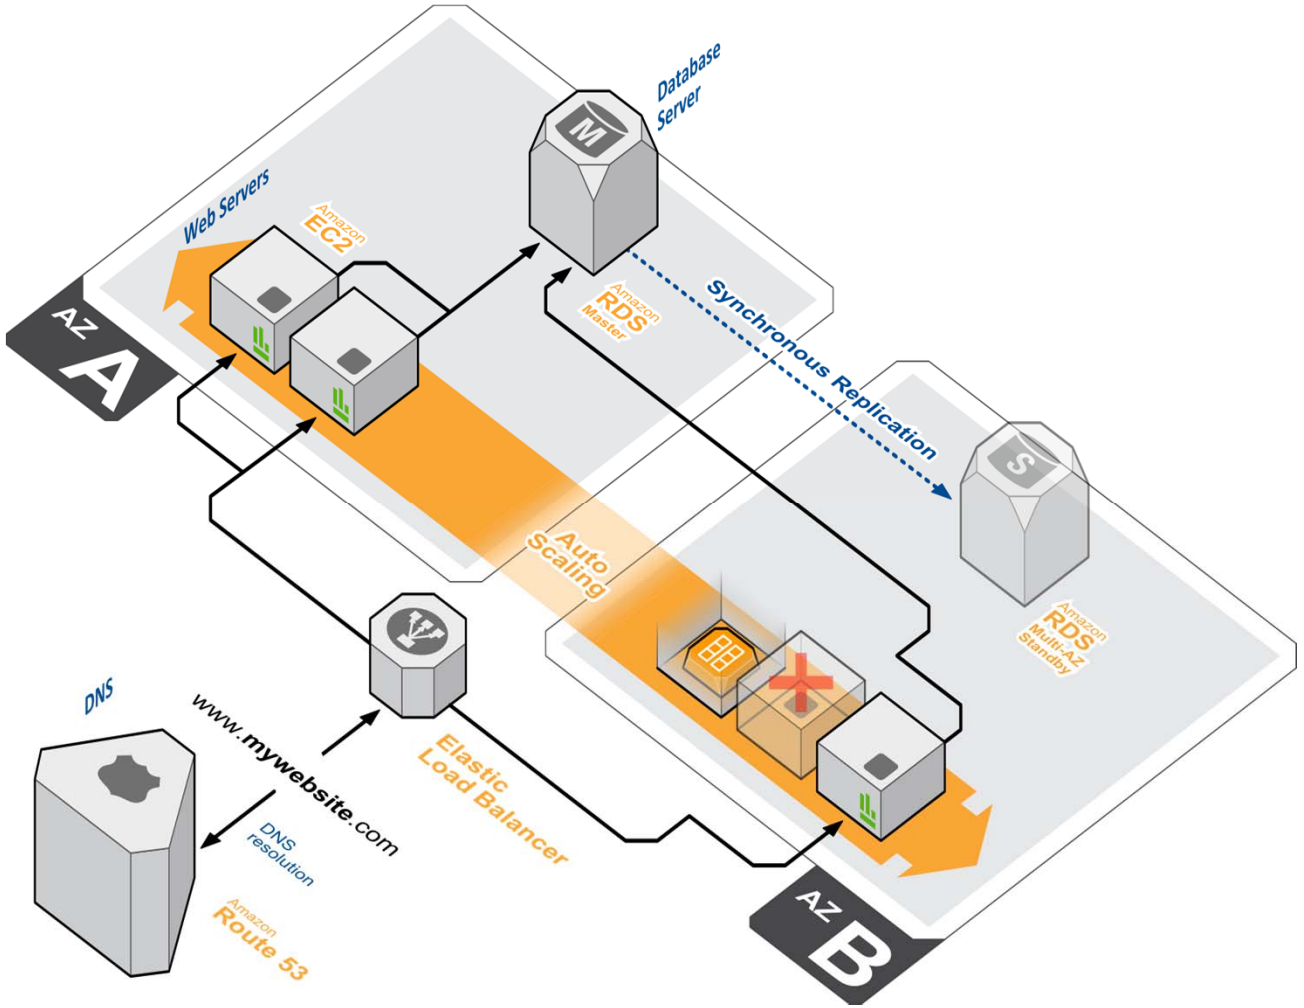
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
DNS
www.mywebsite.com
Elastic Load Balancer
DNS resolution
Amazon Route 53

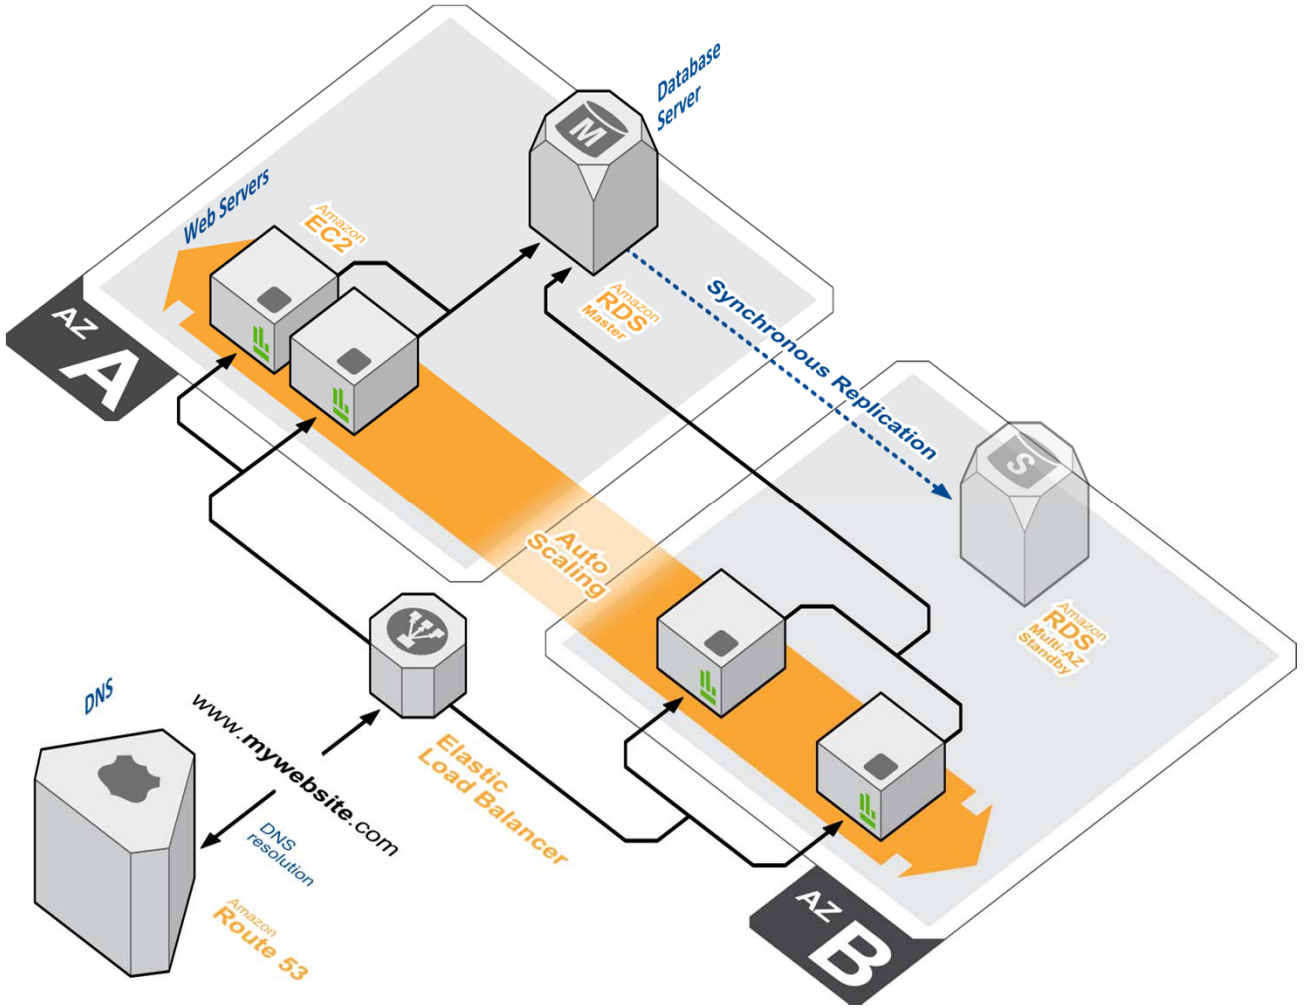
Database Server
Web Servers
Amazon EC2
Amazon RDS Master
Synchronous Replication
AZ A
Auto Scaling
Amazon RDS Multi-AZ Standby
AZ B
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancer

HEALTH CHECKS

+

AUTO SCALING

=

SELF-HEALING

DEGRADED MODE

AMAZON S3

STATIC WEBSITE

+

AMAZON ROUTE53

WEIGHTED RESOLUTION

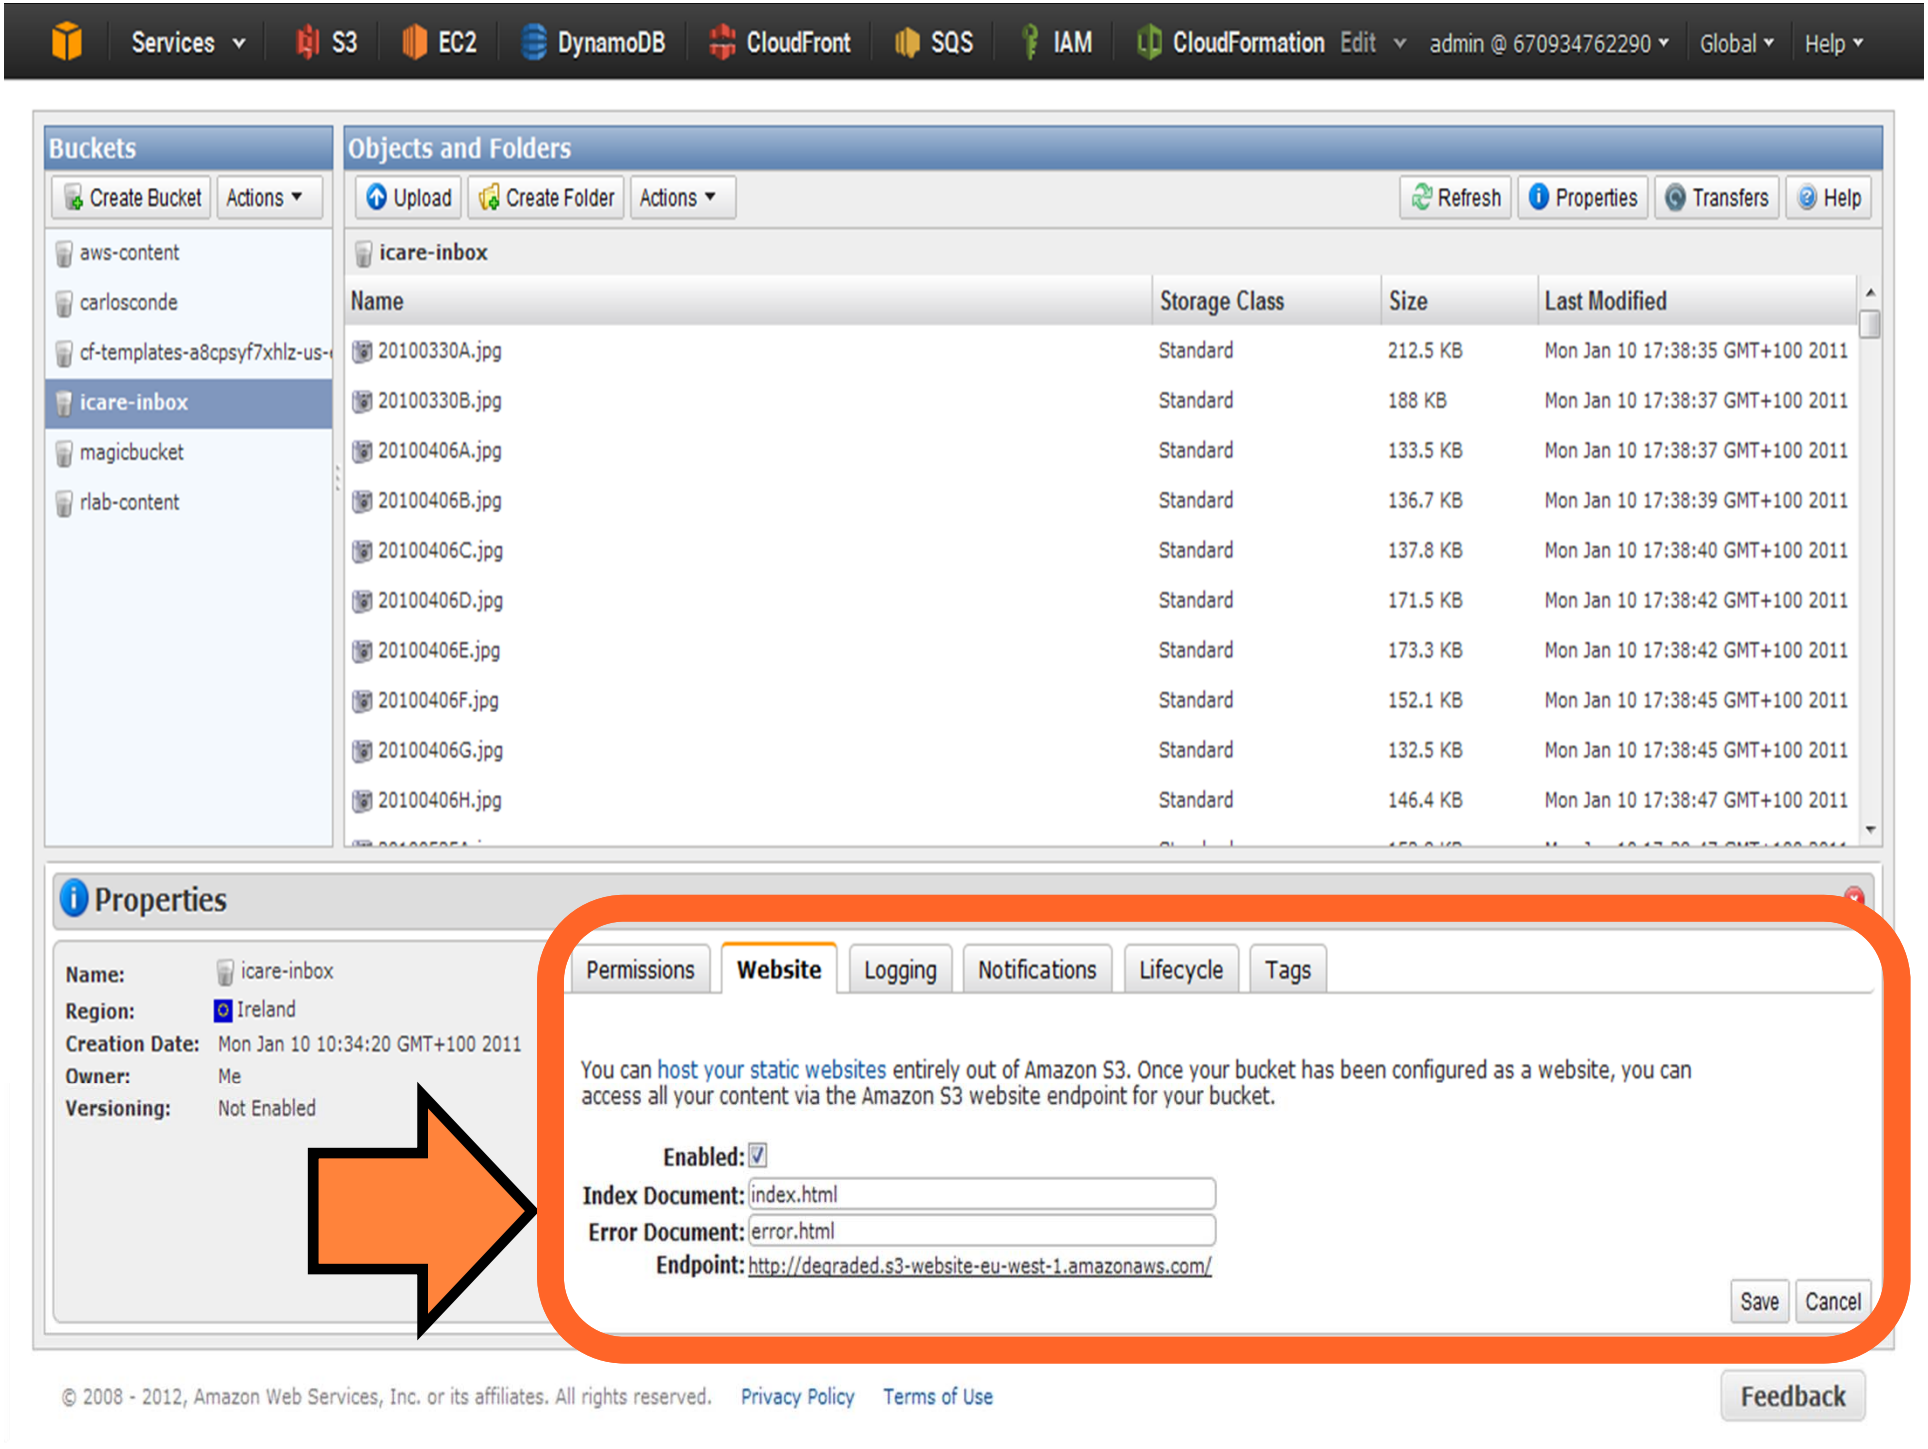
Services
S3
EC2
DynamoDB
CloudFront
SQS
IAM
CloudFormation
Edit
admin @ 670934762290
Global
Help
Buckets
Create Bucket
Actions
aws-content
carlosconde
cf-templates-a8cpsyf7xhlz-us-
icare-inbox
magicbucket
rlab-content
Objects and Folders
Upload
Create Folder
Actions
Refresh
Properties
Transfers
Help
icare-inbox
Name
Storage Class
Size
Last Modified
20100330A.jpg Standard 212.5 KB Mon Jan 10 17:38:35 GMT+100 2011
20100330B.jpg Standard 188 KB Mon Jan 10 17:38:37 GMT+100 2011
20100406A.jpg Standard 133.5 KB Mon Jan 10 17:38:37 GMT+100 2011
20100406B.jpg Standard 136.7 KB Mon Jan 10 17:38:39 GMT+100 2011
20100406C.jpg Standard 137.8 KB Mon Jan 10 17:38:40 GMT+100 2011
20100406D.jpg Standard 171.5 KB Mon Jan 10 17:38:42 GMT+100 2011
20100406E.jpg Standard 173.3 KB Mon Jan 10 17:38:42 GMT+100 2011
20100406F.jpg Standard 152.1 KB Mon Jan 10 17:38:45 GMT+100 2011
20100406G.jpg Standard 132.5 KB Mon Jan 10 17:38:45 GMT+100 2011
20100406H.jpg Standard 146.4 KB Mon Jan 10 17:38:47 GMT+100 2011
Properties
Name: icare-inbox
Region: Ireland
Creation Date: Mon Jan 10 10:34:20 GMT+100 2011
Owner: Me
Versioning: Not Enabled
Permissions
Website
Logging
Notifications
Lifecycle
Tags
You can host your static websites entirely out of Amazon S3. Once your bucket has been configured as a website, you can access all your content via the Amazon S3 website endpoint for your bucket.
Enabled:
Index Document: index.html
Error Document: error.html
Endpoint: http://degraded.s3-website-eu-west-1.amazonaws.com/
Save
Cancel
© 2008 - 2012, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use
Feedback

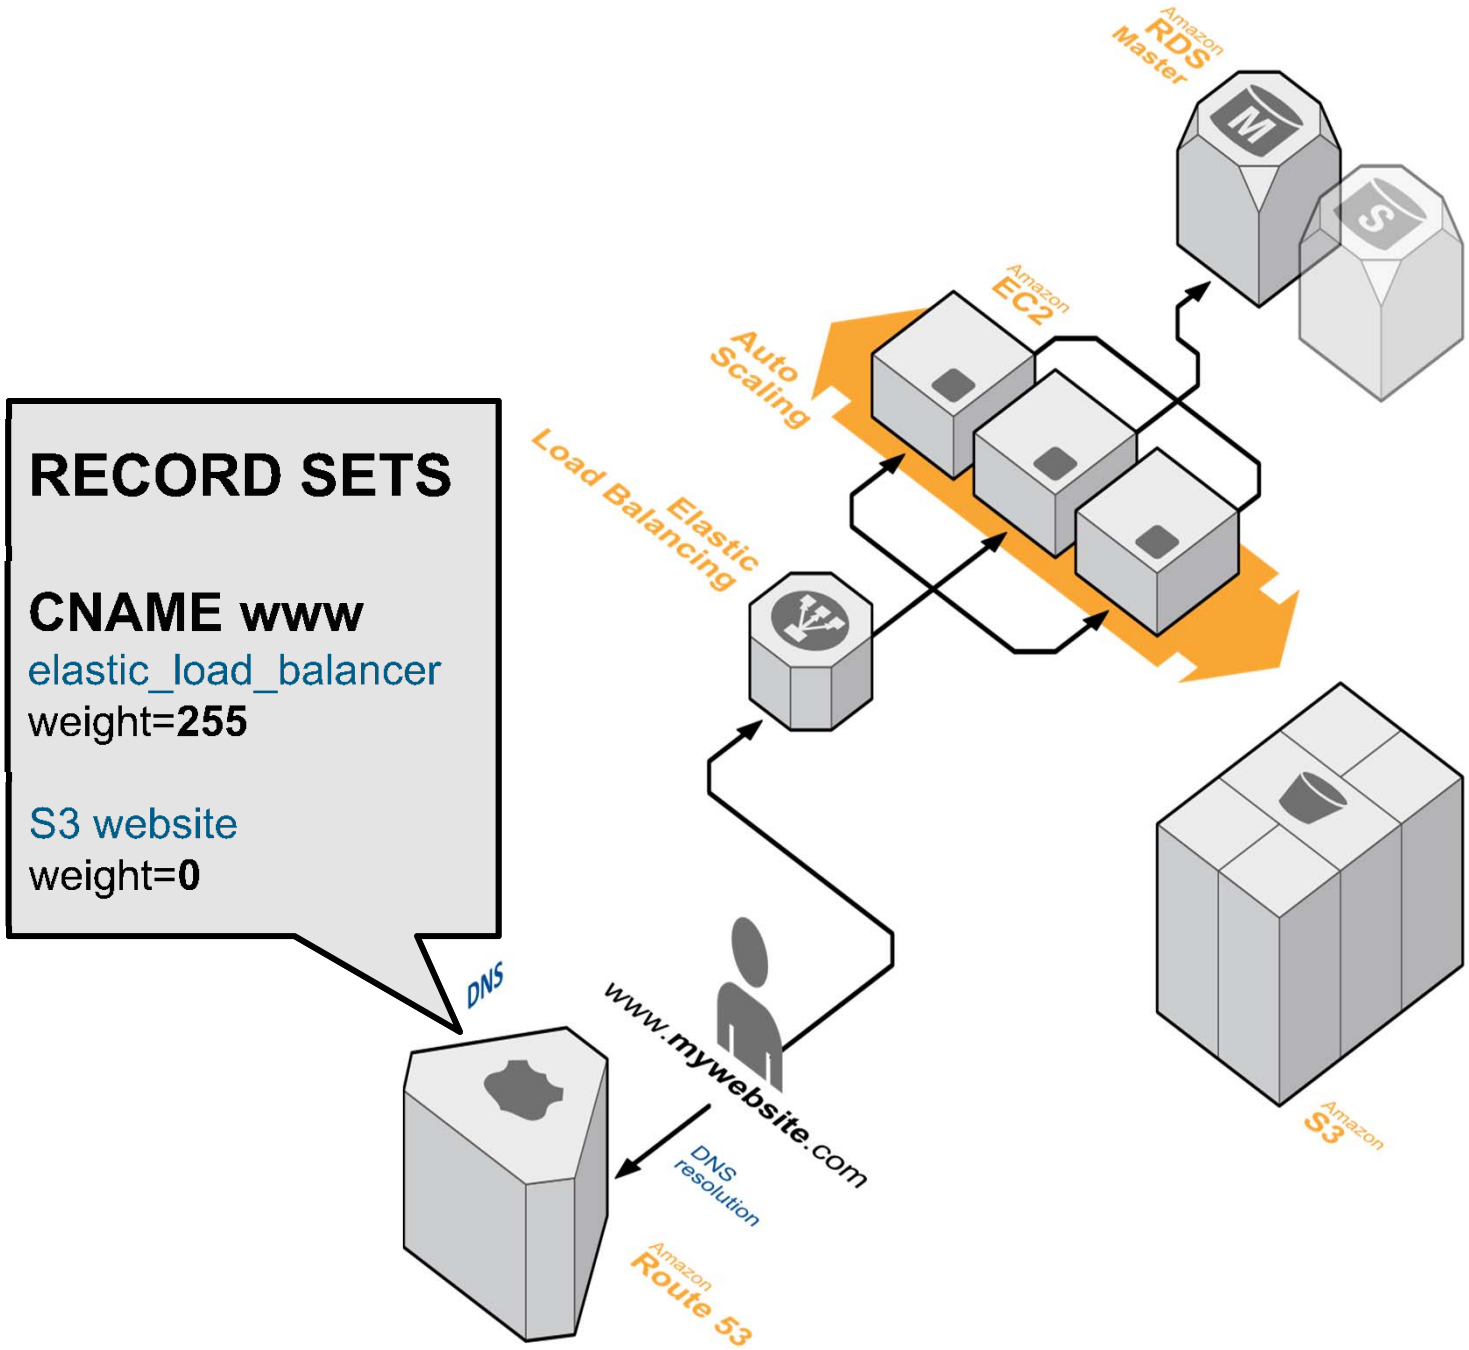
RECORD SETS
CNAME www
elastic_load_balancer
weight=255
S3 website
weight=0
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancing
Auto Scaling
Amazon EC2
Amazon RDS Master
Amazon S3

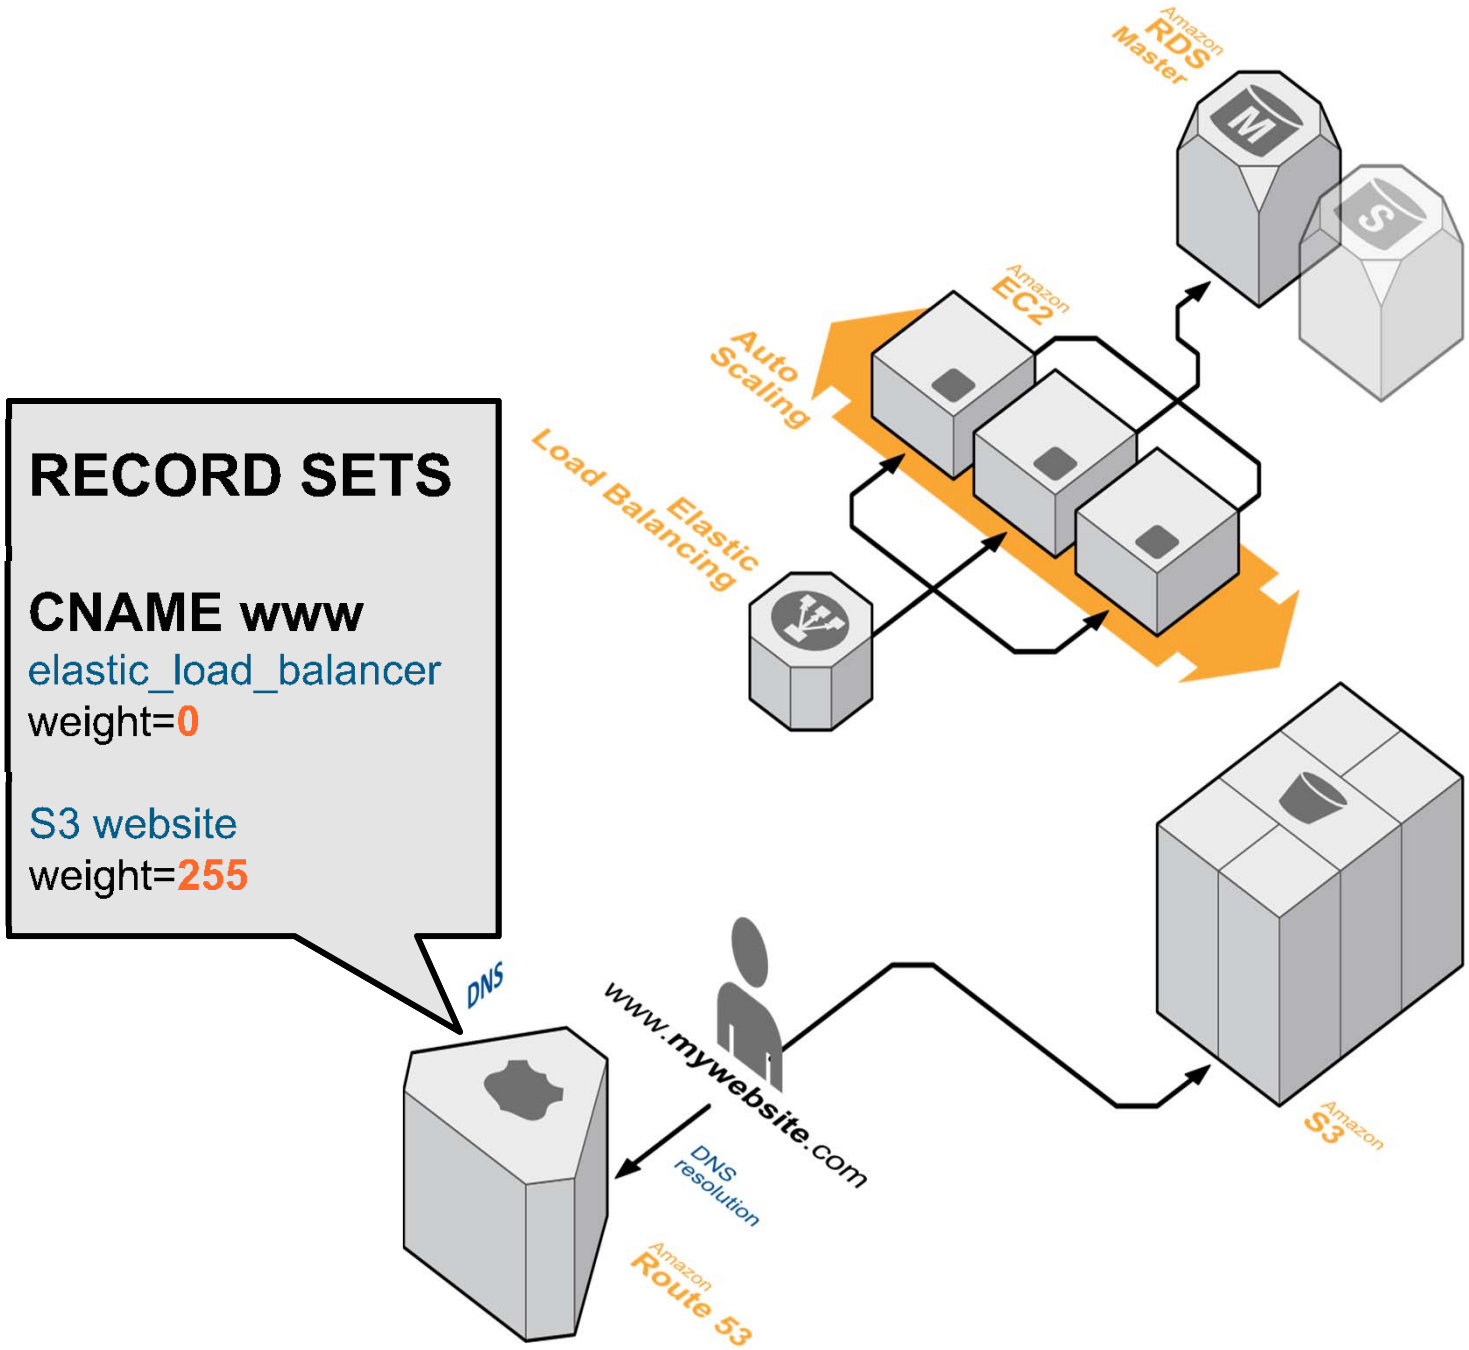
RECORD SETS
CNAME www
elastic_load_balancer
weight=0
S3 website
weight=255
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Amazon S3
Elastic Load Balancing
Auto Scaling
Amazon EC2
Amazon RDS Master

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. **SELF-HEALING**
5. LOOSE COUPLING

#5

LOOSE COUPLING

BUILD LOOSELY COUPLED SYSTEMS

The **looser** they are coupled,
the **bigger** they scale,
the more **fault tolerant** they get…

AMAZON SQS

SIMPLE QUEUE SERVICE

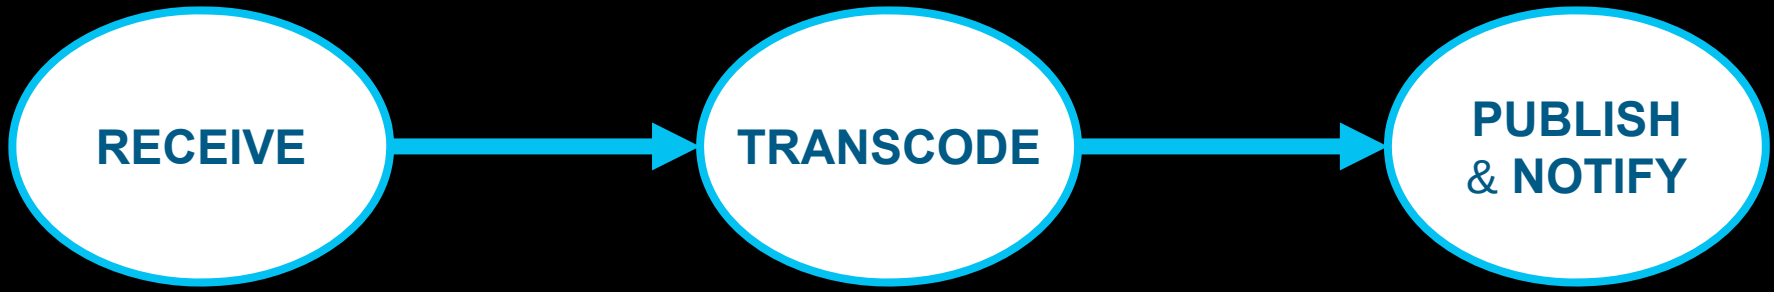
RECEIVE
TRANSCODE
PUBLISH
& NOTIFY

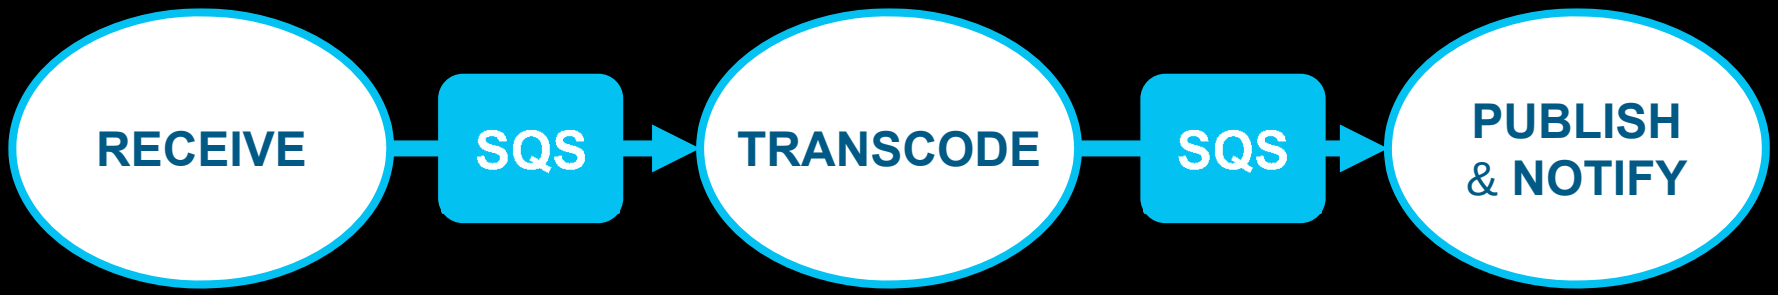
RECEIVE
SQS
TRANSCODE
SQS
PUBLISH
& NOTIFY

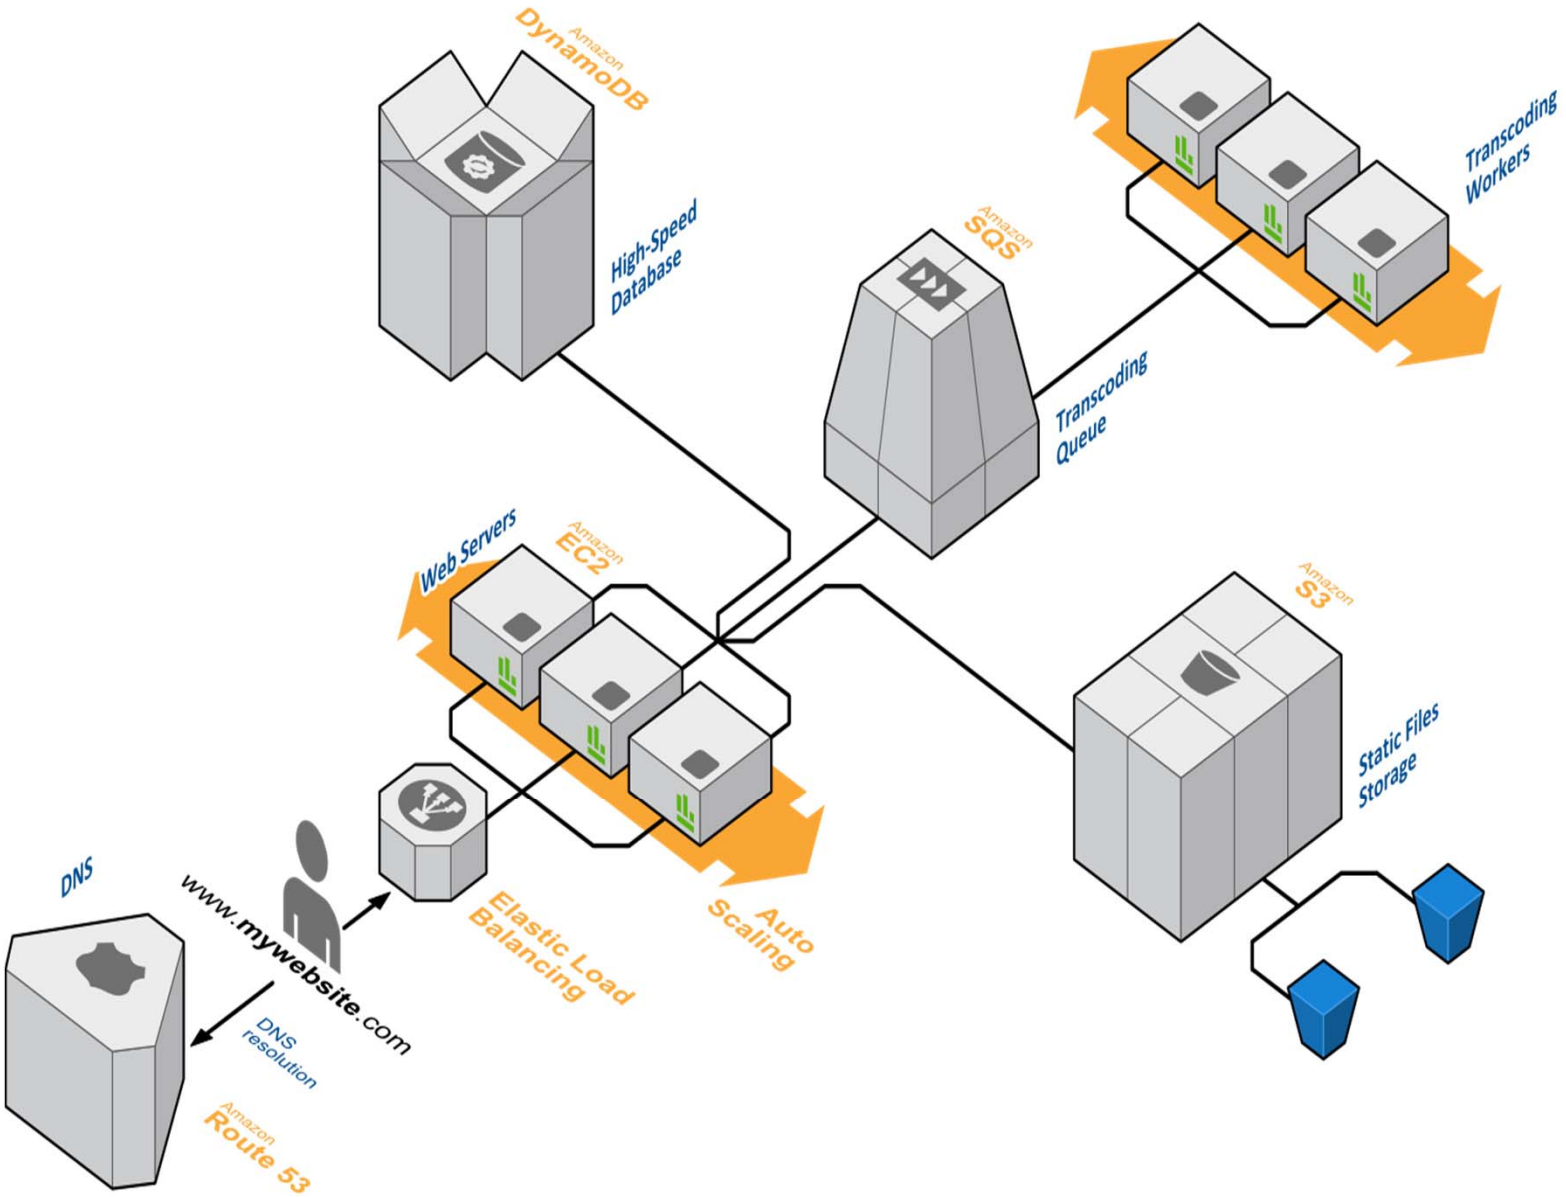
Amazon
DynamoDB
High-Speed
Database
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Web Servers
Amazon
EC2
Amazon
S3
Static Files
Storage
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53
Elastic Load
Balancing
Auto
Scaling

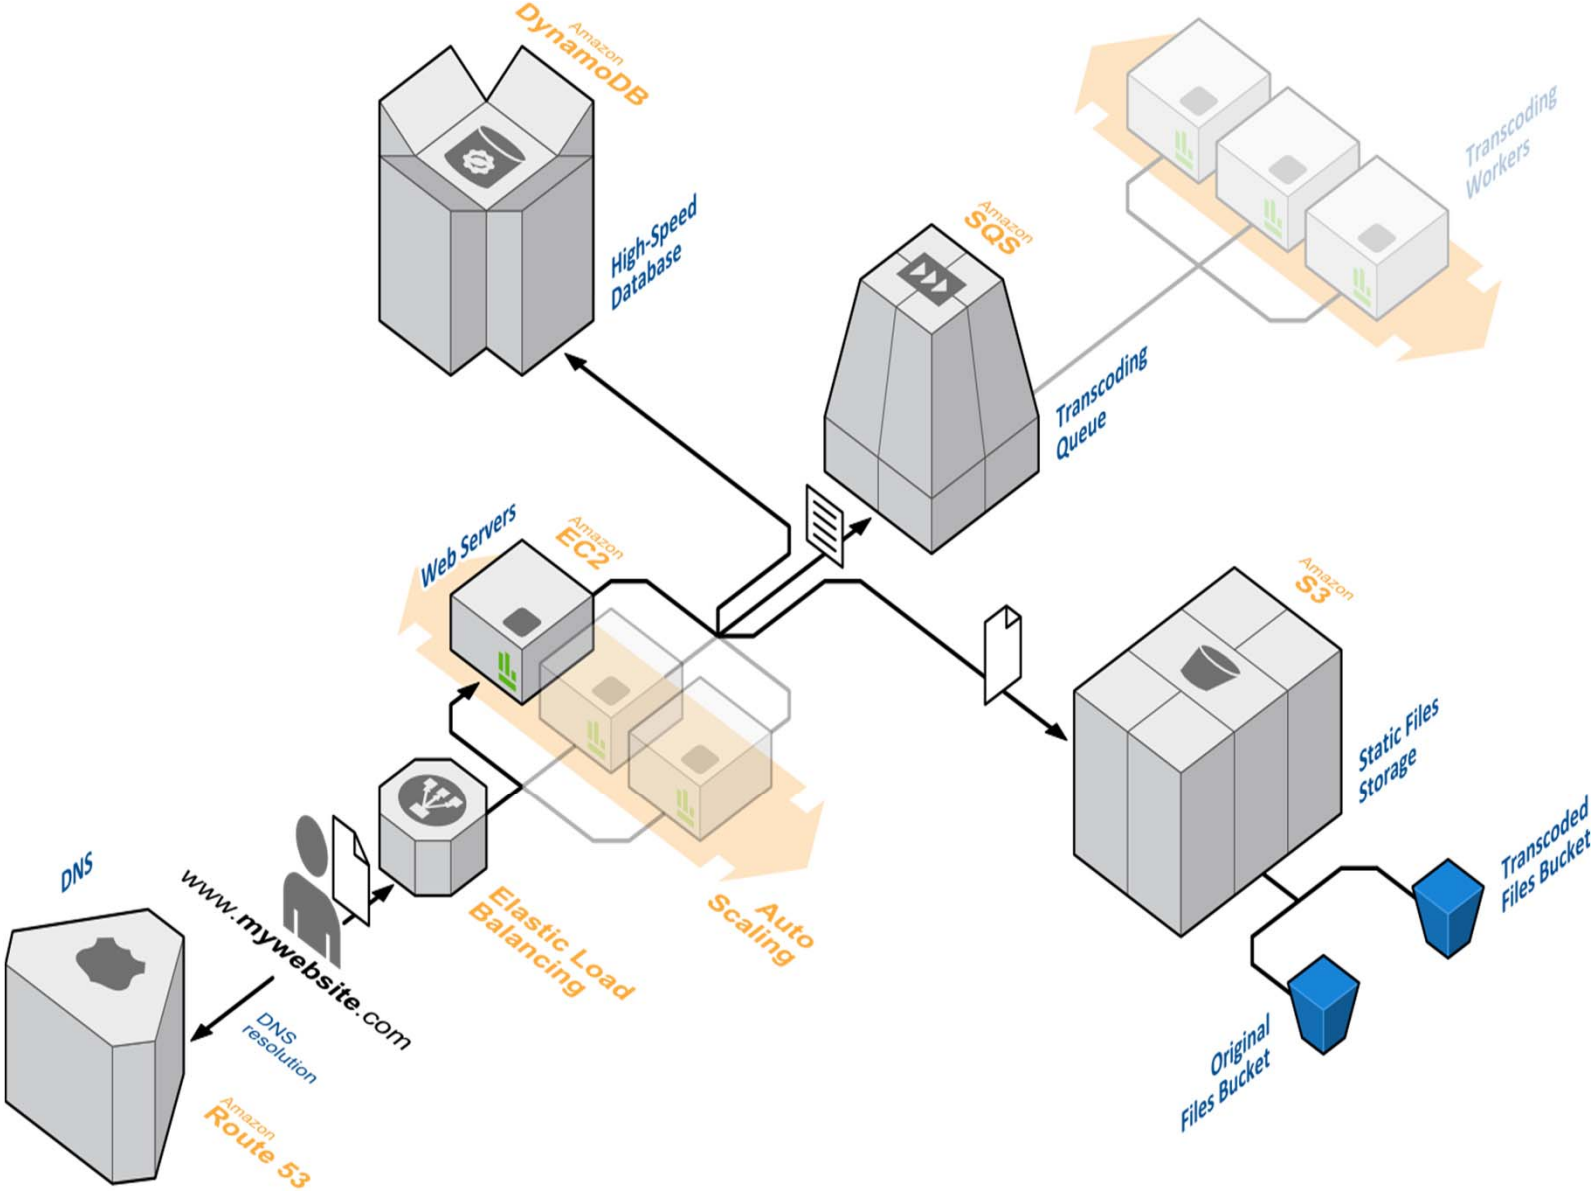
Amazon
DynamoDB
High-Speed
Database
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Web Servers
Amazon
EC2
Amazon
S3
Static Files
Storage
Transcoded
Files Bucket
Original
Files Bucket
Elastic Load
Balancing
Auto
Scaling
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

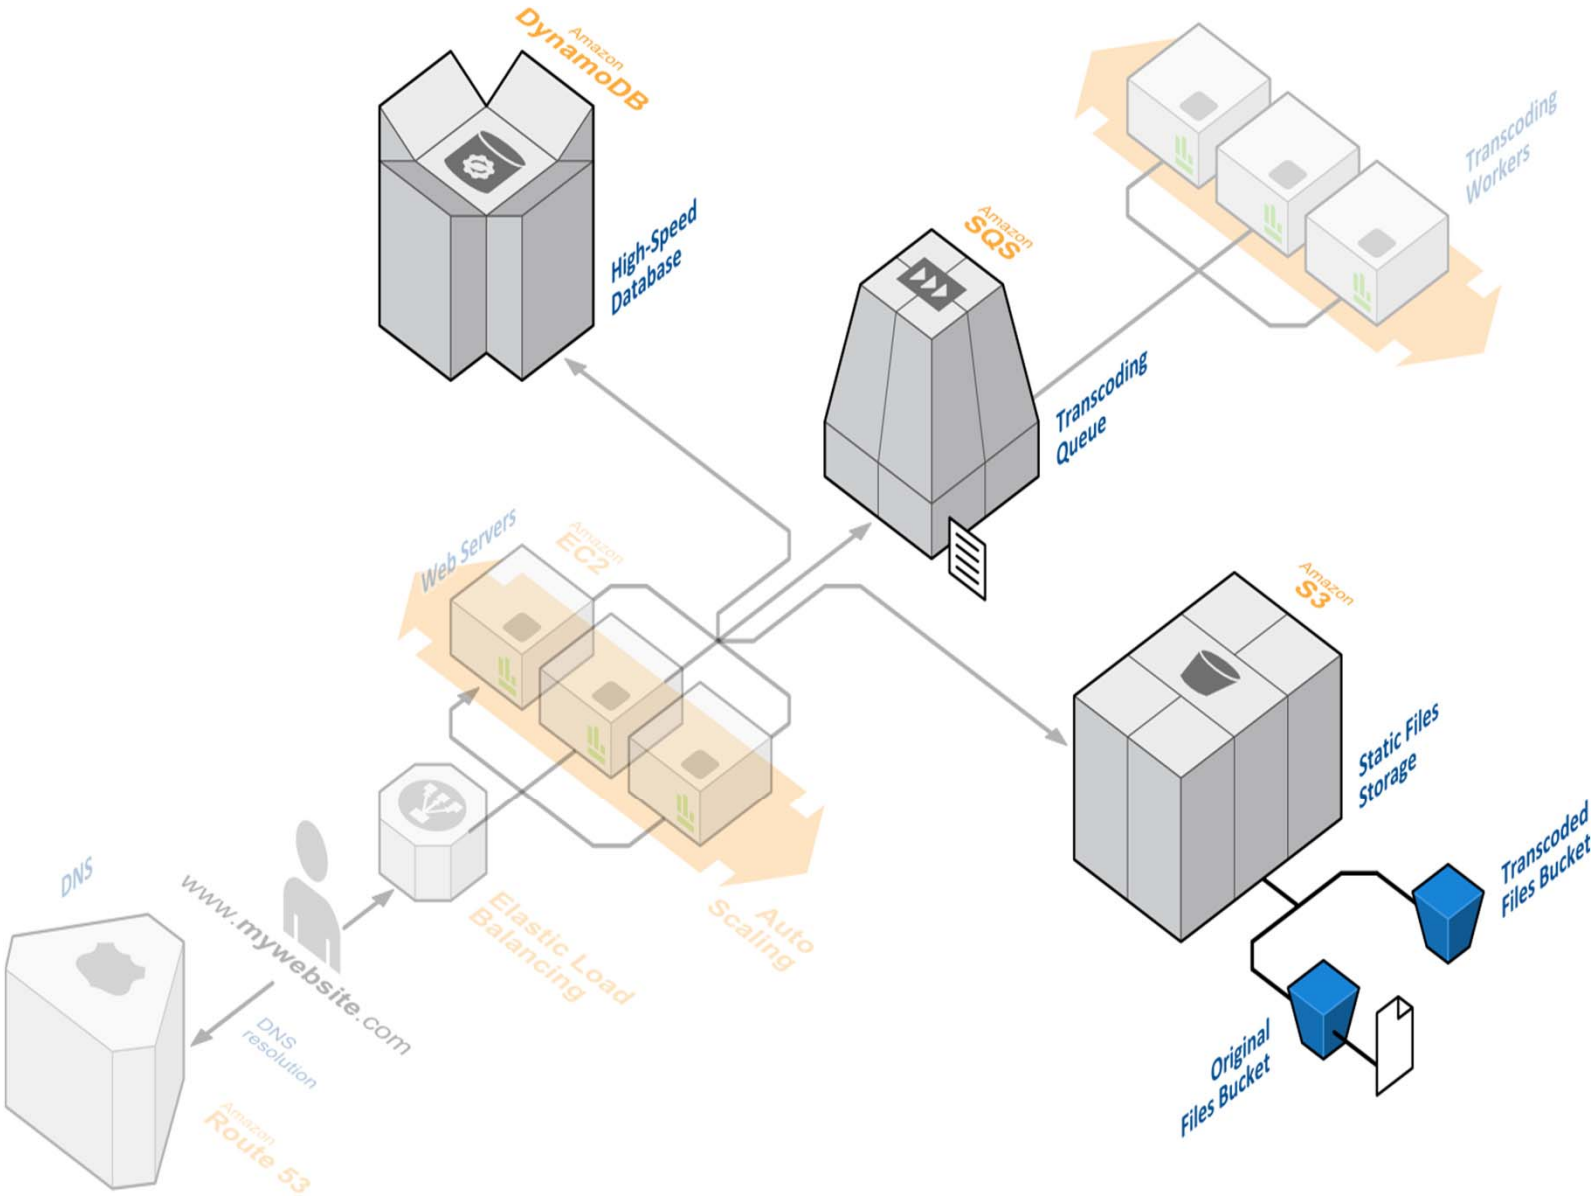
Amazon
DynamoDB
High-Speed
Database
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Web Servers
Amazon
EC2
Elastic Load
Balancing
Auto
Scaling
Amazon
S3
Static Files
Storage
Transcoded
Files Bucket
Original
Files Bucket
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

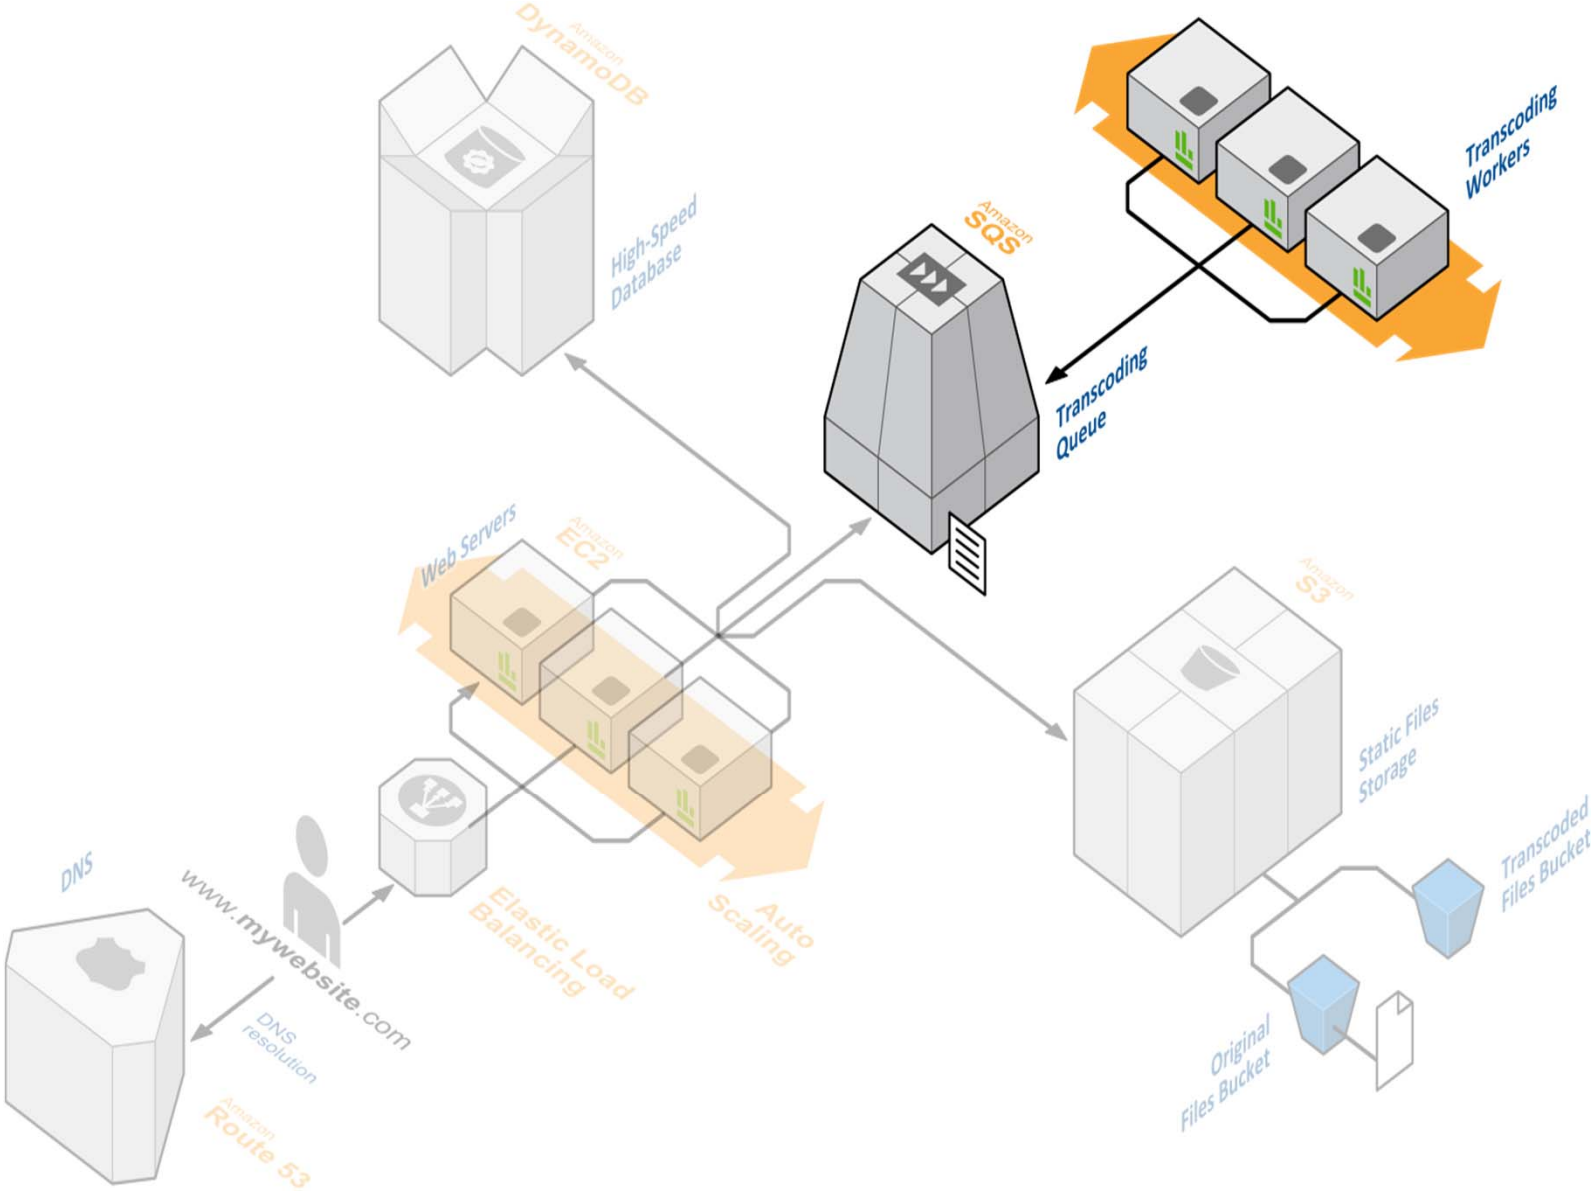
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

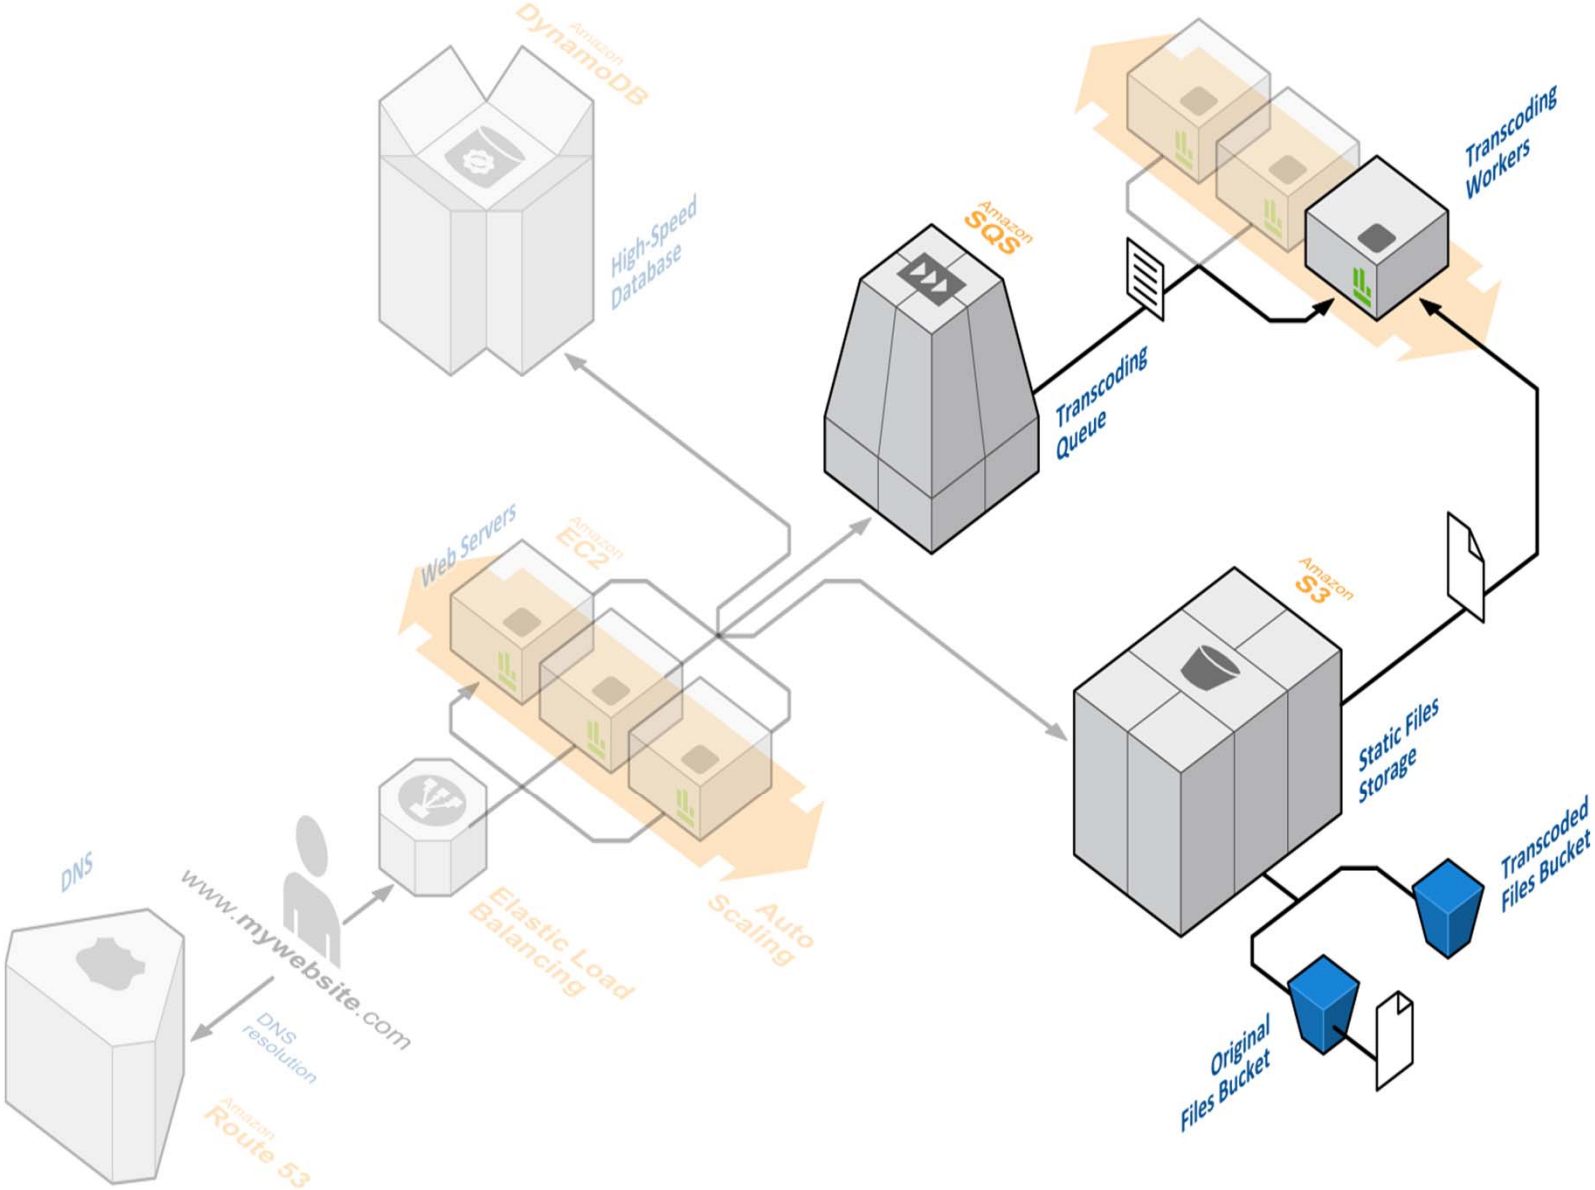
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Auto Scaling
Elastic Load Balancing
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

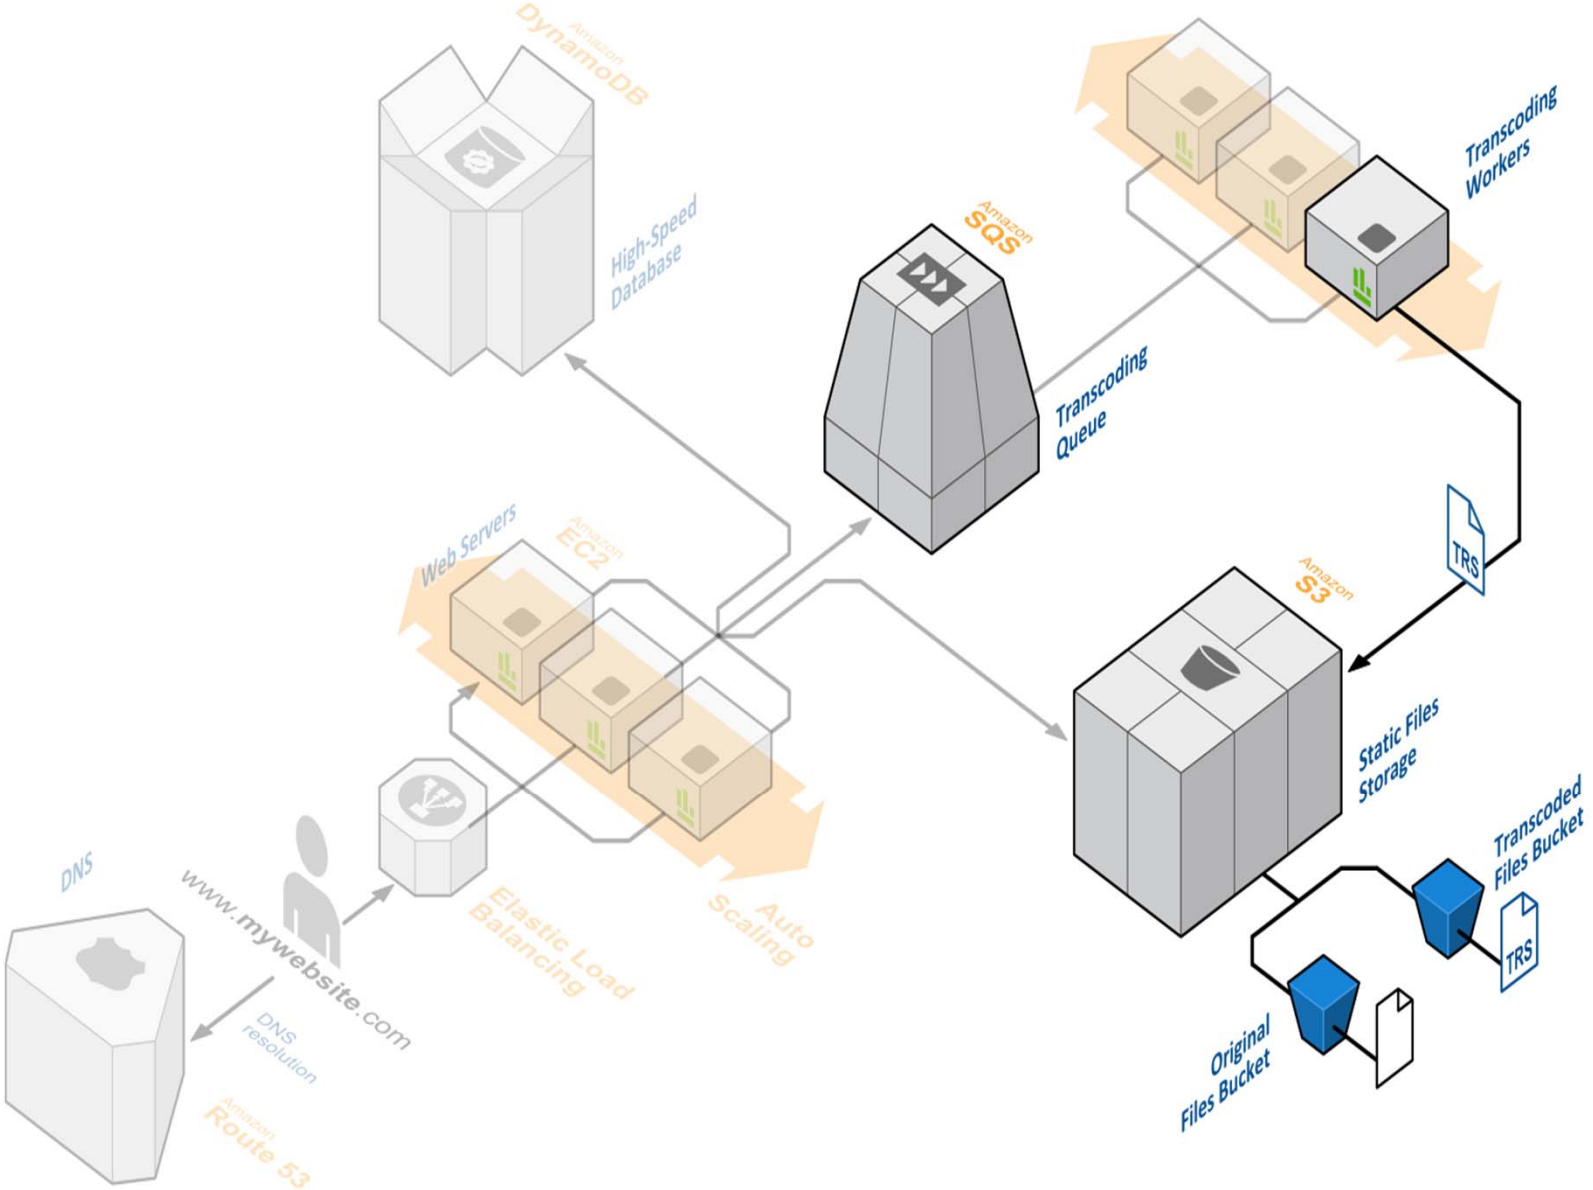
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
TRS
Amazon S3
Static Files Storage
Transcoded Files Bucket
TRS
Original Files Bucket
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

VISIBILITY TIMEOUT

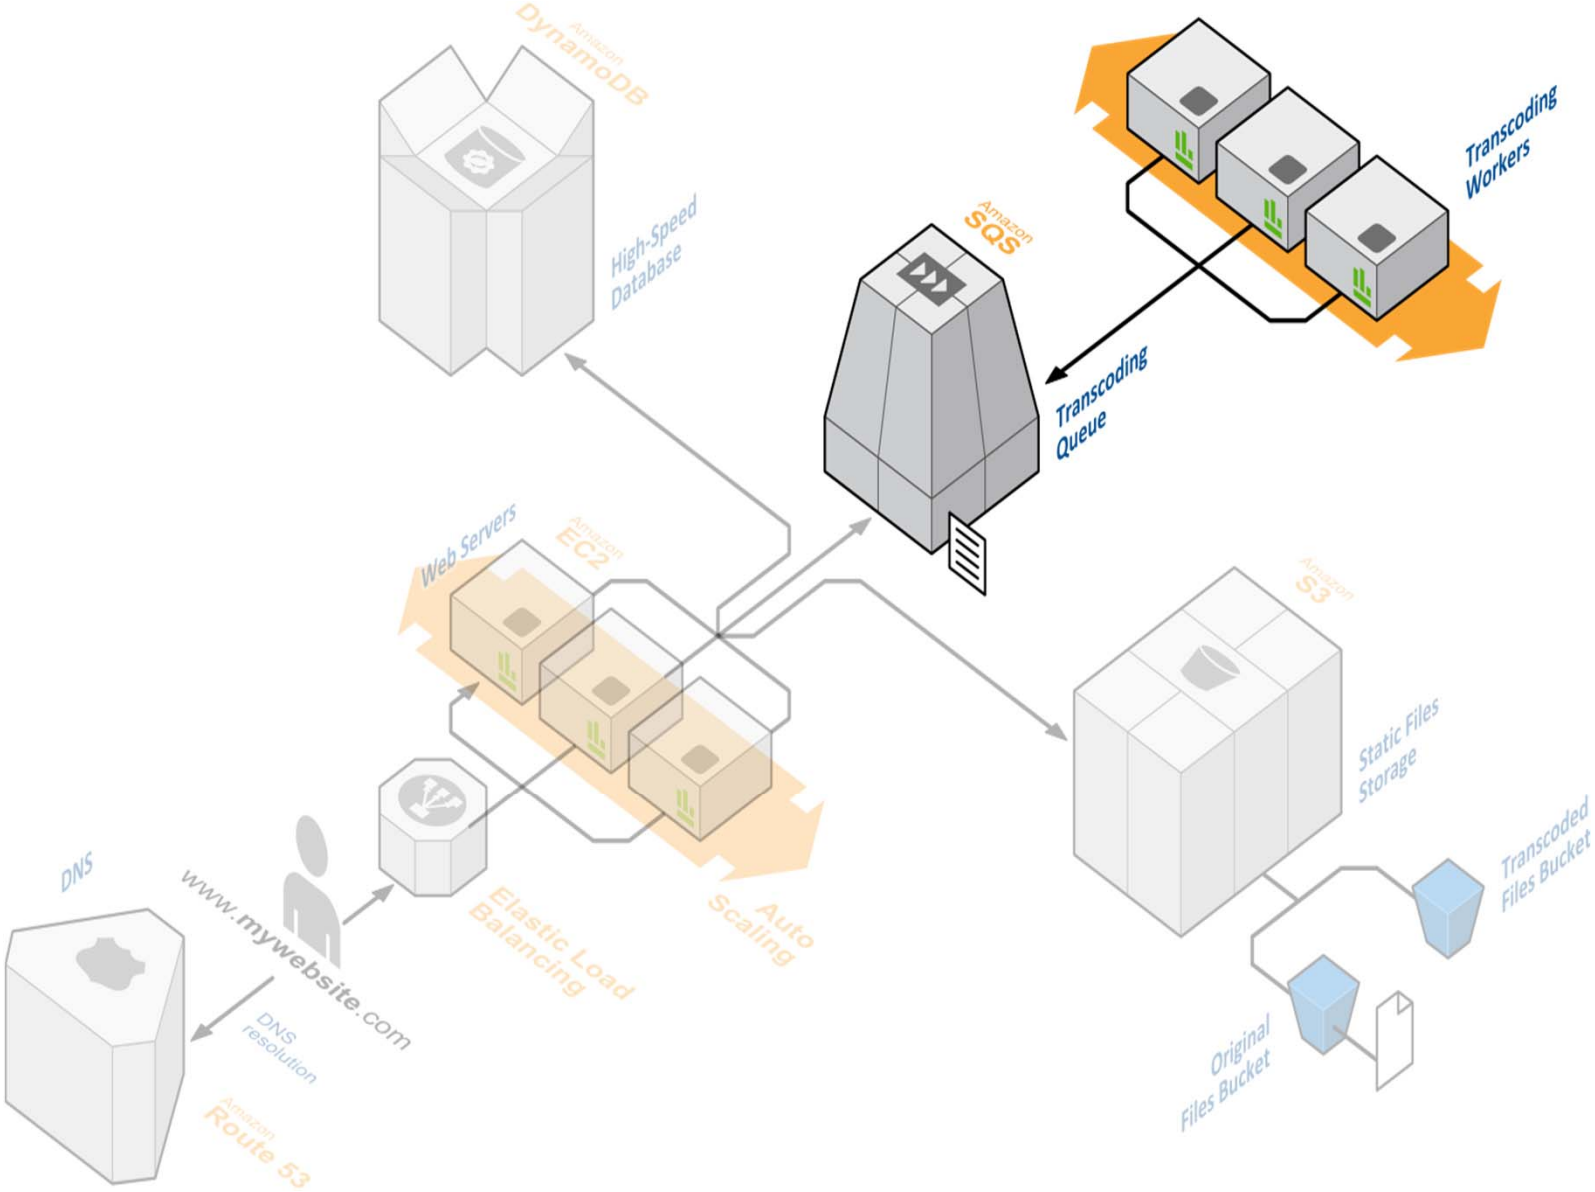
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

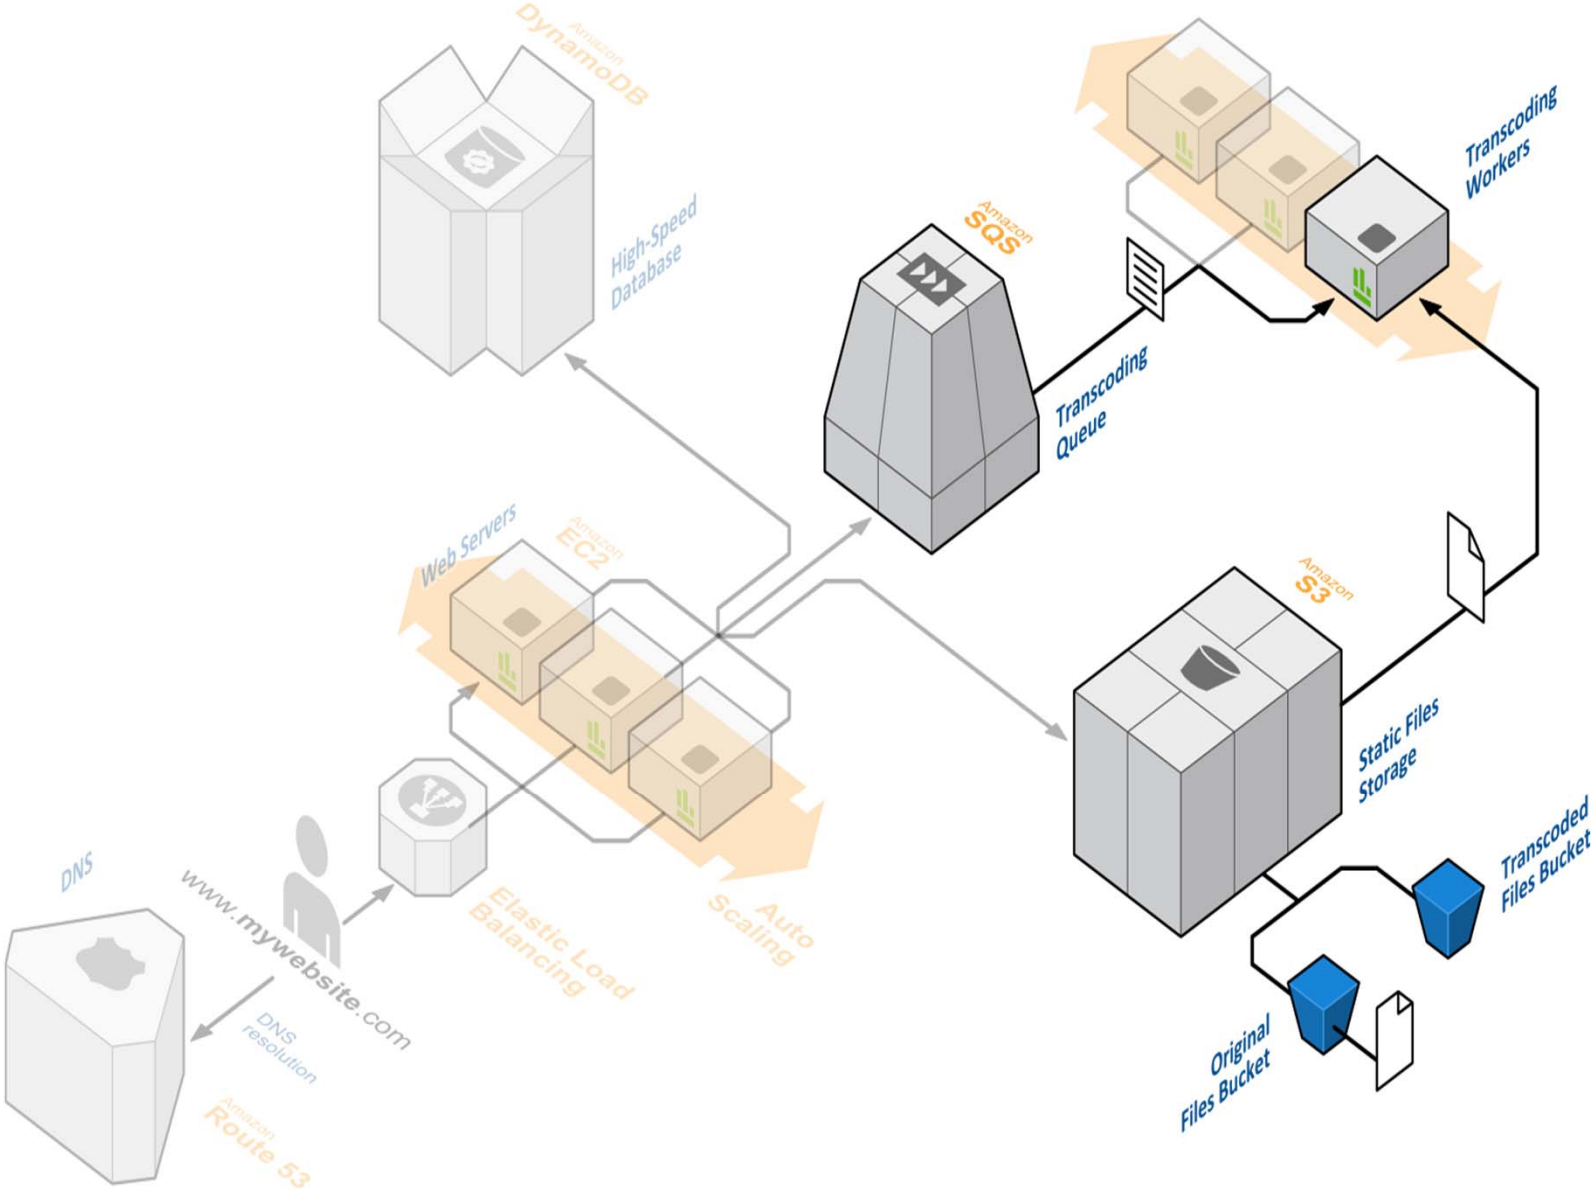
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Auto Scaling
Elastic Load Balancing
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

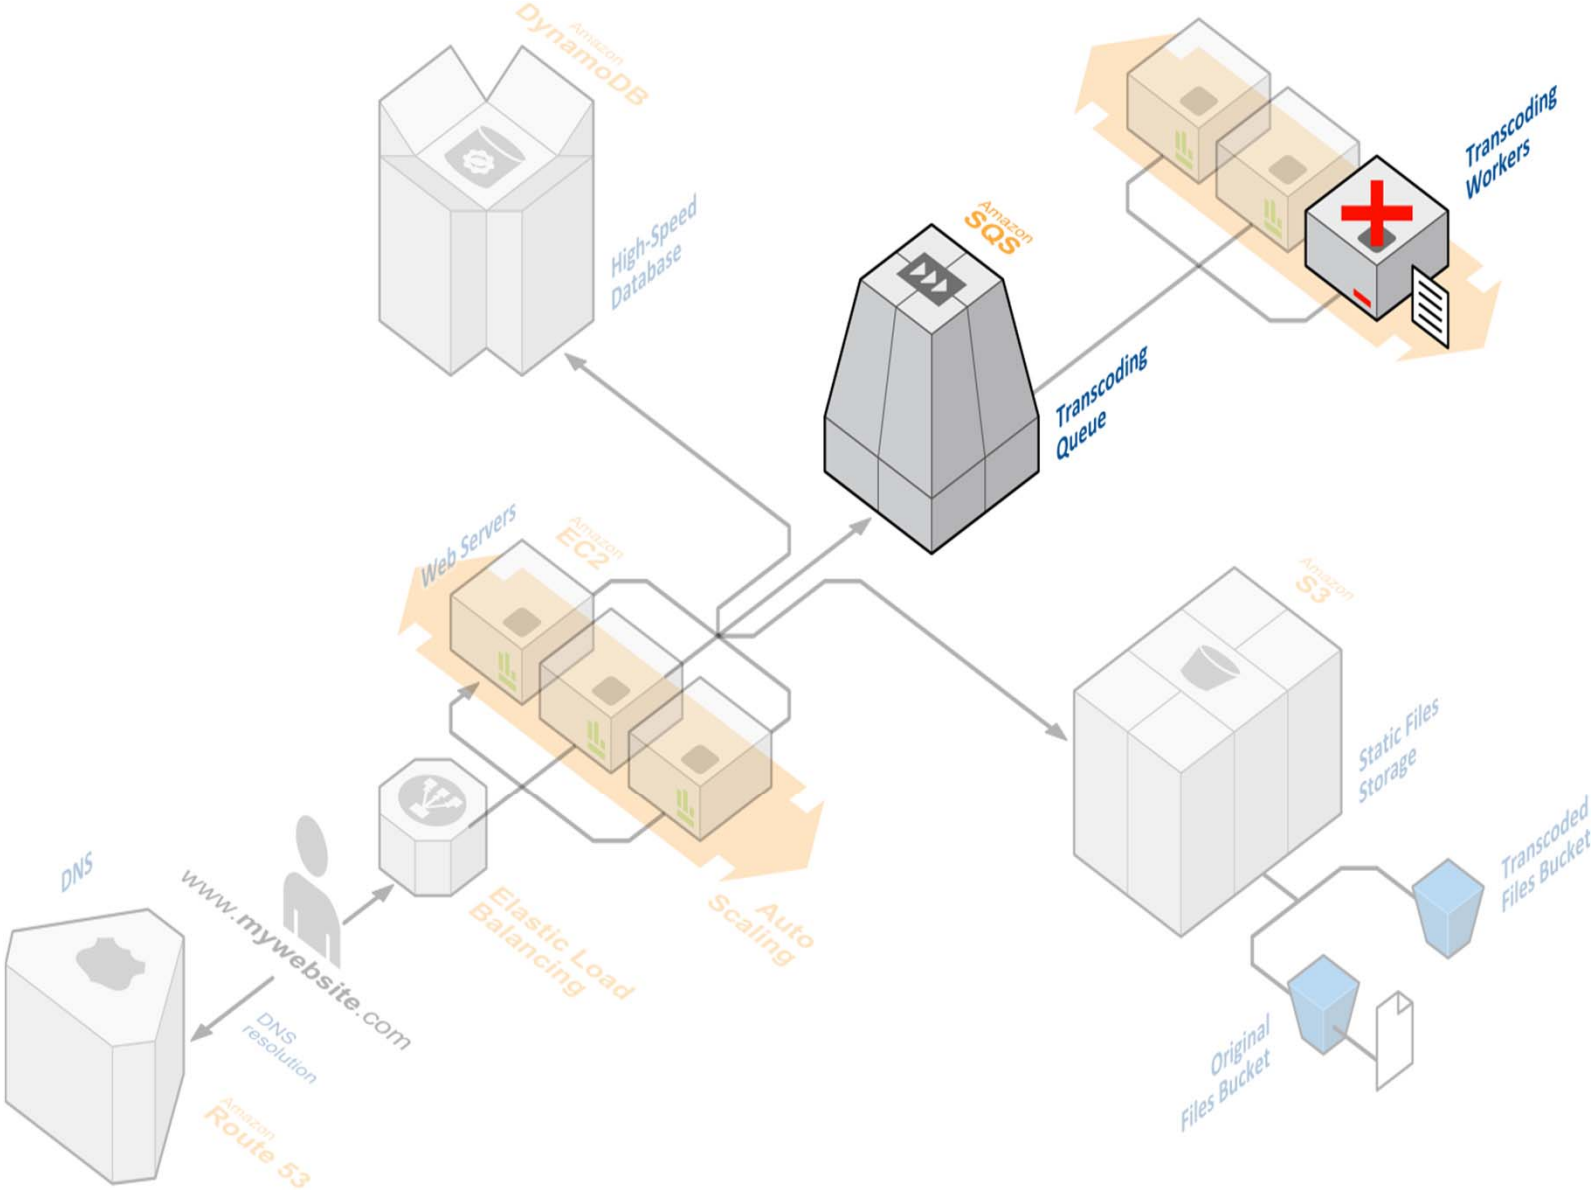
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

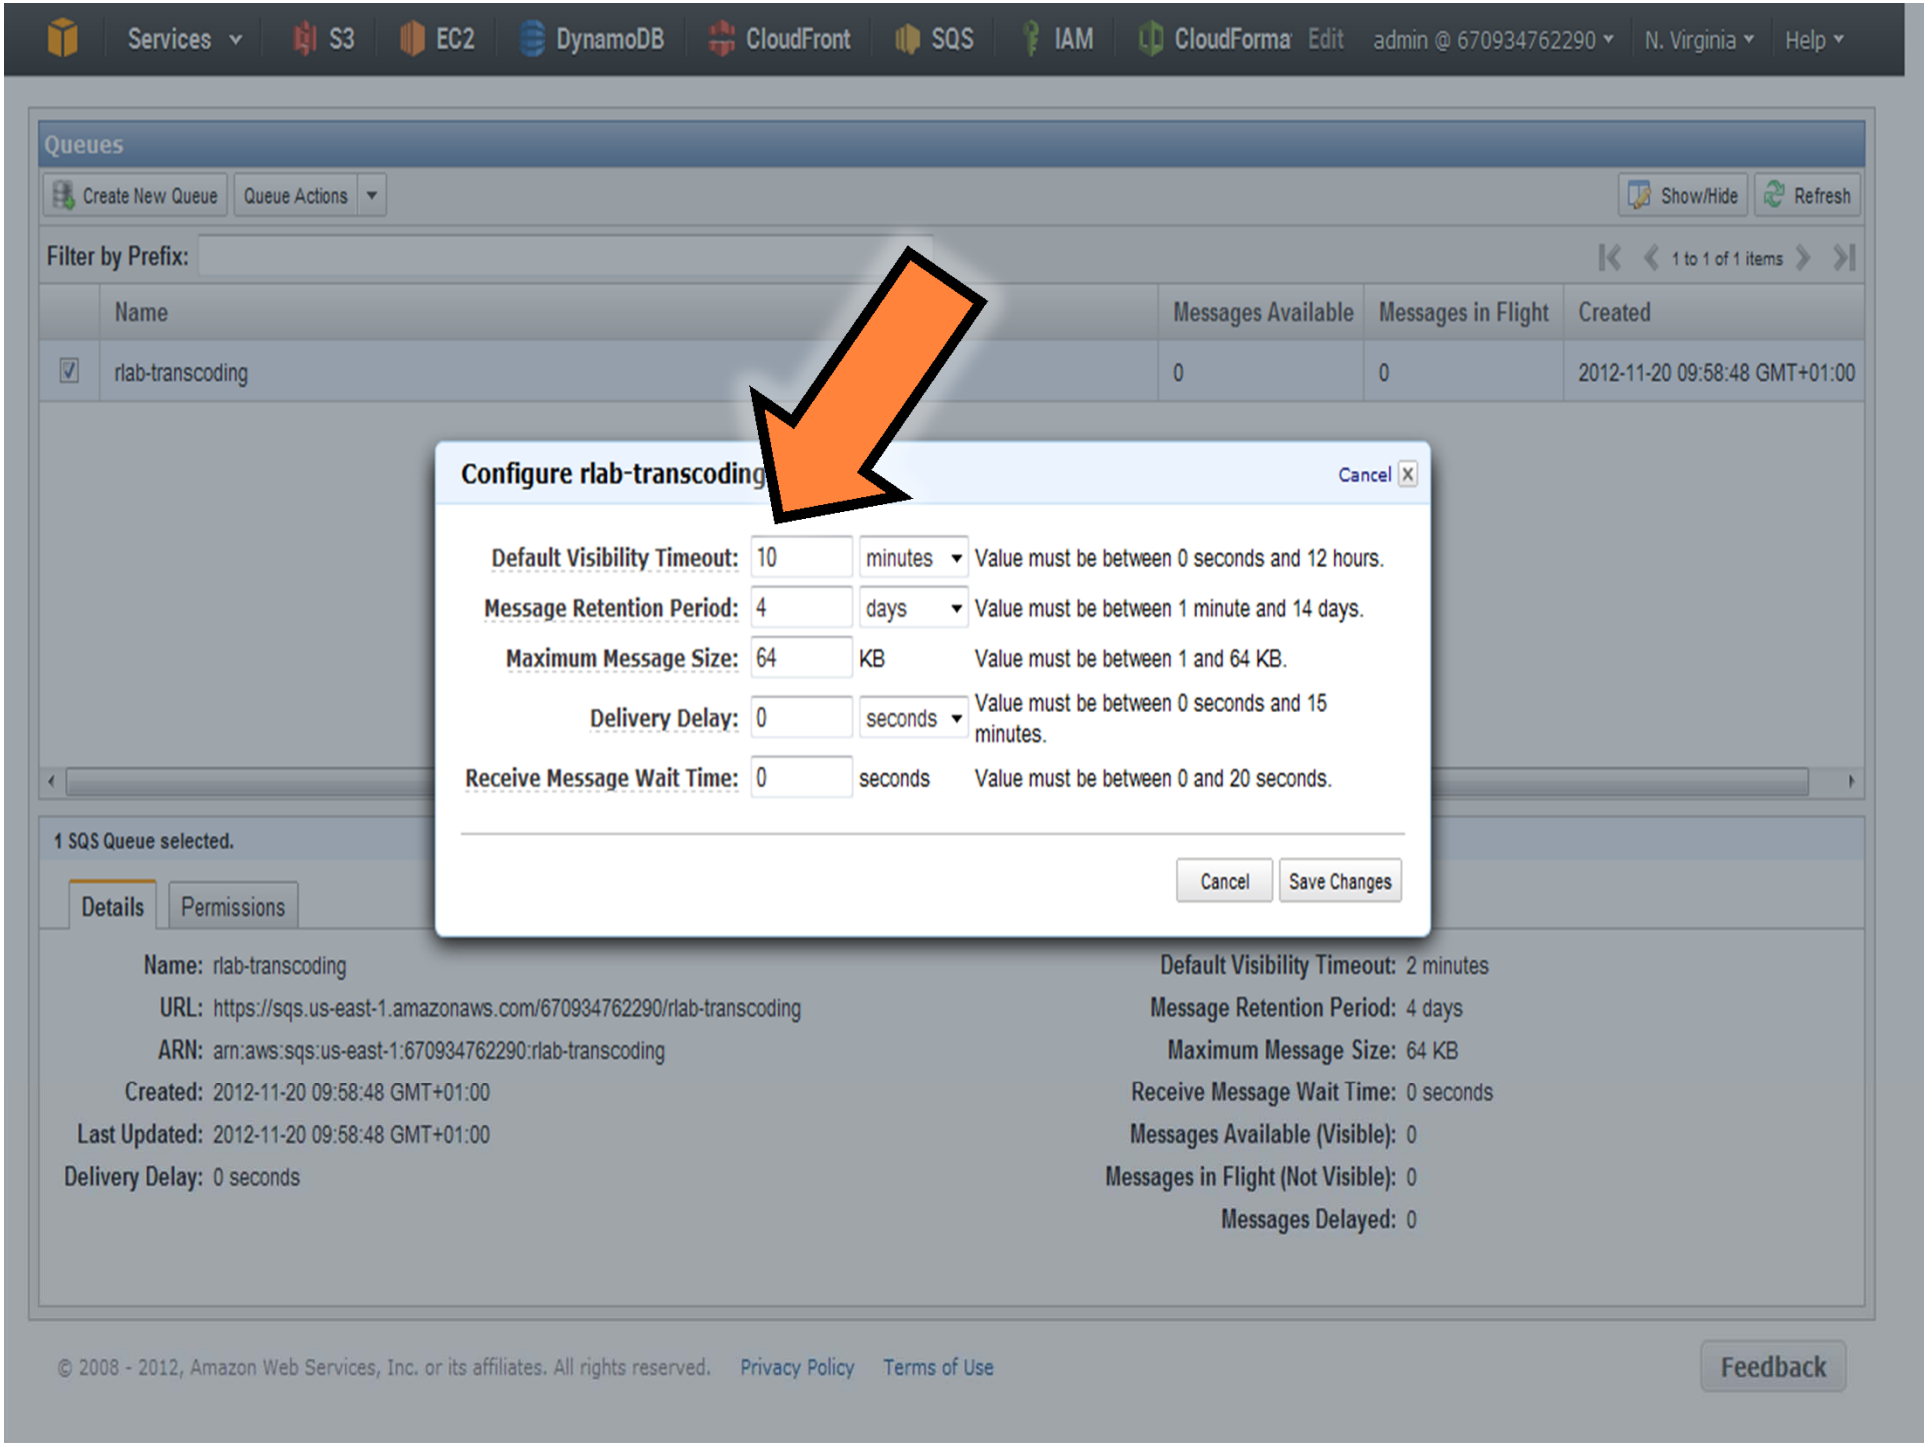
Services
S3
EC2
DynamoDB
CloudFront
SQS
IAM
CloudForma
Edit
admin @ 670934762290
N. Virginia
Help
Queues
Create New Queue
Queue Actions
Show/Hide
Refresh
Filter by Prefix:
1 to 1 of 1 items
Name
Messages Available
Messages in Flight
Created
rlab-transcoding
0
0
2012-11-20 09:58:48 GMT+01:00
Configure rlab-transcoding
Cancel
Default Visibility Timeout:
10
minutes
Value must be between 0 seconds and 12 hours.
Message Retention Period:
4
days
Value must be between 1 minute and 14 days.
Maximum Message Size:
64
KB
Value must be between 1 and 64 KB.
Delivery Delay:
0
seconds
Value must be between 0 seconds and 15 minutes.
Receive Message Wait Time:
0
seconds
Value must be between 0 and 20 seconds.
Cancel
Save Changes
1 SQS Queue selected.
Details
Permissions
Name: rlab-transcoding
URL: https://sqs.us-east-1.amazonaws.com/670934762290/rlab-transcoding
ARN: arn:aws:sqs:us-east-1:670934762290:rlab-transcoding
Created: 2012-11-20 09:58:48 GMT+01:00
Last Updated: 2012-11-20 09:58:48 GMT+01:00
Delivery Delay: 0 seconds
Default Visibility Timeout: 2 minutes
Message Retention Period: 4 days
Maximum Message Size: 64 KB
Receive Message Wait Time: 0 seconds
Messages Available (Visible): 0
Messages in Flight (Not Visible): 0
Messages Delayed: 0
© 2008 - 2012, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use
Feedback

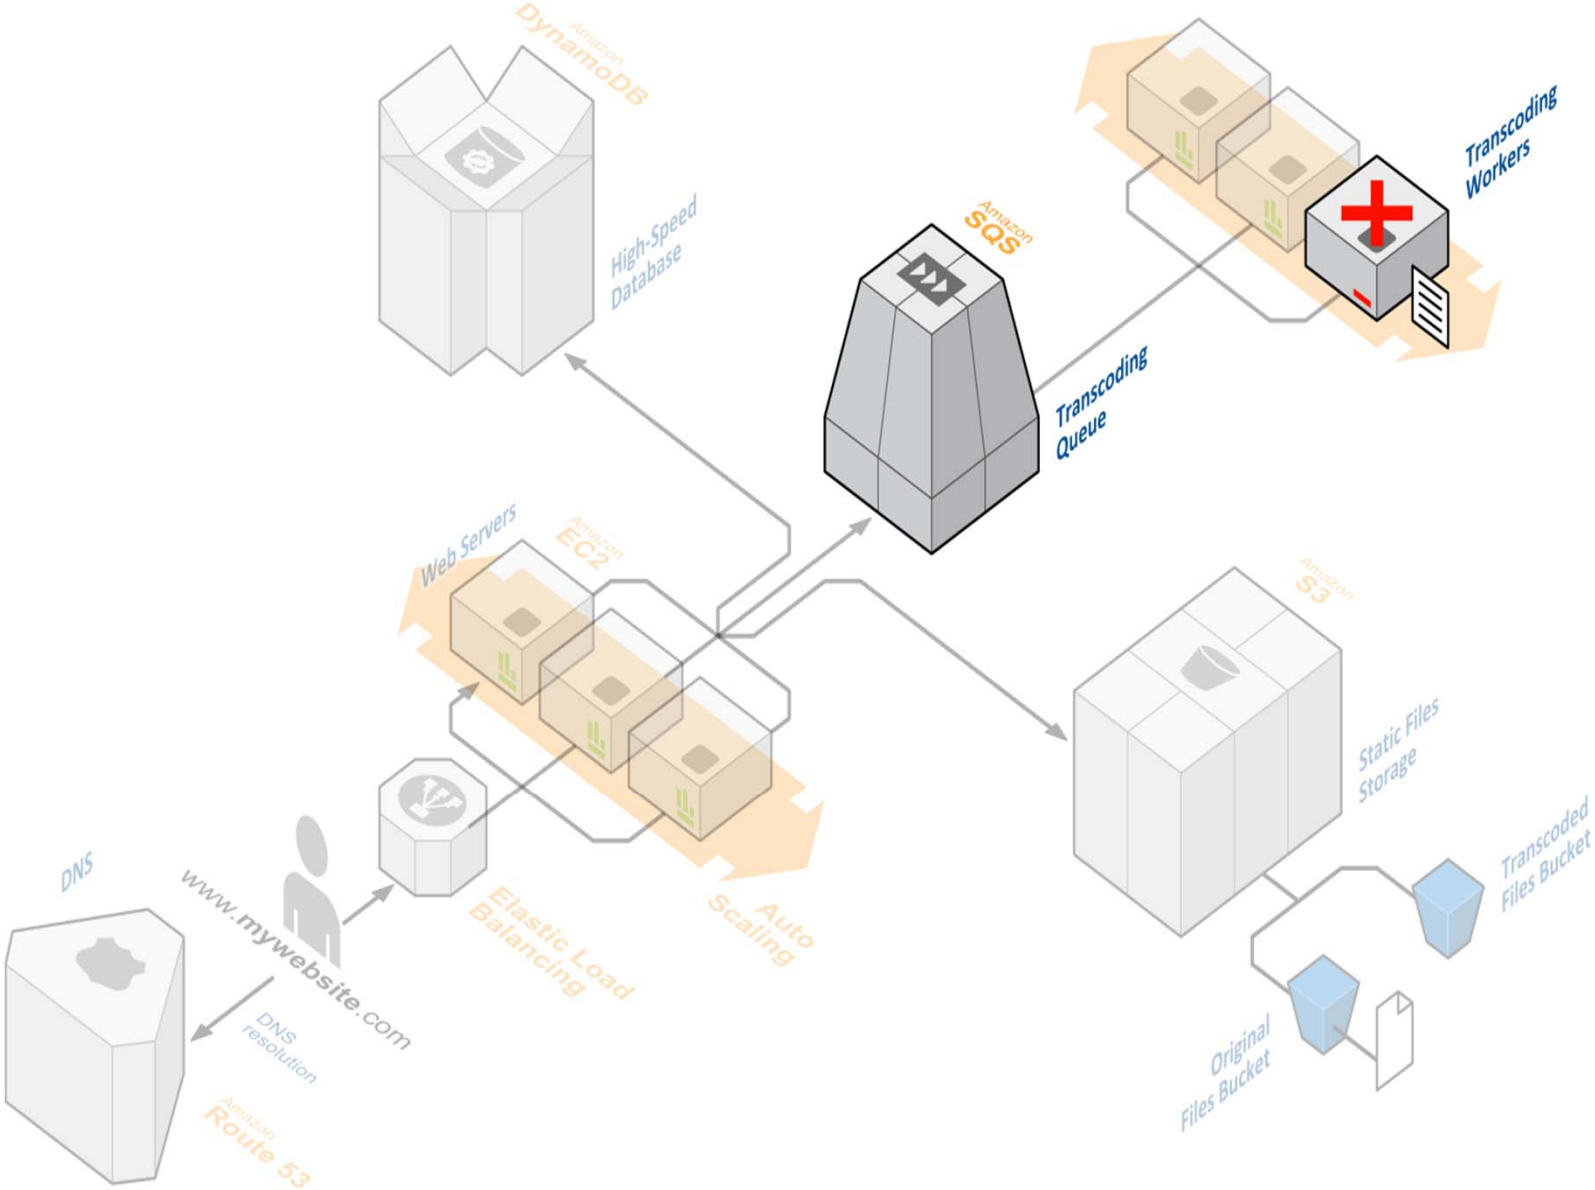
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

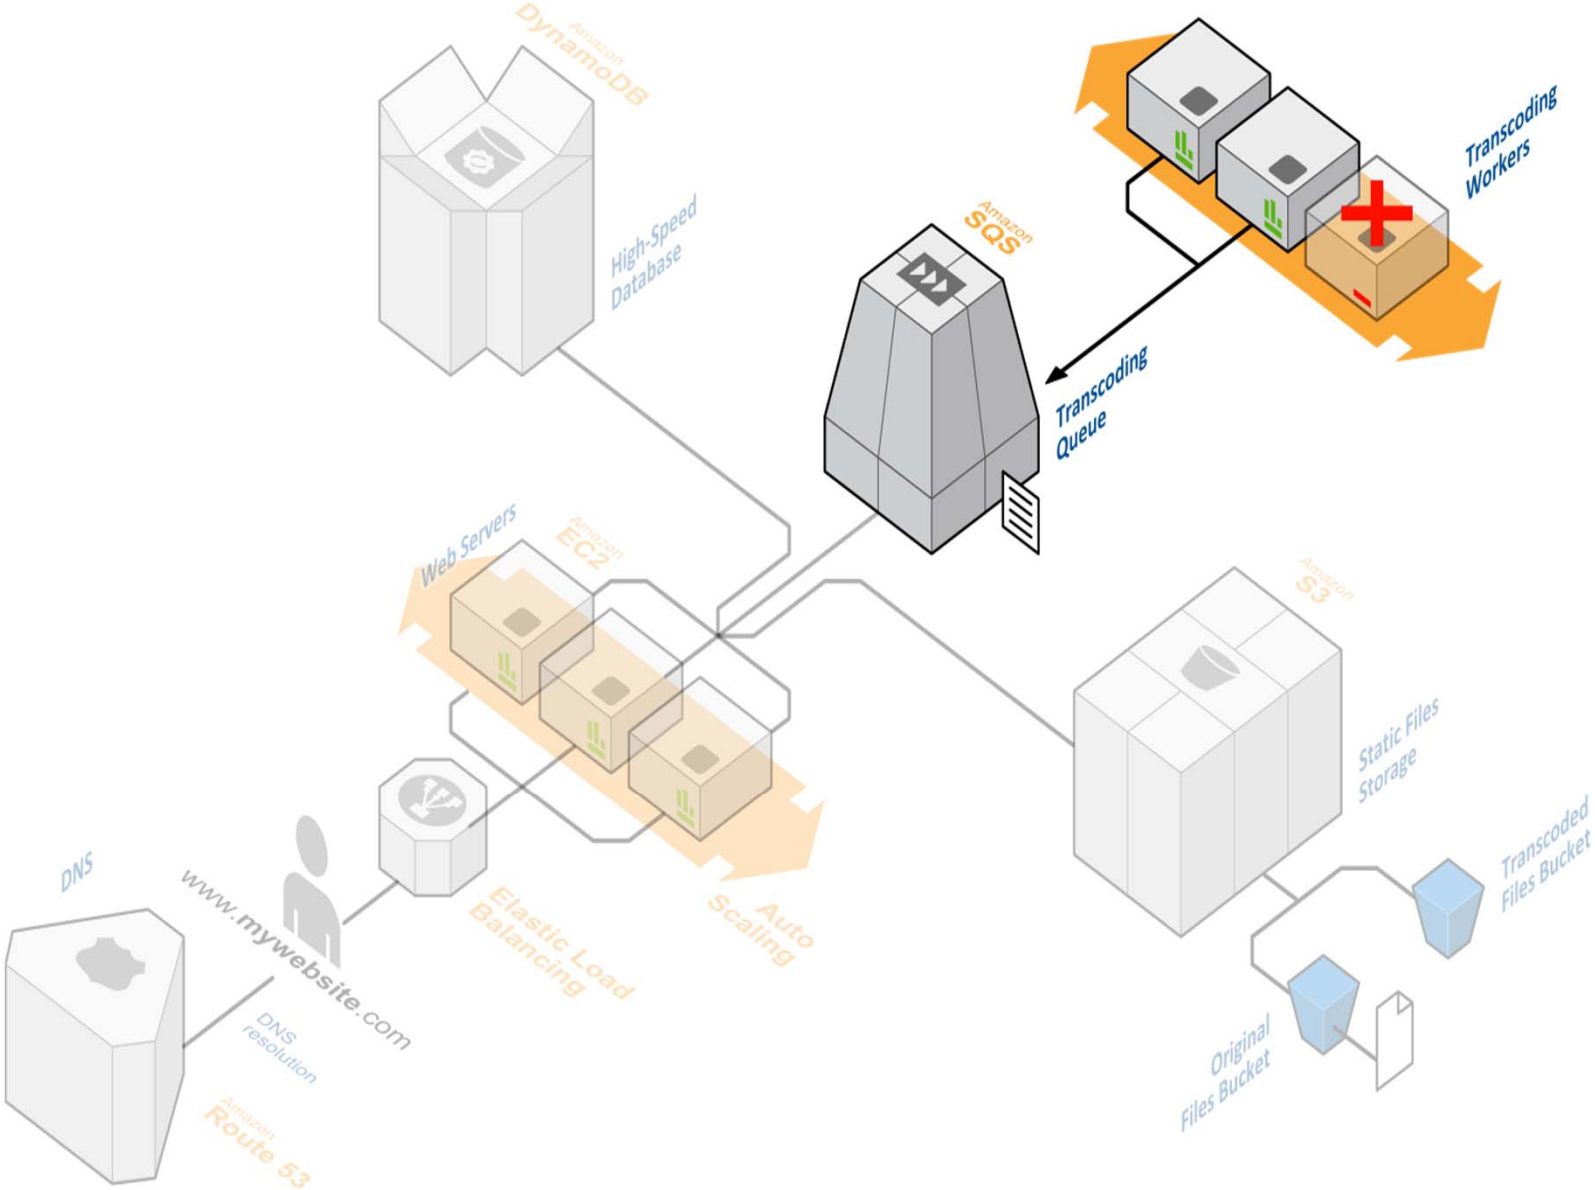
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

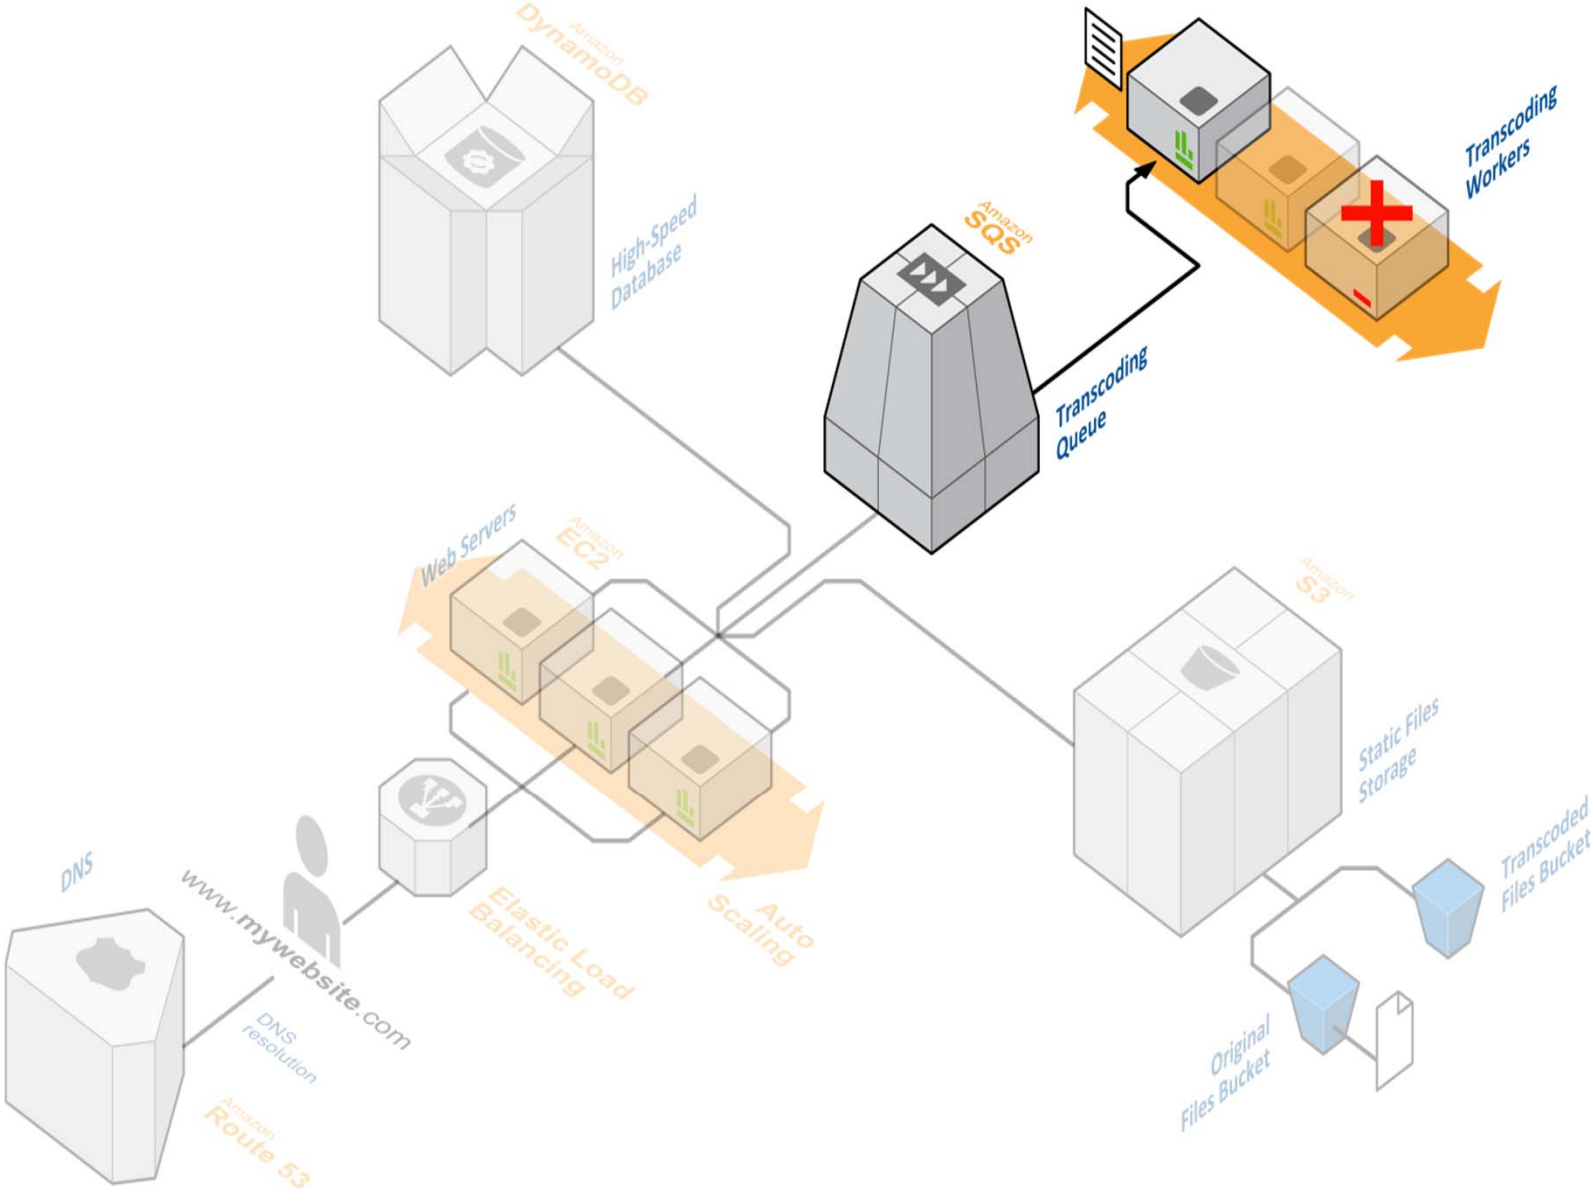
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

BUFFERING

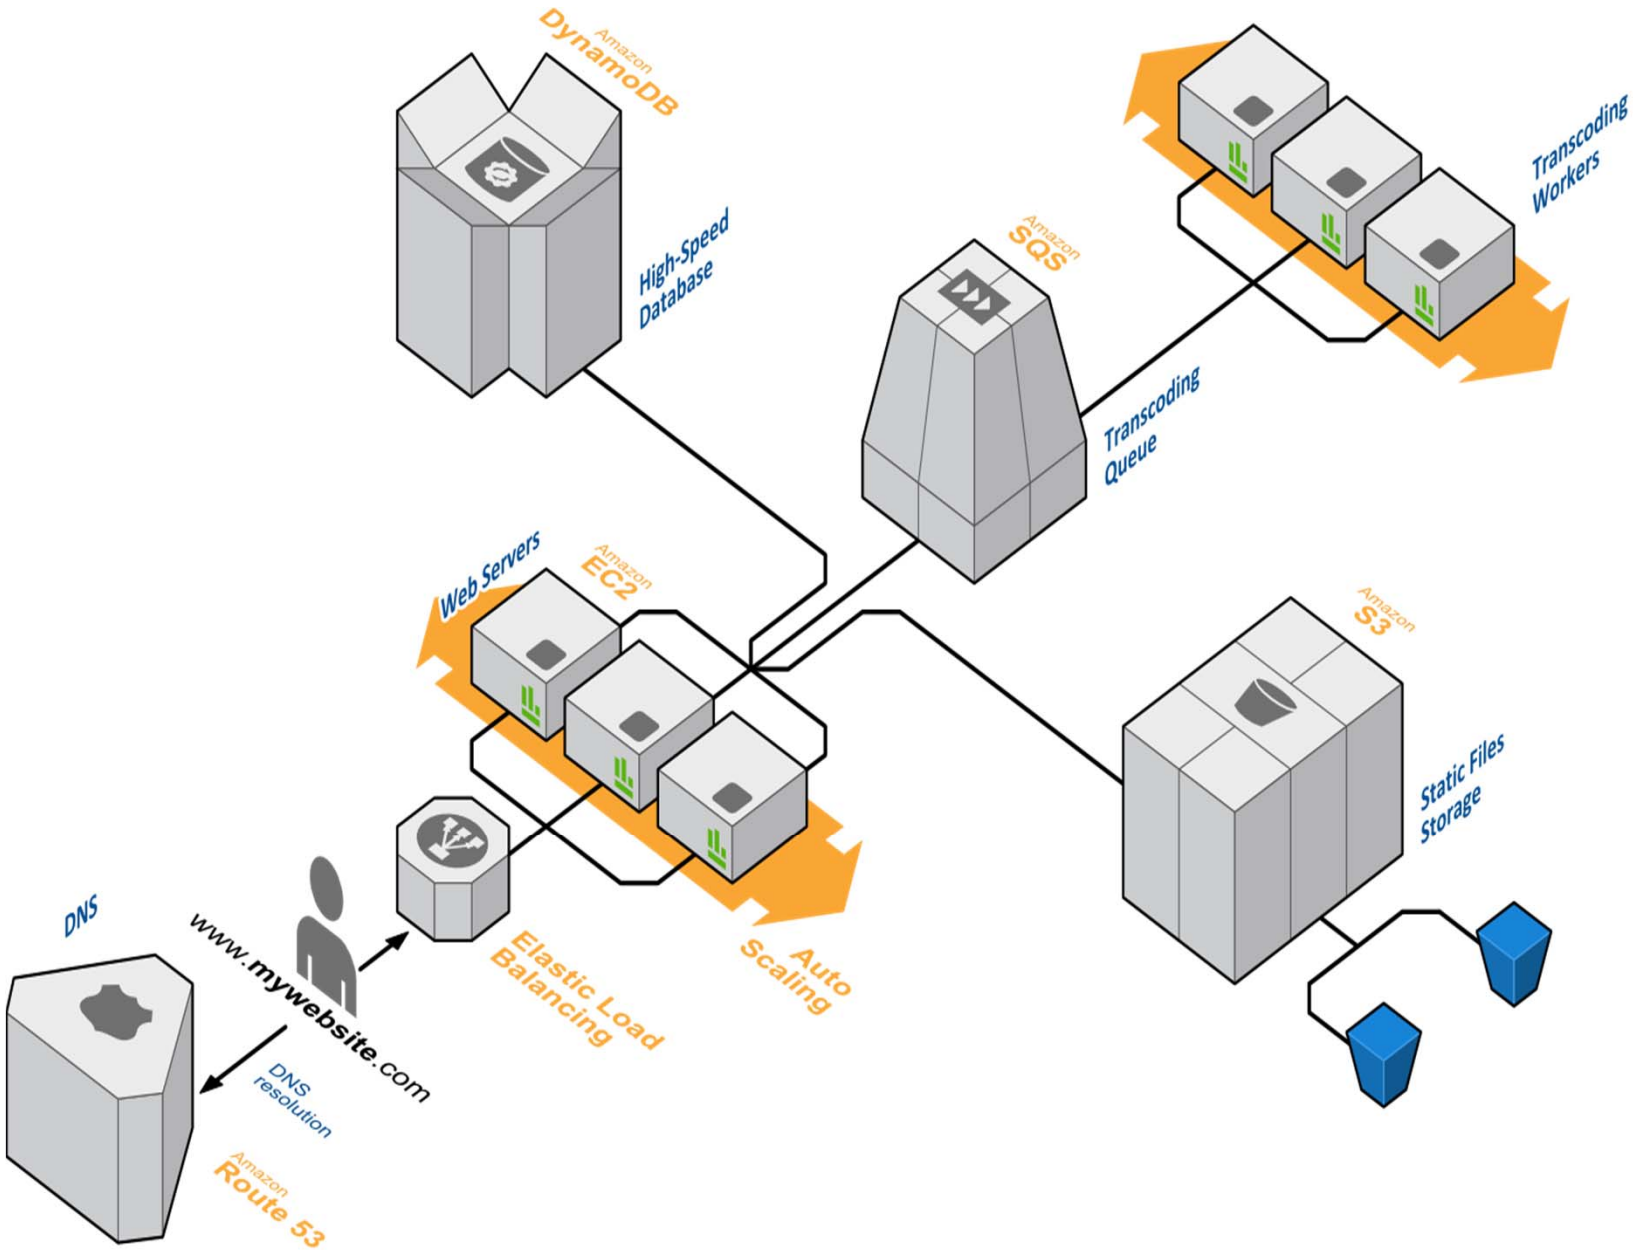
Amazon
DynamoDB
High-Speed
Database
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Web Servers
Amazon
EC2
Amazon
S3
Static Files
Storage
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53
Elastic Load
Balancing
Auto
Scaling

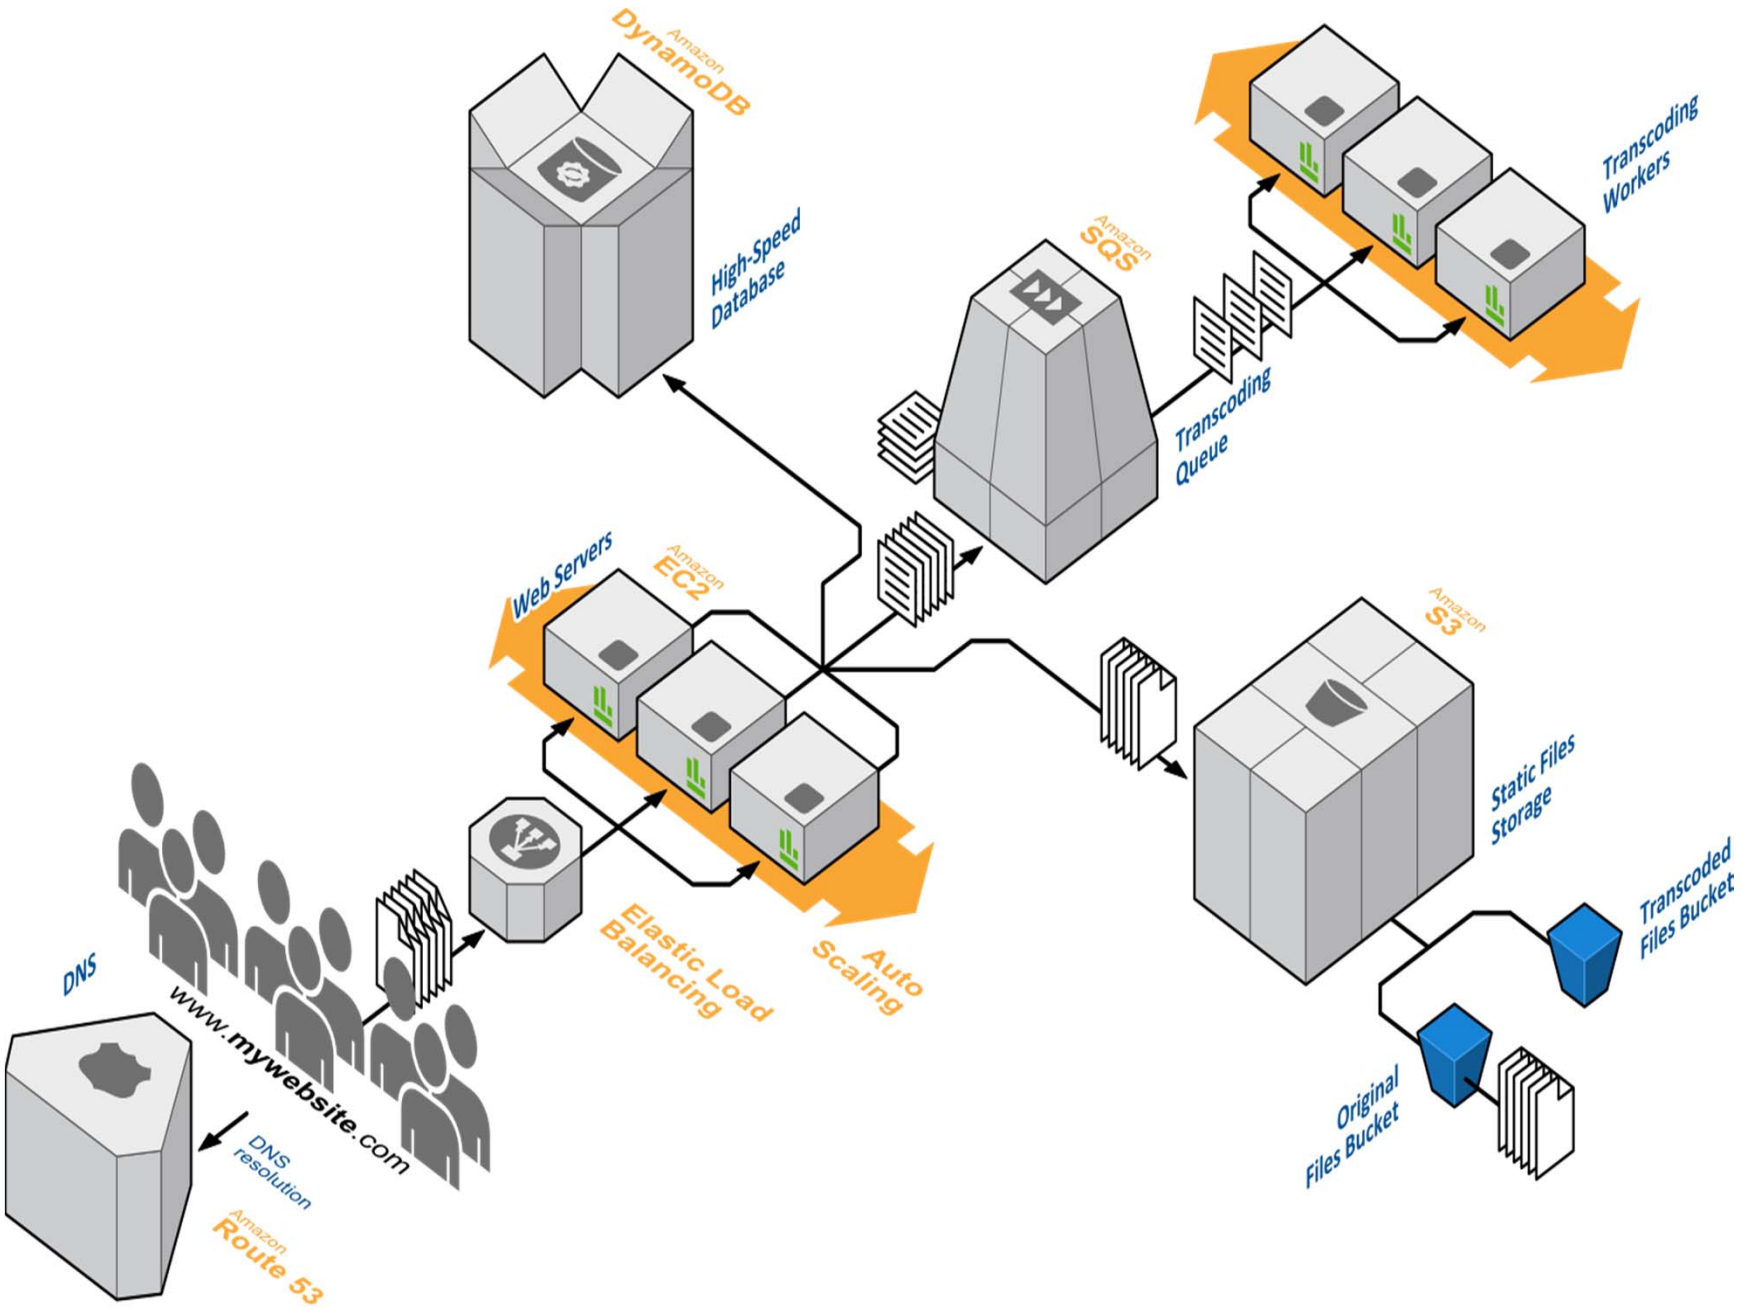
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Web Servers
Amazon EC2
Elastic Load Balancing
Auto Scaling
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

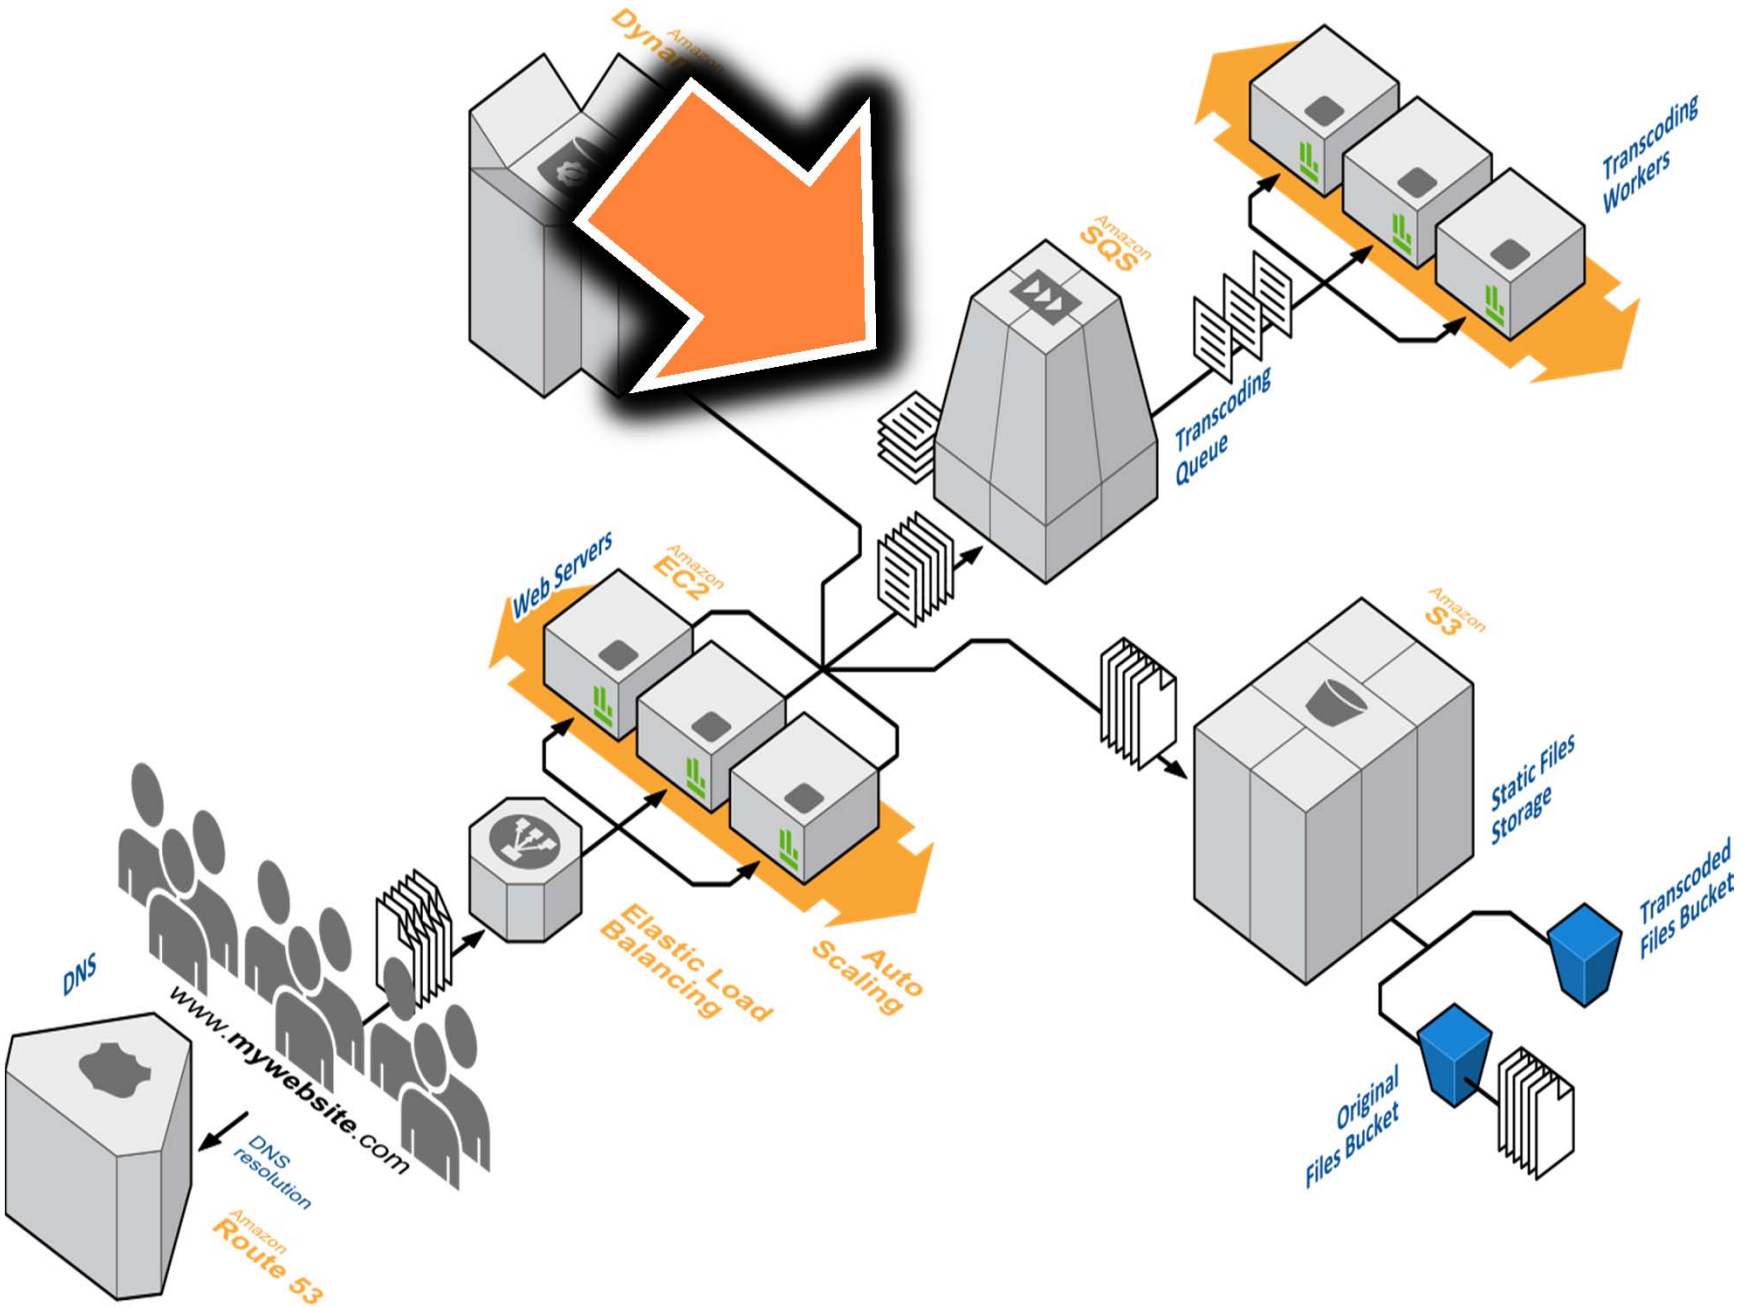
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Web Servers
Amazon
EC2
Elastic Load
Balancing
Auto
Scaling
Amazon
S3
Static Files
Storage
Transcoded
Files Bucket
Original
Files Bucket
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

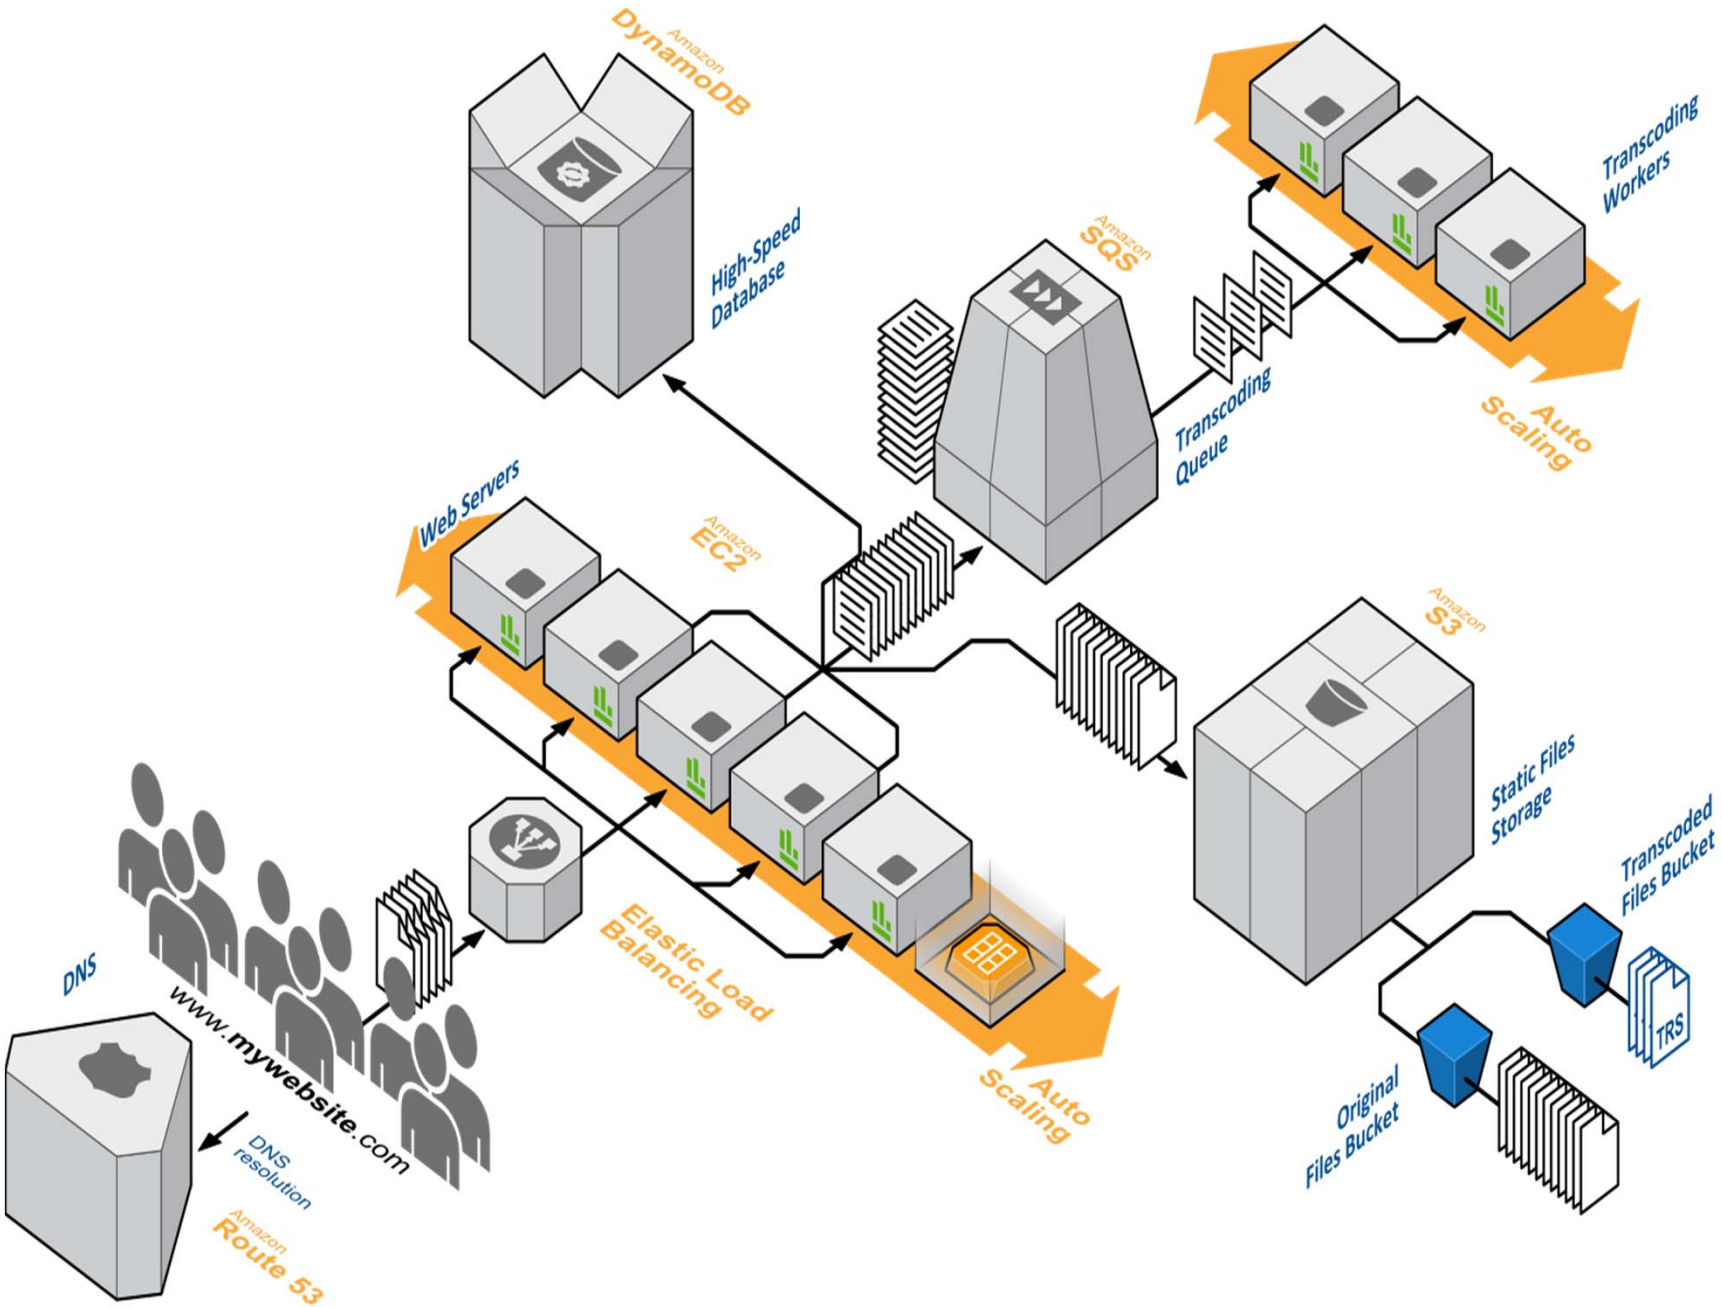
Amazon
DynamoDB
High-Speed
Database
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Auto
Scaling
Web Servers
Amazon
EC2
Amazon
S3
Static Files
Storage
Transcoded
Files Bucket
TRS
Original
Files Bucket
Elastic Load
Balancing
Auto
Scaling
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

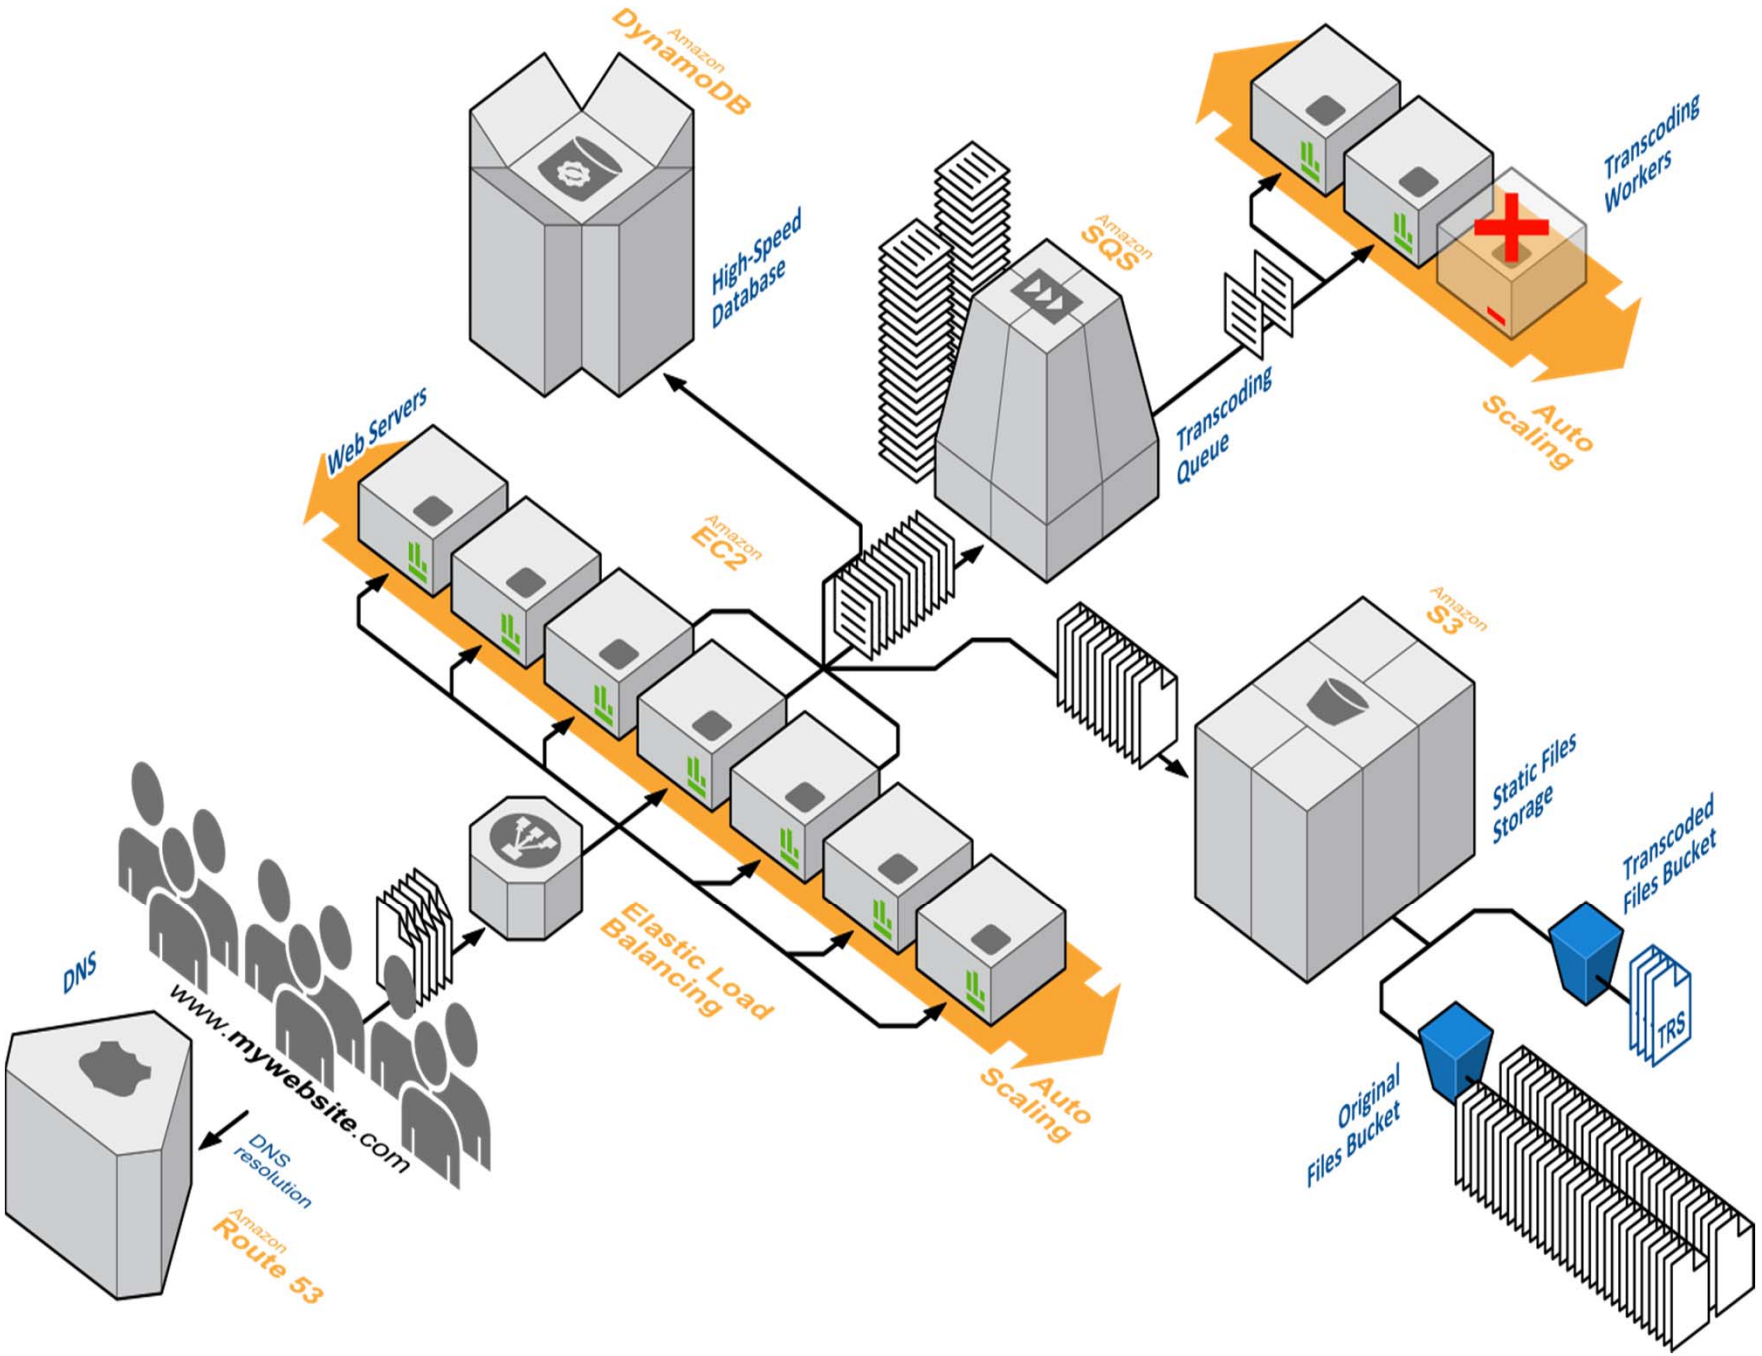
Amazon
DynamoDB
High-Speed
Database
Amazon
SQS
Transcoding
Queue
Transcoding
Workers
Auto
Scaling
Web Servers
Amazon
EC2
Amazon
S3
Static Files
Storage
Transcoded
Files Bucket
TRS
Original
Files Bucket
Elastic Load
Balancing
Auto
Scaling
DNS
www.mywebsite.com
DNS
resolution
Amazon
Route 53

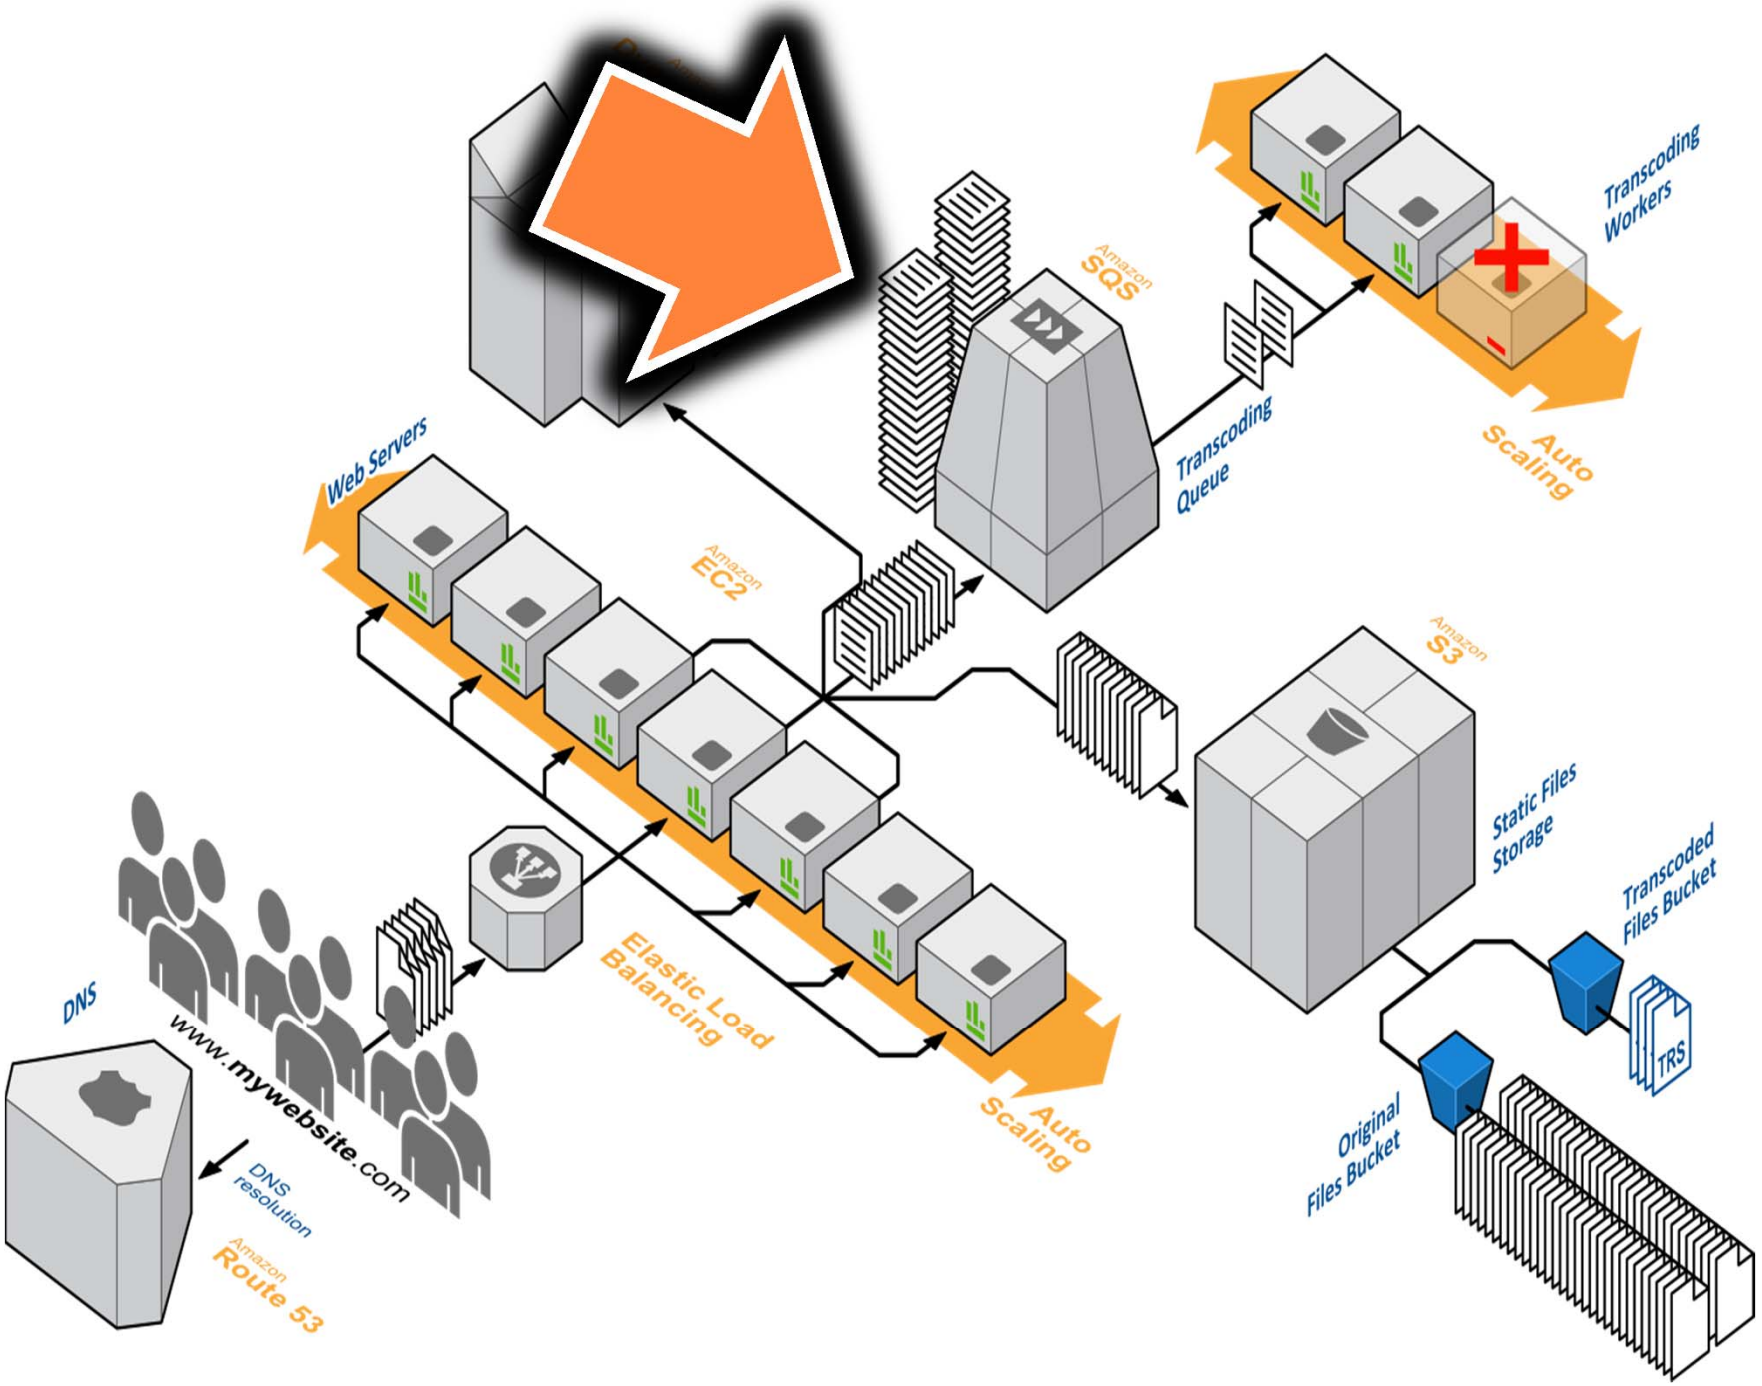
Transcoding Workers
Amazon SQS
Auto Scaling
Transcoding Queue
Web Servers
Amazon EC2
Amazon S3
Static Files Storage
Transcoded Files Bucket
TRS
Original Files Bucket
Elastic Load Balancing
Auto Scaling
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

CLOUDWATCH METRICS FOR AMAZON SQS

+

AUTO SCALING

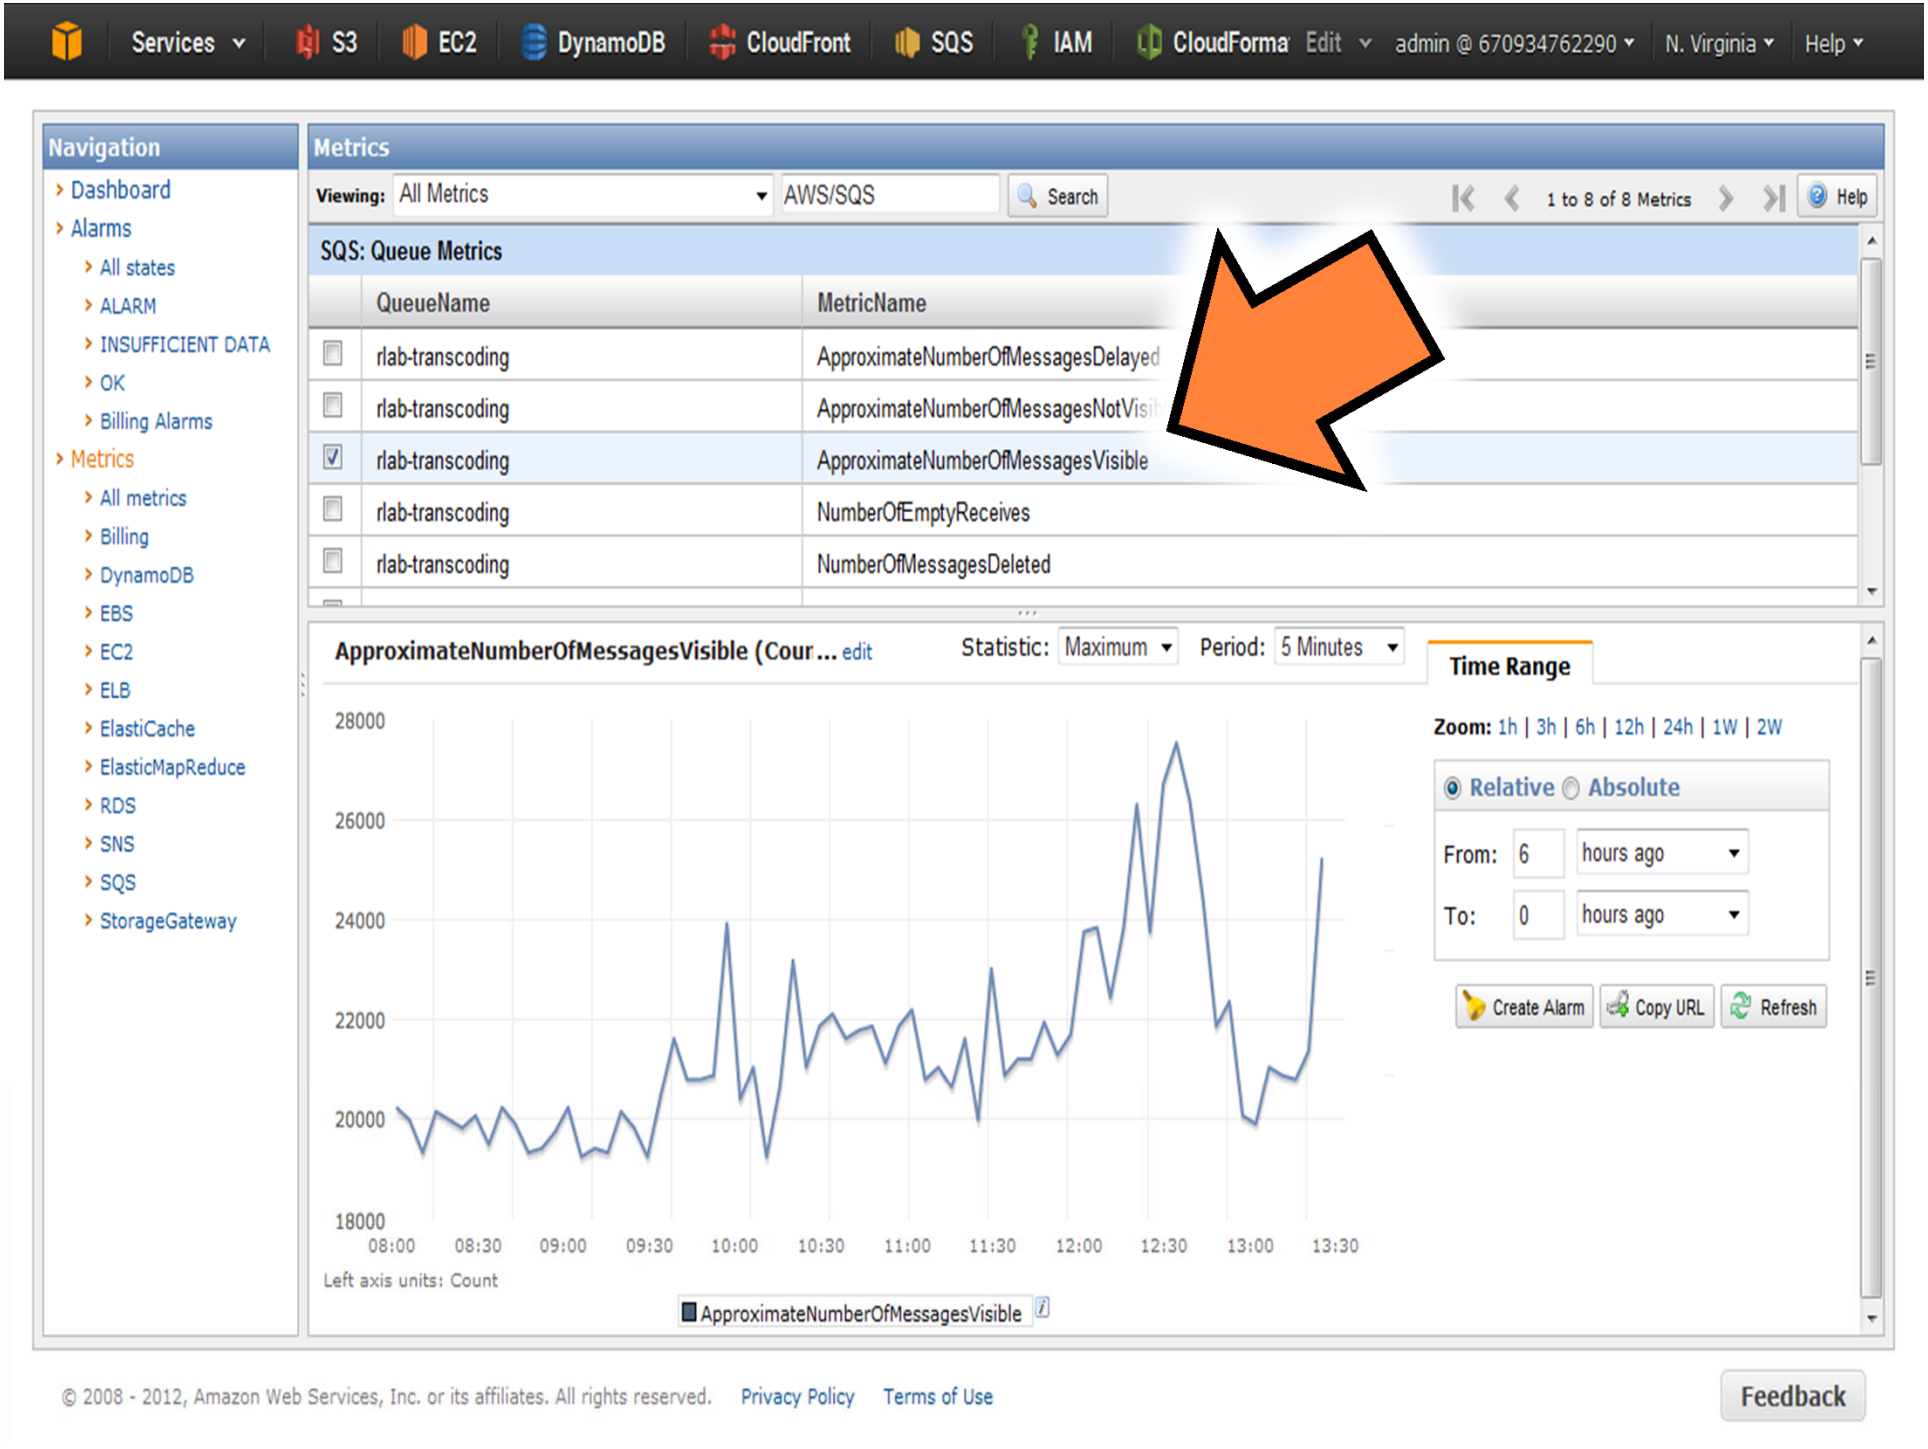
Services
S3
EC2
DynamoDB
CloudFront
SQS
IAM
CloudForma
Edit
admin @ 670934762290
N. Virginia
Help
Navigation
Dashboard
Alarms
All states
ALARM
INSUFFICIENT DATA
OK
Billing Alarms
Metrics
All metrics
Billing
DynamoDB
EBS
EC2
ELB
ElastiCache
ElasticMapReduce
RDS
SNS
SQS
StorageGateway
Metrics
Viewing: All Metrics
AWS/SQS
Search
1 to 8 of 8 Metrics
Help
SQS: Queue Metrics
QueueName
MetricName
rlab-transcoding
ApproximateNumberOfMessagesDelayed
rlab-transcoding
ApproximateNumberOfMessagesNotVisible
rlab-transcoding
ApproximateNumberOfMessagesVisible
rlab-transcoding
NumberOfEmptyReceives
rlab-transcoding
NumberOfMessagesDeleted
ApproximateNumberOfMessagesVisible (Cour... edit
Statistic: Maximum
Period: 5 Minutes
Time Range
Zoom: 1h | 3h | 6h | 12h | 24h | 1W | 2W
Relative
Absolute
From: 6 hours ago
To: 0 hours ago
Create Alarm
Copy URL
Refresh
28000
26000
24000
22000
20000
18000
08:00
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
Left axis units: Count
ApproximateNumberOfMessagesVisible
© 2008 - 2012, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use
Feedback

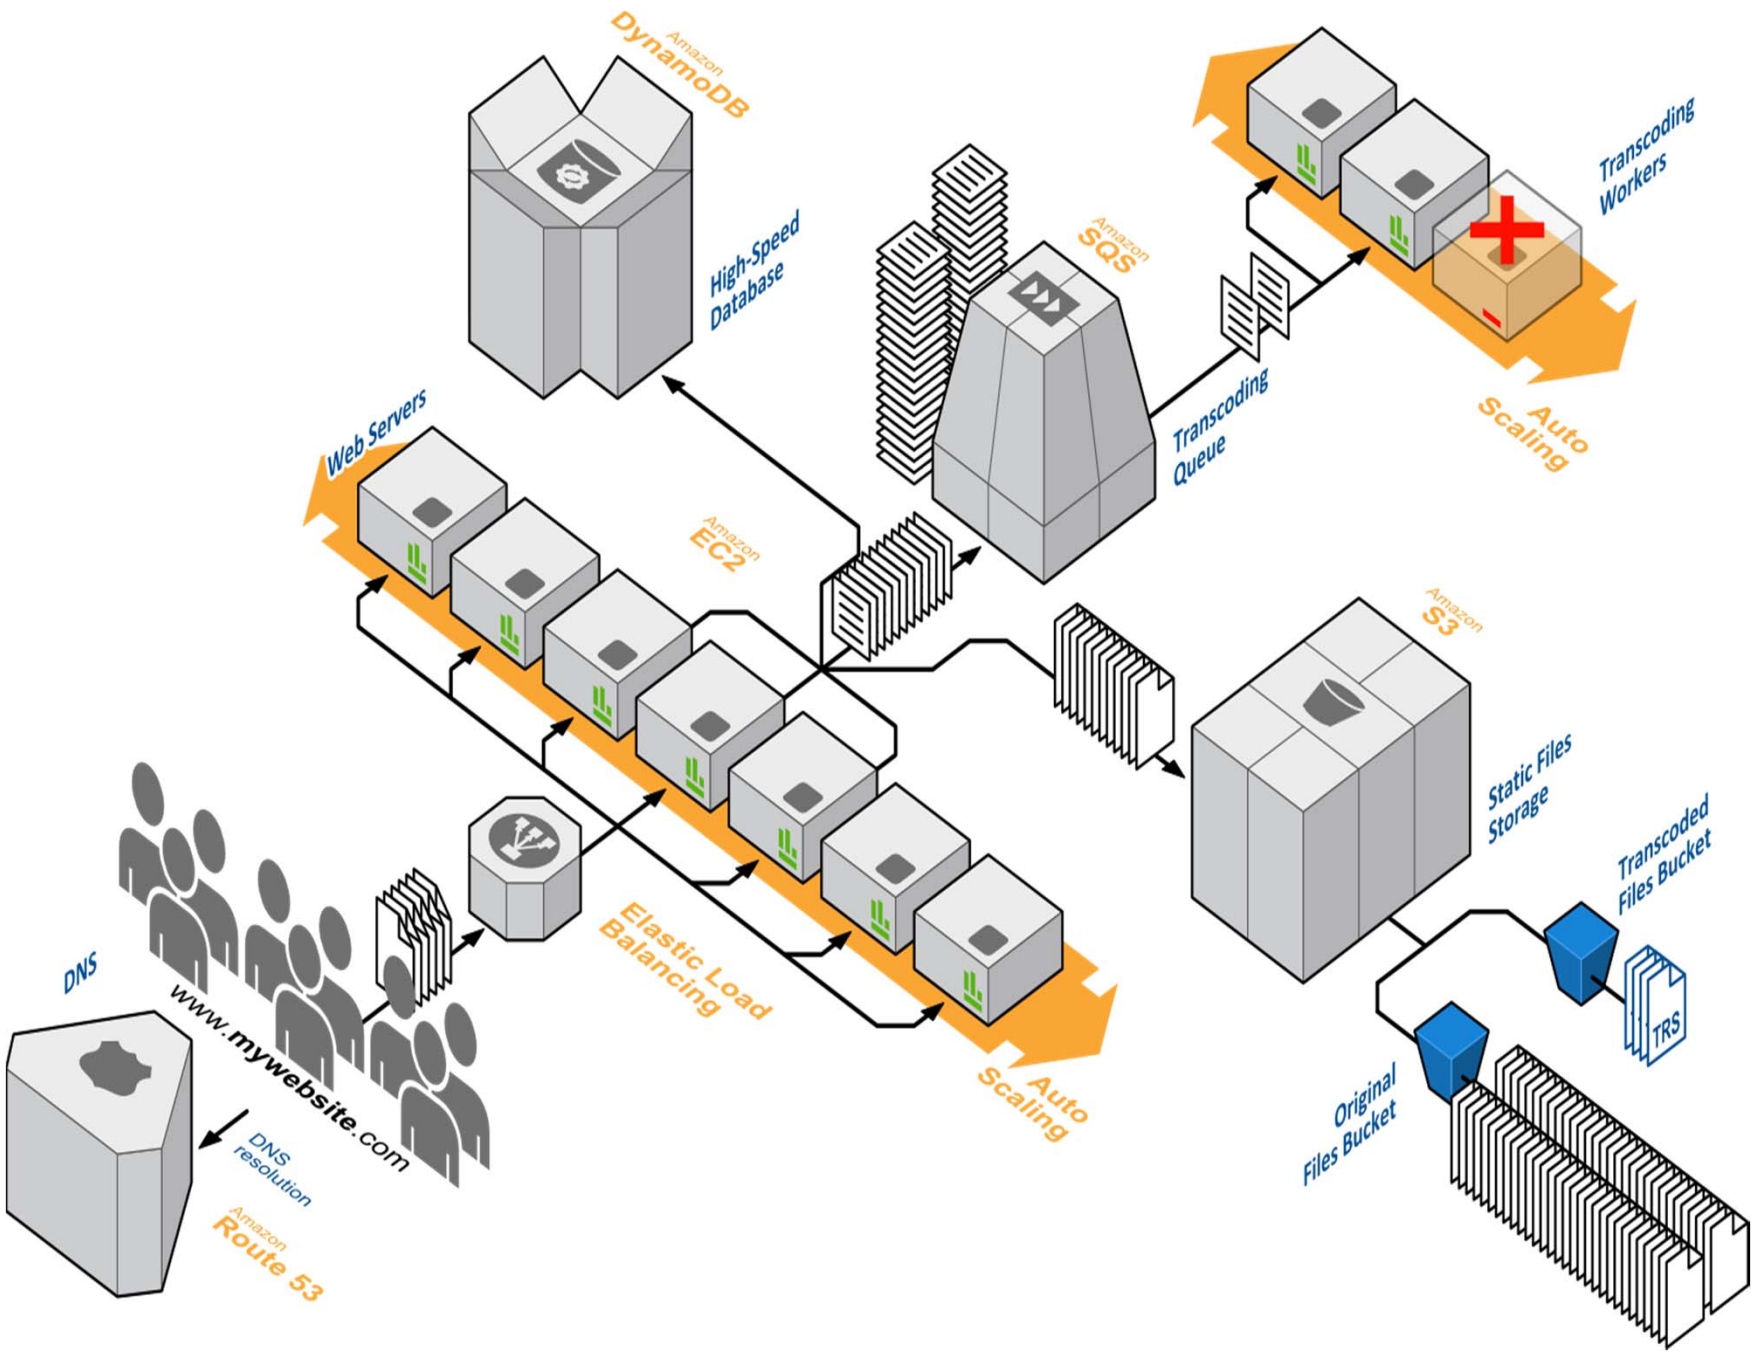
Amazon DynamoDB
High-Speed Database
Amazon SQS
Transcoding Queue
Transcoding Workers
Auto Scaling
Web Servers
Amazon EC2
Amazon S3
Static Files Storage
Transcoded Files Bucket
TRS
Original Files Bucket
Elastic Load Balancing
Auto Scaling
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

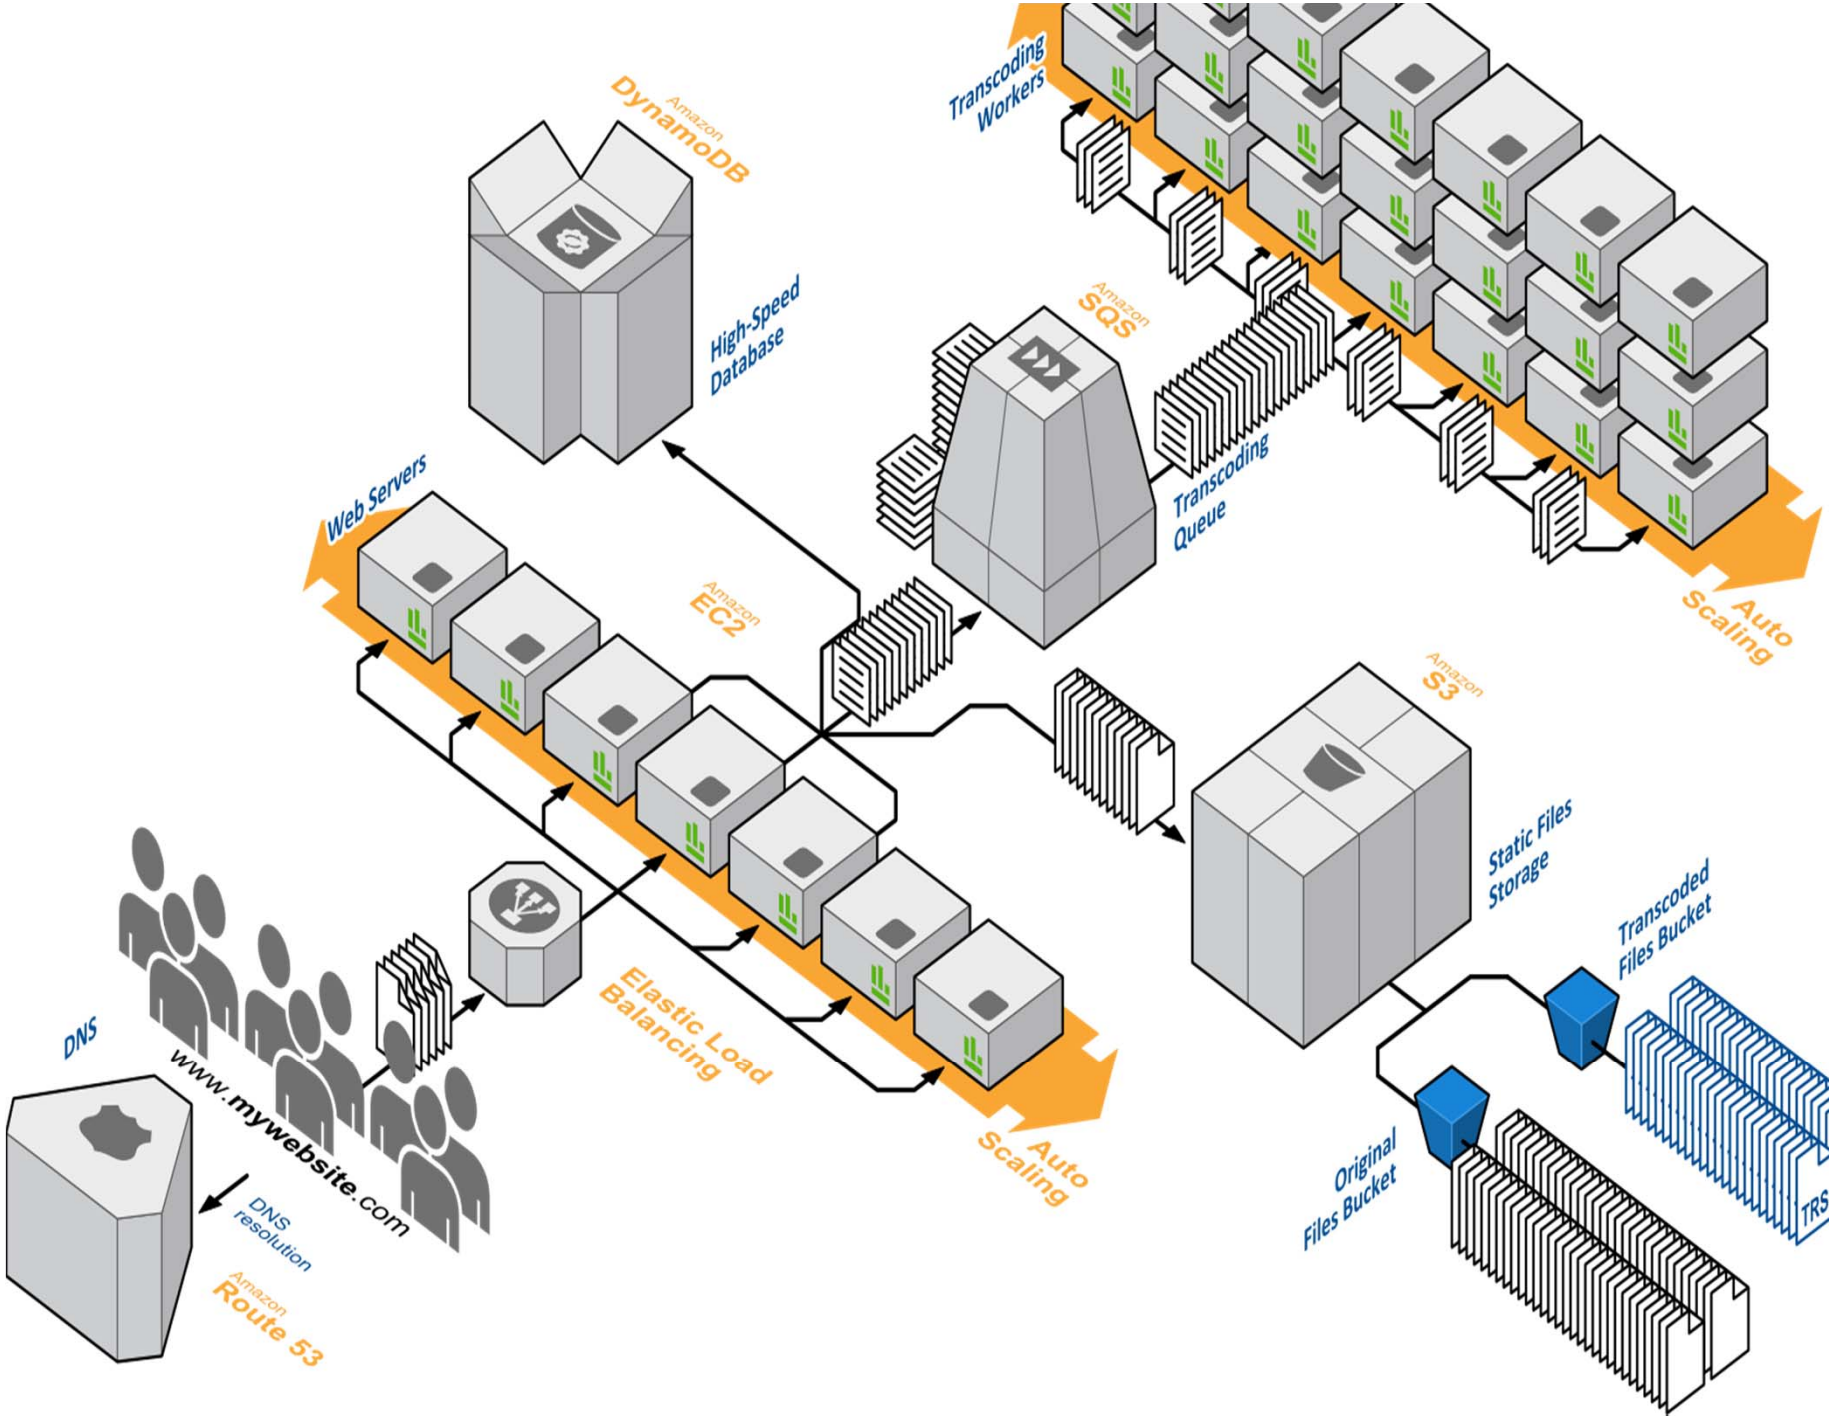
Amazon DynamoDB
High-Speed Database
Transcoding Workers
Amazon SQS
Transcoding Queue
Auto Scaling
Web Servers
Amazon EC2
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
TRS
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancing
Auto Scaling

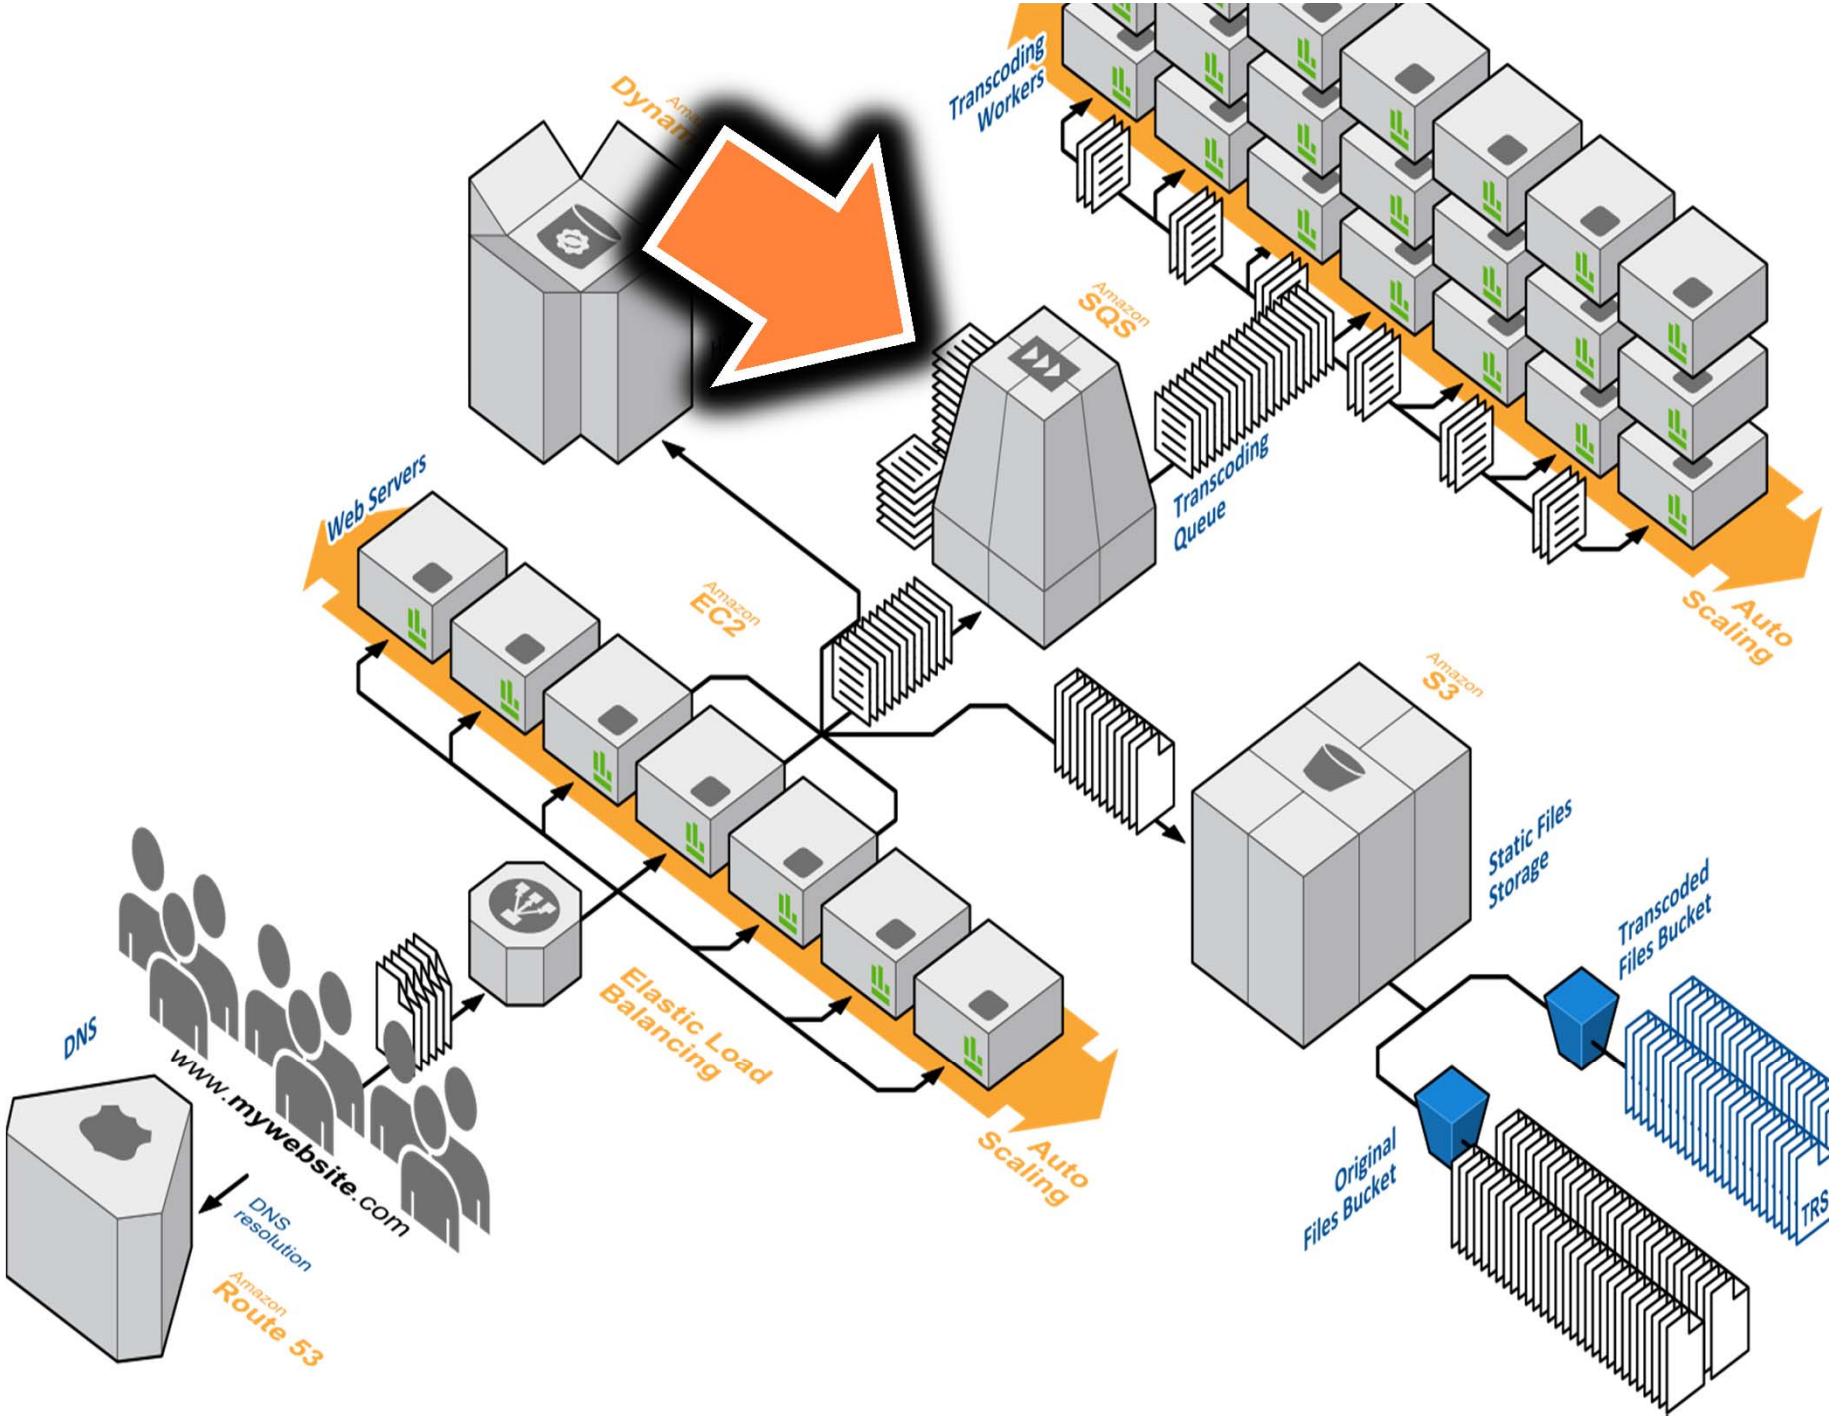
Transcoding Workers
Amazon SQS
Transcoding Queue
Auto Scaling
Web Servers
Amazon EC2
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
TRS
Elastic Load Balancing
Auto Scaling
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53

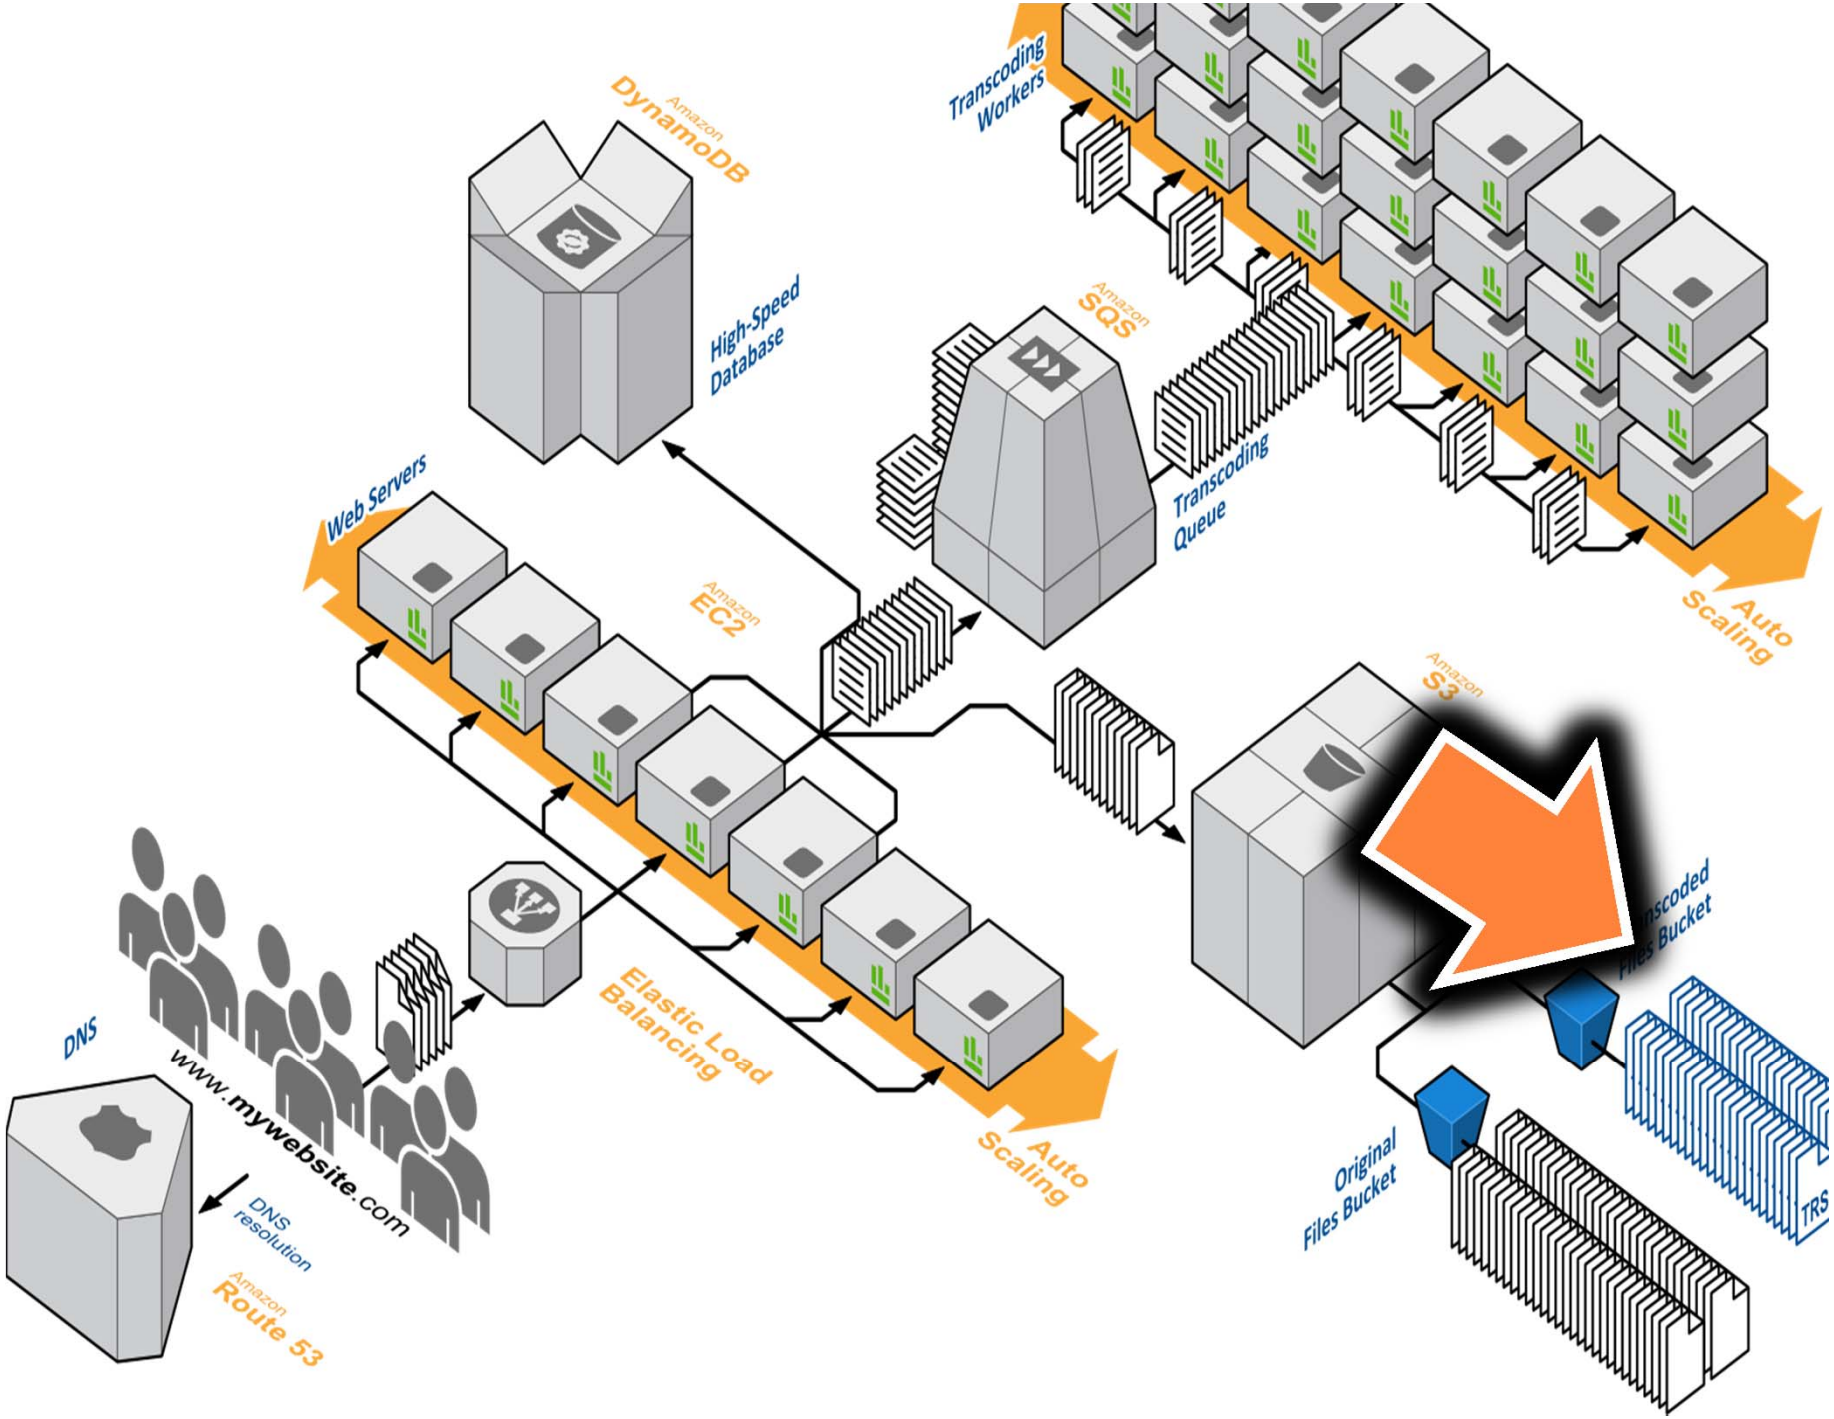
Amazon DynamoDB
High-Speed Database
Transcoding Workers
Amazon SQS
Transcoding Queue
Auto Scaling
Web Servers
Amazon EC2
Amazon S3
Elastic Load Balancing
Auto Scaling
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Original Files Bucket
TRS

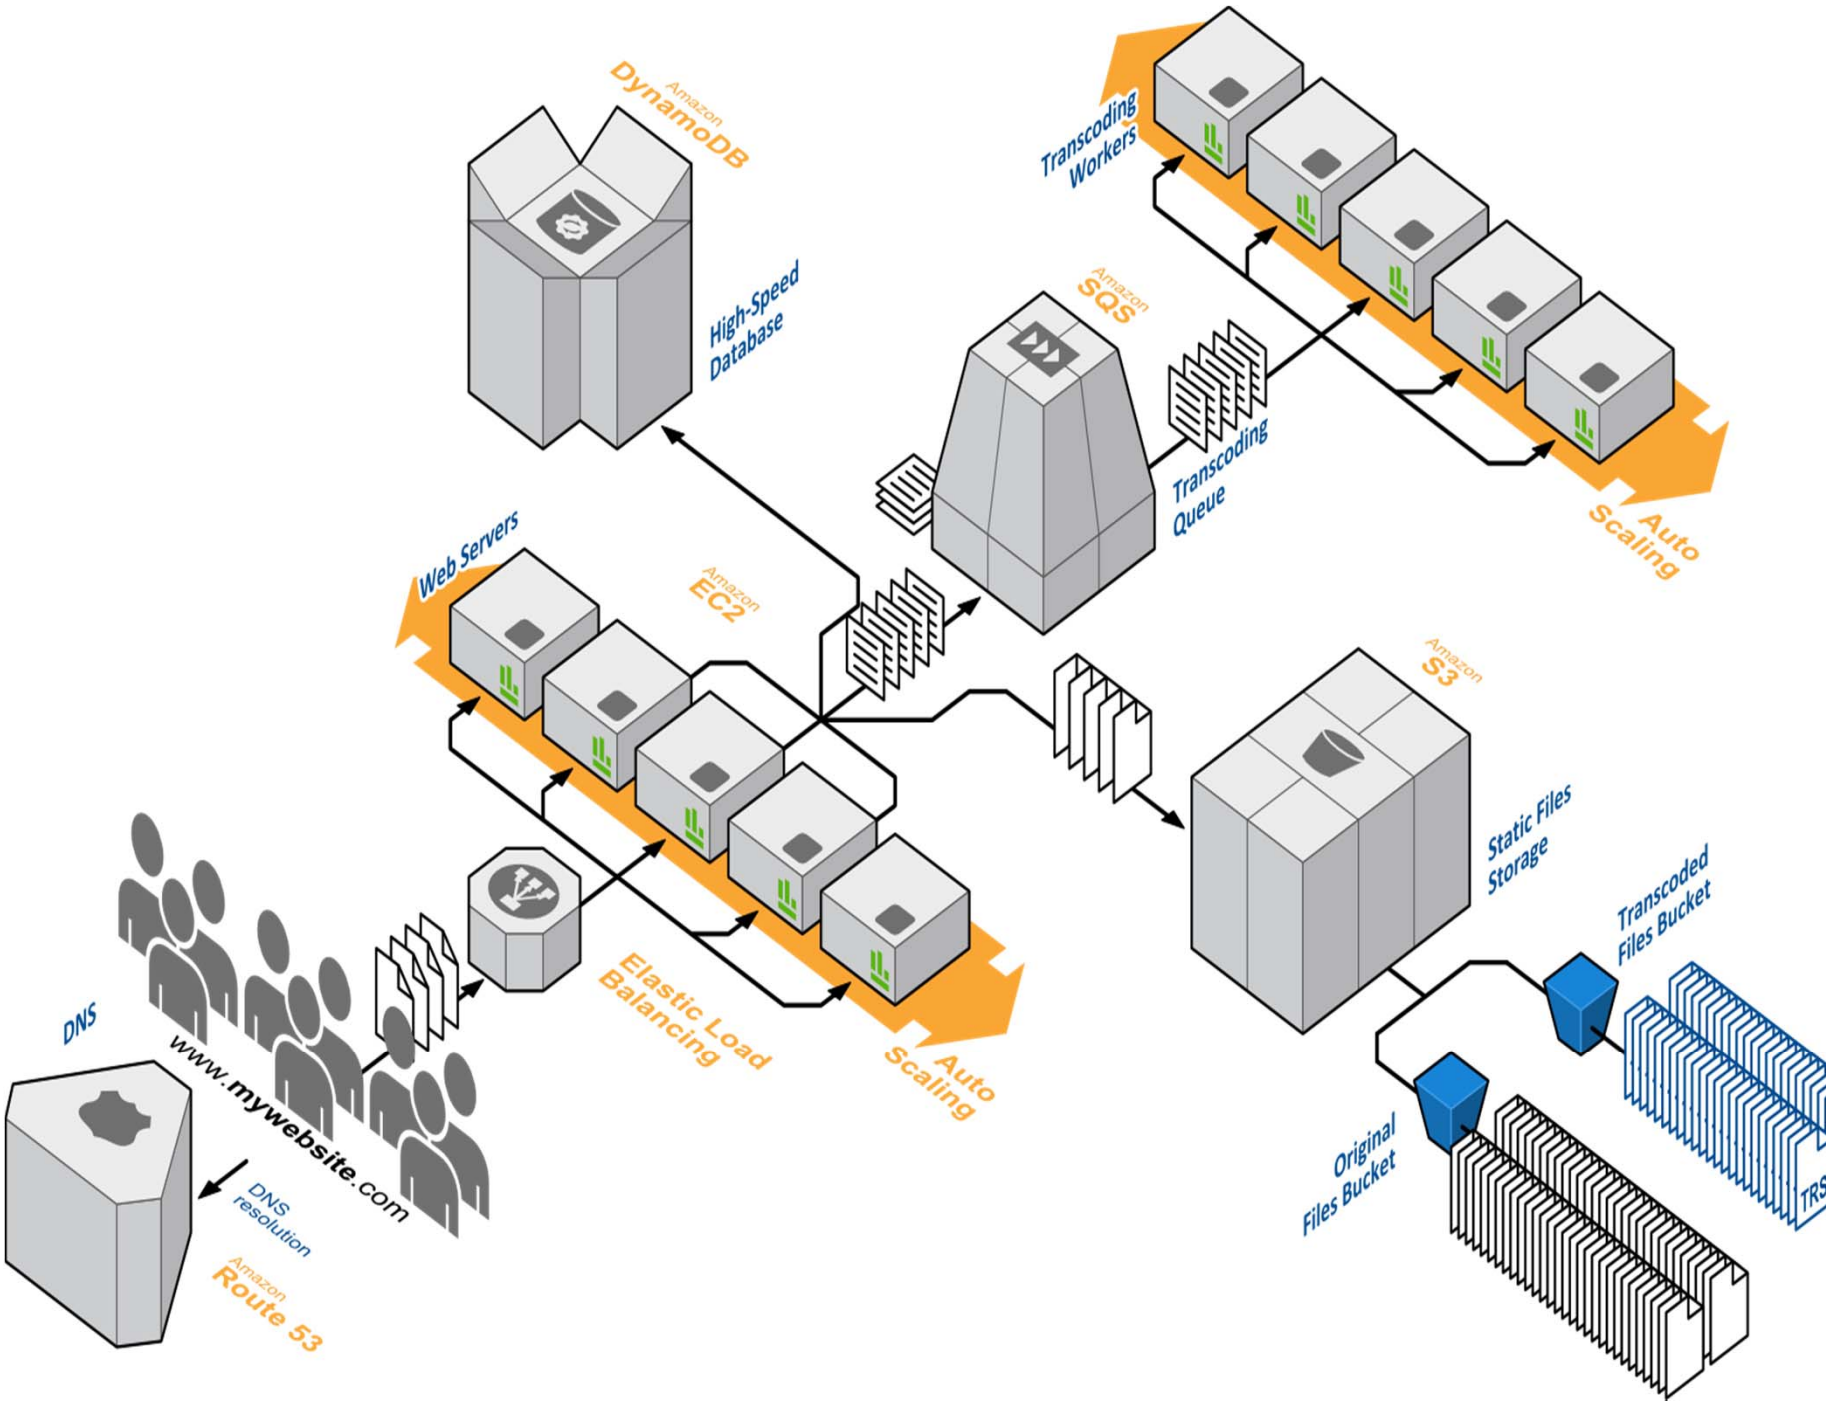
Amazon DynamoDB
High-Speed Database
Transcoding Workers
Amazon SQS
Transcoding Queue
Auto Scaling
Web Servers
Amazon EC2
Amazon S3
Static Files Storage
Transcoded Files Bucket
Original Files Bucket
DNS
www.mywebsite.com
DNS resolution
Amazon Route 53
Elastic Load Balancing
Auto Scaling

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY
3. SCALING
4. SELF-HEALI
5. LOOSE COU

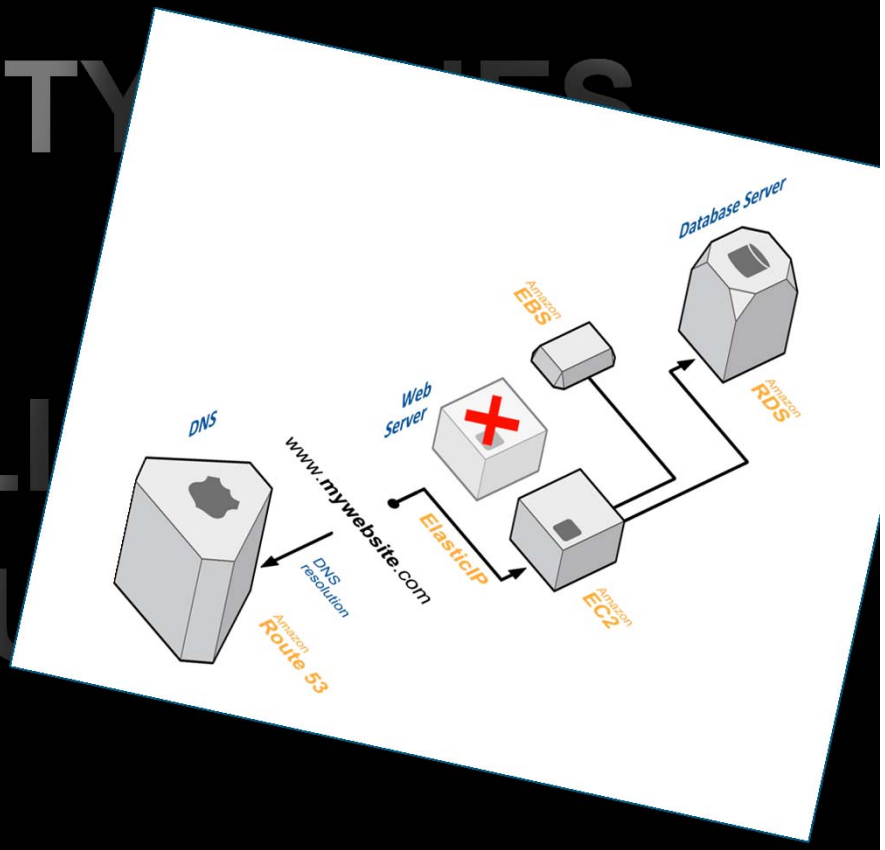

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILITY ZONES
3. SCAL
4. SELF
5. LOO G

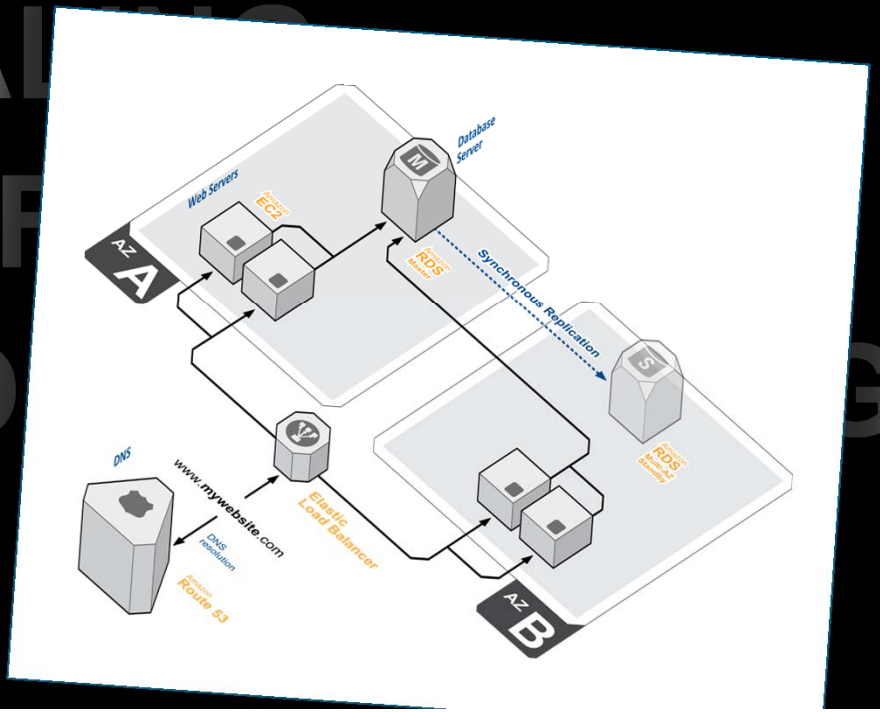

1. DESIGN FOR FAILURE
2. MULTIPLE AVAILABILIT
3. **SCALING**
4. SELF-HEALI
5. LOOSE COUP ING

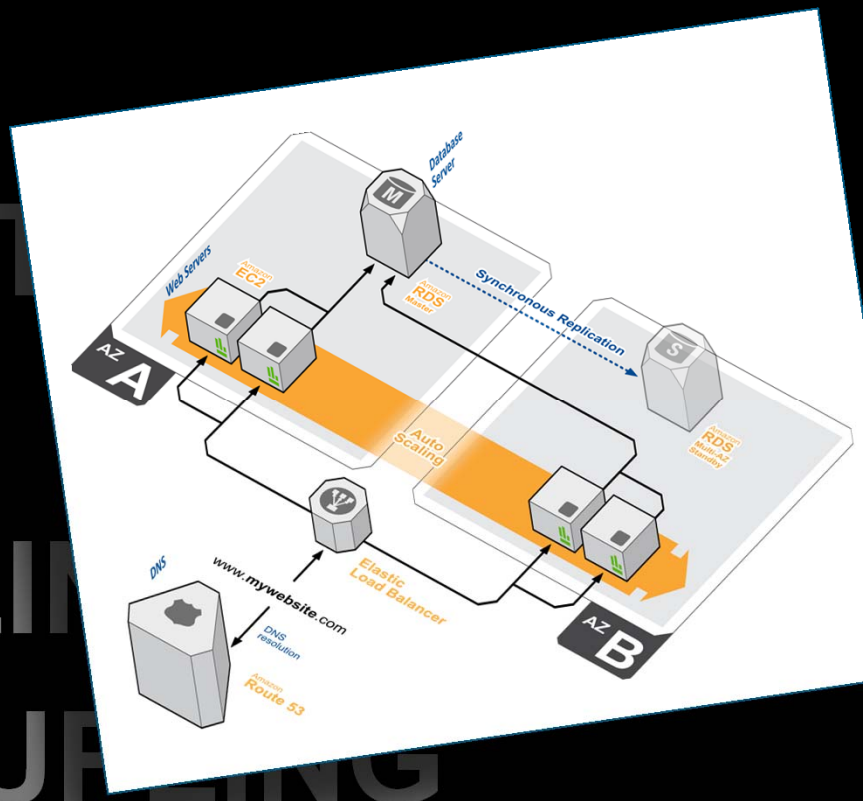

1. DESIGN FOR
2. MULTIPLE AVAILABILIT
3. SCALING
4. **SELF-HEALING**
5. LOOSE COUPLING

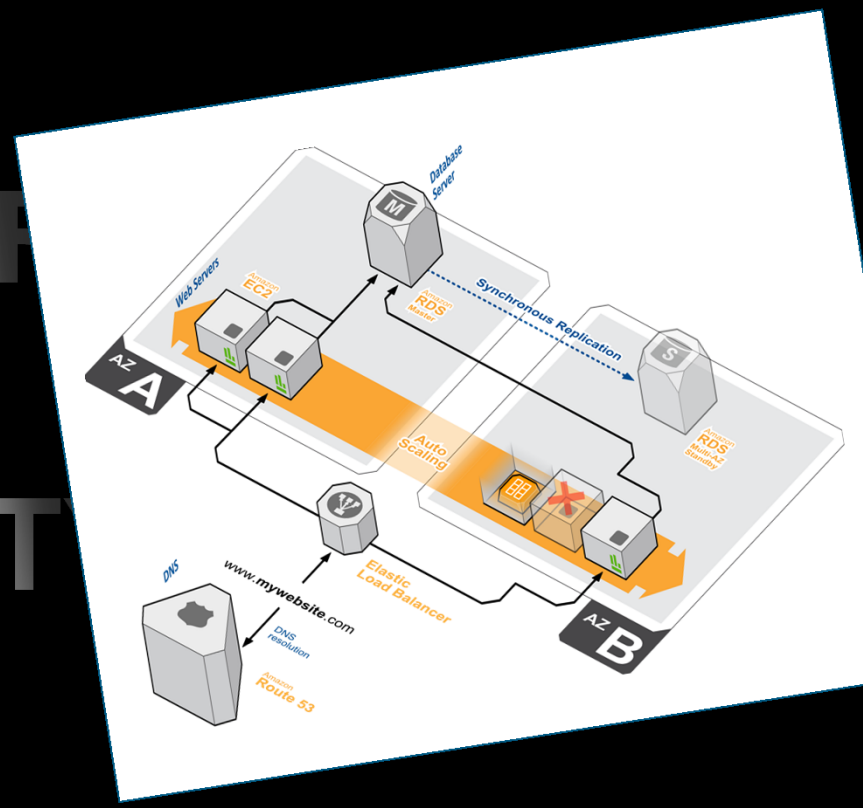

1. DESIGN FOR
2. MULTIPLE AVAILABILI
3. SCALING
4. SELF-HEALING
5. LOOSE COUPLING

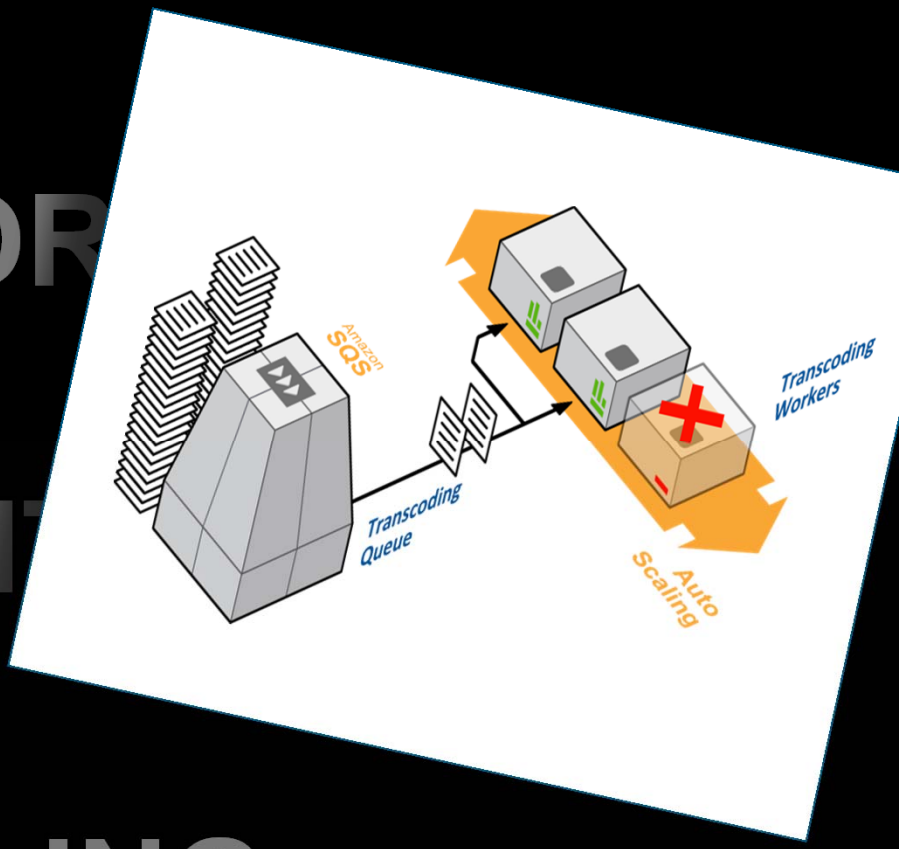

YOUR GOAL

Applications should
continue to function

IT'S ALL ABOUT
CHOICE

BALANCE COST & HIGH AVAILABILITY

Thank You!

Feedback ID 1927 ;-)

AWS is Hiring

Solutions Architects

London – Ireland – Luxembourg – Berlin – Munich – Paris

Attila Narin

attila@amazon.com

linkedin.com/in/attilanarin