

# Server Load Balancer

---

jnlin

# Introduction

---

- ❑ More users, more resources needed
  - CPU, RAM, HDD ...
  
- ❑ Scale Up & Scale Out
  - One powerful server to service more users; or
  - Multiple servers to service more users
  
- ❑ Pros & Cons ?
  
- ❑ C10K Problem

# Introduction

---

## ❑ High Availability

- A characteristic of a system, which aims to ensure an agreed level of operational performance, usually uptime, for a higher than normal period.


## ❑ Availability (per year)

- 99%: 3.65days
- 99.9%: 8.77 hours (3 nines)
- 99.99%: 52.60 minutes (4 nines)
- 99.999%: 5.26 minutes (5 nines)

# High Availability

## ❑ Principles

- Elimination of single points of failure.
- Reliable crossover.
  - Reliable configuration / topology change
- Detection of failures as they occur.



## ❑ Graceful Degradation

- the ability of a computer, machine, electronic system or network to maintain limited functionality even when a large portion of it has been destroyed or rendered inoperative.

# Load Balancing

---

## ❑ Client Side

- e.g: DNS round-robin
- Pros & Cons

## ❑ Server Side

- Server Load Balancer

# Server Load Balancer (1)

---

- ❑ Provide “Scale-Out” and HA features
  
- ❑ Share loading among all backend nodes with some algorithms
  - Static Algorithms: does not take into account the state of the system for the distribution of tasks.
  - Dynamic Algorithms

# Server Load Balancer (2)

---

- ❑ Layer 4 or Layer 7
  - Layer 4 Switch
  
- ❑ Distribution Algorithms
  - Round-robin
  - Random
  - Ratio
  - Hash Table
  - Least-connections
  - Persistence
    - Session-ID (e.g. HTTP Cookie)

# Server Load Balancer (3)

## ❑ Persistence (Stickiness)

- “The Server” in OLG
- How to handle information that must be kept across the multiple requests in a user's session.

## ❑ Session ID?

- Cookie
- IP Address
- TCP Connection

## ❑ Pros & Cons ?





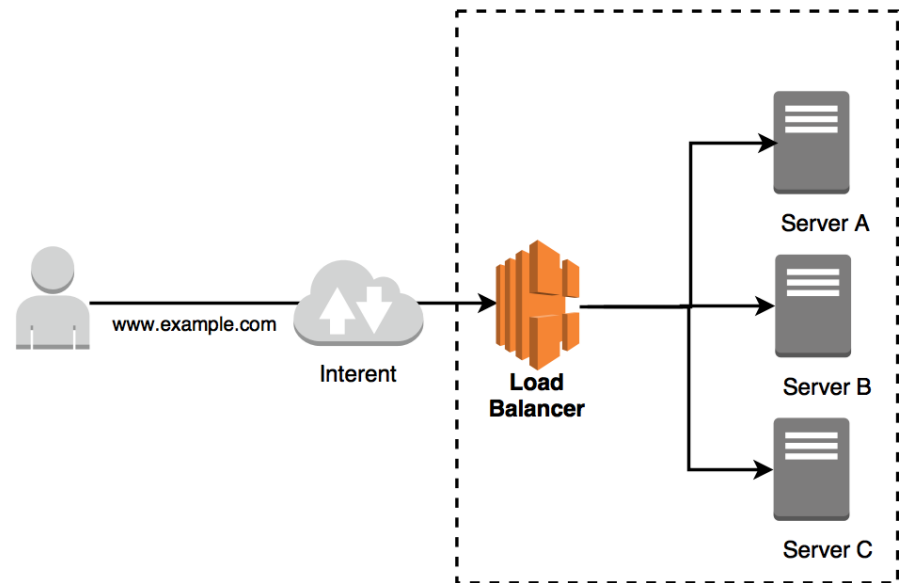
# Server Load Balancer (4)

## ❑ SSL offloading (SSL/TLS termination)

- Pros?

## ❑ Problems of Server Load Balancer

- SPoF
- Capacity Limit
- Latency



# Haproxy

---

❑ <http://www.haproxy.org>

❑ Reliable & High Performance TCP/HTTP Load Balancer

- Layer 4 (TCP) and Layer 7 (HTTP) load balancing
- SSL/TLS termination
- Gzip compression
- Health checking
- HTTP/2

# Haproxy - Installation

---

- ❑ `pkg install haproxy`
- ❑ You can also build it from ports
- ❑ Config file: `/usr/local/etc/haproxy.conf`

# Haproxy - Configuration

```
1  global
2    daemon
3    log 127.0.0.1 local0
4    log 127.0.0.1 local1 notice
5    maxconn 4096
6    tune.ssl.default-dh-param 2048
7
8  defaults
9    log                global
10   retries             3
11   maxconn             2000
12   timeout connect    5s
13   timeout client     50s
14   timeout server     50s
15
16  listen stats
17   bind 127.0.0.1:9090
18   balance
19   mode http
20   stats enable
21   stats auth admin:admin
```

# Haproxy - Configuration

```
22
- 23   frontend www_csie_nctu
24     bind 140.113.208.102:80
25     mode http
26     use_backend www_csie_nctu_server
27
28   frontend cscs_csie_nctu
29     bind 140.113.208.103:80
30     mode http
31     use_backend cscs_csie_nctu_server
32
33   frontend game_server
34     bind 140.113.208.104:9876
35     mode tcp
36
37   backend www_csie_nctu_server
38     balance roundrobin
39     mode http
40     option forwardfor
41     http-request set-header X-Forwarded-Port %[dst_port]
42     http-request add-header X-Forwarded-Proto https if { ssl_fc }
43     server ww1 192.168.99.1:80
44     server ww2 192.168.99.2:80
```

# Haproxy - Configuration

---

```
backend cscs_csie_nctu_server
  balance roundrobin
  mode http
  option httpchk HEAD /health_check.php HTTP/1.1\r\nHost:\ cscs.cs.nctu.edu.tw
  option forwardfor
  http-request set-header X-Forwarded-Port %[dst_port]
  http-request add-header X-Forwarded-Proto https if { ssl_fc }
  server ww1 192.168.99.101:80 check fall 3 rise 2
  server ww2 192.168.99.102:80 check fall 3 rise 2
```

# Haproxy Configuration

---

- ❑ global
  - log
  - chroot
  - uid / gid
  - pidfile

# Haproxy Configuration

---

- ❑ defaults
  - log
  - option
  - retries
  - timeout



# Haproxy Configuration

## □ listen

- stats

← → ↻ 192.168.10.10:1936/haproxy?stats

## HAProxy

### Statistics Report for pid 7076 on tecadmin.net

> General process information

pid = 7076 (process #1, nbproc = 1)  
 uptime = 0d 0h00m32s  
 system limits: memmax = unlimited; ulimit-n = 90017  
 maxsock = 90017; maxconn = 45000; maxpipes = 0  
 current conns = 1; current pipes = 0/0  
 Running tasks: 1/5

■ active UP      ■ backup UP  
■ active UP, going down      ■ backup UP, going down  
■ active DOWN, going up      ■ backup DOWN, going up  
■ active or backup DOWN      ■ not checked

Display option:  
 • [Hide 'DOWN' servers](#)  
 • [Disable refresh](#)  
 • [Refresh now](#)  
 • [CSV export](#)

External resources:  
 • [Primary site](#)  
 • [Updates \(v1.3\)](#)  
 • [Online manual](#)

Note: UP with load-balancing disabled is reported as "NOLB".

#### stats

	Queue			Session rate			Sessions				Bytes		Denied		Errors		Warnings		Server										
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle	
Frontend				1	2	-	1	2	10	4		1 372	26 971	0	0	0	0					OPEN							
Backend	0	0		0	1		0	1	10	1	0	1 372	26 971	0	0		1	0	0	0	0	32s UP	0	0	0			0	

#### http\_tecadmin\_net

	Queue			Session rate			Sessions				Bytes		Denied		Errors		Warnings		Server										
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle	
Frontend				0	0	-	0	0	2 000	0		0	0	0	0	0	0					OPEN							
server1	0	0	-	0	0		0	0	512	0	0	0	0	0	0	0	0	0	0	0	0	32s UP	1	Y	-	0	0	0s	-
server2	0	0	-	0	0		0	0	512	0	0	0	0	0	0	0	0	0	0	0	0	32s UP	1	Y	-	0	0	0s	-
Backend	0	0		0	0		0	0	2 000	0	0	0	0	0	0	0	0	0	0	0	0	32s UP	2	2	0		0	0s	

#### https\_tecadmin\_net

	Queue			Session rate			Sessions				Bytes		Denied		Errors		Warnings		Server										
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle	
Frontend				0	0	-	0	0	2 000	0		0	0	0	0	0	0					OPEN							
server1	0	0	-	0	0		0	0	512	0	0	0	0	0	0	0	0	0	0	0	0	32s UP	1	Y	-	0	0	0s	-
server2	0	0	-	0	0		0	0	512	0	0	0	0	0	0	0	0	0	0	0	0	32s UP	1	Y	-	0	0	0s	-
Backend	0	0		0	0		0	0	2 000	0	0	0	0	0	0	0	0	0	0	0	0	32s UP	2	2	0		0	0s	

# Haproxy Configuration

---

## ❑ frontend

- bind
- mode
- option
- use\_backend

# Haproxy Configuration

---

## ❑ backend

- balance
  - roundrobin, leastconn, hdr(param)
- mode
- http-request
- server
  - check
  - fall
  - rise
  - inter
  - cookie

# Haproxy - run

---

- ❑ /etc/rc.conf.local
  - haproxy\_enable="YES"
- ❑ /usr/local/etc/rc.d/haproxy start
- ❑ Question: how to setup a backup node for haproxy?

# Haproxy - Reference

---

- ❑ <http://cbonte.github.io/haproxy-dconv/2.1/configuration.html>

# Envoy

---

- ❑ <https://www.envoyproxy.io>
- ❑ Developed by Lyft (a ride-sharing company like Uber) and opensourced in 2017
  - Apache License 2.0
- ❑ Features
  - Dynamic APIs for configuration
  - Service Discovery
  - gRPC / MongoDB / HTTP support

# Envoy - Installation

---

- ❑ Broken in FreeBSD now
  - You can install it on Linux instead
  
- ❑ <https://www.getenvoy.io>
  - Debian: <https://www.getenvoy.io/install/envoy/debian/>
  - Ubuntu: <https://www.getenvoy.io/install/envoy/ubuntu/>
  - Centos: <https://www.getenvoy.io/install/envoy/centos/>

# Envoy - Configuration

```
1  static_resources:
2    listeners:
3      - address:
4          # Tells Envoy to listen on 0.0.0.0:80
5          socket_address:
6            address: 0.0.0.0
7            port_value: 80
8          filter_chains:
9            # Any requests received on this address are sent through this chain of filters
10         - filters:
11             # If the request is HTTP it will pass through this HTTP filter
12           - name: envoy.http_connection_manager
13             typed_config:
14               "@type": type.googleapis.com/envoy.config.filter.network.http_connection_manager.v2.HttpConnectionManager
15               codec_type: auto
16               stat_prefix: http
17               access_log:
18                 name: envoy.file_access_log
19                 typed_config:
20                   "@type": type.googleapis.com/envoy.config.accesslog.v2.FileAccessLog
21                   path: /dev/stdout
```



# Envoy - Configuration

```
22     route_config:
23       name: search_route
24       virtual_hosts:
25         - name: backend
26           domains:
27             - "*"
28           routes:
29             - match:
30                 prefix: "/"
31                 route:
32                   # Send request to an endpoint in the Bing cluster
33                   cluster: backend_server
34       http_filters:
35         - name: envoy.router
36           typed_config: {}
37     clusters:
38     - name: backend_server
39       connect_timeout: 1s
40       # Instruct Envoy to continuously resolve DNS asynchronously
41       type: logical_dns
42       dns_lookup_family: V4_ONLY
43       lb_policy: round_robin
```

# Envoy - Configuration

```
44     load_assignment:
45       cluster_name: backend_server
46       endpoints:
47         - lb_endpoints:
48           - endpoint:
49             address:
50               socket_address:
51                 address: 192.168.77.1
52                 port_value: 80
53           - endpoint:
54             address:
55               socket_address:
56                 address: 192.168.77.2
57                 port_value: 80
58     admin:
59       access_log_path: "/dev/stdout"
60       address:
61         socket_address:
62           address: 0.0.0.0
63           port_value: 15000
```

# Envoy - Configuration

---

- ❑ YAML file format
  
- ❑ Basic concept is same as haproxy
  - Listen (frontend) address
  - Backend addresses
  - Healthy Checks
    - <https://www.envoyproxy.io/learn/health-check>
  - Routes

# Envoy - Run

---

❑ `envoy -c config.yaml`

# Envoy - Reference

---

- ❑ <https://www.envoyproxy.io/docs/envoy/latest/>
- ❑ <https://blog.getambassador.io/envoy-vs-nginx-vs-haproxy-why-the-open-source-ambassador-api-gateway-chose-envoy-23826aed79ef>