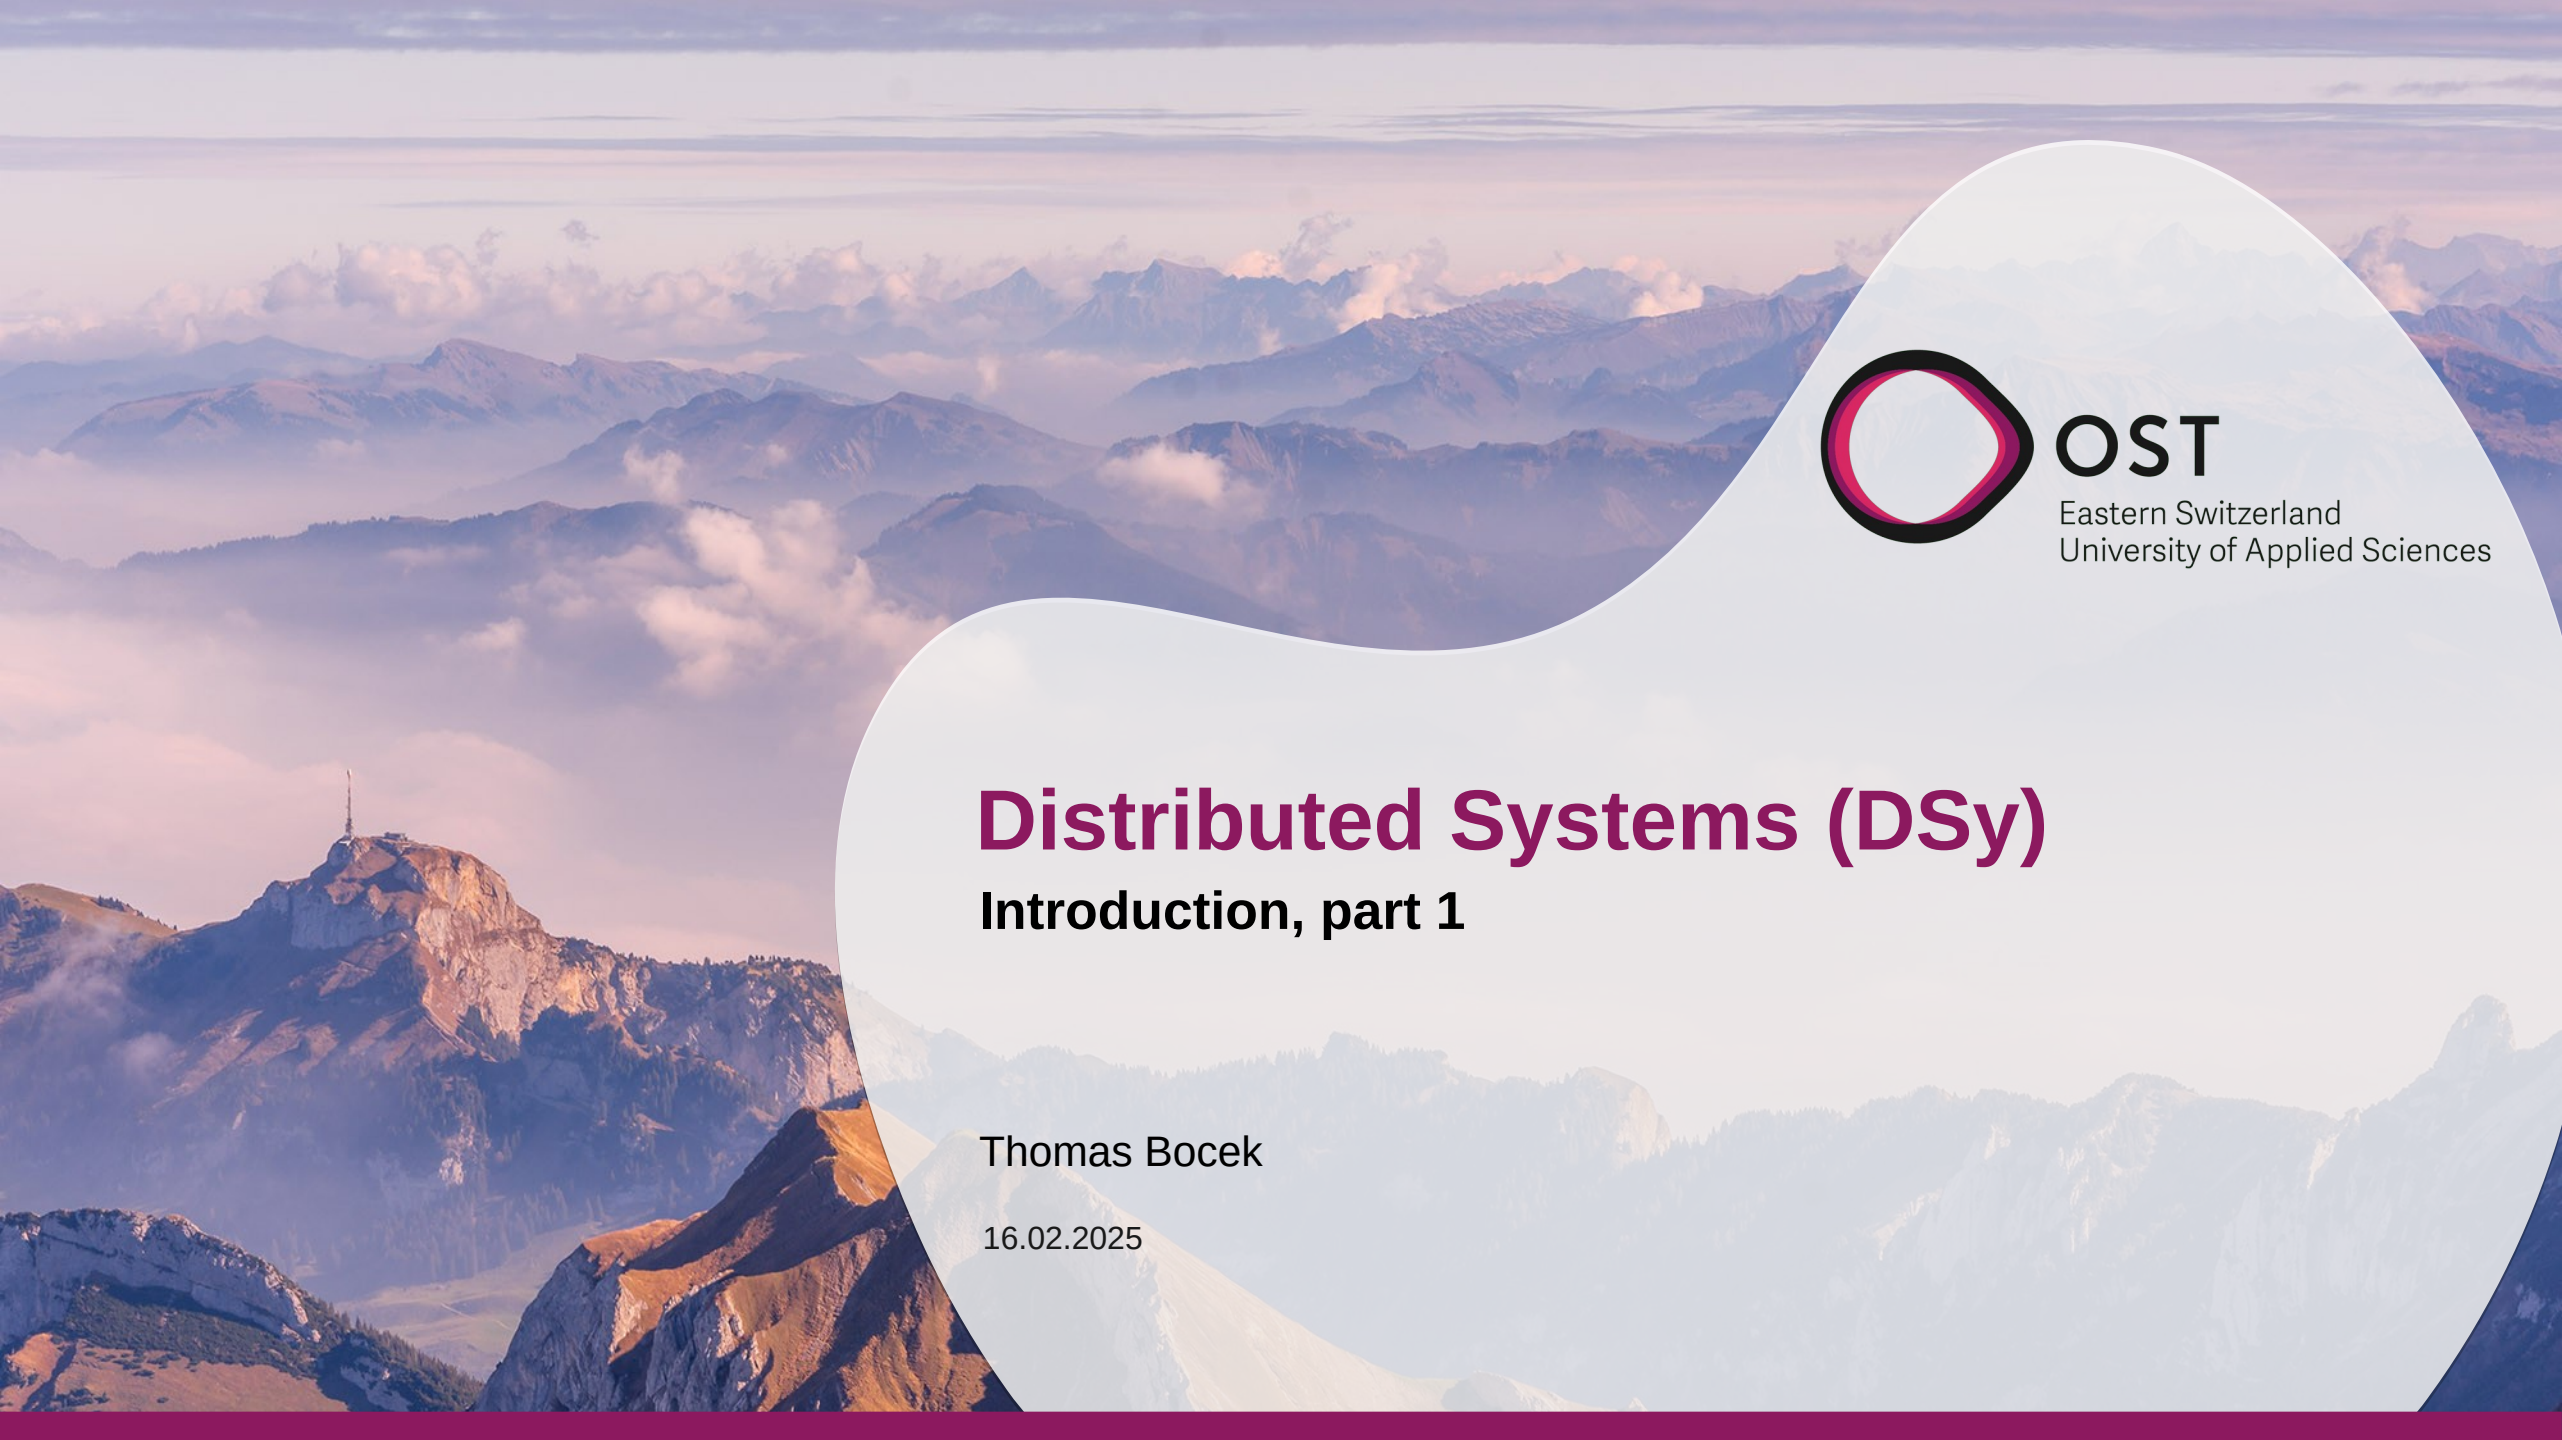

OST
Eastern Switzerland
University of Applied Sciences

Distributed Systems (DSy)

Introduction, part 1

Thomas Bocek

16.02.2025

Learning Goals

- Distributed systems add complexity. Avoid complexity!
- Why do we need distributed systems?
  1) Scaling (if one machine is not enough)
  2) Location (to move closer to the user)
  3) Fault-tolerance (HW will fail eventually)

Distributed Systems Motivation

- Why Distributed Systems
  - Scaling
    - Vertical (scale up), more memory, faster CPU
    - Horizontal (scale out), more machines
    - Apple has 75’000 Apache Cassandra nodes storing 10 petabytes of data in 2015 [source]

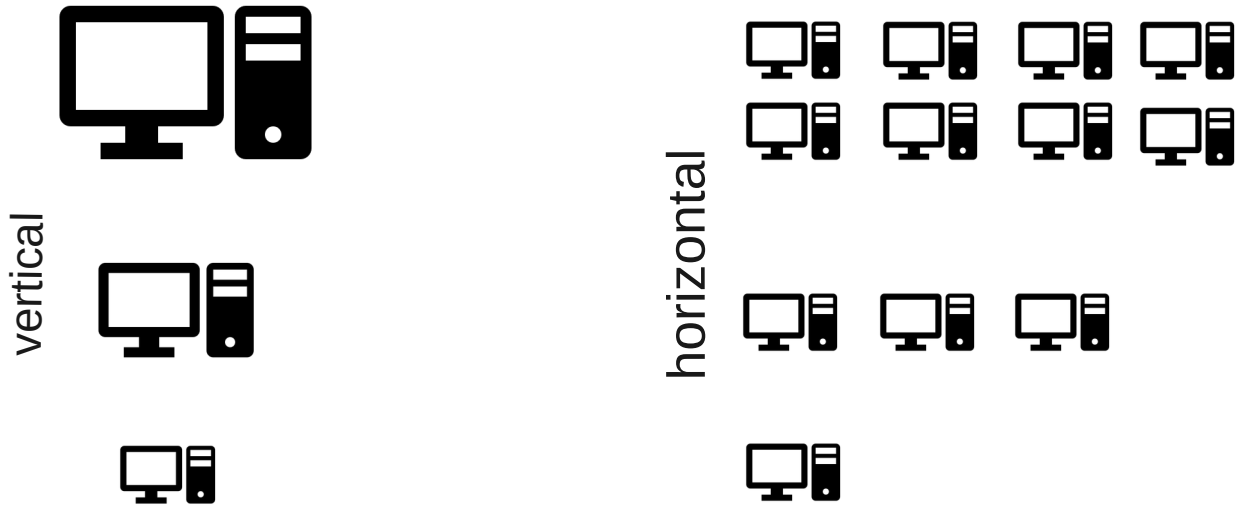

- Scaling Vertical
  - Example Ryzen 7 1700X, Ryzen 9, 3900X, Ryzen 9, 5950
  - Single Core
    - Cinebench R23:
      - 1700X: 981 (2017) [link]
      - 3900X: 1302 (2019) [link]
      - 5950X: 1644 (2020)
  - In 3 years, ~66% faster

OST

Distributed Systems Motivation

- Machine Learning
  - Current trend: scale horizontally
  - NVIDIA H200 Tensor Core GPU [link], 141GB
  - AMD Instinct MI325X Accelerators [link], 288GB
  - High param LLM with large context size
    - 1 User = 1 Card
    - Building models, fine-tuning → faster with multiple cards
    - Storing data → scalable storage
  - ML with vertical scaling is not (yet) feasible
- Economics
  - Initially scaling vertically is cheaper, until you max out HW
  - Current servers are fast: Rasperry Pi4 [link]
  - Raspberry Pi5 with llama.cpp / llama 3.2 [link] [link]

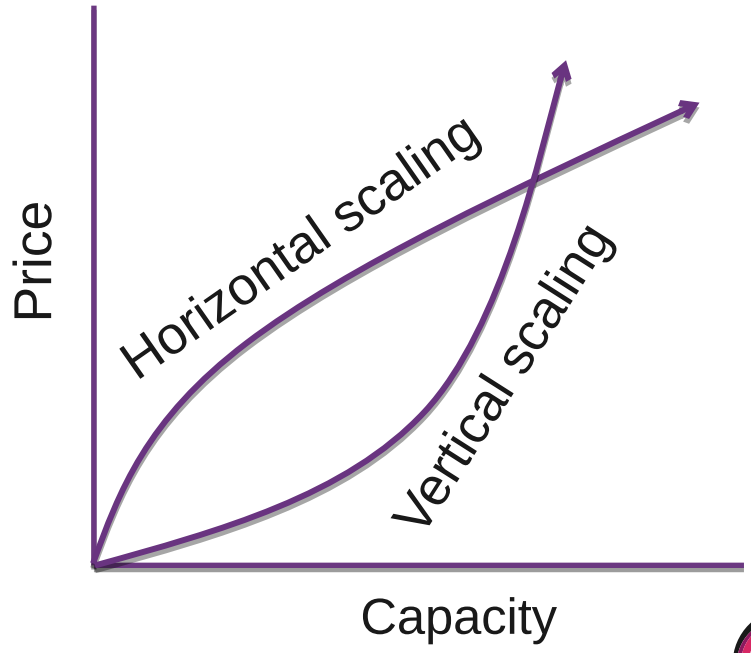

Distributed Systems Motivation

Horizontal Scaling

+ Lower cost with massive scale

+ Easier to add fault-tolerance

+ Higher availability

- Adaption of software required

- More complex system, more components involved

Vertical Scaling

+ Lower cost with small scale

+ No adaption of software required

+ Less complexity

- HW limits for scaling

- Risk of HW failure causing outage

- More difficult to add fault-tolerance

Vertical Scaling Performance

- Moore’s Law – nr. of transistors doubles every 2 years (other predictions, doubling chip performance every 18 month)
- Dead in 2025?
- Forbes 1995: "The price per transistor will bottom out sometime between 2003 and 2005. From that point on, there will be no economic point to making transistors smaller. So Moore's Law ends in seven years."

  AMD EPYC, 96 cores, 2socket server → 384 threads
- Apple M2 Ultra ~134b

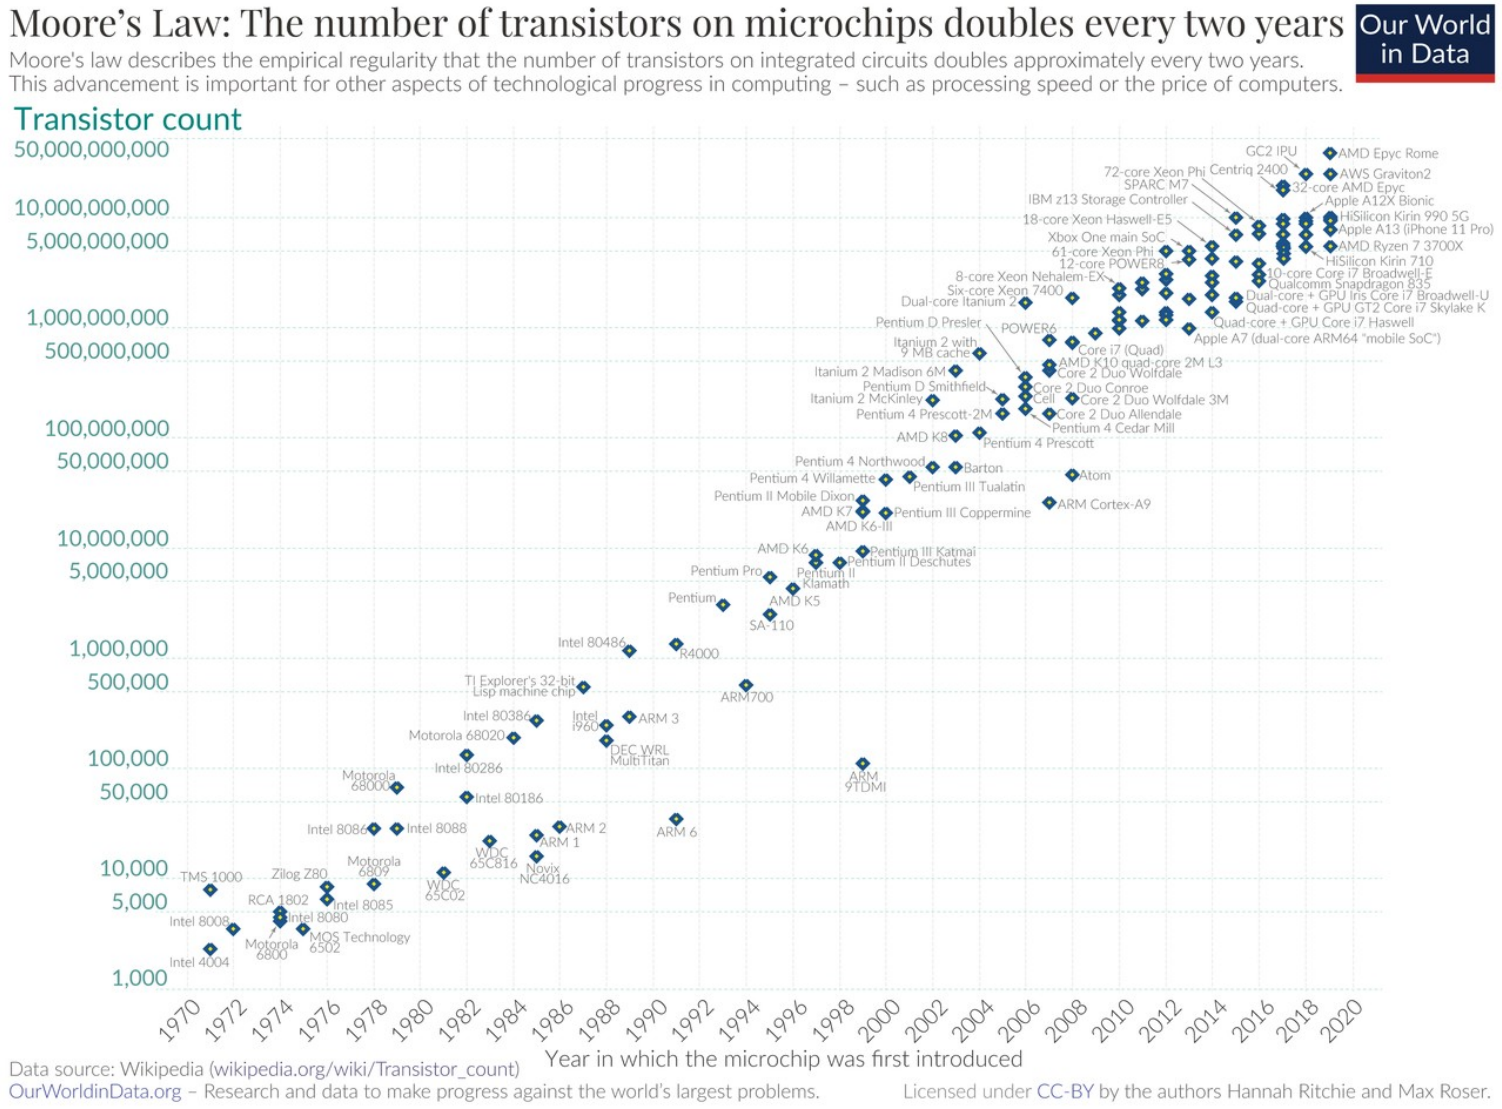

Vertical Scaling Performance

- Nielsen's Law: a high-end user's connection speed grows by 50% per year
- Bandwidth grows slower than computer power
  - Telecoms companies are conservative
  - Users are reluctant to spend much money on bandwidth
  - The user base is getting broader
- Optimize for bandwidth not for CPU
- Zmap complete scan of the IPv4 address space in under 5 minutes
- Init7: Fiber7-X2 25/25 Gbit ~64CHF/month

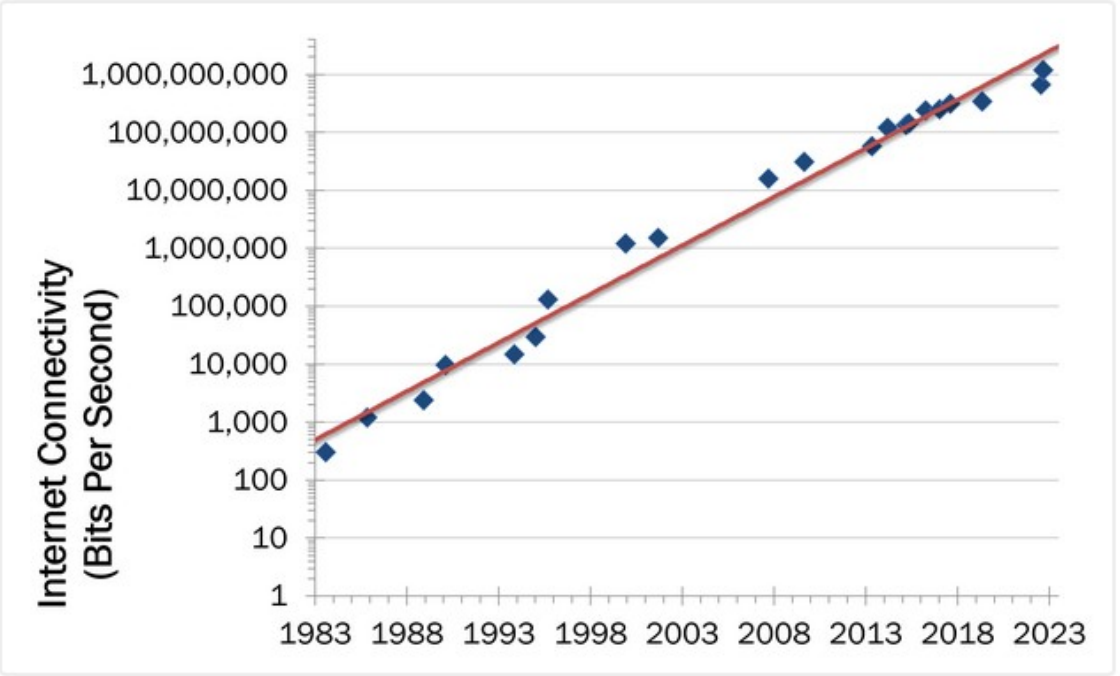

https://www.nngroup.com/articles/law-of-bandwidth/

		Annualized Growth Rate	Compound Growth Over 10 Years
Nielsen's law	Internet bandwidth	50%	57×
Moore's law	Computer power	60%	100×

OST

Vertical Scaling Performance

- Kryder’s Law: disk density doubling every 13 month
- «Soon hard drives will migrate into phones, still cameras, PDAs, cars and everyday appliances» https://www.scientificamerican.com/article/kryders-law/ , Aug. 2005
- User behavior changed
  - SSD, speed is important
- Cloud – Dropbox, Spotify
  - Streaming

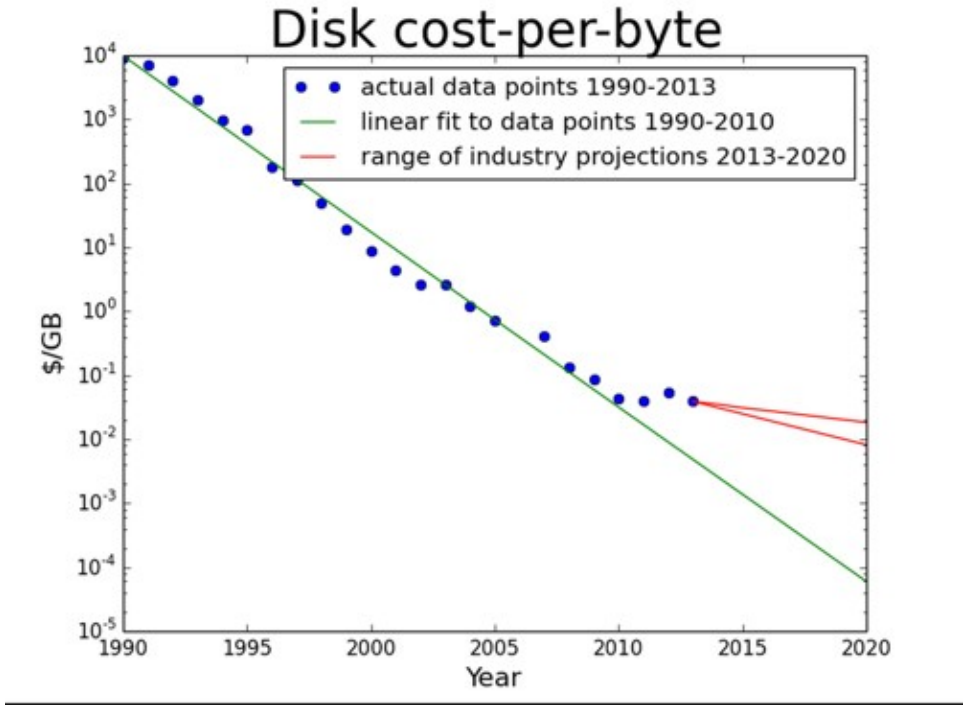

http://blog.dshr.org/2016/05/the-future-of-storage.html

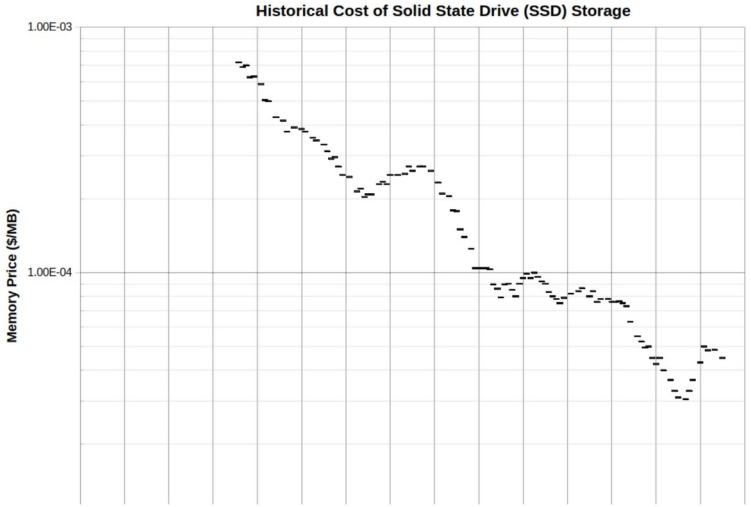

Source: https://jcmit.net/flashprice.htm

OST

Vertical Scaling Performance

- Vertical scaling
  - HW today is fast!
    - Database benchmark with a fast machine in 2020 (96 cores, 384GB RAM, 4 x NVMe SSD)
    - 70k TPS
- Best principle for small and simple applications!
- RPi 4 [link]
- Simple website with a few DB calls is not HW intensive
  - But: ML, Gaming (cloud gaming) are HW intensive

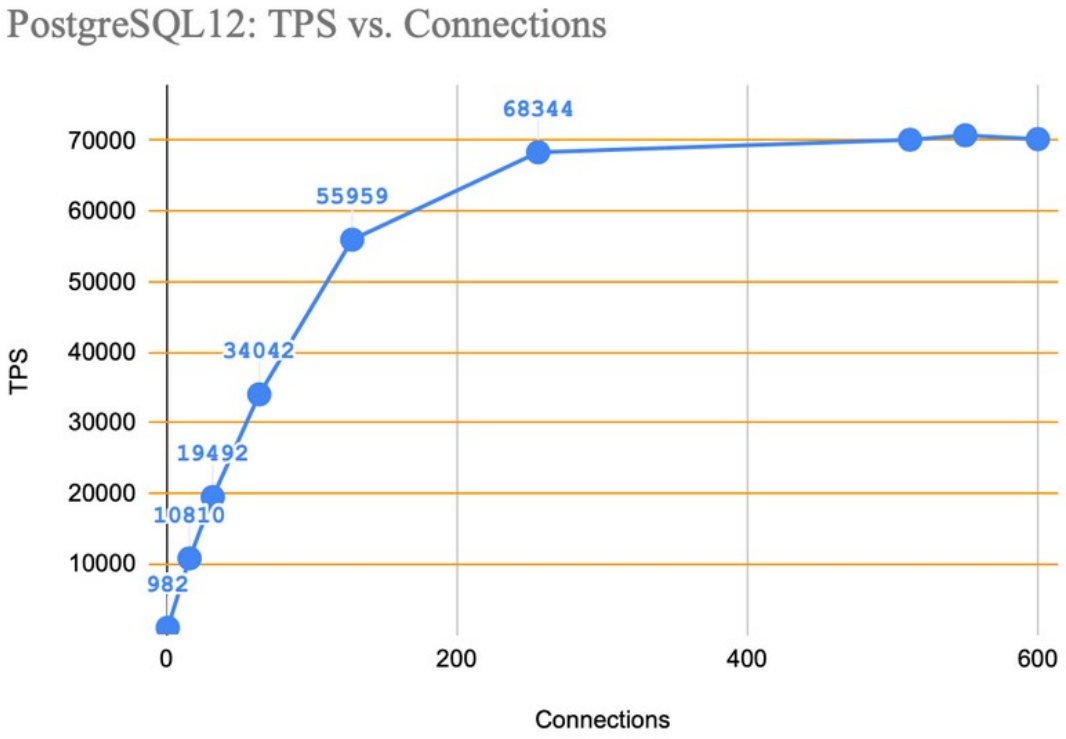

https://www.enterprisedb.com/blog/pgbench-performance-benchmark-postgresql-12-and-edb-advanced-server-12

OST

Vertical Scaling Performance

- Example: Let's Encrypt
- 21.01.2021: The Next Gen Database Servers Powering Let's Encrypt [link]
  - Providing certificates for 275m websites
  - “A database is at the heart of how Let’s Encrypt manages certificate issuance” - 1 single MariaDB
  - “We run the CA against a single database in order to minimize complexity” – Some read operations at replicas, one server for writes
  - 2x Xeon 24-cores running at 90%
  - Upgrade to 2x64 Epyc, on 15.09, running at 25%
    - Query 3 times faster
    - SATA → NVMe - IO from 500MB/s to 3 GB/s

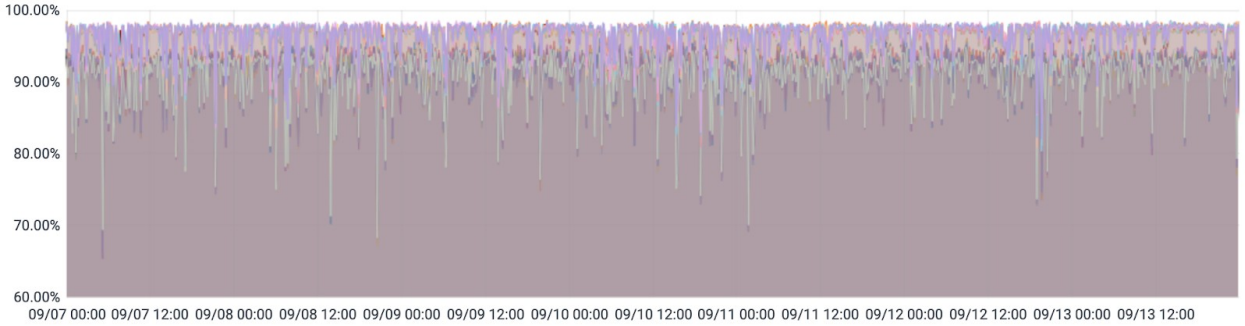

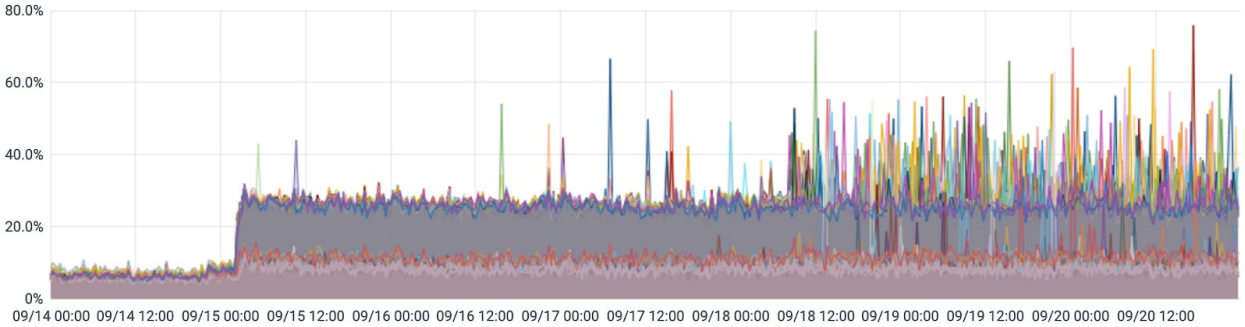

OST

Wirth's law (1995)

- "Software is getting slower more rapidly than hardware becomes faster." [link]
- Hardware get faster (e.g., Moore's Law)
- Software complexity and resource demands grow even faster - while widely observed, this is not scientifically measured like Moore's Law (On Bloat)
- Result: Modern software often feels slower despite running on much faster hardware - this is anecdotal experience many developers share

- Examples of bloat: Modern web pages often larger than the original DOOM game, electron apps consuming gigabytes of RAM
- Will it run DOOM? [link]
  - Doom Ledger Application [link]

OST