

DOCK

George Ioannidis, gi@in-two.com and Alexandru Stan, as@in-two.com

GOALS


- Understand better how a hybrid deployment (Docker containers inside VMs) performs for MyMeedia
- H1. There are no significant performance drawbacks for a hybrid deployment of the MyMeedia application.
- H2. Using the hybrid approach it is possible to optimise the resource utilisation and scalability of MyMeedia deployments (under conditions of increased load and concurrent requests).

CHALLENGES

- MyMeedia is a cloud service for collecting, organising, discovering and presenting multimedia stream and feed
- Heavy duty media processing (videos, audio, photos, documents)
- Ease of deployment as a hosted service across operating systems



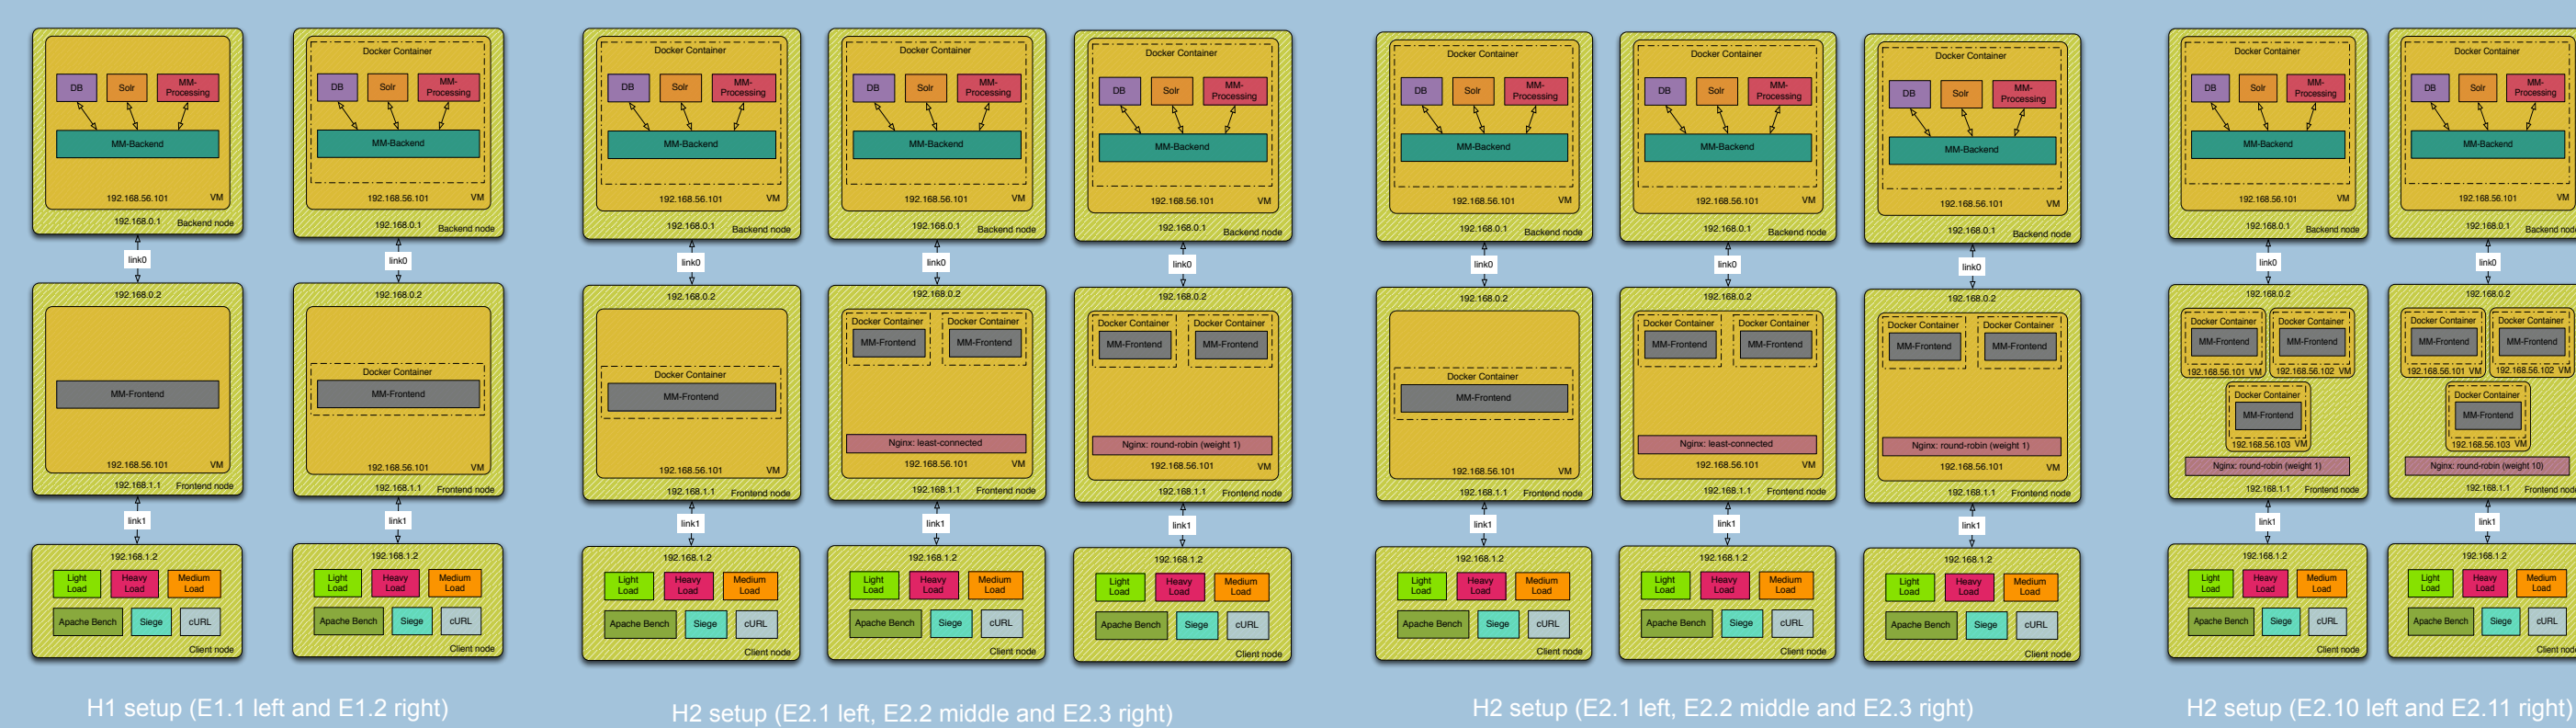
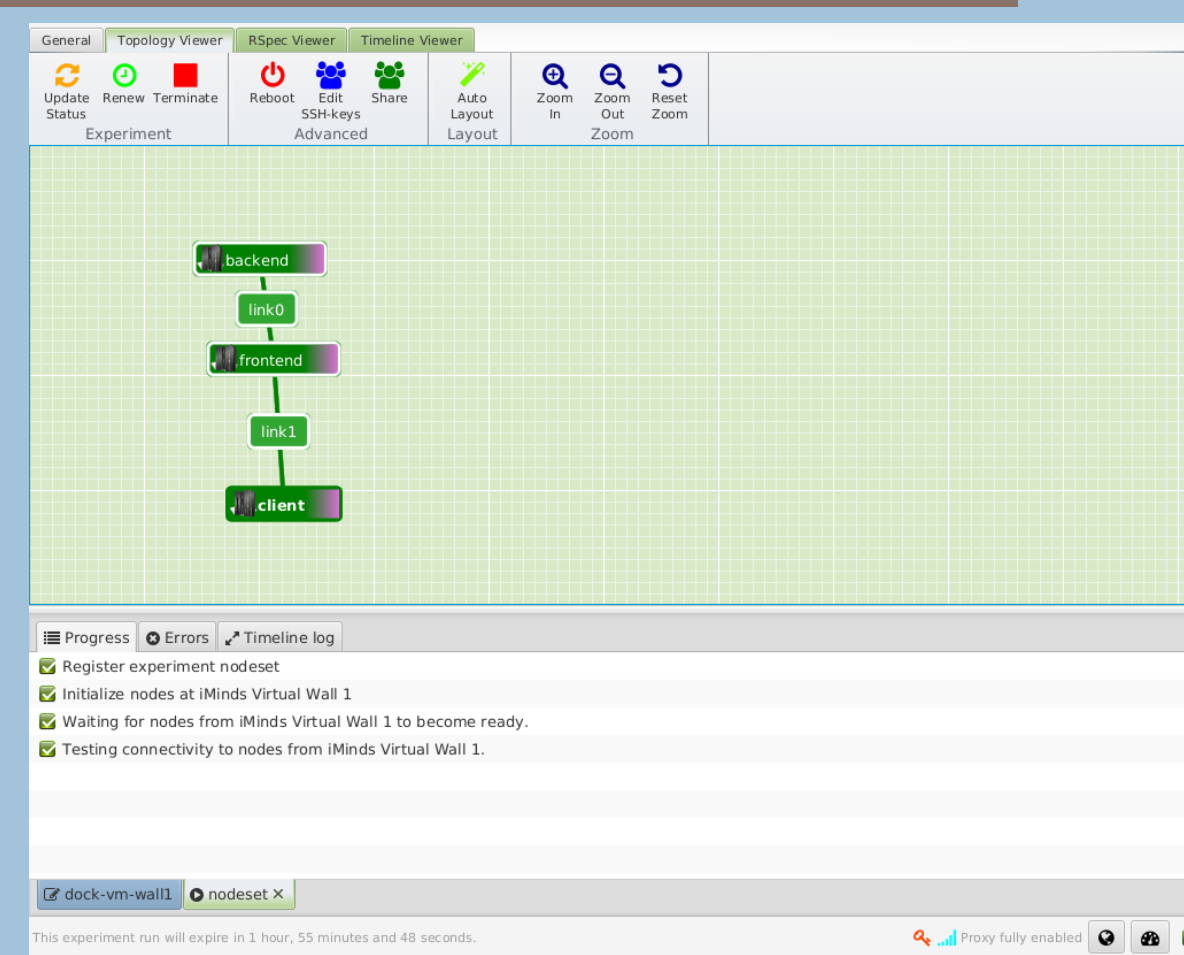
Software people since 2005
<https://in-two.com>



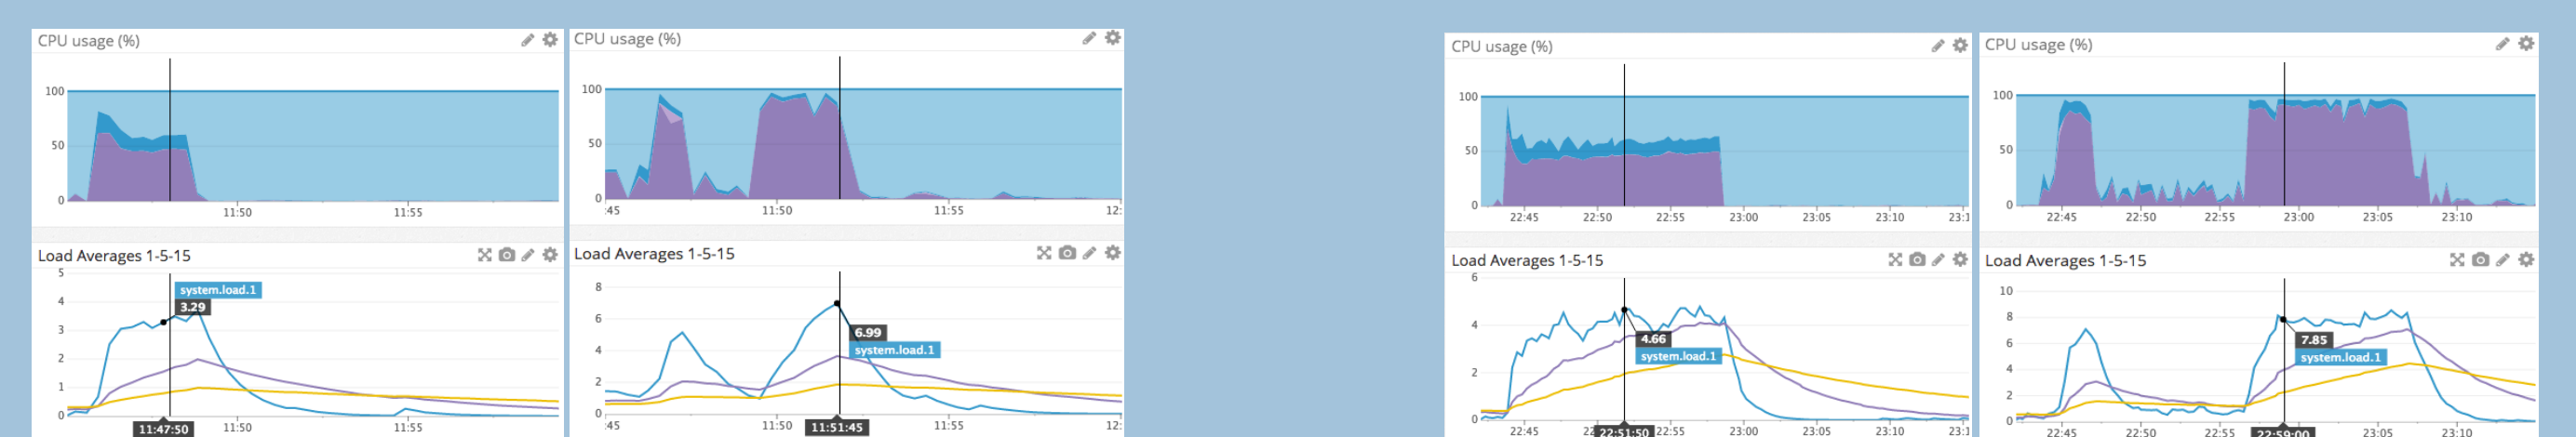
<https://mymedia.com>

DEMO SETUP

- Backend (database, indexing, MM-Processing, MM-Backend)
- Frontend (1..3 MM-Frontend)
- Client (perform stress tests and simulate load)



RESULTS




Light Load (top) [E1.2 (L)]
AB: 1000 reqs / concurrency 10
Siege: 100 reqs / repeated 10x
Upload and processing of 25 photos, 8 audio and 8 video files, 8 .doc and 8 .pdf files

Frontend Node (left)
1 VM (8GB RAM, 4 CPU cores)
1 MM-Frontend; inside Docker

Backend Node (right)
1 VM (8GB RAM, 4 CPU cores)
DB, Solr, MM-Processing, MM-Backend; inside Docker

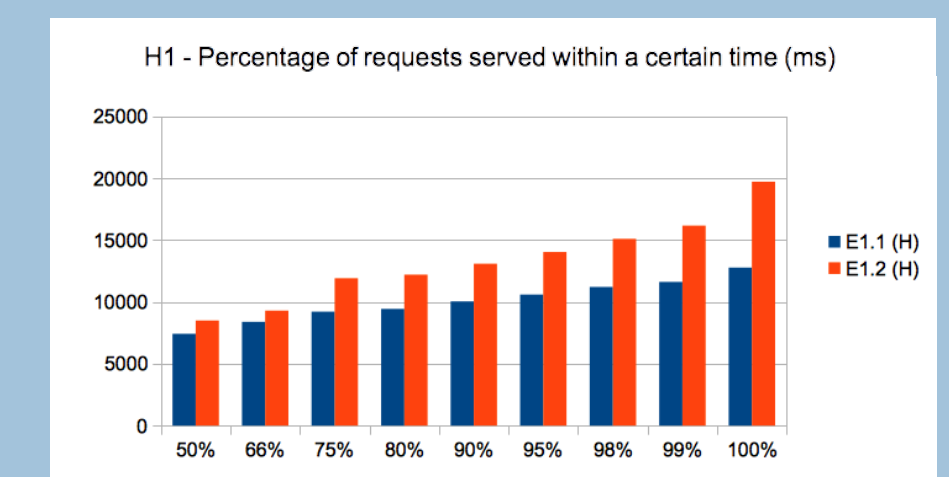
Heavy Load (bottom) [E1.2 (H)]
AB: 4000 reqs / concurrency 40
Siege: 100 reqs / repeated 40x
Upload and processing of 100 photos, 32 audio and 32 video files, 32 .doc and 32 .pdf files



Light Load (top)
50% of reqs are served within 2000 ms

For 66% to 90% of reqs no considerable lag in serving them

VM only setup performs generally better, but even more for the last 5% of reqs



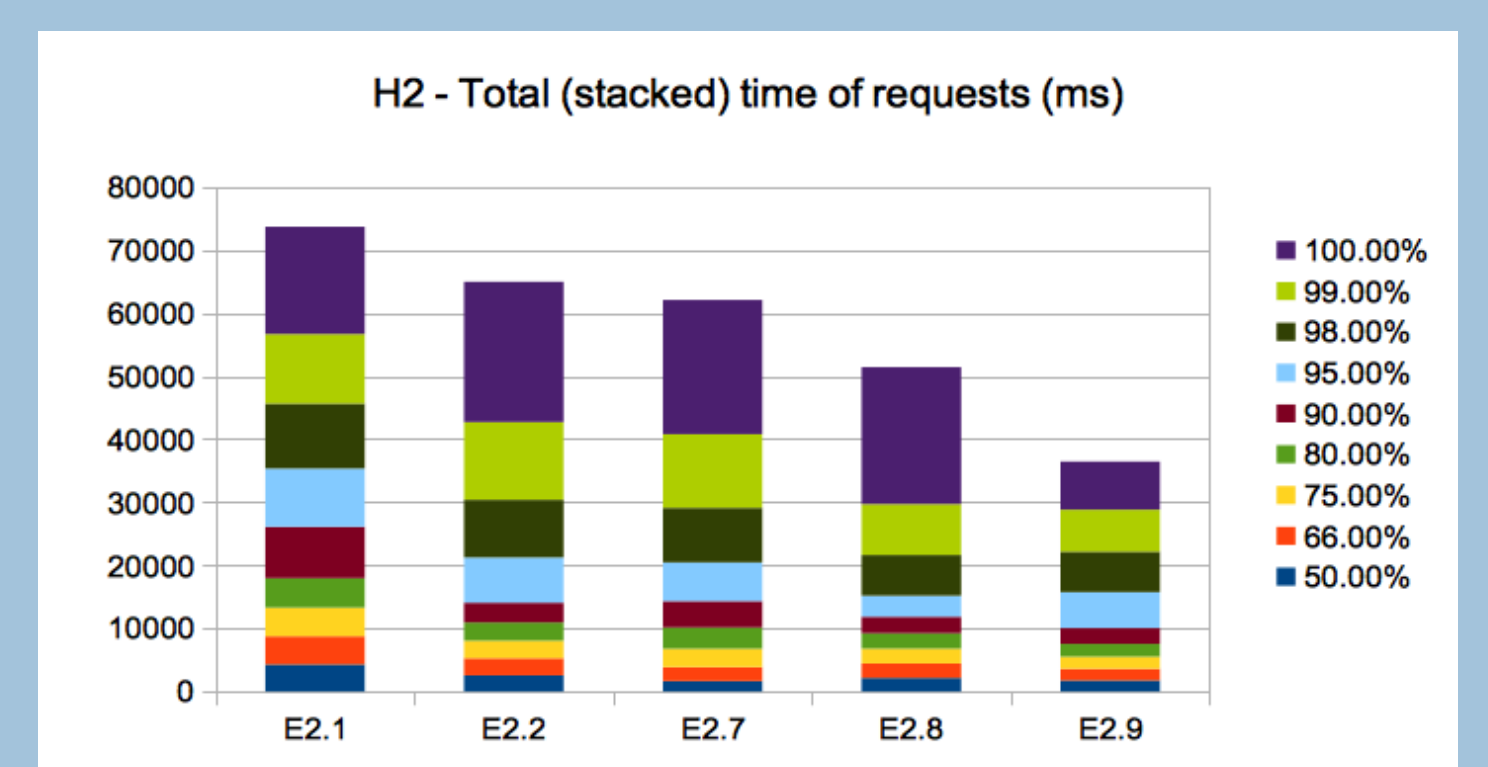

Heavy Load (bottom)
50% of reqs are served within 8000ms

For 66% of reqs no considerable lag in serving them

VM only setup performs generally better, but even more for the last 25% of reqs

MORE RESULTS

- For more than 3 hybrid VMs it is better to use a least-connected load-balancing mode when enough computing resources are available.
- In the case computing resources are scarce, a round-robin load-balancing mode performs better.
- The amount of computing power per VM can be decreased for a hybrid deployment of the MM-Frontend (E2.9 setup was 3 VMs each with 3GB RAM and 2 CPU cores, while E2.7 was 1 VM with 9GB RAM and 6 CPU cores)
- It is much more efficient for the MM-Frontend to scale horizontally (adding more computing nodes) instead of scaling vertically (adding more computing power).



CONCLUSIONS

- Considerable reduction of 50% in response time
- Improved User Experience
- Improved search engine visibility
- Applied lessons learnt in production at MyMeedia
- Increased confidence in experimenting with set-up topologies and in deploying applications



POST MORTEM

“Thanks to the experiment we conducted within Fed4FIRE we determined how Docker containers can be best used in conjunction with Virtual Machines, which ultimately helped us double the performance of our MyMeedia (<https://mymedia.com>) service for media management and publishing.