

CLOUD-EDGE COMPUTING

The background features a series of white, wavy, parallel lines that create a sense of motion and depth. On the right side, there is a network of blue nodes of varying sizes connected by thin white lines, suggesting a data network or cloud infrastructure.

UNIVERSITY
OF TWENTE.

DIGITAL SOCIETY
INSTITUTE




WORDPRESS

www.distributed-systems.net

UNIVERSITY
OF TWENTE.

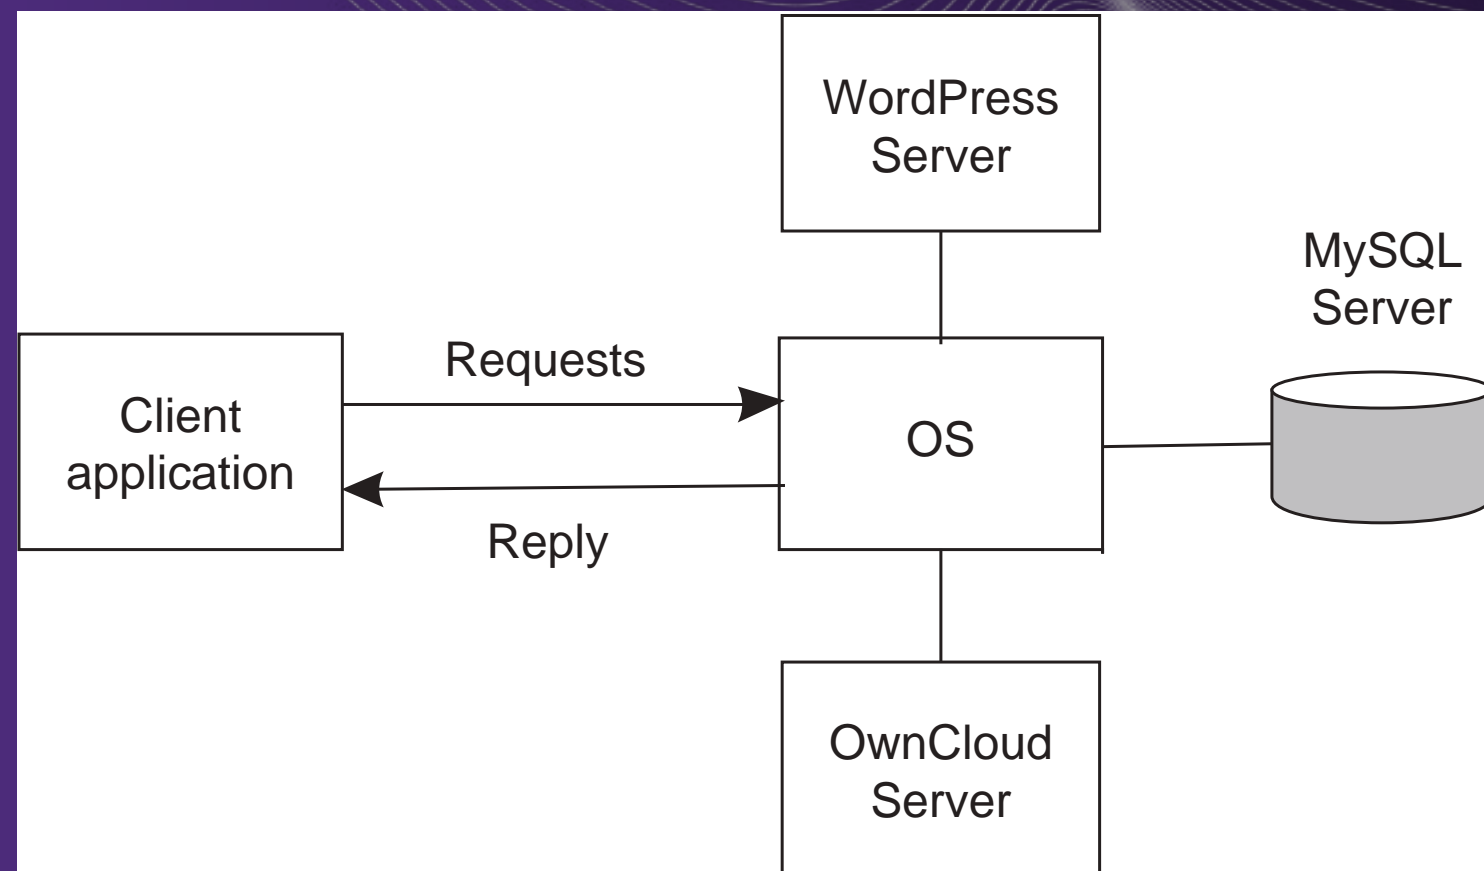
DIGITAL SOCIETY
INSTITUTE



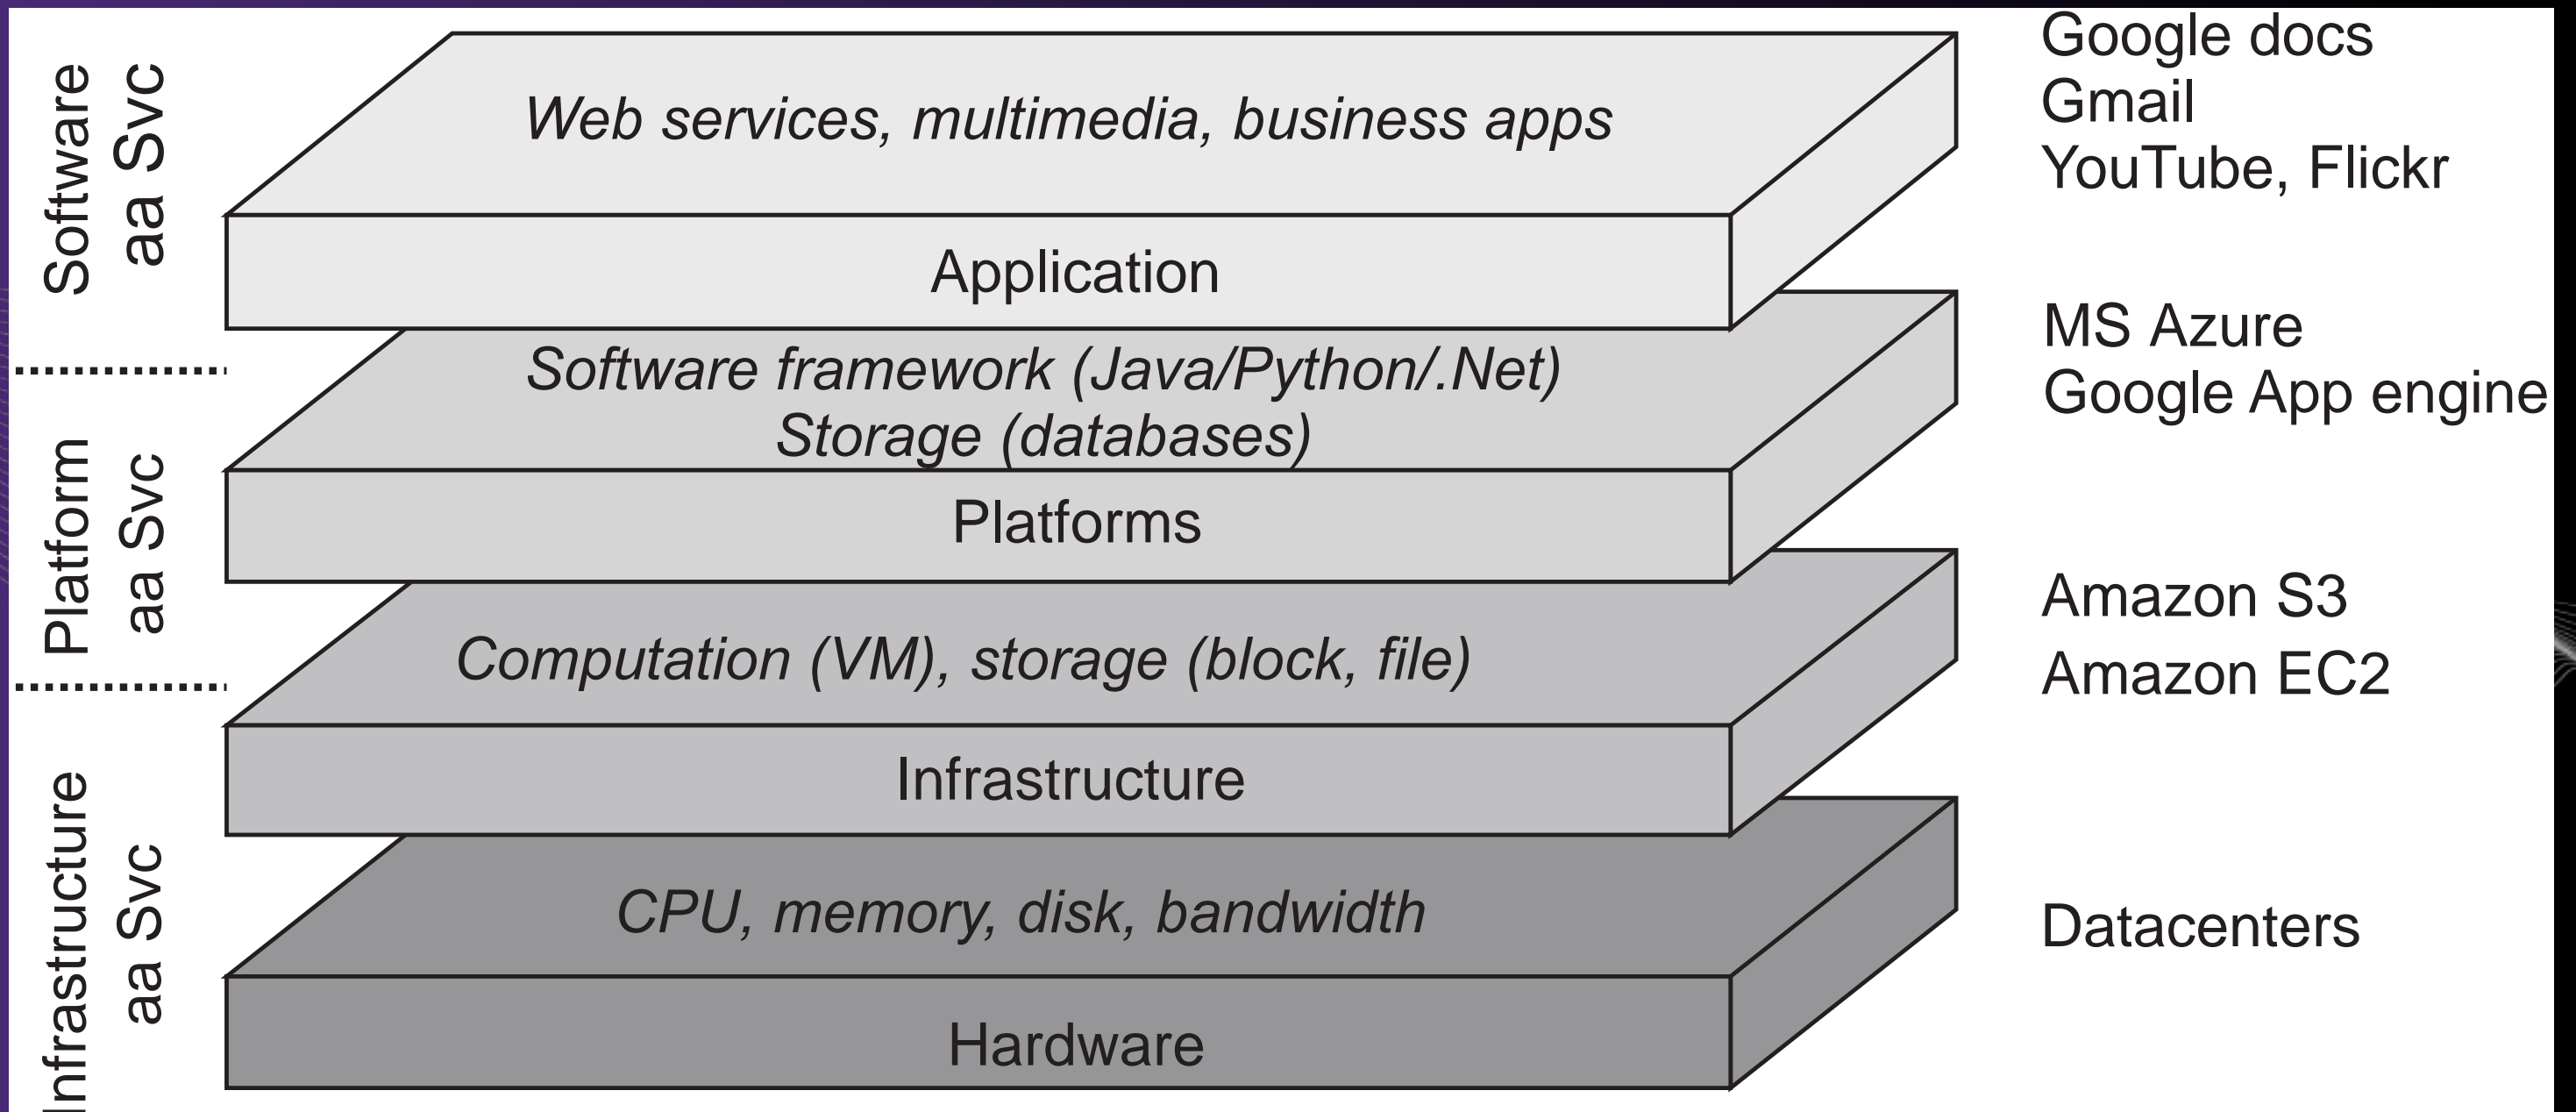
owncloud.van-steen.net

A VERY SIMPLE SETUP


- A single physical machine with standard UNIX (Ubuntu 18.04)
- Servers operate as separate processes (mysqld) or collection of scripts executed by common daemon (apache).
- Web-based service is selected by Apache's virtualhost, which spawns a new process to run associated scripts.





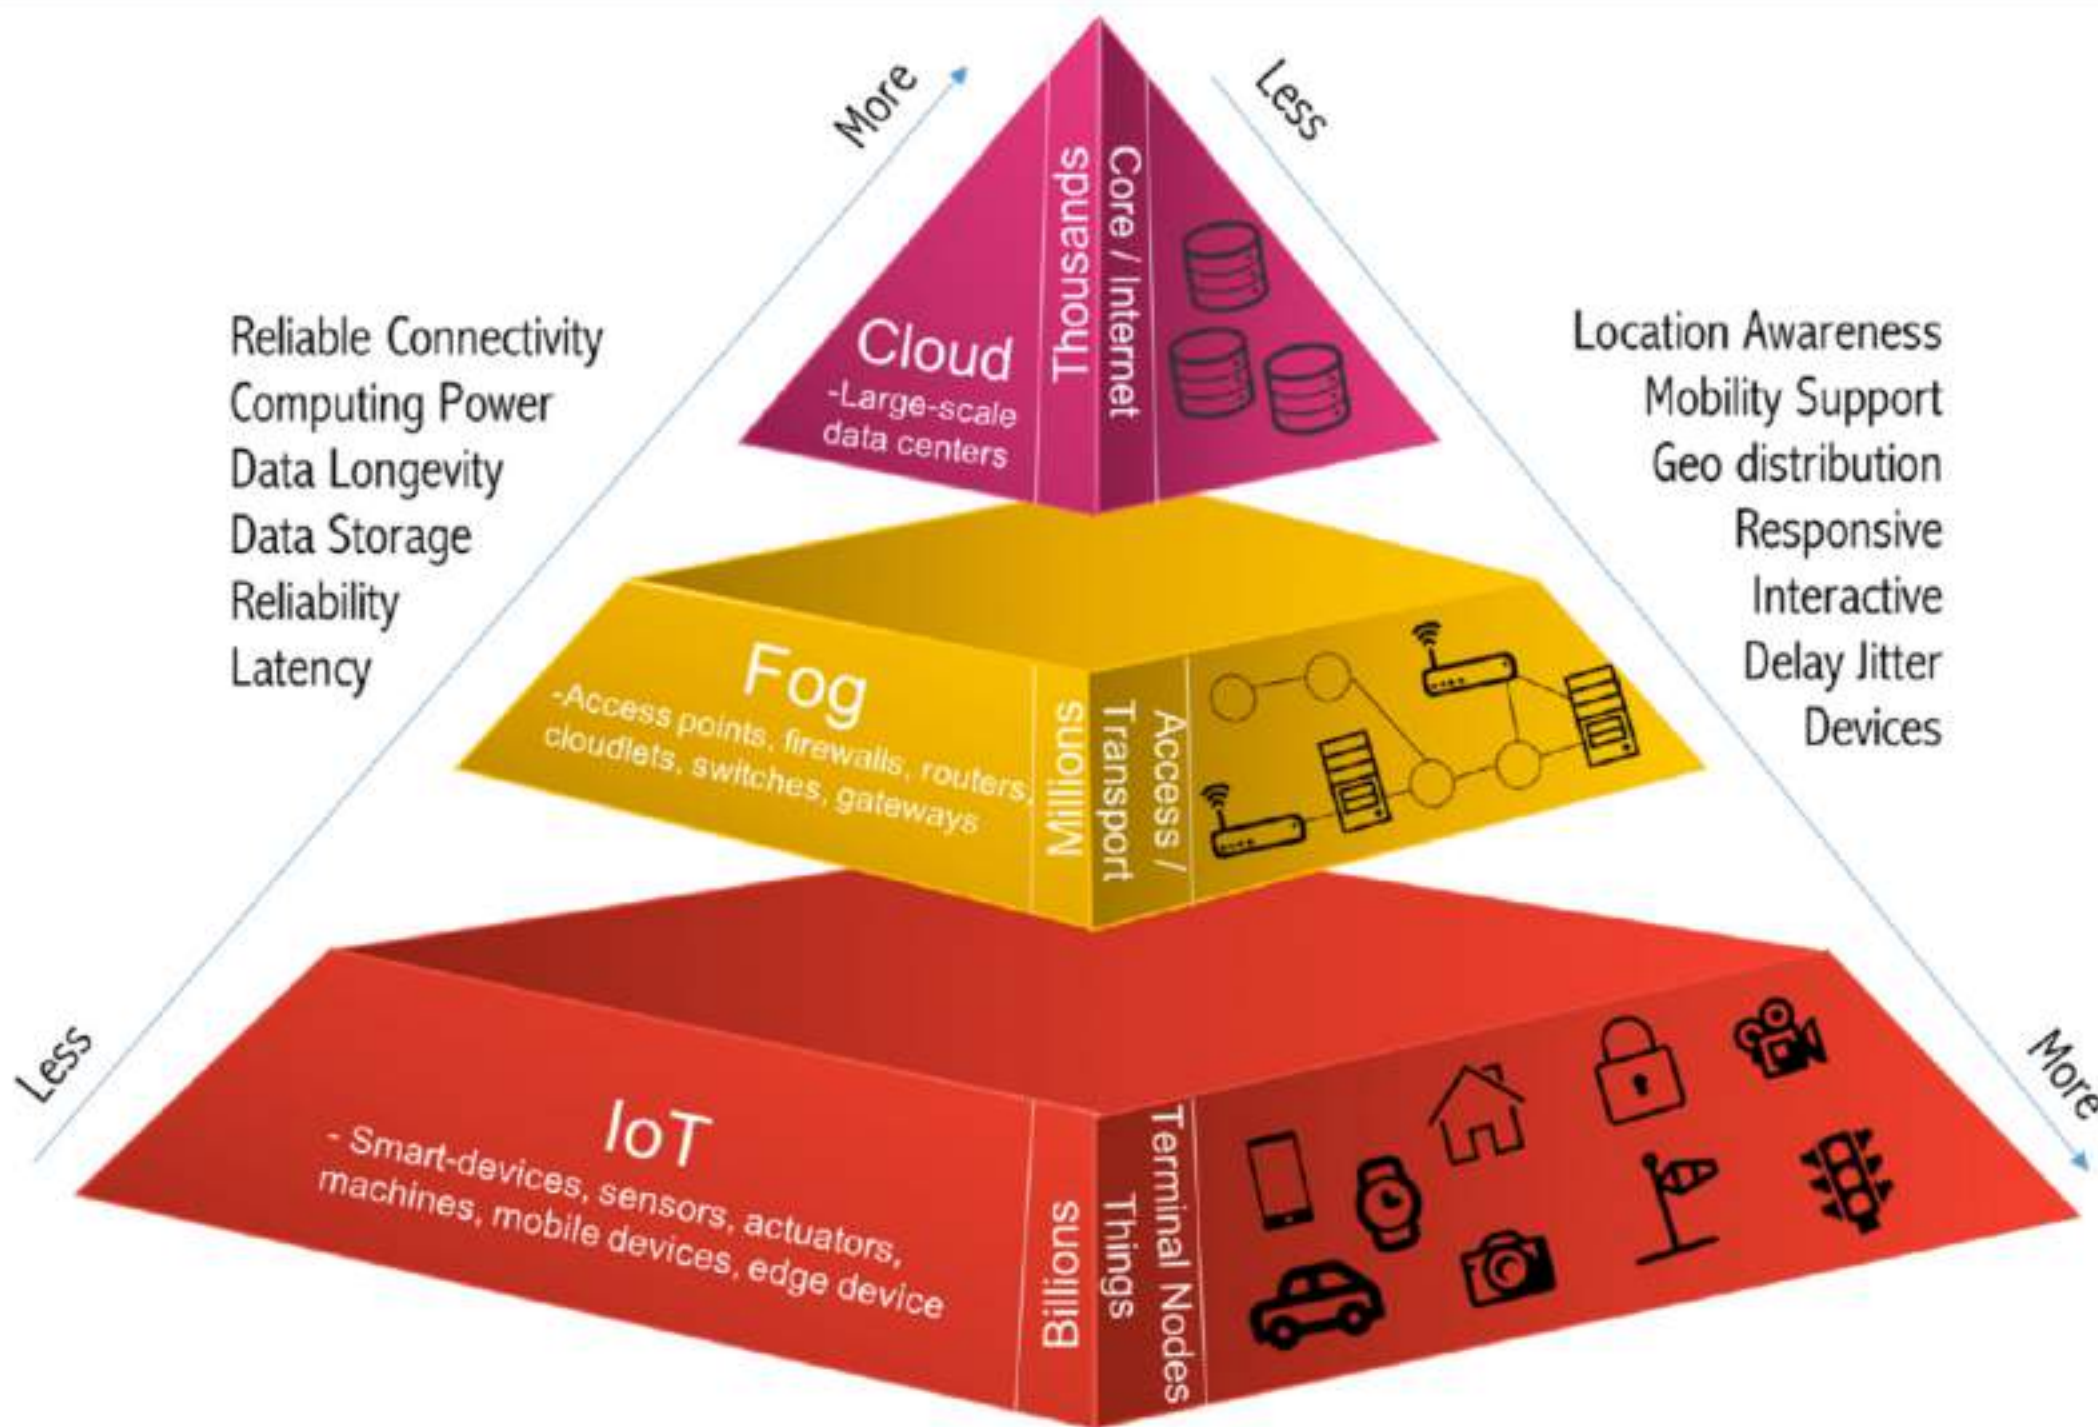



ORCHESTRATION




UNIVERSITY
OF TWENTE.

DIGITAL SOCIETY
INSTITUTE





Monolithic Architecture




Microservice Architecture

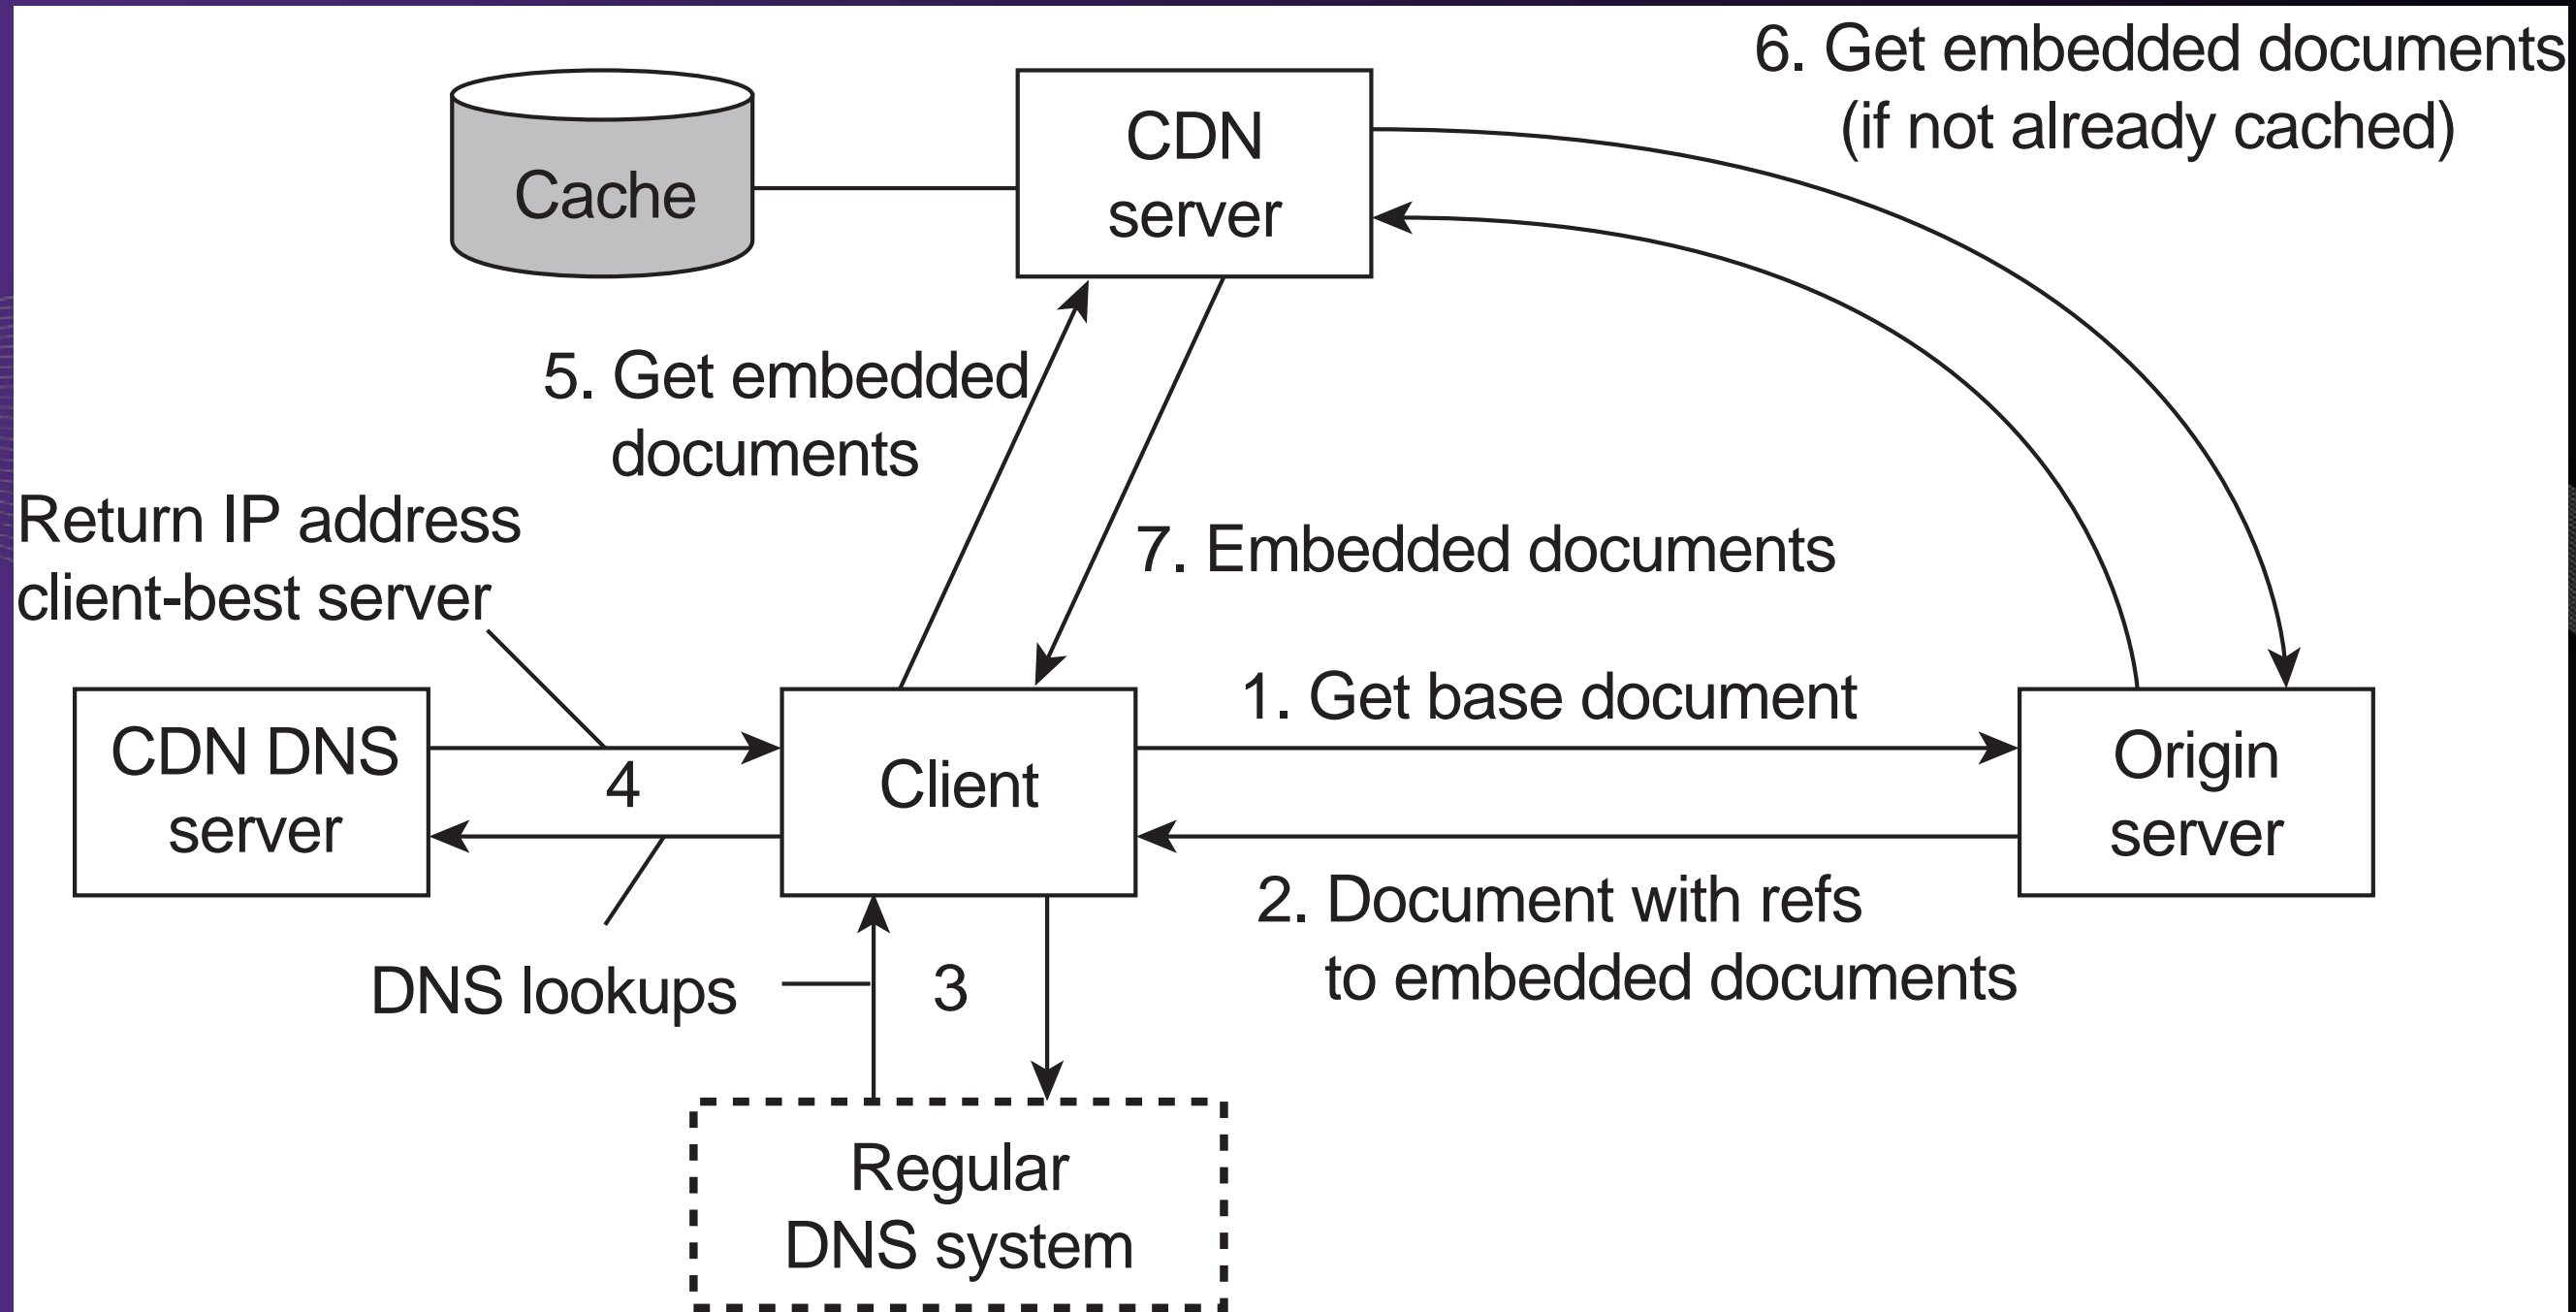
QUESTION FOR MANY RESEARCHERS

- How can you ensure that your distributed application, consisting of possibly many microservices, is conveniently distributed across all those nodes in the cloud and on the edge?

ANSWER: (SCALABLE) ORCHESTRATION


- Essentially a global system that monitors and decides where to place computations (or microservices).
- Do we need it?







ABOUT (HIDDEN) ASSUMPTIONS

- Many of us often pose problems and solutions based on assumptions that have not been made explicit and that may turn out to be false.
- **Assumption**: when dealing with edge devices we need to carefully decide which computations (or microservices) we place at the edge, in turn leading to a need for orchestration.
- **Hypothesis**: edge computing is generally not driven by the fact that it's easier to offload computations to the edge, but from the fact that there are compelling reasons not to offload computations to the cloud.


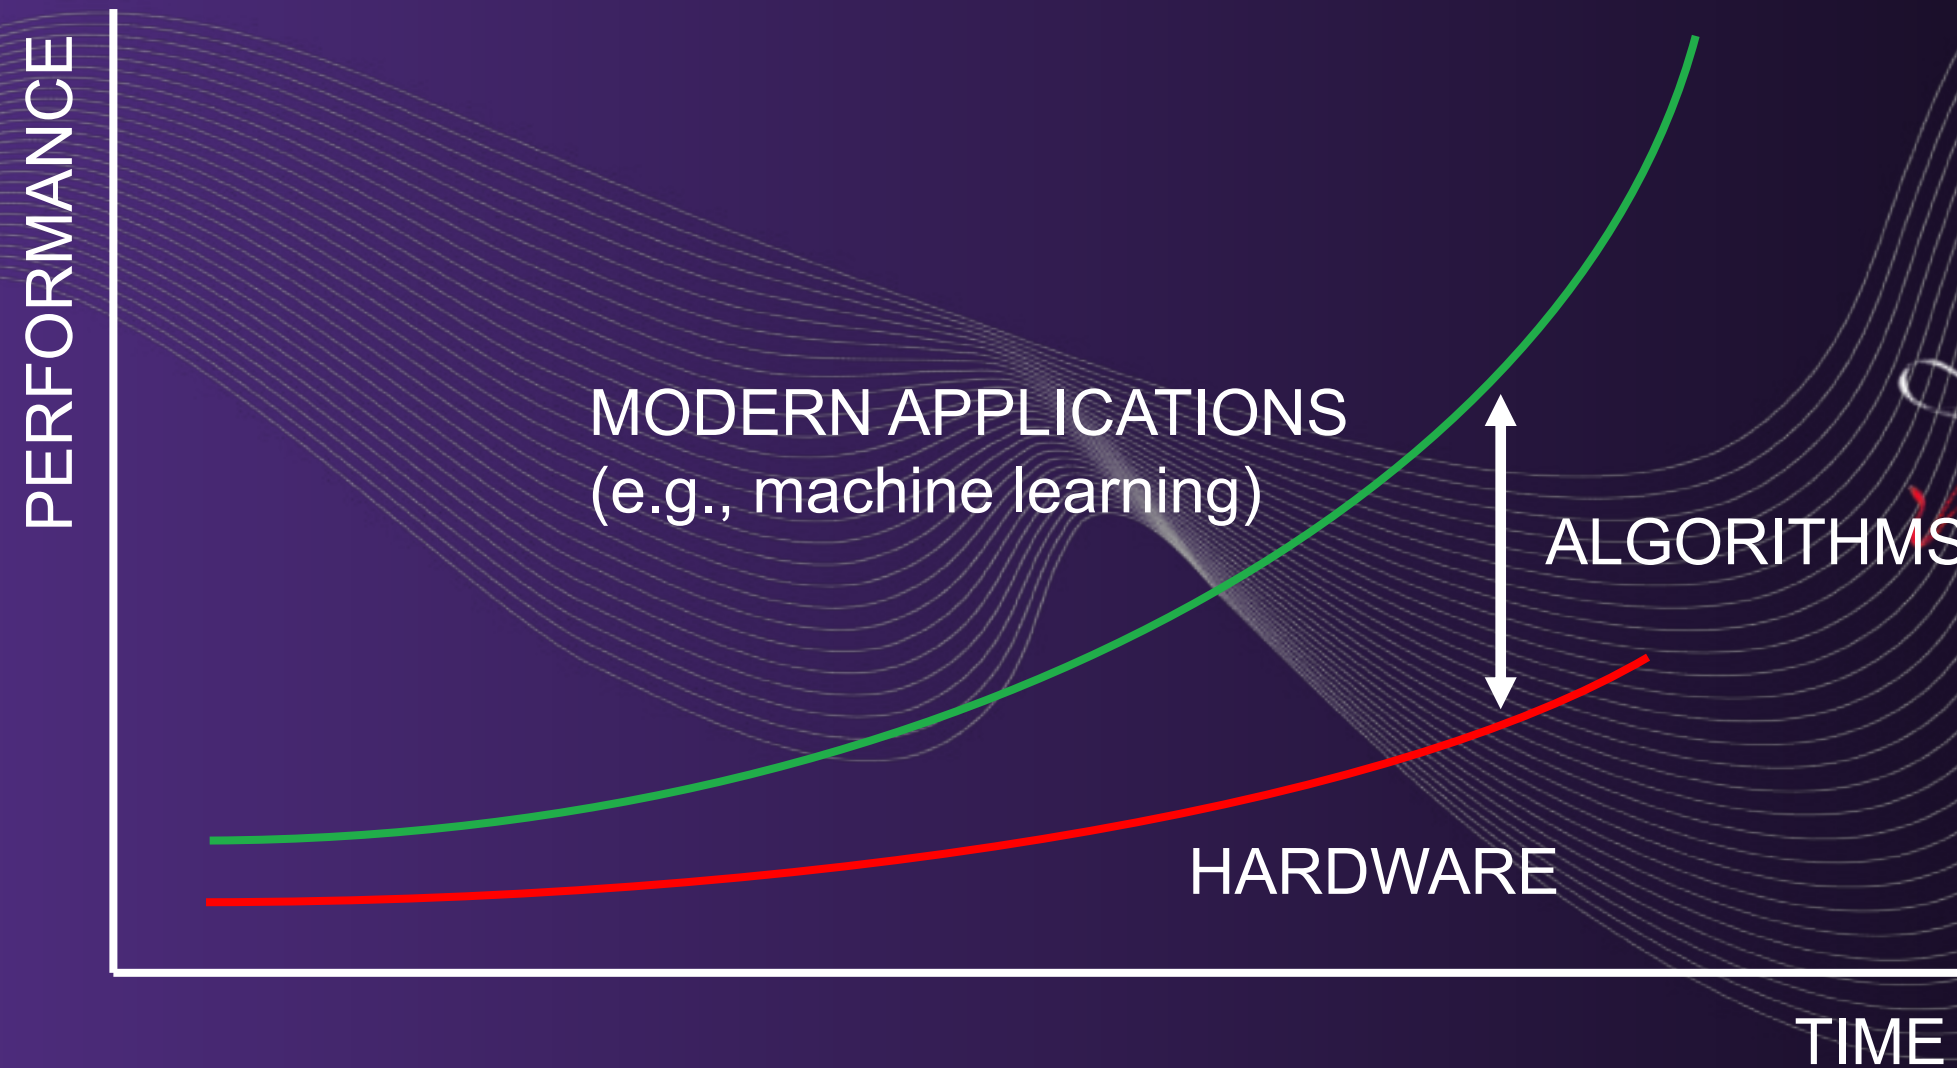


ABOUT (HIDDEN) ASSUMPTIONS

- **Assumption:** resources are so limited that we need to carefully optimize their usage. This is particularly the case for the Internet of Things as the number of things are growing exponentially.
- **Fact:** Computational, networking, and storage resources have grown incredibly fast the past decades. We often don't even know what to do with them.
- **Hypothesis:** we can't beat Mother Nature, so we need to concentrate on situations where she is bugging us. Bandwidth we can control, latency we cannot, although even that is not as true as it would seem.


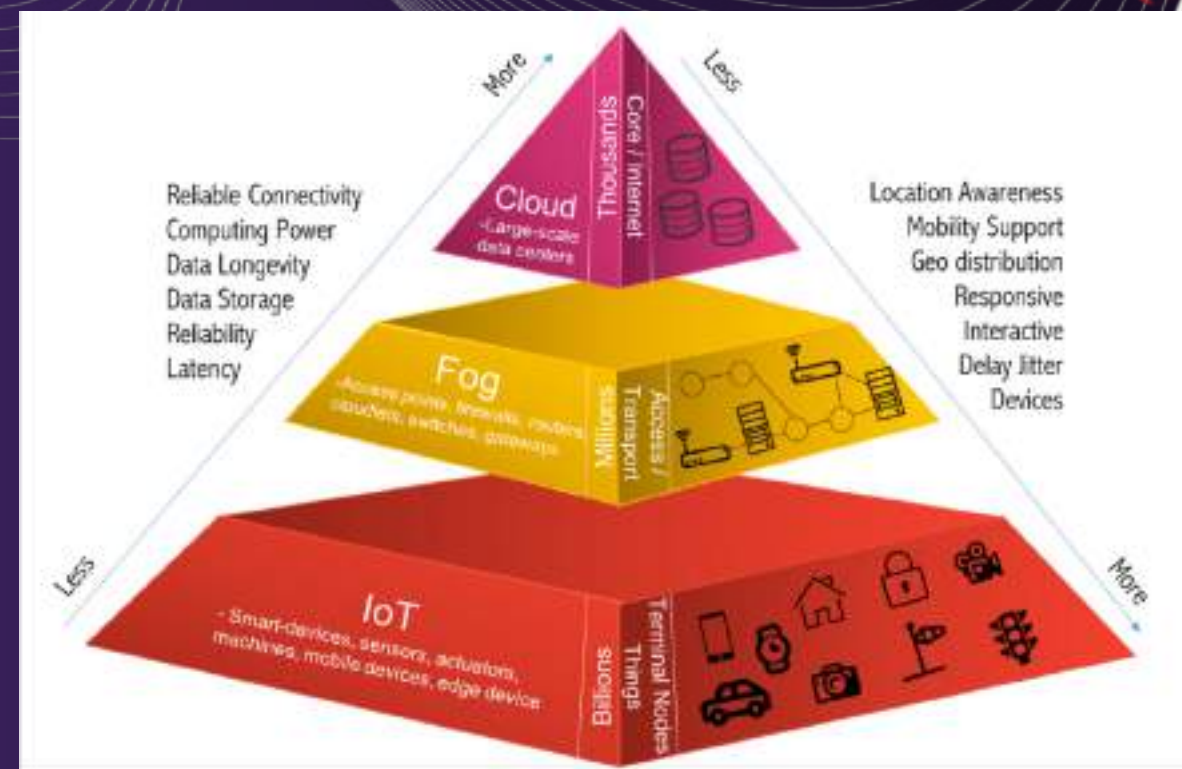


ABOUT (HIDDEN) ASSUMPTIONS



MORE (HIDDEN) ASSUMPTIONS

- We need to keep data closeby (we can't trust the cloud)
- Decentralization is better centralized approaches
- We're missing out on resource optimization (there is research on sharing resources of and among mobile phones)



WORK TO DO

The background features a series of thin, white, wavy lines that flow from the left side towards the right. On the right side, there is a complex 3D geometric structure composed of numerous triangular and quadrilateral faces, rendered in shades of green and white. The structure appears to be a stylized, abstract representation of a building or a network of interconnected nodes.

UNIVERSITY
OF TWENTE.

DIGITAL SOCIETY
INSTITUTE

STARTING POINTS & SUGGESTIONS

- A Comprehensive Survey on Fog Computing: State-of-the-Art and Research Challenges.
C. Mouradian, D. Naboulsi, S. Yangui, R.H. Glitho, M.J. Morrow, P.A. Polakos.
IEEE Communications Surveys & Tutorials, vol 20 (1), 2018.
- Contains lots of references for further reading and inspiration.
- Edge Computing: A Survey.
W.Z. Khan, E. Ahmed, S. Hakak, I. Yaqoob, A. Ahmed.
Future Generation Computer Systems, vol 97, 2019.
- Yet another survey (less comprehensive), but also lists open research questions.
- Distributed Systems book
H1, H2.3, H3.2, H3.4, H7.2, H7.4, H7.6