



**OST**

Eastern Switzerland  
University of Applied Sciences

# Distributed Systems (DSy)

## Introduction - Scaling

Thomas Bocek

14.02.2026

# 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
- 2025, Google: largest known Kubernetes cluster, with 130,000 nodes [source] [source]

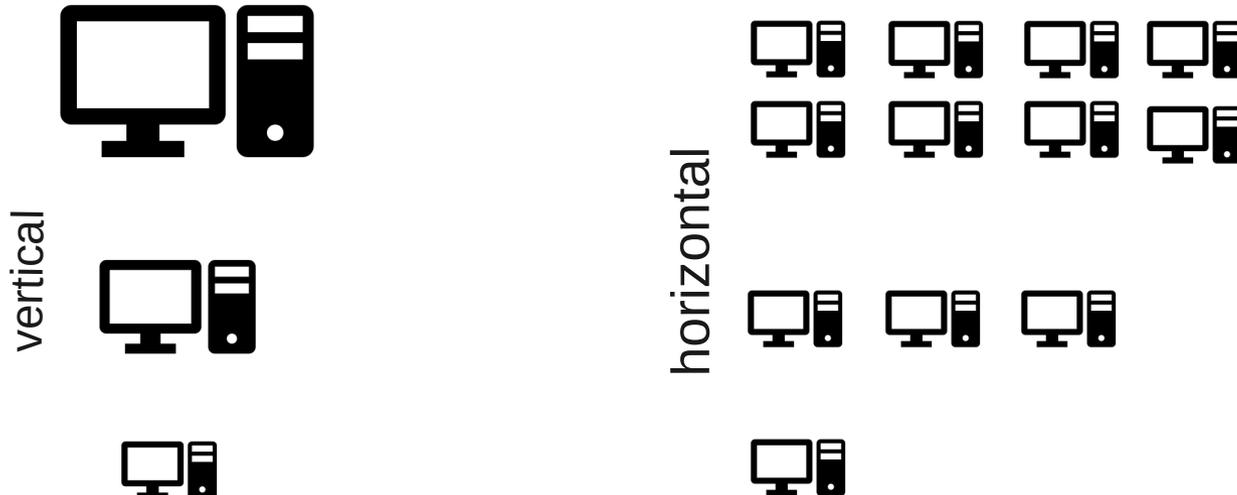
- Scaling Vertical

- Single Core

- Cinebench R23:

- 1700X: 981 (2017) [link]
- 3900X: 1302 (2019) [link]
- 5950X: 1644 (2020)
- 9950X: 2202 (2024) [link]

- In 7 years, more than 2x as fast

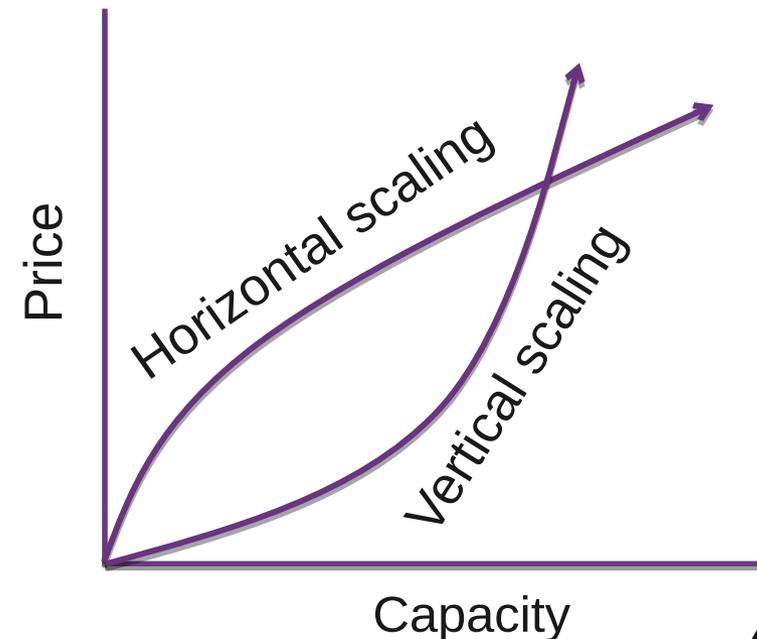


# Single Core – Multi Core

- Amdahl's Law ([Gene Amdahl, 1967](#))
  - $S = 1 / ((1 - P) + P/N)$
  - $S$  = Speedup,  $P$  = parallelizable fraction (0–1),  
 $N$  = number of parallel units
- Parallelization applies at every level
  - Cores
  - CPUs
  - Servers
  - Clusters
  - Datacenters
- Example:  $P = 95\%$ 
  - Max speedup = 20x (even with  $\infty$  processors)
    - 10 processors → 6.9x
    - 100 → 16.8x
    - 1000 → 19.6x
- The sequential part is always the bottleneck
- Real-world: DB locks limit throughput regardless of server count

# Distributed Systems Motivation

- Machine Learning
  - Current trend: scale horizontally
  - NVIDIA Blackwell GPU [\[link\]](#), 192GB, [NVL72](#)
  - AMD Instinct MI350X Accelerators [\[link\]](#), 288GB, MI400X [announced](#) with 432GB.
  - High param LLM with large context size
  - ML with vertical scaling is not (yet) feasible
- Economics
  - Initially scaling vertically cheaper, until max out HW
  - Current RAM [prices](#), [increasing](#)
  - Current servers are fast: Raspberry 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

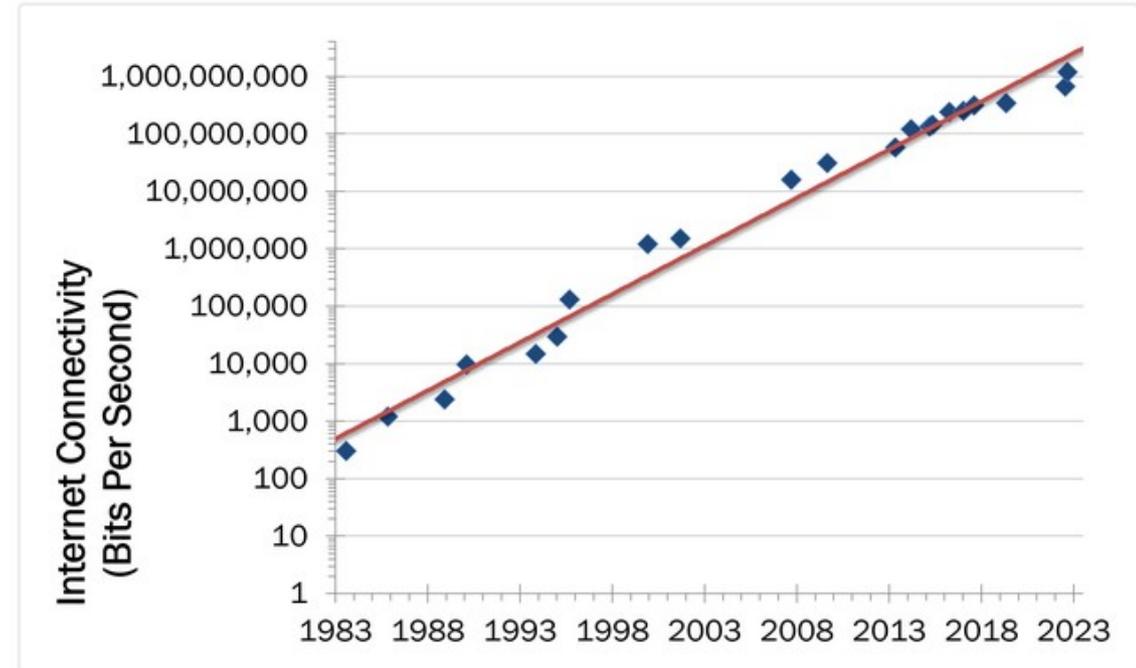
## Vertical Scaling

- + 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

- 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

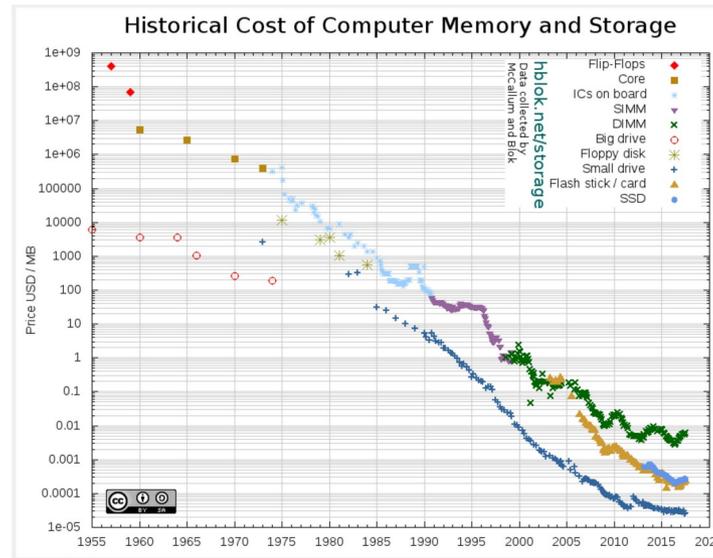


<https://www.nngroup.com/articles/law-of-bandwidth/>

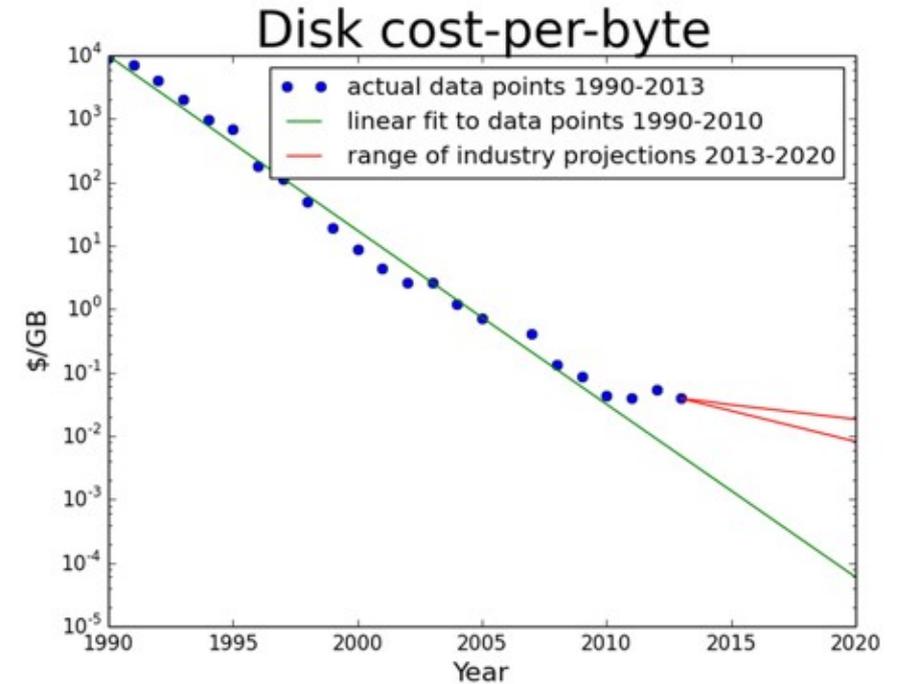
		Annualized Growth Rate	Compound Growth Over 10 Years
Nielsen's law	Internet bandwidth	50%	57×
Moore's law	Computer power	60%	100×

# Vertical Scaling Performance

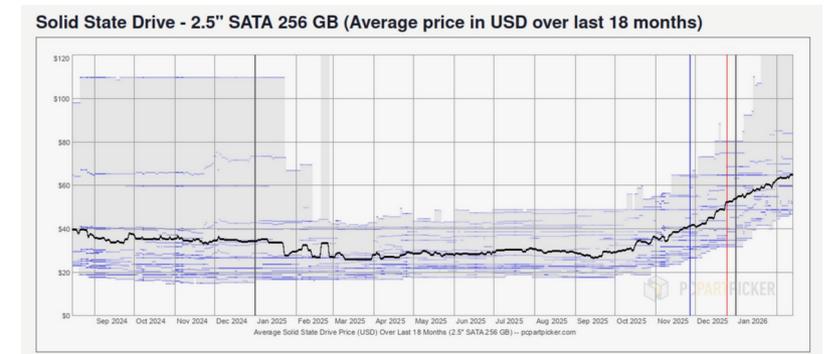
- Kryder's Law: disk density doubling every 13 month
- «Soon hard drives will migrate into phones, still cameras, PDAs, cars and everyday appliances»  
<https://www.scientificamerican.com/article/kryders-law/> , Aug. 2005
- User behavior changed
  - SSD, speed is important
- Cloud – Dropbox, Spotify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



Source: <https://hblok.net/blog/storage/>



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

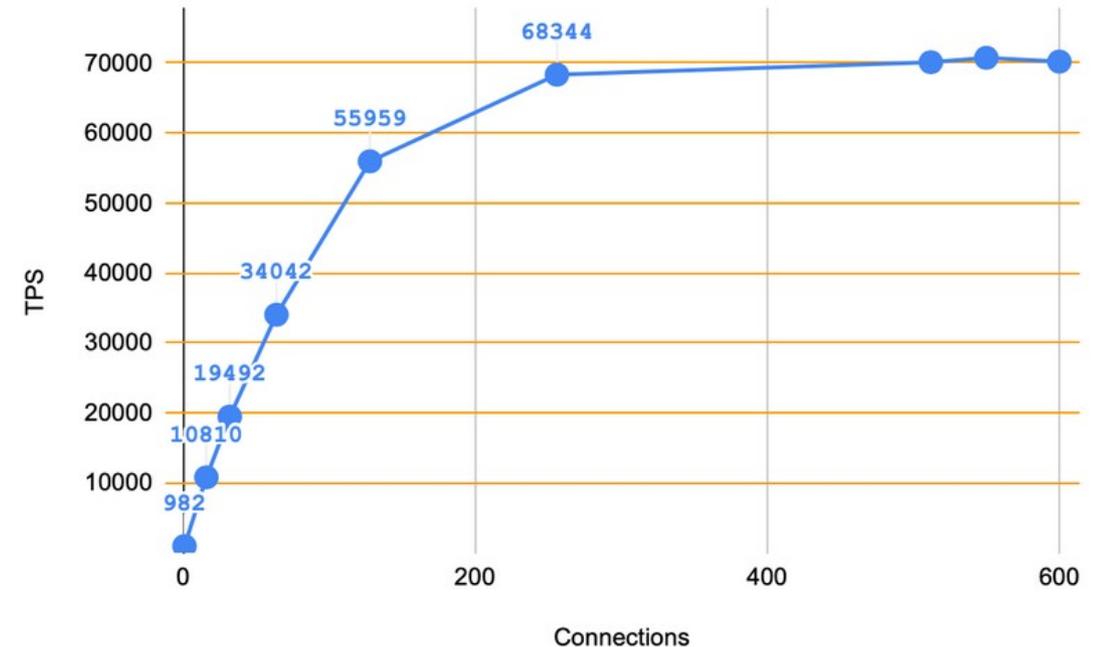


Source: <https://pcpartpicker.com/trends/price/internal-hard-drive/>

# 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

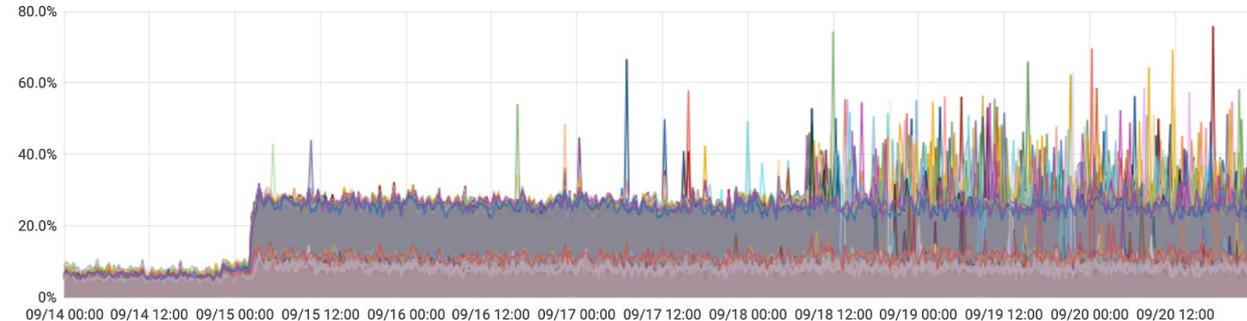
PostgreSQL12: TPS vs. Connections



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

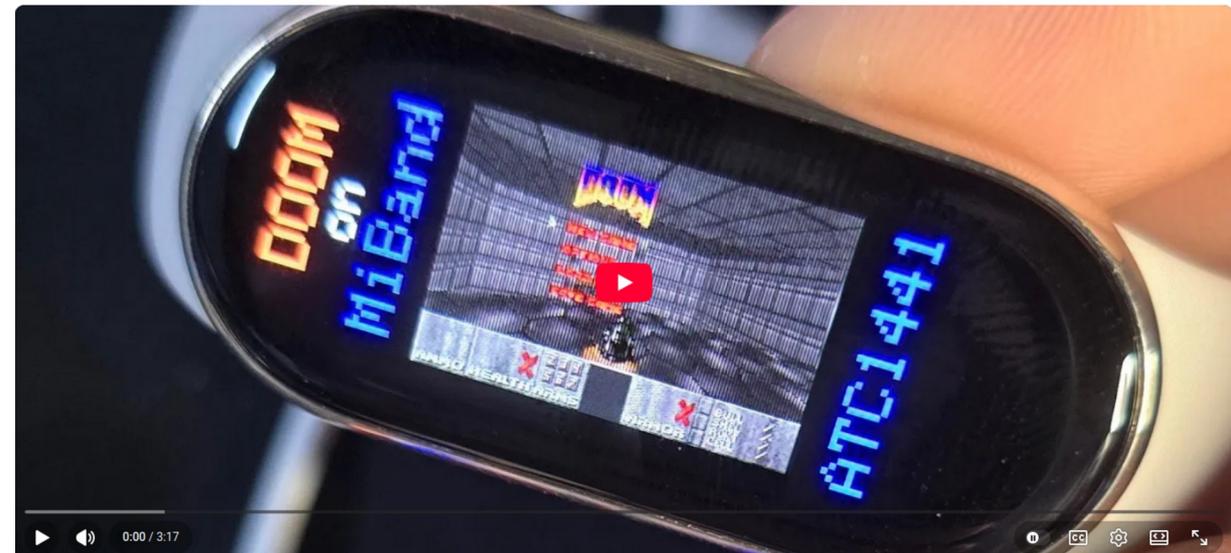
# Vertical Scaling Performance

- 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



# 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](#)]



# Energy: The New Scaling Constraint

- Global data center power: ~415 TWh (2024), projected 945 TWh by 2030 [[source](#)]
- AI-optimized servers: 21% of DC power (2025), 44% by 2030, growing 5× faster than conventional [[source](#)]
- US data centers: 183 TWh ~4% of national electricity [[source](#)]
- One GB200 NVL72 rack = 120 kW ≈ 40 households [[source](#)]
- Bitcoin network: ~204 Twh/year [[source](#)]
- AI-specific hardware will surpass Bitcoin soon
- ChatGPT query: ~0.3 Wh, Google Gemini: 0.24 Wh [[source](#), [source](#)]
  - + 100 queries/day = 30 Wh < 0.3% Swiss household consumption ≈ 15 min Netflix
  - + Efficiency improving
  - + Time savings, indirect energy savings
  - Rebound effect, OpenAI ~2.5B queries/day
  - Infrastructure needs to be built anyway
- Recommendation: Use consciously and targeted