

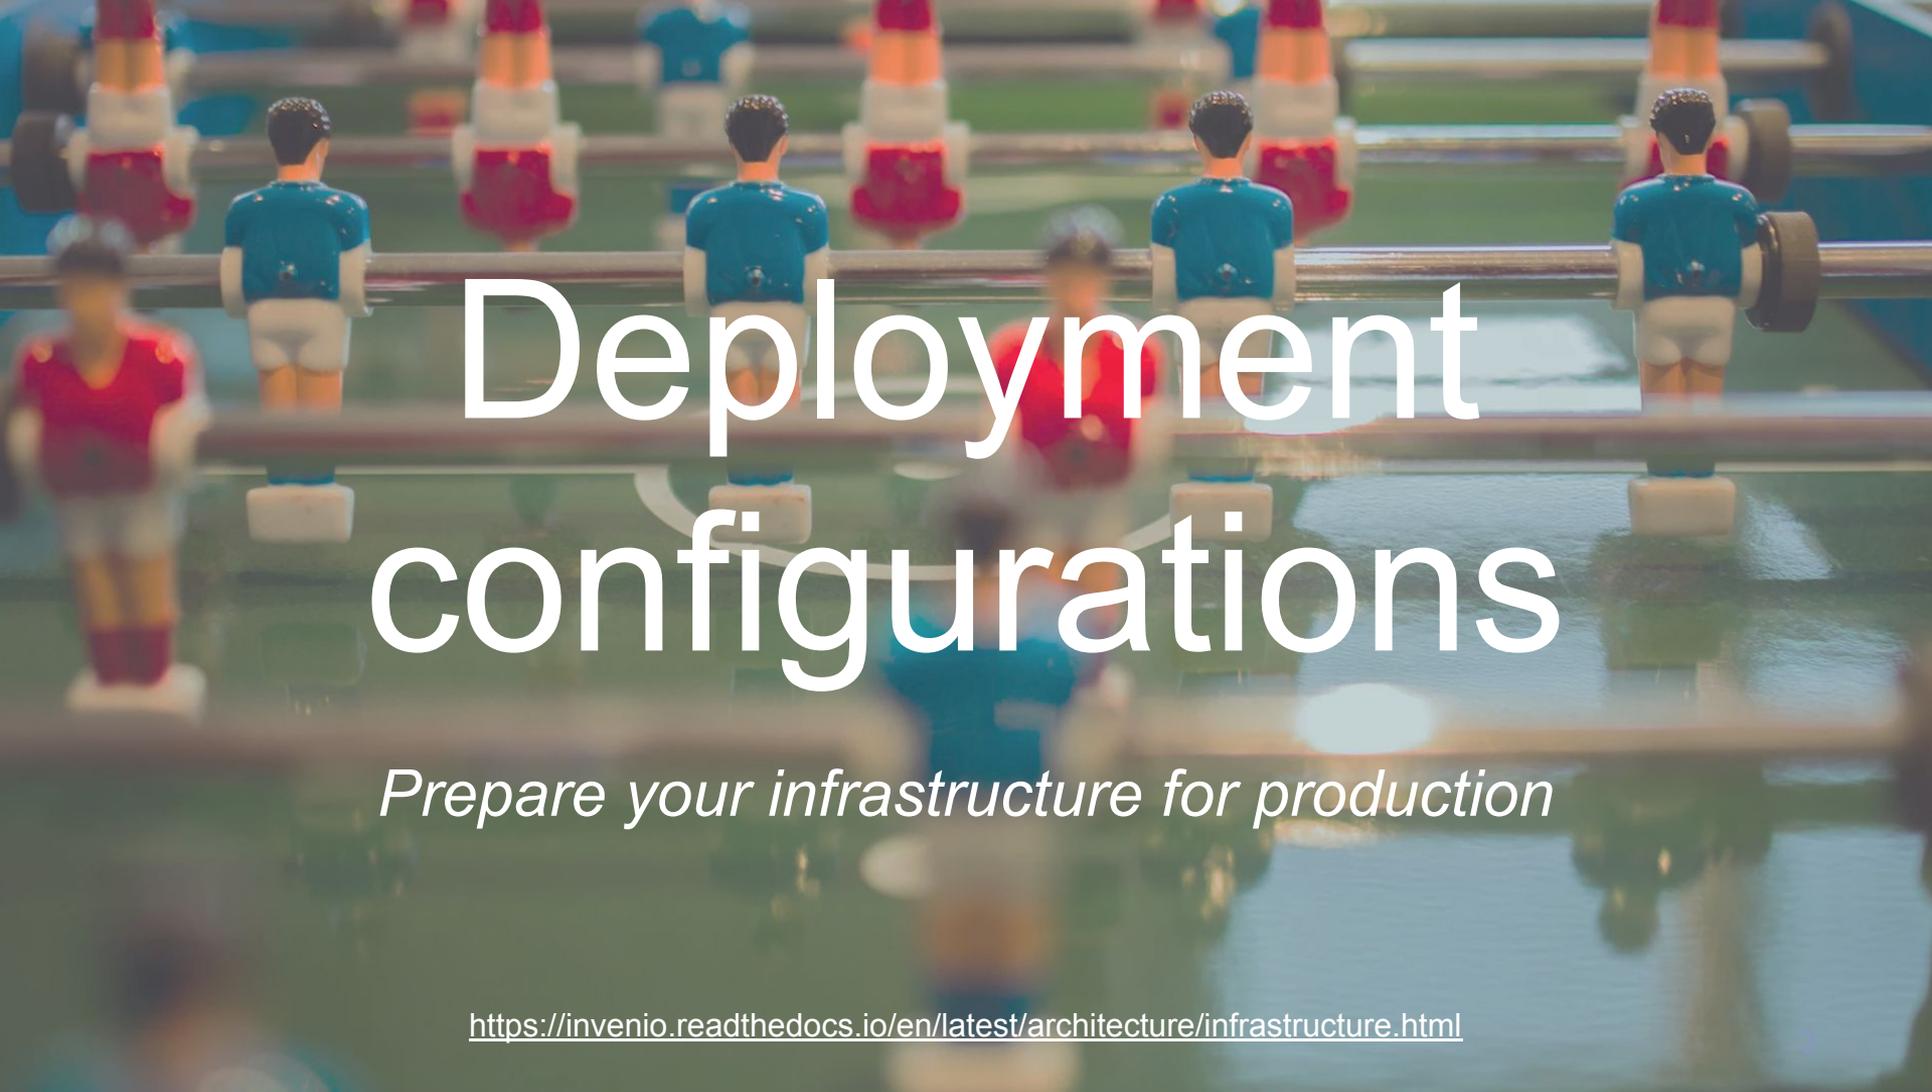
Deployment and monitoring

Nicola Tarocco and **our cool team**

CERN

Outline

- Deployment configurations
- Dimensions
- Monitoring
- Logging

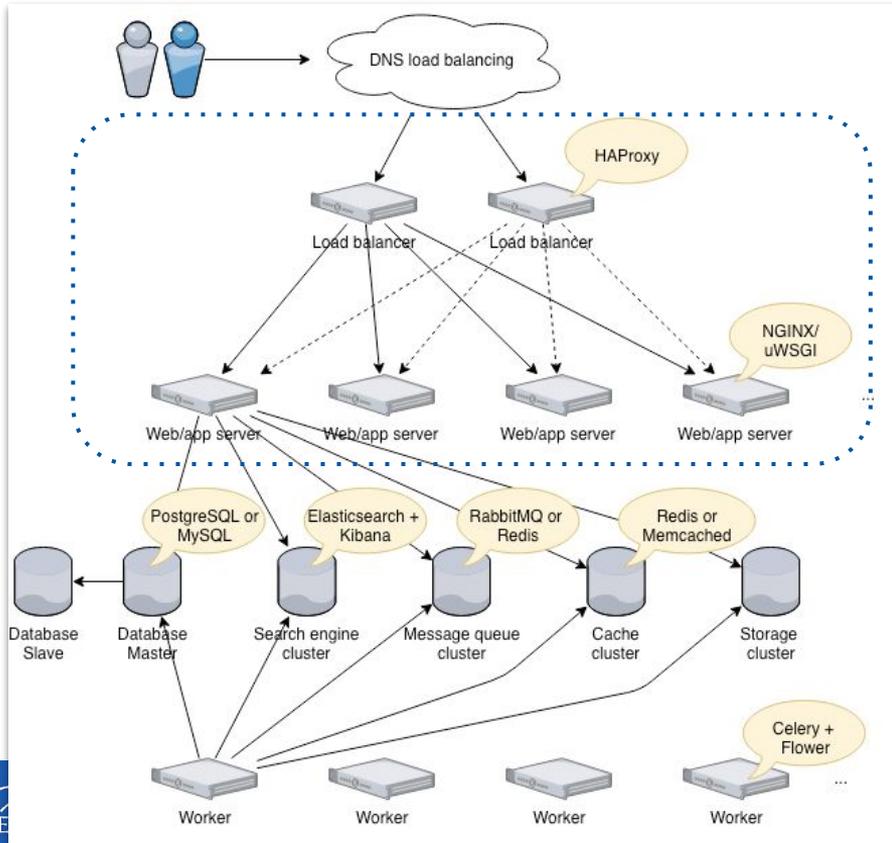


Deployment configurations

Prepare your infrastructure for production

<https://invenio.readthedocs.io/en/latest/architecture/infrastructure.html>

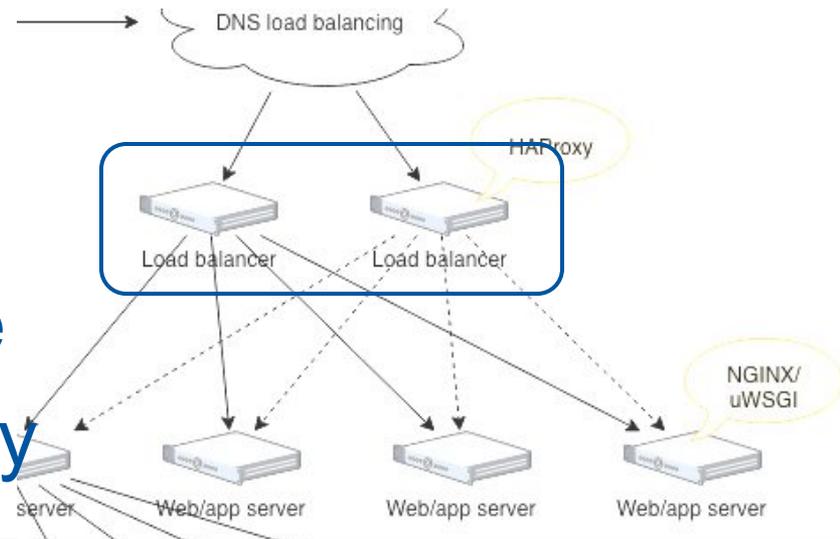
Infrastructure overview



- *Load balancers:* HAProxy, Nginx or others.
- *Web servers:* Nginx, Apache or others.
- *Application servers:* uWSGI, Gunicorn or mod_wsgi.
- *Distributed task queue:* Celery
- *Database:* PostgreSQL, MySQL or SQLite.
- *Search engine:* Elasticsearch (v5 and v6).
- *Message queue:* RabbitMQ, Redis or Amazon SQS.
- *Cache system:* Redis or Memcache.
- *Storage system:* Local, S3, XRootD, WebDAV and more.

Load balancer: HAProxy

- Traffic distribution depending on its “type”
 - /static, /<path>
- URL Rewriting
- SSL Termination
- Downtime static website
- More than 1: redundancy



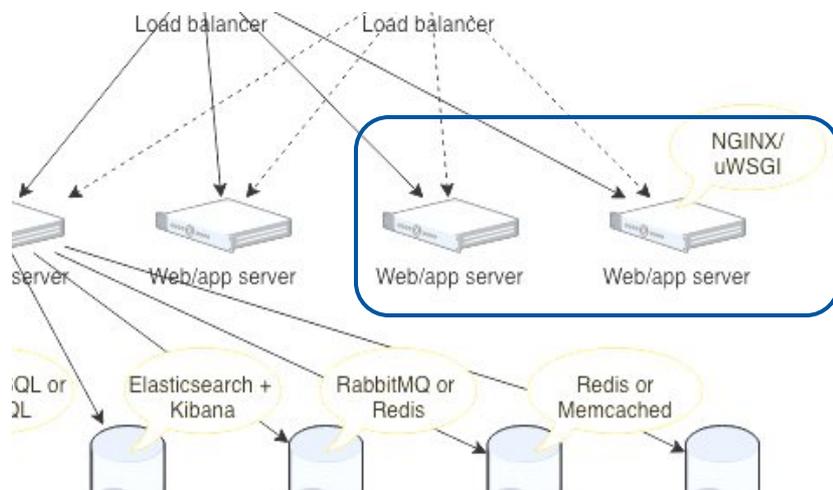
HAProxy configuration

- Frontend
- Backend
- Headers: X-Forwarded-For

[Go to documentation](#)

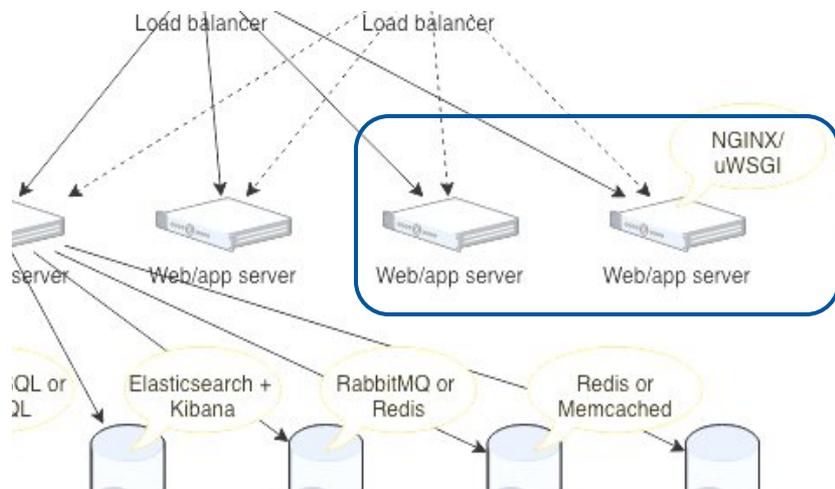
Web app: uWSGI (behind Nginx)

- PEP 333 (TL;DR)
- Invenio UI and REST
- processes and threads



Web app: Nginx

- ngx http uwsgi module: translate HTTP -> WSGI
- Reverse proxy to distribute traffic
 - ui, api, static, files
- protocols and ciphers
- logging
- headers: X-Request-ID



A row of six vintage metal wrenches hanging on a wooden rack. The wrenches are arranged horizontally, with their heads pointing to the left. Each wrench has a different size and design, with some having blue-painted handles. The background is a wooden surface with a visible grain.

Dimensions

Processes and Threads

How many servers do we need? How do I configure them?

Let's start with the database

- Depends on DB: Postgresql? MySQL? Servers?
- General: configure max connections
- DB specific: check documentation

Example: 200 max connections

```
$ less /var/lib/pgsql/data/postgresql.conf
max_connections = 200

$ mysql
> SET GLOBAL max_connections = 200;
```

Web nodes: uWSGI

- From the [uWSGI documentation](#) ([interesting discussion here](#))

*“There is no magic rule for setting the number of processes or threads to use. It is very much application and system dependent. Simple math like $\text{processes} = 2 * \text{cpucore}$ s will not be enough. You need to experiment with various setups and be prepared to constantly monitor your apps. `uwsgi`top could be a great tool to find the best values.”*

- **Processes:**
 - 1 process -> 1 Invenio app
 - start with $2 * \text{CPU Cores}$, and then measure and adjust
- **Threads: low, it really depends** (take into account GIL, multithreading, etc...)

Example: 8 processes, 2 threads

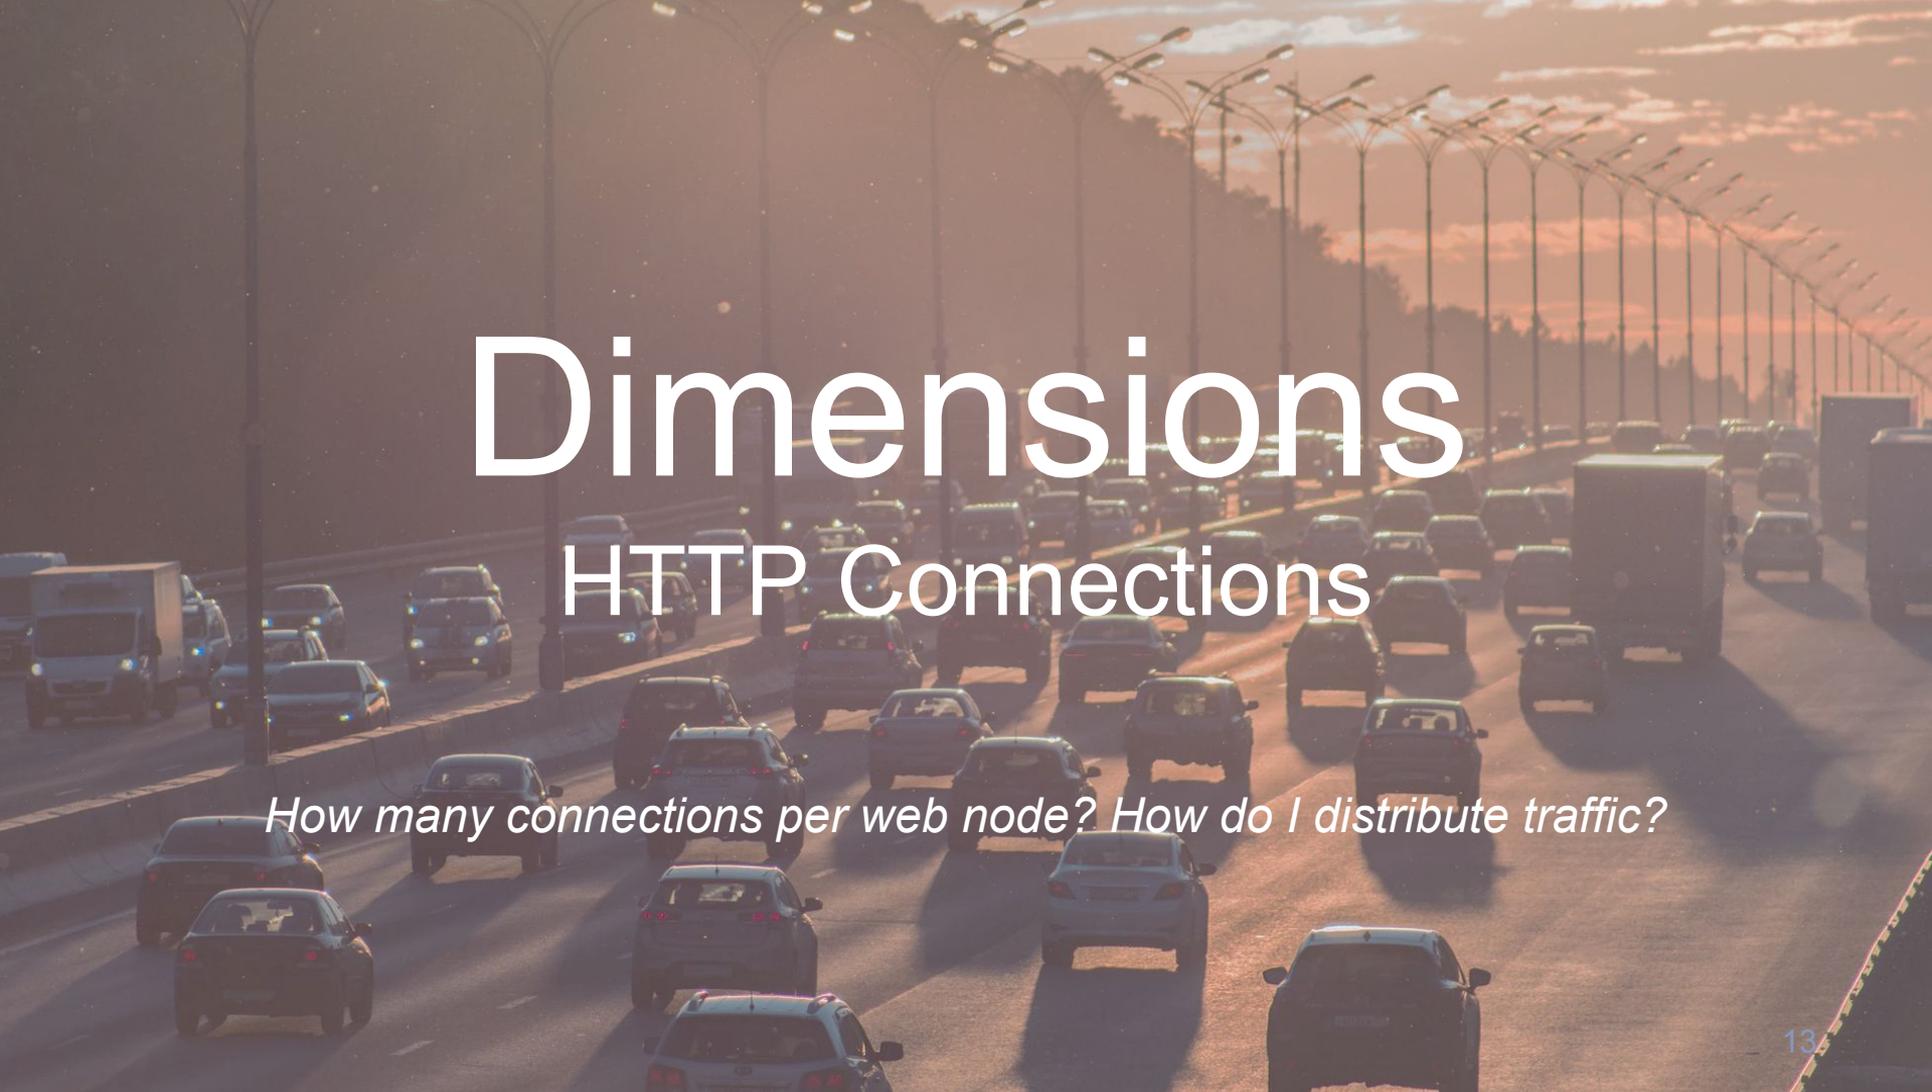
Max processes per web node

- SQLAlchemy: pool of db connections `pool_size` (default 5) per process
- Each uWSGI process -> 5 db connections
- 8 process per web node -> $5 * 8 = 40$ connections
- **200** max connections on db / **40** per web node = ???

5 web nodes max

BUT

- You have to count all Invenio apps (your workers, celery apps)!

A wide-angle, high-angle shot of a multi-lane highway during sunset. The sky is a mix of orange, yellow, and blue. The road is filled with cars and trucks, moving away from the viewer. The perspective is from an elevated position, looking down at the traffic. The text is overlaid on the center of the image.

Dimensions

HTTP Connections

How many connections per web node? How do I distribute traffic?

A collection of vintage drafting tools is arranged on a wooden workbench. In the foreground, a pair of large, dark metal shears with a circular hole in the handle lies on the left. To its right, a pair of dividers and a pair of compasses are positioned. A long, thin metal ruler with numerical markings is placed vertically on the right side. A large set square is also visible. In the background, a detailed topographical map with contour lines and text is spread across the table. The word "Measure!" is overlaid in the center in a large, white, sans-serif font.

Measure!

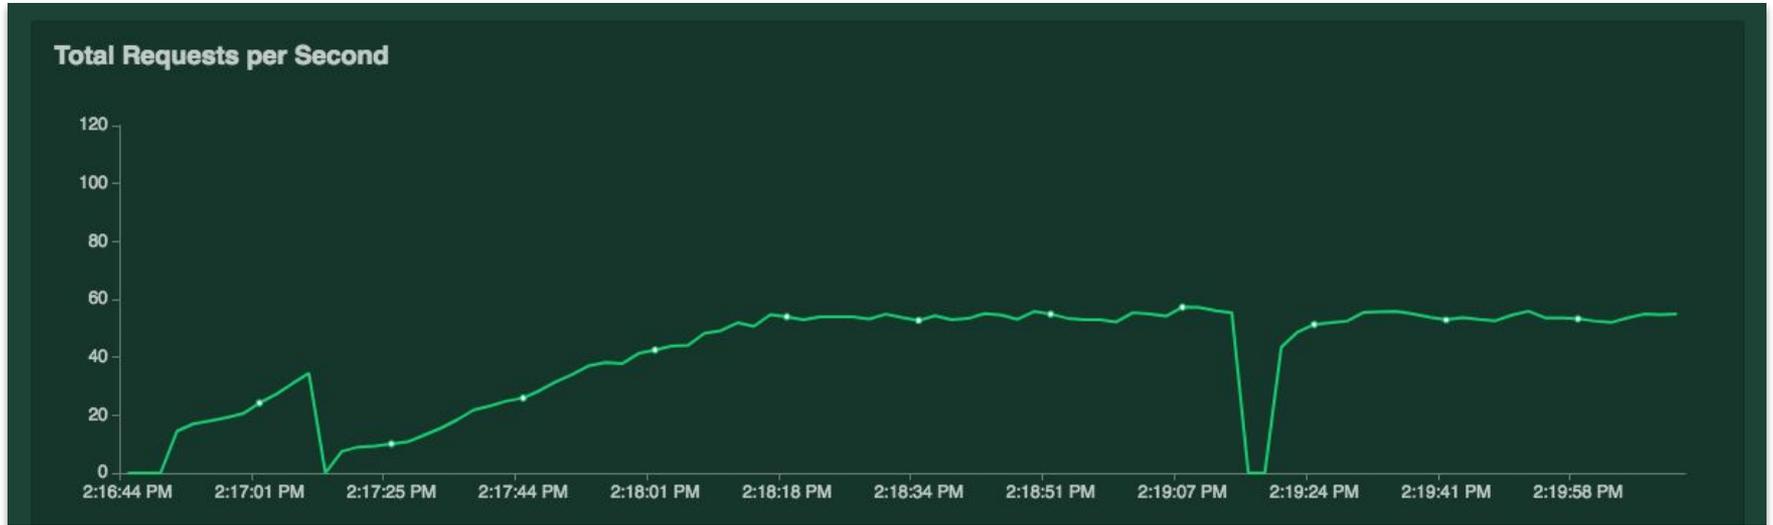
uWSGI: Start small

1. Predict/Estimate the n. of request/sec you need
2. Take one web node
3. Configure number of uWSGI processes
4. Benchmark by increasing number of requests

```
locust --host=https://mydomain.org/ --no-web -c 1000 -r 100
```

uWSGI: Iterate

Keep tuning processes/memory until you reach your target



We found out how many
connections per web nodes.

And now?

HAProxy: distribute load

Since we now know how many connections per node, we can distribute the load

```
backend ssl_invenio_app
...
server web-backup 127.0.0.1:444 backup check check-ssl fall 2 inter 5000 maxconn 255 rise 1 ssl verify none weight 2
server web1 frontend:443 check check-ssl fall 2 inter 5000 maxconn 40 rise 1 ssl verify none weight 2
server web2 frontend:443 check check-ssl fall 2 inter 5000 maxconn 40 rise 1 ssl verify none weight 2

backend ssl_invenio_files
...
server web5 frontend:443 check check-ssl fall 2 inter 5000 maxconn 100 rise 1 ssl verify none weight 2
server web6 frontend:443 check check-ssl fall 2 inter 5000 maxconn 100 rise 1 ssl verify none weight 2

backend ssl_invenio_static
...
server web3 frontend:443 check check-ssl fall 2 inter 5000 maxconn 255 rise 1 ssl verify none weight 2
server web4 frontend:443 check check-ssl fall 2 inter 5000 maxconn 255 rise 1 ssl verify none weight 2
```

Recap

1. Max db connections
2. Processes/threads per web node
3. Max connections per web node
4. Load balancer config

A utility pole stands against a clear blue sky, equipped with several monitoring devices. At the top is a large, white, dome-shaped sensor. Below it, a horizontal rectangular device is mounted. Further down, a cluster of three smaller, white, rectangular meters is attached to the pole. Each of these meters has a small blue label with the word 'meter' printed on it. The pole is a dark grey or black color.

Monitoring

Health of your production infrastructure

Load balancer: HAProxy

<http://127.0.0.1:8080>

HAProxy

Statistics Report for pid 20749

> General process information

pid = 20749 (process #1, nbproc = 1)
 uptime = 93d 18h25m07s
 system limits: memmax = unlimited; ulimit-n = 8057
 maxsock = 8057; maxconn = 4000; maxpipes = 0
 current conns = 3; current pipes = 0; conn rate = 1/sec
 Running tasks: 1/38; idle = 100 %

■ active UP
■ active UP, going down
■ active DOWN, going up
■ active or backup DOWN
■ active or backup DOWN for maintenance (MAINT)
■ active or backup SOFT STOPPED for maintenance
■ backup UP
■ backup UP, going down
■ backup DOWN, going up
■ not checked
 Note: "NO LB"/"DRAIN" = UP with load-balancing disabled.

		Queue			Session rate			Sessions						Bytes			Denied	
		Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	
in-cds_cdslabs_http		0	68	-	0	67	3 000	16 535						14 945 003	20 964 447	0		

		Queue			Session rate			Sessions						Bytes			Denied	
		Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	
in-cds_cdslabs_https		0	43	-	2	47	3 000	801 526						330 327 397 112	2 265 959 457 193			

		Queue			Session rate			Sessions						Bytes			Denied	
		Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	
in-cds_cdslabs_elists		1	2	-	1	2	3 000	37						108 078	11 547 457	0		

		Queue			Session rate			Sessions						Bytes			Denied	
		Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	
cda_cdslabs_app		0	0	-	0	0	255	0	0	40m8s	12 904	12 904	40m8s	3 738 807	5 206 625	0		
videos-cds-lb2.cern.ch-backup		0	0	-	0	16	40	12 904	12 904	19m52s	12 908	12 908	19m52s	3 680 512	5 211 161			
videos-cds-web1.cern.ch		0	0	-	0	19	40	12 903	12 903	39m44s	12 903	12 903	39m44s	3 768 591	5 205 159			
videos-cds-web2.cern.ch		0	0	-	0	16	40	12 904	12 904	9m5s	12 904	12 904	9m5s	3 658 629	5 203 619			
videos-cds-web3.cern.ch		0	0	-	0	17	40	51 619	51 619	9m5s	51 619	51 619	9m5s	14 846 539	20 826 564			
videos-cds-web4.cern.ch		0	0	-	0	68	300											
Backend		0	0	-	0	8	300											

		Queue			Session rate			Sessions						Bytes			Denied	
		Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	
cda_cdslabs_files		0	0	-	0	1	256	8	8	25d10h	8	8	25d10h	1 199	3 486			
videos-cds-web5.cern.ch		0	0	-	0	1	256	7	7	28d20h	7	7	28d20h	897	3 052			
videos-cds-web6.cern.ch		0	0	-	0	2	300	15	15	25d10h	15	15	25d10h	2 096	6 538	0	0	
Backend		0	0	-	0	2	300											



Web app: uWSGI

- uWSGI Stats Server
- uwsgitop

```
uwsgi-2.0.14 - Wed Mar 22 12:17:30 2017 - req: 1195253 - RPS: 1 - lq: 0 - tx: 616.0G  
node: 127.0.0.1 - cwd: /opt/zenodo - uid: 1001 - gid: 1001 - masterpid: 28932
```

WID	%	PID	REQ	RPS	EXC	SIG	STATUS	AVG	RSS	VSZ	TX	ReSpwn	HC	RunT	LastSpwn
7	10.1	28968	120274	0	2528	0	idle	776ms	0	0	52.0G	1	0	1278m	11:45:37
10	10.1	28971	120126	0	2559	0	idle	1ms	0	0	57.0G	1	0	1394m	11:45:37
2	10.0	28946	120067	0	2564	0	idle	1ms	0	0	54.0G	1	0	1447m	11:45:37
1	10.0	28942	119980	1	2526	0	idle	1ms	0	0	57.0G	1	0	1546m	11:45:37
8	10.0	14707	119859	0	2548	0	idle	2ms	0	0	49.0G	2	0	1552m	10:13:16
5	10.0	28958	119785	0	2616	0	idle	846ms	0	0	89.0G	1	0	1655m	11:45:37
3	10.0	28952	119741	0	2500	0	busy	5ms	0	0	61.0G	1	0	1589m	11:45:37
4	10.0	8635	119062	0	2643	0	idle	1ms	0	0	67.0G	2	0	2230m	12:07:51
9	9.9	28970	118862	0	2623	0	busy	1ms	0	0	63.0G	1	0	2199m	11:45:37
6	9.8	28967	117497	0	2505	0	busy	12ms	0	0	62.0G	1	0	3222m	11:45:37

ElasticSearch: Kibana

<http://127.0.0.1:5601>

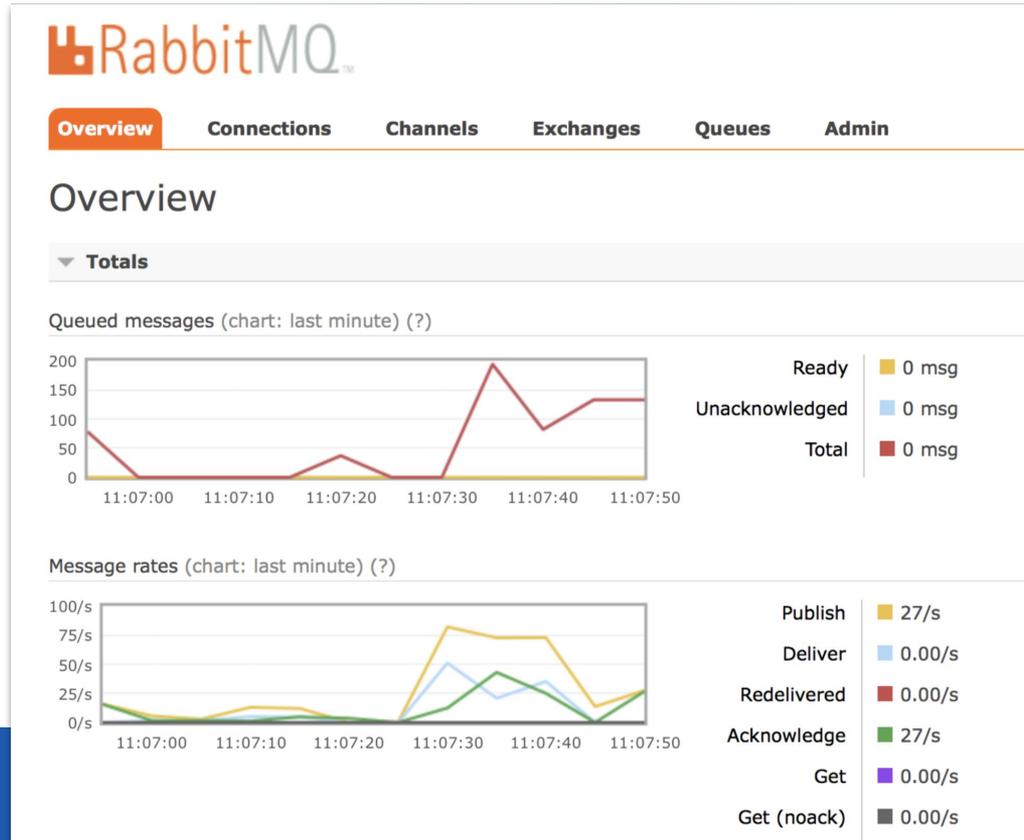
Elasticsearch index status

Sort by
name

Index name	health	# replicas	# shards	# docs	size (including replicas)				
<input type="text" value="index name ..."/>	<input type="text" value="health..."/>	<input type="text" value="replicas ..."/>	<input type="text" value="shards ..."/>	<input type="text" value="docs ..."/>	<input type="text" value="size ..."/>				
.kibana_cdsbooksqa	green	1	1	1	6.8kb				
cdsbooks-documents-document-v1.0.0	green	1	5	19	259.3kb				
cdsbooks-internal_locations-internal_location-v1.0.0	green	1	5	8	64.4kb				
cdsbooks-items-item-v1.0.0	green	1	5	49	489.6kb				
cdsbooks-loans-loan-v1.0.0	green	1	5	99	337.4kb				
cdsbooks-locations-location-v1.0.0	green	1	5	1	12.6kb				
Selected: 6		Selected: 6		Selected: 26		Selected: 177		Selected: 1.1mb	
Total: 6		Total: 6		Total: 26		Total: 177		Total: 1.1mb	

Job queue

<http://127.0.0.1:15672> - guest:guest



Job queue

<http://127.0.0.1:15672>

Celery Flower **Dashboard** Tasks Broker Monitor Docs Code

Active: 0 Processed: 19781 Failed: 759 Succeeded: 18687 Retried: 361

Shut Down

	Worker Name	Status	Active	Processed	Failed	Succeeded	Retried	Load Average
<input type="checkbox"/>	celery@zenodo-task7.cern.ch	Online	0	4929	202	4633	93	0.04, 0.03, 0.05
<input type="checkbox"/>	celery@zenodo-task5.cern.ch	Online	0	4954	255	4607	101	0.08, 0.08, 0.05
<input type="checkbox"/>	celery@zenodo-task6.cern.ch	Online	0	4953	133	4759	80	0.0, 0.01, 0.05
<input type="checkbox"/>	celery@zenodo-task4.cern.ch	Online	0	4945	169	4688	87	0.0, 0.01, 0.05

Monitoring



Alerting: UpTimeRobot

The screenshot displays the UpTimeRobot dashboard. At the top left is the UpTimeRobot logo. On the top right, there are links for 'Upgrade', 'Dashboard', and 'My Settings'. Below the logo is a dark navigation bar with a '+ Add New Monitor' button and a search bar. Underneath, there are options for '(Bulk Actions)', '(Export Monitors - Shorten Monitor Names)', 'Sort Monitors', and a time filter set to 'Last 24 Hours'. The main content area shows a list of monitors with their status percentages and progress bars. To the right, there is an 'Account Dashboard' section with a link to 'details about the account' and a note that it's a place to find details about monitors. Below that is a 'Quick Stats' panel showing 'UP MONITORS' with a large green circle containing the number '8' and a link to 'show "up" monitors'.

Uptime	Protocol	Monitor Name	Status
99.67%	http	CDS Videos - Streaming	OK
100%	http	CDS Videos - Search	OK
100%	http	CDS Videos - Record	OK
100%	http	CDS Videos - EOS	OK
94.76%	http	CDS Legacy - Search	OK
100%	http	CDS Legacy - MediaArchive	OK
97.87%	http	CDS Legacy - Home	OK
82.84%	http	CDS Legacy [SSL] - Home	OK



Logging

Understand what happened

Request tracing

How do we find the traces of a specific request in each server?

How can we get the request IP?

Load balancer: HAProxy

```
backend ssl_app
  balance leastconn
  http-check disable-on-404
  option http-server-close
  option forwardfor except 127.0.0.0/8
  option httpchk HEAD /ping HTTP/1.0
  server web1 frontend:443 check check-ssl fall 2 inter 20000 maxconn 30 rise 1 ssl verify none weight 2

backend ssl_static
  balance leastconn
  http-check disable-on-404
  option http-server-close
  option forwardfor except 127.0.0.0/8
  option httpchk HEAD /ping HTTP/1.0
  server web1 frontend:443 check check-ssl fall 2 inter 5000 maxconn 255 rise 1 ssl verify none weight 2
```

Web app: Nginx

Nginx log config

Infrastructure IPs

```
172.22.0.12 my-site_lb_1
172.22.0.11 my-site_frontend_1
172.22.0.10 my-site_web-ui_1
172.22.0.7 my-site_worker_1
172.22.0.9 my-site_web-api_1
172.22.0.6 my-site_flower_1
172.22.0.8 my-site_kibana_1
172.22.0.4 my-site_mq_1
172.22.0.5 my-site_db_1
172.22.0.3 my-site_cache_1
172.22.0.2 my-site_es_1
```

```
log_format trace '$remote_addr - [$time_local] "$request" '
                '$status $body_bytes_sent "$http_referer"
                "$http_user_agent" '
                "$http_x_forwarded_for" $request_id '
                '$msec $request_time '
                '$upstream_http_x_session_id $upstream_http
```

Log example

```
172.22.0.12 - [15/Mar/2019:14:56:27 +0000] "GET /api/records
200 209 "https://127.0.0.1/search?q="
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10.14; rv:65.0) Gecko
"172.22.0.1" 00fdb97e91fe18238733045e9c23a331
1552661787.662 2.242
84e4ee86efbf0e53_5c8bb8cc -
```

Web app: Nginx

Nginx passes
params to the
uWSGI server

```
location /api/files {
    gzip off;
    uwsgi_pass api_server;
    include uwsgi_params;
    uwsgi_buffering off;
    uwsgi_request_buffering off;
    uwsgi_param Host $host;
    uwsgi_param X-Forwarded-For $proxy_add_x_forwarded_for;
    uwsgi_param X-Forwarded-Proto $scheme;
    # Pass request id to api server
    uwsgi_param X-Request-ID $request_id;
    # X-Session-ID / X-User-ID is read by nginx and included in the logs,
    # however we don't want to expose them to clients so we are hiding them.
    uwsgi_hide_header X-Session-ID;
    uwsgi_hide_header X-User-ID;
    # Max upload size for files is set to 50GB (configure as needed).
    client_max_body_size 50G;
}
```

Web app: uWSGI

- invenio-app can read the header X-Forwarded-For and the request_id
- It will contain the list of all IPs (HAProxy, Nginx, etc...)

Error reporting

- invenio-logging
- Sentry

Invenio Logging

- Log to file
- Configure logging level
- Configure file rotation
- Configure Sentry integration

Sentry

Home | **cds / videos** | Issues | Overview | User Feedback | Releases

Star Project | Project Settings

Overview

1 hour | 1 day | 1 week

March 15, 2019

TRENDING ISSUES	EVENTS	USERS
KeyError cds.modules.records.serializers.v... 'duration' ⌚ an hour ago — a year old	16k	2.3k
GET ⌚ 2 hours ago — a year old	10.3k	0
TypeError /opt/cds/src/cds/cds/modules/r... sequence item 4: expected string or Unicode, ... ⌚ 3 hours ago — 9 months old	316	8
AssertionError invenio_files_rest.storage... ⌚ 12 hours ago — a month old	1.9k	0
RequestError /records/ TransportError(400, u'search_phase_executio... ⌚ 20 hours ago — 2 months old	256	22

NEW ISSUES	EVENTS	USERS
PIDDeletedError cds.modules.deposit.ap... ⌚ 8 days ago — 8 days old	1	0
PIDDeletedError /hooks/receivers/<stri... ⌚ 8 days ago — 8 days old	52	1
KeyError cds.modules.webhooks.receivers... '_tasks' ⌚ 9 days ago — 9 days old	1	0
KeyError /hooks/receivers/<string:receive... '_tasks' ⌚ 9 days ago — 9 days old	5	2
RequestError cds.modules.stats.views in g... TransportError(400, u'<html>\r\n<head><tit... ⌚ 9 days ago — 10 days old	154	1



Ready to GO LIVE?



Questions?