MASSIVIZING COMPUTER SYSTEMS

MASSIVIZING SERVERLESS COMPUTING:
THE SCIENCE, DESIGN, AND ENGINEERING OF SERVERLESS ECOSYSTEMS

@Large Research
Massivizing Computer Systems

http://atlarge.science

bit.ly/MassivizingServerless22

Serverless computing =
Extreme automation + fine-grained, utilization-based billing

Sponsored by:
NWO

@Alosup

VU VRIJE UNIVERSITEIT AMSTERDAM

Contributions from the MCS team. Many thanks!
Many thanks to our collaborators, international working groups, authors of all images included here.
Also to Pedro García López, WOSC for invitation!

Prof.dr.ir. Alexandru
IOSUP

US IN 1 MINUTE

WE'RE MASSIVIZING COMPUTER SYSTEMS!

VU AMSTERDAM < SCHIPHOL < THE NETHERLANDS < EUROPE

Amsterdam
founded 10th century
pop: 850,000

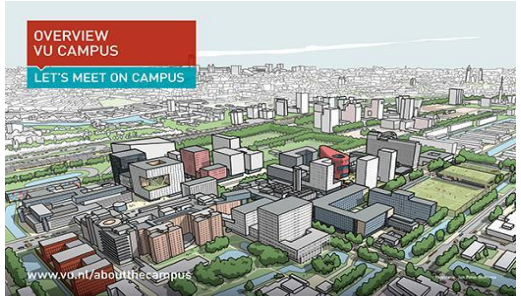

VU
founded 1880
pop: 23,500

The Netherlands

Europe

VU UNIVERSITY AMSTERDAM

WHO AM I?
PROF. DR. IR. ALEXANDRU IOSUP

- Education, my courses:
  - > Honours Programme, Computer Org. (BSc)
  - > Distributed Systems, Cloud Computing (MSc)
- Research, 15 years in DistribSys:
  - > Massivizing Computer Systems
- About me:
  - > Worked in 7 countries, NL since 2004
  - > I like to help… I train people in need
  - > VU University Research Chair + Group Chair
  - > NL ICT Researcher of the Year
  - > NL Higher-Education Teacher of the Year
  - > NL Young Royal Academy of Arts & Sciences
  - > Knighted in 2020

VU

http://atlarge.science

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

CURRENT TEAM

WE ARE HIRING
A NEW ASST. PROF.!

- Professor
- Assistant Prof.
- Teacher
- Visitor/P.-doc
- Ph.D. student
- Early Scientist

ALUMNI

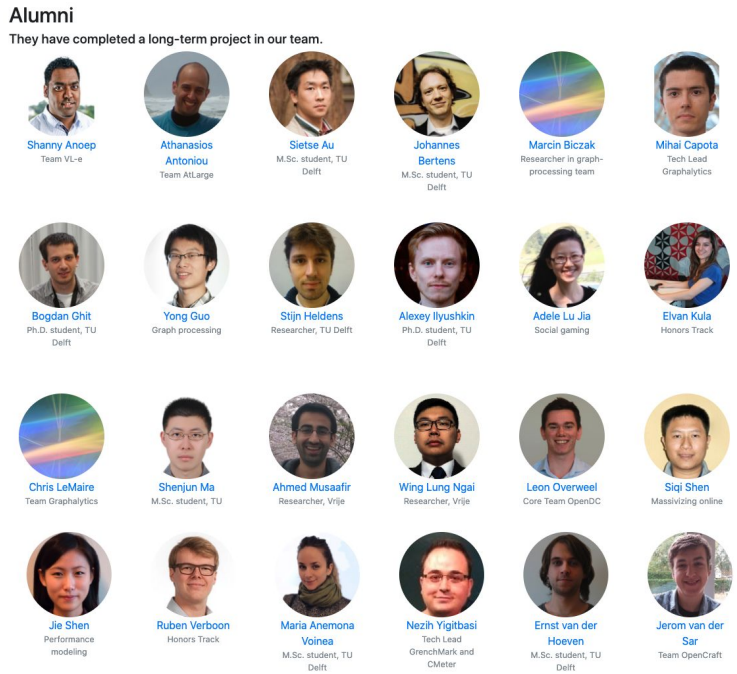

VISITORS

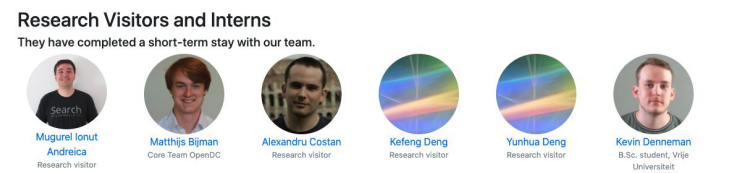

WE ARE A FRIENDLY, DIVERSE GROUP, OF DIFFERENT RACES AND ETHNICITIES, GENDERS AND SEXUAL PREFERENCES, AND VIEWS OF CULTURE, POLITICS, AND RELIGION. YOU ARE WELCOME TO JOIN!

MASSIVIZING COMPUTER SYSTEMS: OUR MISSION

http://atlarge.science/about.html

1. Improve the lives of millions through impactful research.

2. Educate the new generation of top-quality, socially responsible professionals.

3. Make innovation available to society and industry.

THIS IS THE GOLDEN AGE OF COMPUTER ECOSYSTEMS

1

THIS IS THE GOLDEN AGE OF MASSIVE COMPUTER ECOSYSTEMS

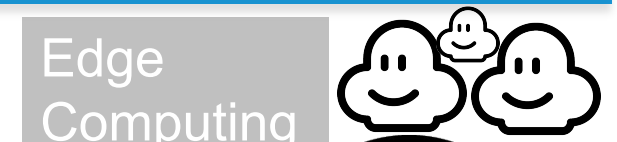

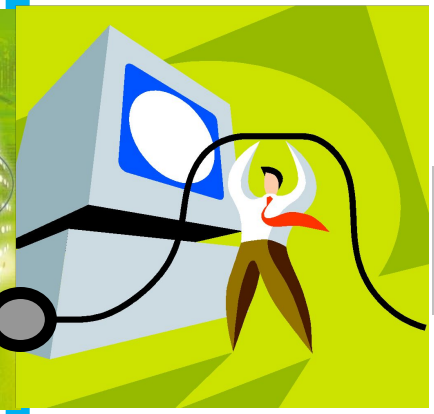

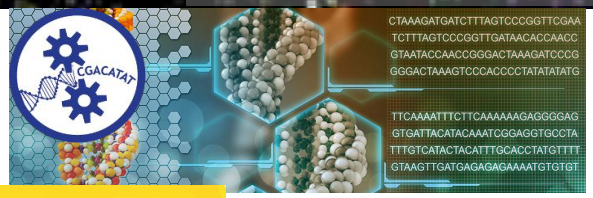

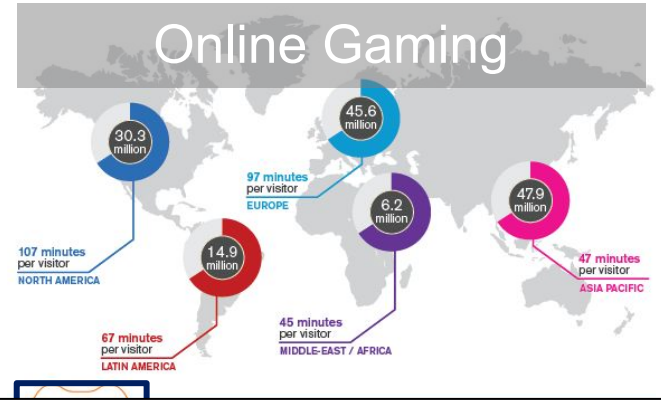

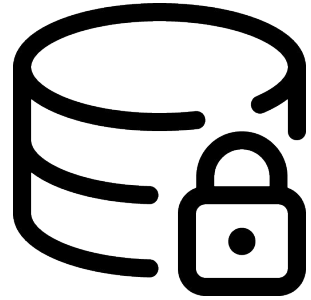

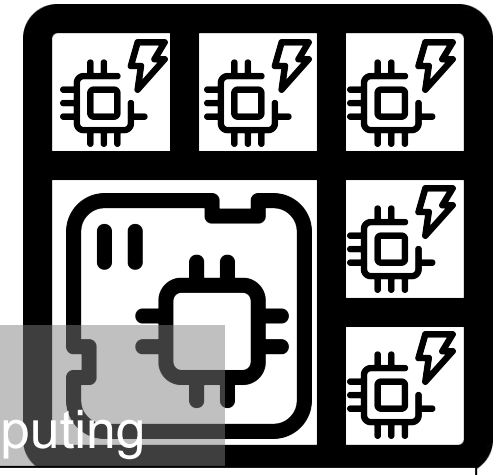

Daily Life

Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online] Hesselman, Grosso, Kuipers, et al. (2020) A Responsible Internet to increase Trust in the Digital World, JNSM [Online]

THE ECONOMIC IMPACT OF MASSIVE COMPUTER ECOSYSTEMS

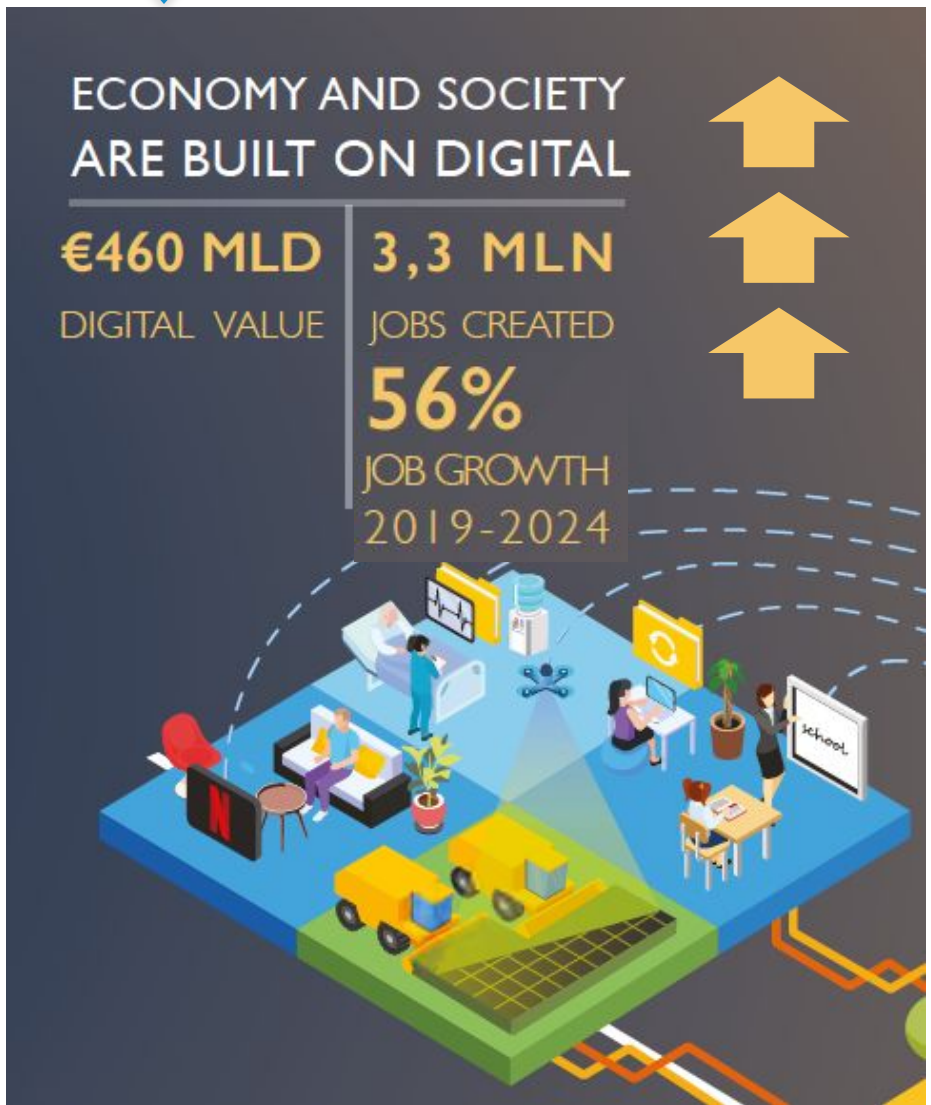

DIVERSE SERVICES FOR ALL

EVERY €1 → €15 ADDED VALUE

Impacting ≥60% of the NL GDP (1 trillion EUR/y)

Attracting ≥20% of all foreign direct investments in NL

Sources: Iosup et al., Massivizing Computer Systems, ICDCS 2018 [Online] / Dutch Data Center Association, 2020 [Online] / Growth: NL Gov't, Flexera, Binx 2020. Gartner 2019. IA 2017.

BUT WE CANNOT TAKE THIS TECHNOLOGY FOR GRANTED

2

(So, this is why I am giving this talk)

PHENOMENON: FAILURES IN CLOUD SERVICES

UNCOVERING THE PRESENCE OF FAILURES

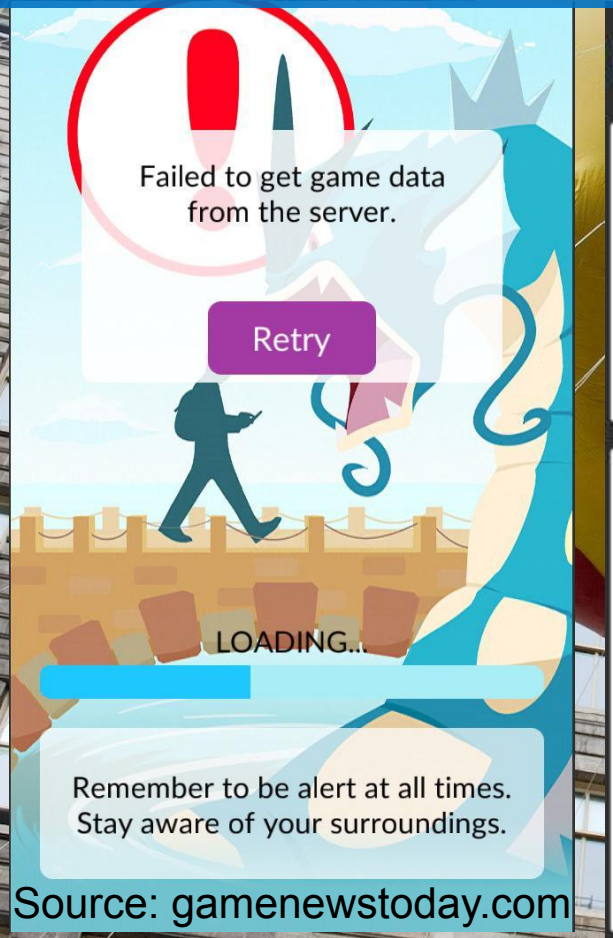

Source: gamenewstoday.com

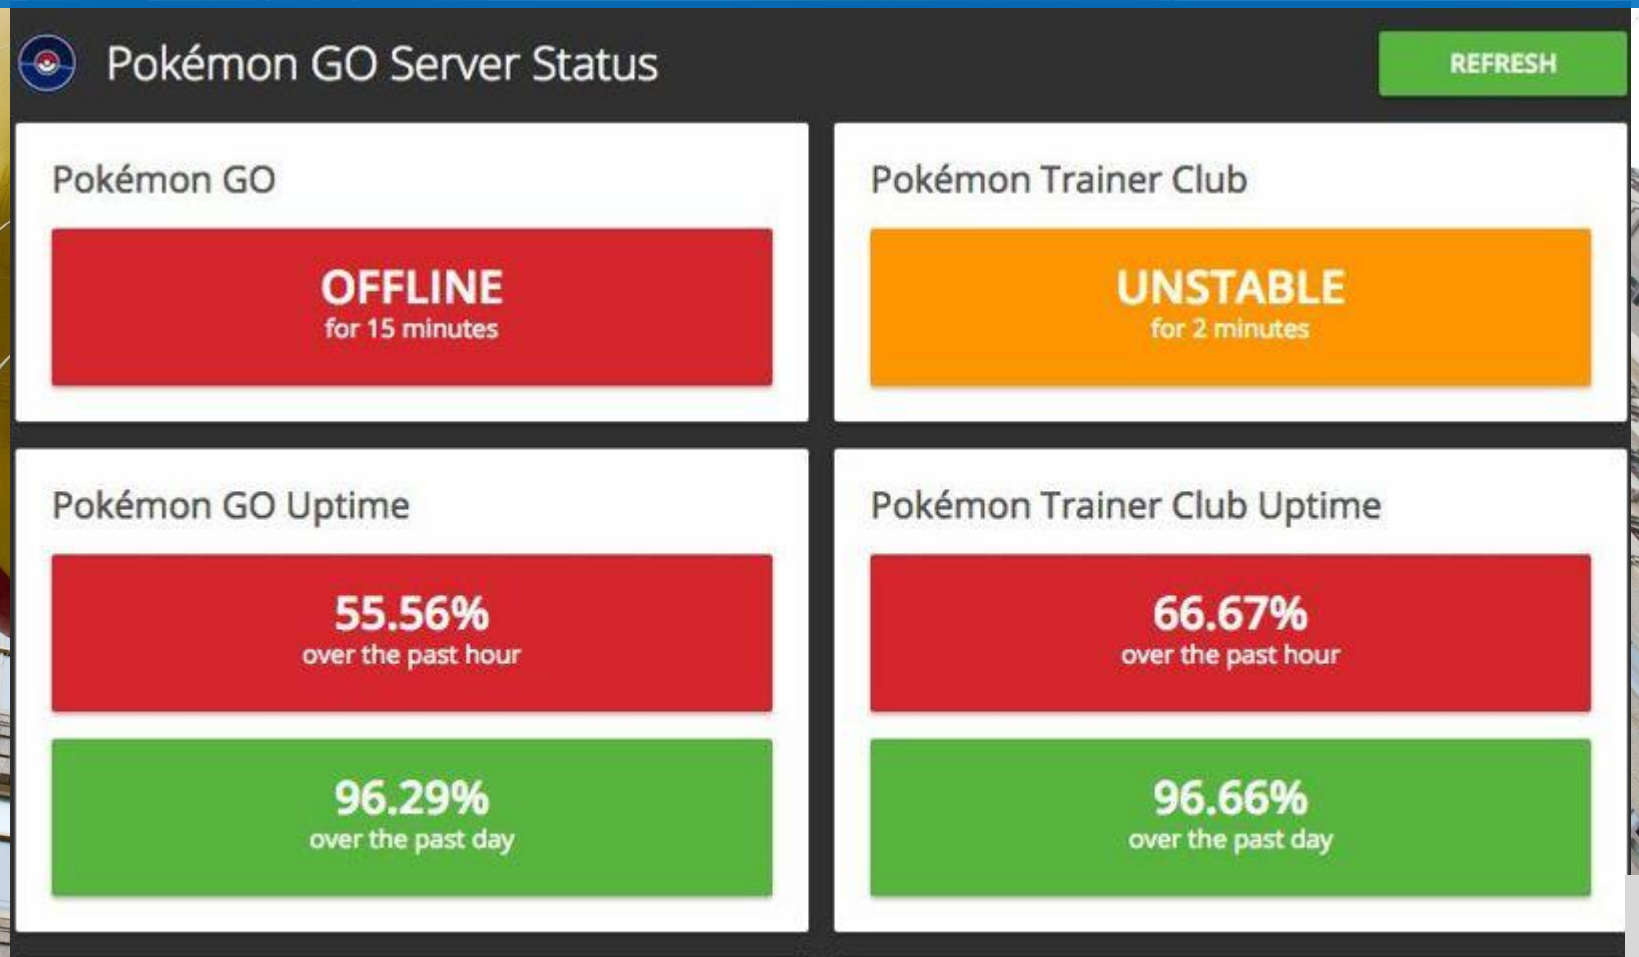

PHENOMENON: PERFORMANCE IN CLOUD SERVICES

UNCOVERING THE PRESENCE OF PERFORMANCE ISSUES, EVEN LEADING TO CRASHES

Polygon

Source: http://bit.ly/EveOnline21Crash

NEWS

Players in Eve Online broke a world record — and then the game itself

Developers said they're not 'able to predict the server performance in these kinds of situations'

By Charlie Hall | @Charlie_L_Hall | Jan 5, 2021, 2:54pm EST

Source: Razorien/CCP Games

PHENOMENON: CLOUD DATACENTER SUSTAINABILITY

UNCOVERING THE USE OF ENERGY AND WATER, THE IMPACT ON CLIMATE

Power consumption of datacenters:
>1% of global electricity

Source: Nature, 2018 [Online]

Water consumption of datacenters in the US:
>625Bn. l/y (0,1%)

Source: Energy Technologies Area, 2016 [Online]

Power consumption of datacenters in the Netherlands:
1→3% of national electricity

Source: NRC, 2019 [Online]

Other greenhouse emissions:
Largely unknown

Source: Nature Climate Change, 2020 [Online]

Source: NASA Earth Observatory

THIS TALK, IN A NUTSHELL

!

Serverless =

1. Extreme automation
2. Fine-grained reporting / utilization-based billing

Serverless can only be achieved through complex, smart computer ecosystems (operational simplicity is for the user)

3

A TYPICAL ECOSYSTEM: SERVICE, DATACENTER, SCHEDULER

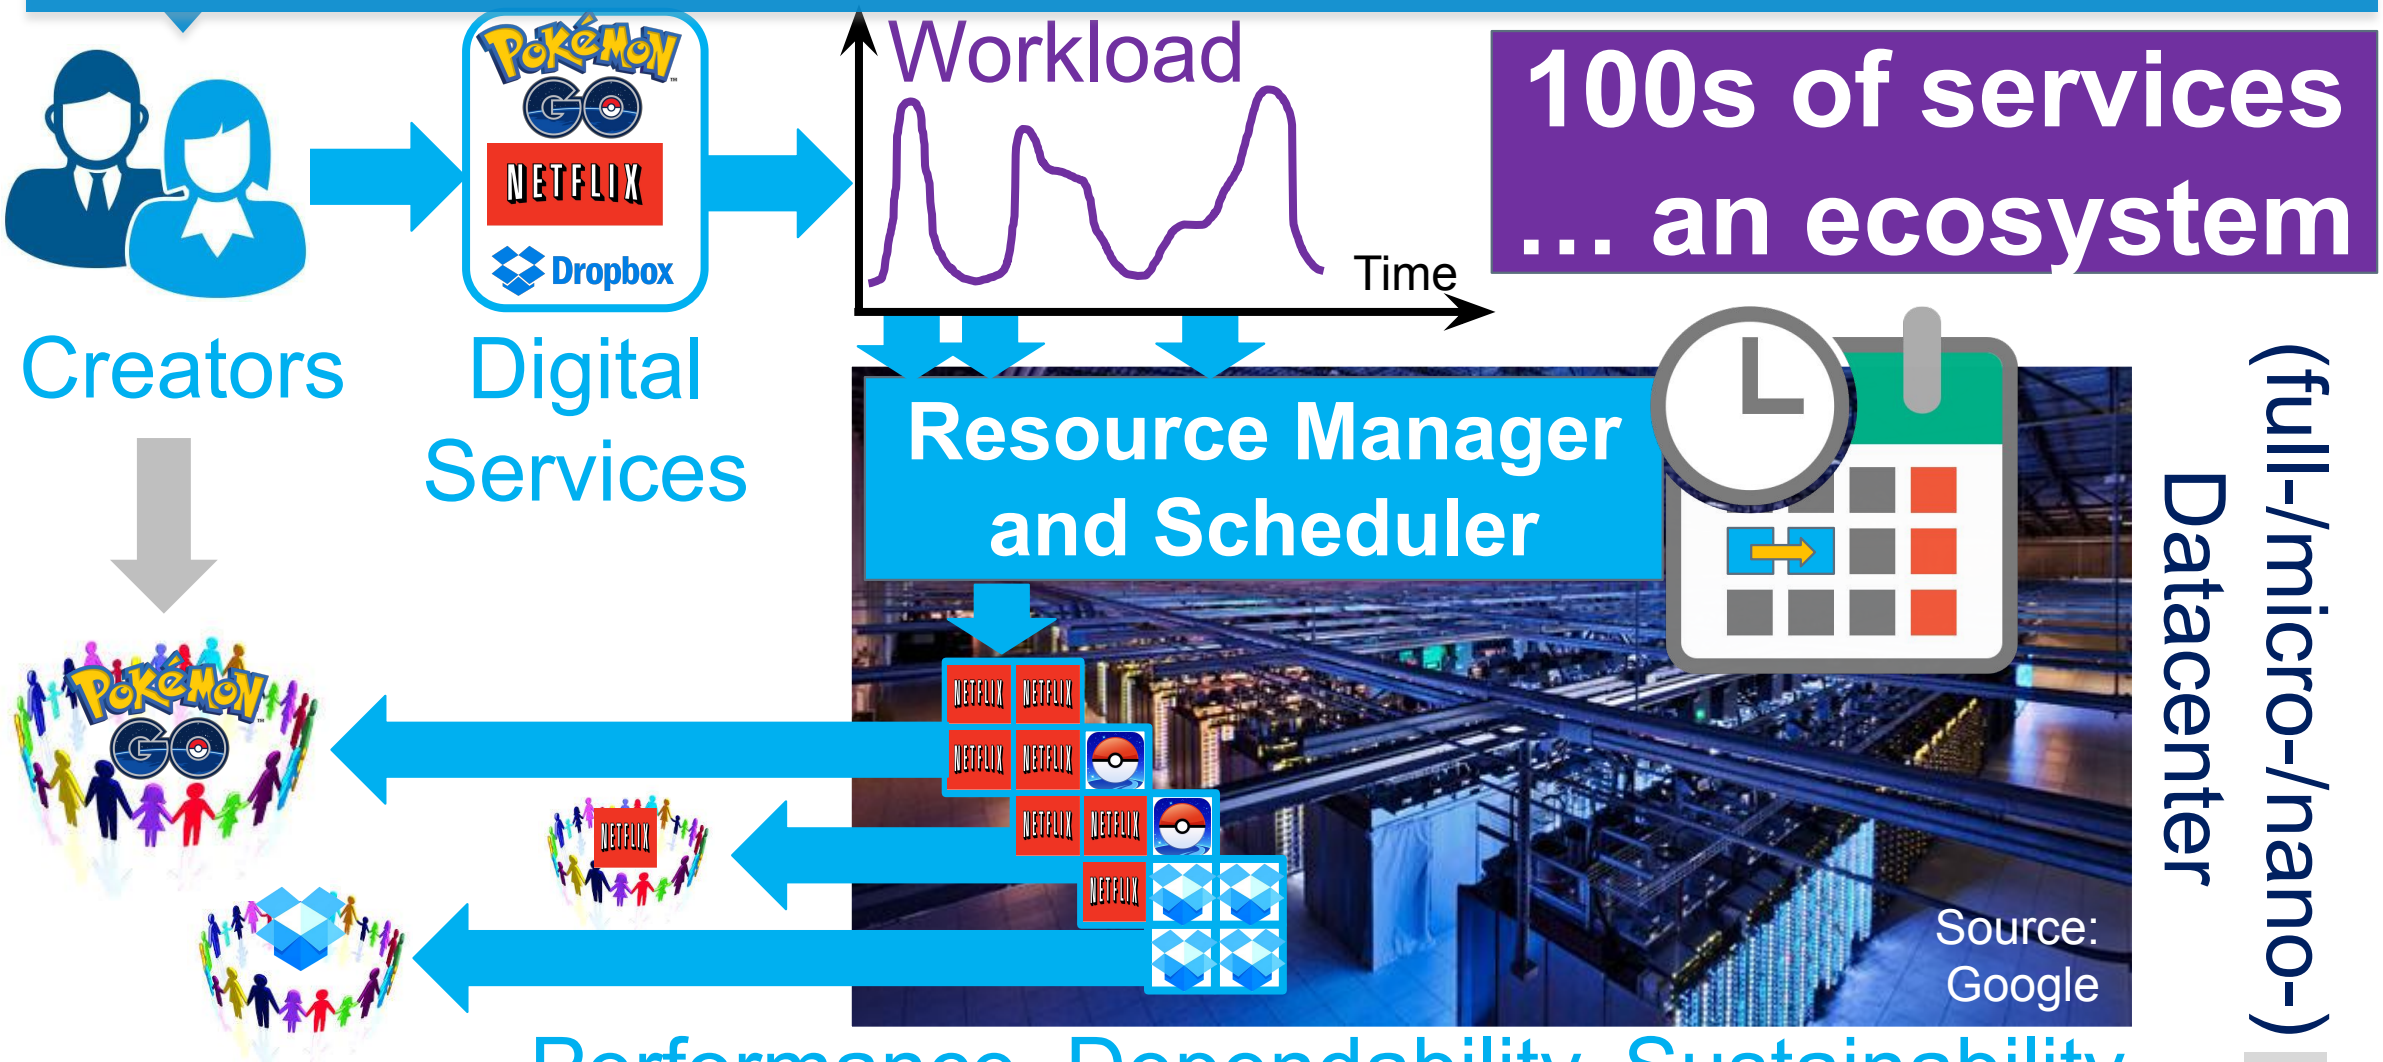

ECOSYSTEM = SERVICES + COMPUTING + SMARTS + GOALS

Creators

Digital Services

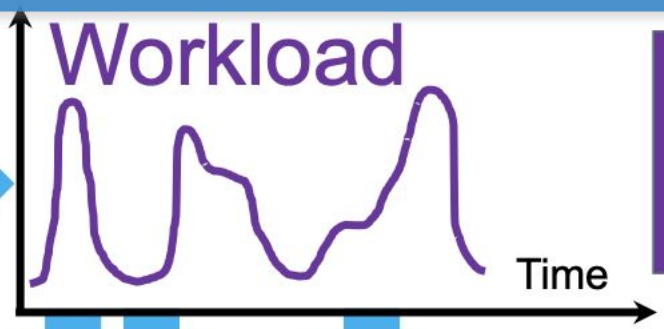

100s of services … an ecosystem

Resource Manager and Scheduler

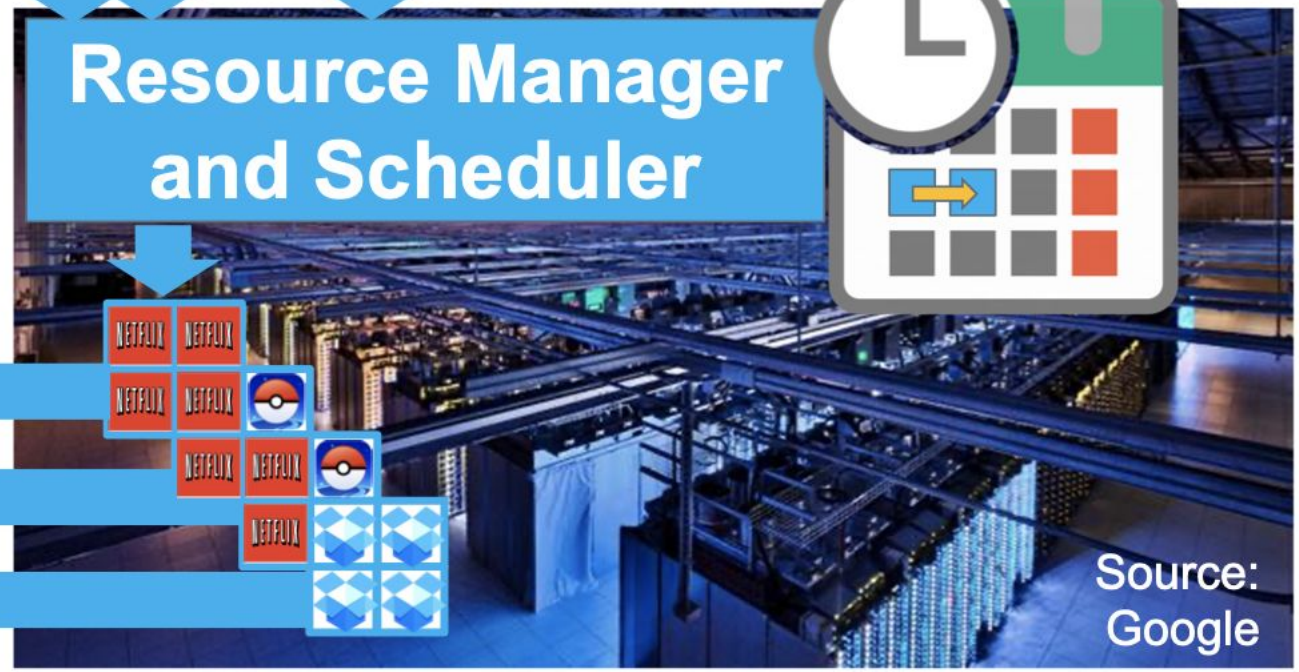

Computing Continuum

Extreme Automation, Performance, Dependability, Sustainability

50+ PLATFORMS … EMERGENT FEATURES

3A

THE COMPLEXITY CHALLENGE

REFERENCE ARCHITECTURE OF FAAS PLATFORMS

Workflow Composition Layer

- W1 Workflow Registry
- W2 Workflow Engine
- W3 Workflow Scheduler
- W4 Workflow Execution Store

Function Management Layer

- F1 Function Registry
- F2 Function Builder
- F3 Function Deployer
- F4 Function Instance
- F5 Function Router
- F6 Function Autoscaler

Resource Orchestration Layer

- R1 Naming Service
- R2 Resource Manager
- R3 Resource Scheduler
- R4 Node Agent

Business Concerns

Operational Concerns

Erwin van Eyk

[van Eyk et al. (2019) Serverless is More: From PaaS to Present Cloud Computing, IEEE Internet Computing] [Online]

SERVERLESS AI/ML/DL OPERATIONS

3B

ISSUES: COMPLEXITY, NON-TECHNICAL

IOSUP ET AL. 2021

Actual ML app is a very small part!

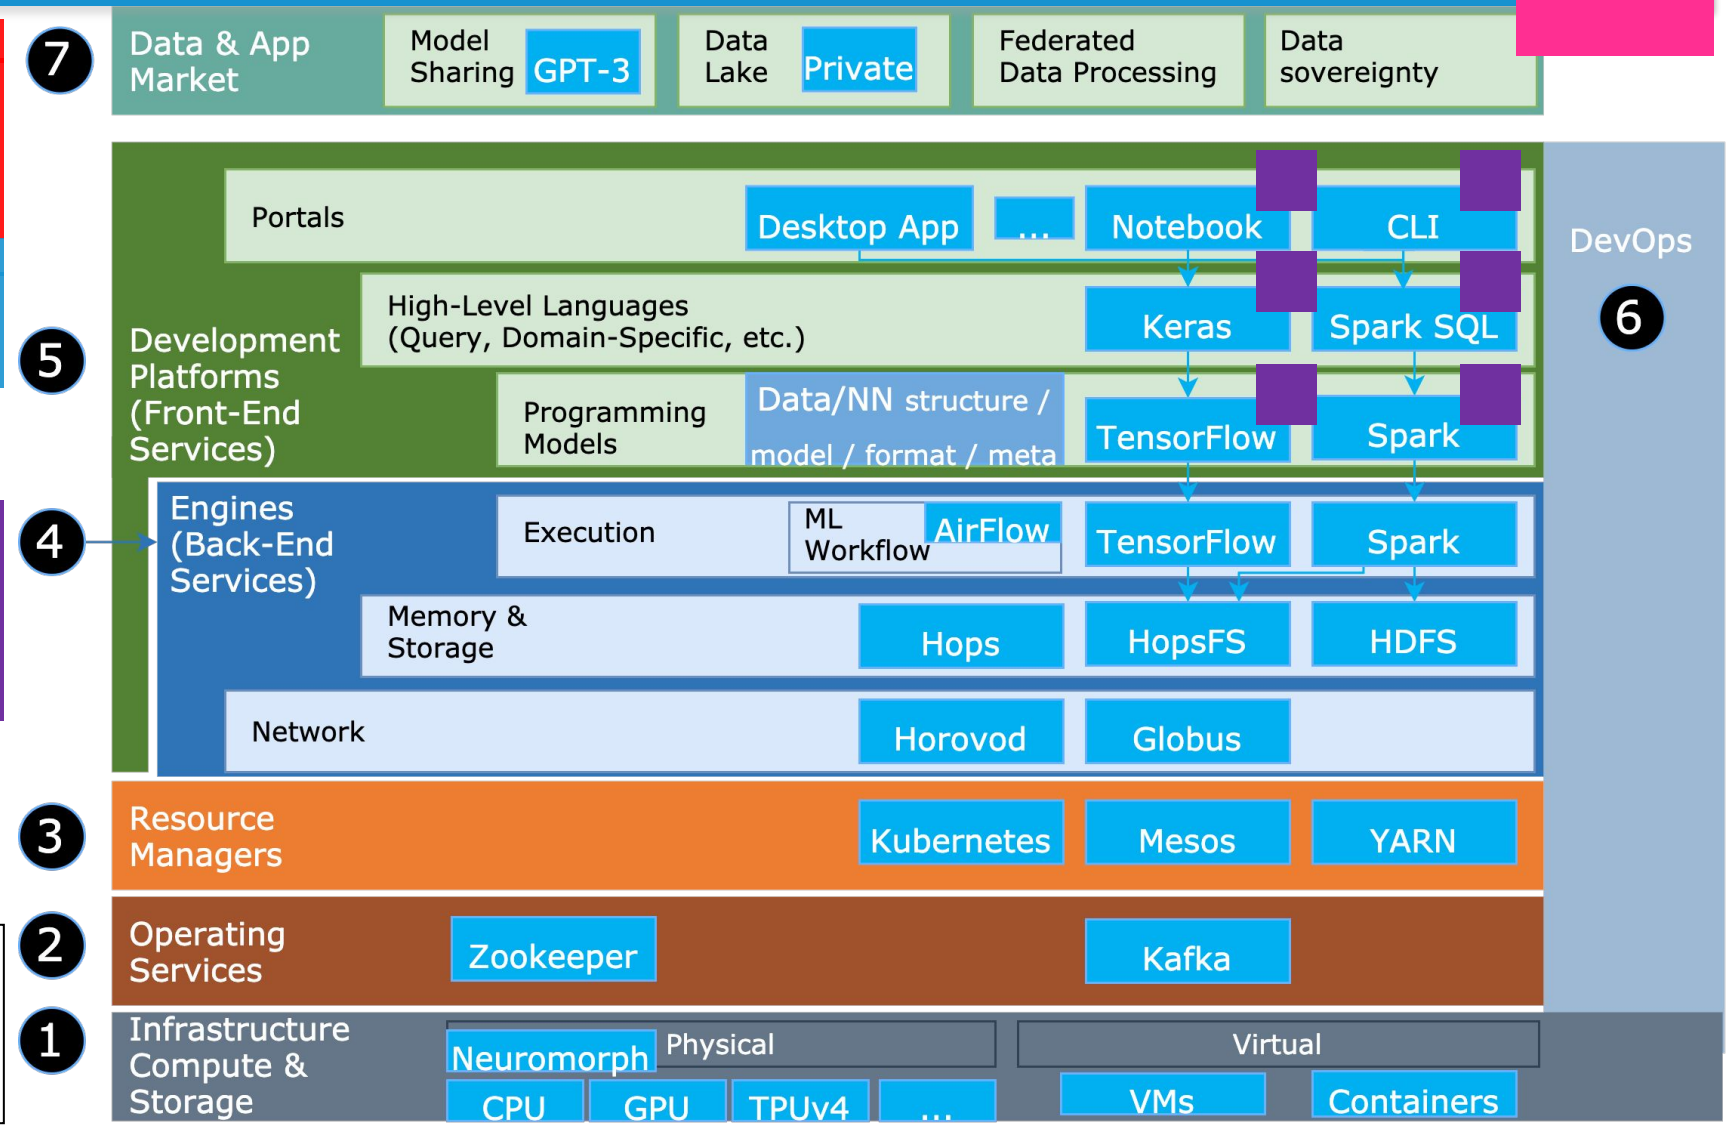

Adapted from:

Sakr, Bonifati, Voigt, Iosup, et al. (2021) The Future Is Big Graphs! CACM

SERVERLESS … WHAT COULD BE THE BENEFITS?

3C

TOO COSTLY TO CONDUCT REAL-WORLD EXPERIMENTS, SO WE BUILT A SIMULATOR

OpenDC
simulator

Learn more:
opendc.org

- Short-term resource management
- Long-term capacity planning
- Sophisticated model
- Support for many kinds of workloads and resources
- Validated for various scenarios
- Work with major NL hoster
- Used in training

Fabian Mastenbroek

and more...

4

Serverless should be concerned with modern non-functional properties, observed continuously, addressed short- and long-term

HOW TO ACHIEVE FINE-GRAINED BILLING AND UTILIZATION-BASED BILLING?

4A

THE COMPLEXITY CHALLENGE

REFERENCE VIEW ON OPERATIONAL METRICS

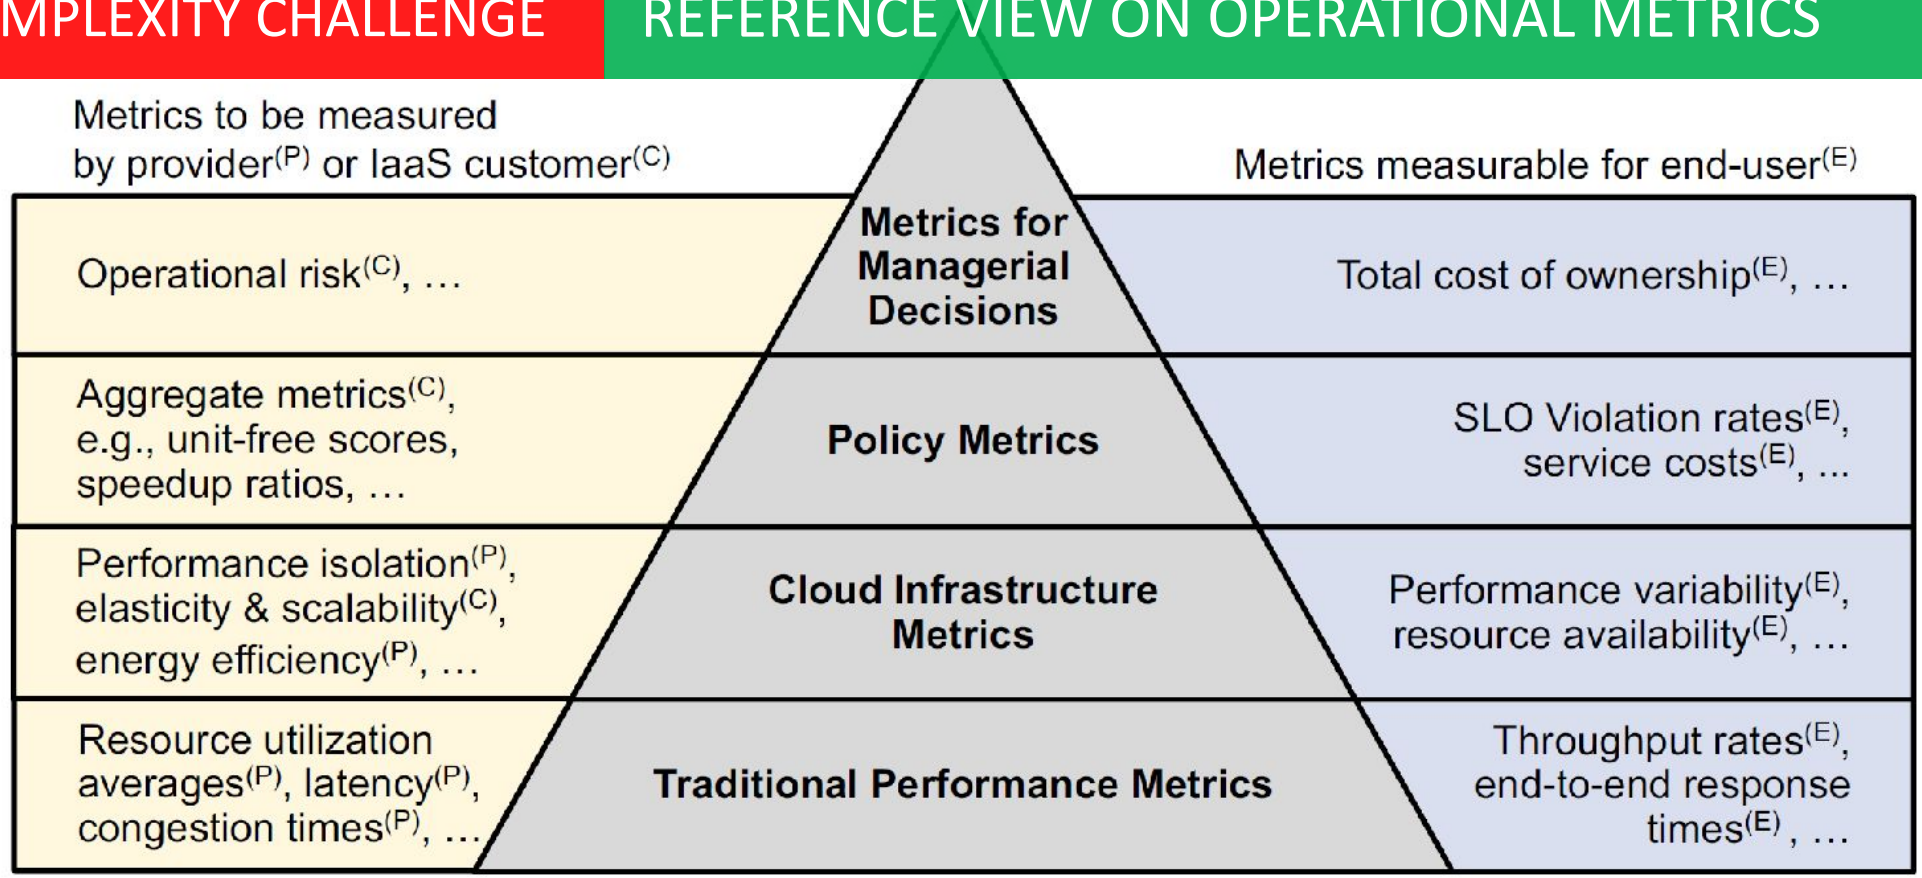

N. Herbst, E. Van Eyk, C. L. Abad, A. Iosup, et al. (2018) Quantifying Cloud Performance and Dependability: Taxonomy, Metric Design, and Emerging Challenges. TOMPECS 3(4): 19:1-19:36

HOW TO AUTOMATE X ACROSS THE ECOSYSTEM?

4B

IT'S OPERATIONS!

REFERENCE VIEW ON OPERATIONAL TECHNIQUES

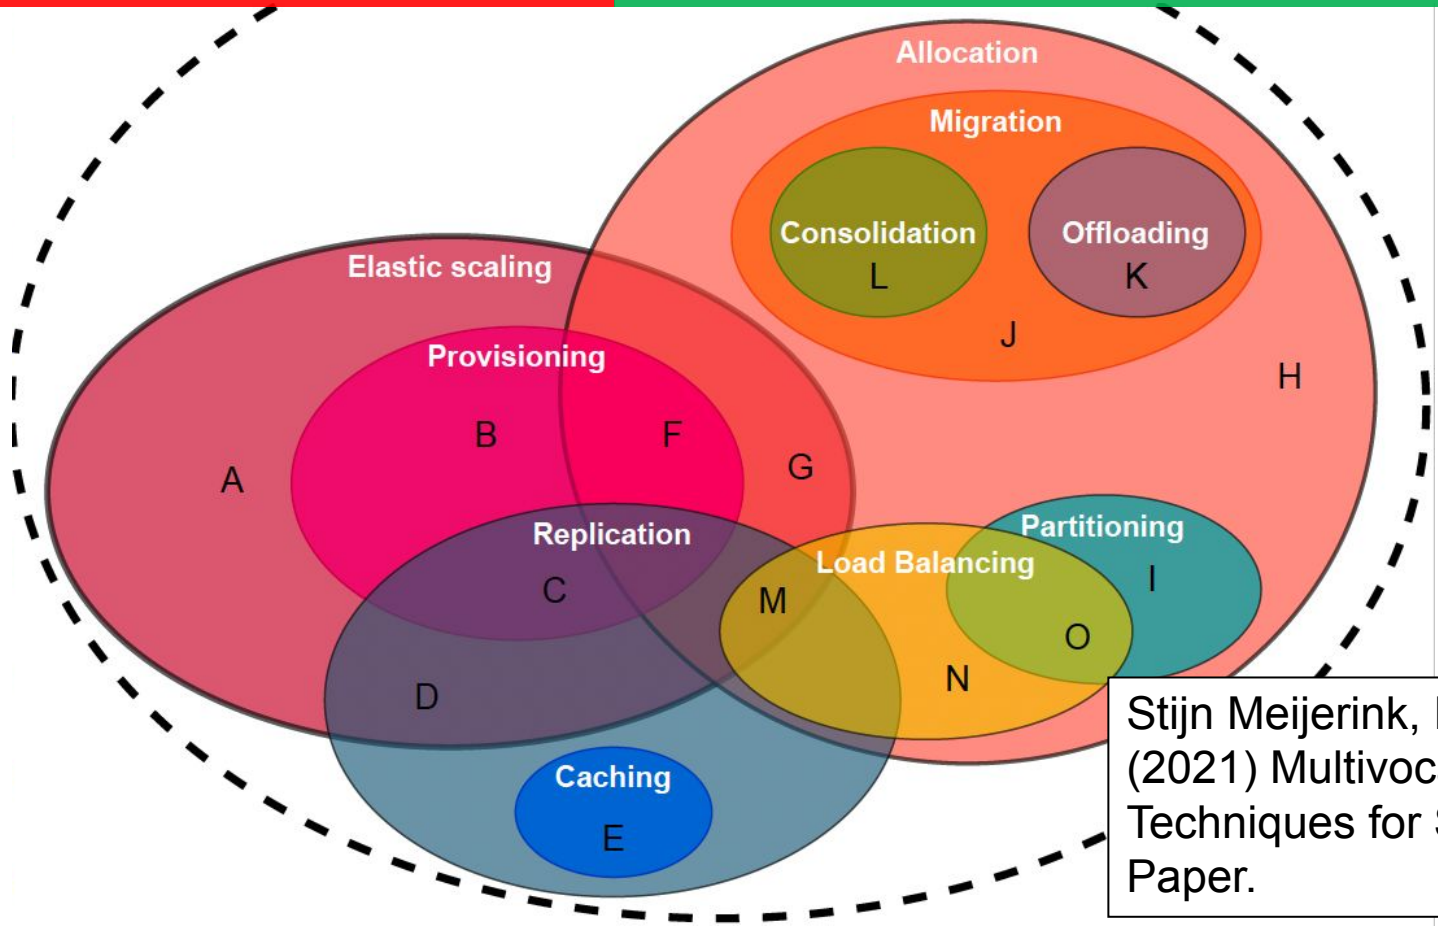

Stijn Meijerink, Erwin van Eyk, Alexandru Iosup (2021) Multivocal Survey of Operational Techniques for Serverless Computing. White Paper.

VU

SERVERLESS STREAMING WORKFLOWS

Erwin van Eyk

4b

DESIGN AND ENGINEERING: SERVERLESS ARCHITECTURE, API, SCHEDULER

One of the first serverless workflow management engine, part of Fission.io

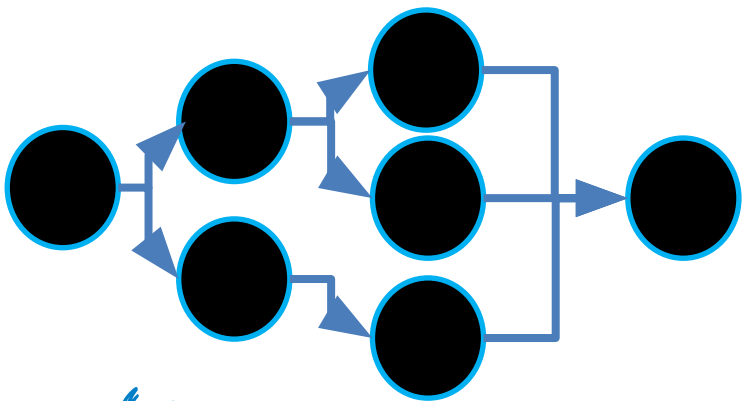

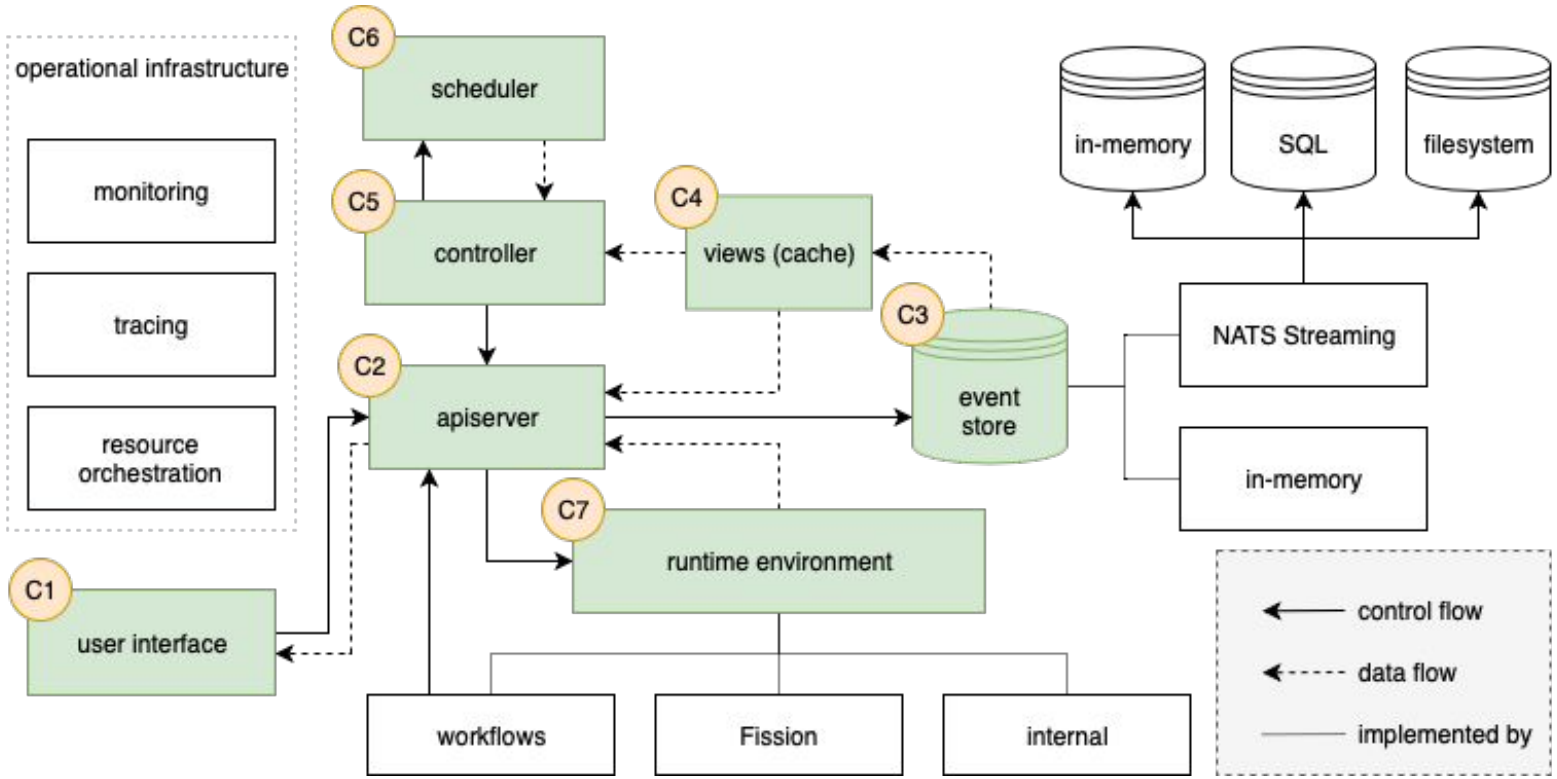

SERVERLESS STREAMING WORKFLOWS

Erwin van Eyk

4b

DESIGN AND ENGINEERING: SERVERLESS ARCHITECTURE, API, SCHEDULER

Fission Workflows delivers good performance, which also lowers cost

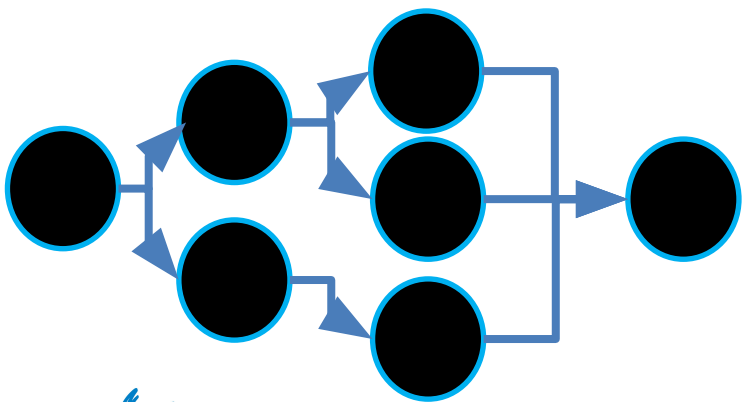

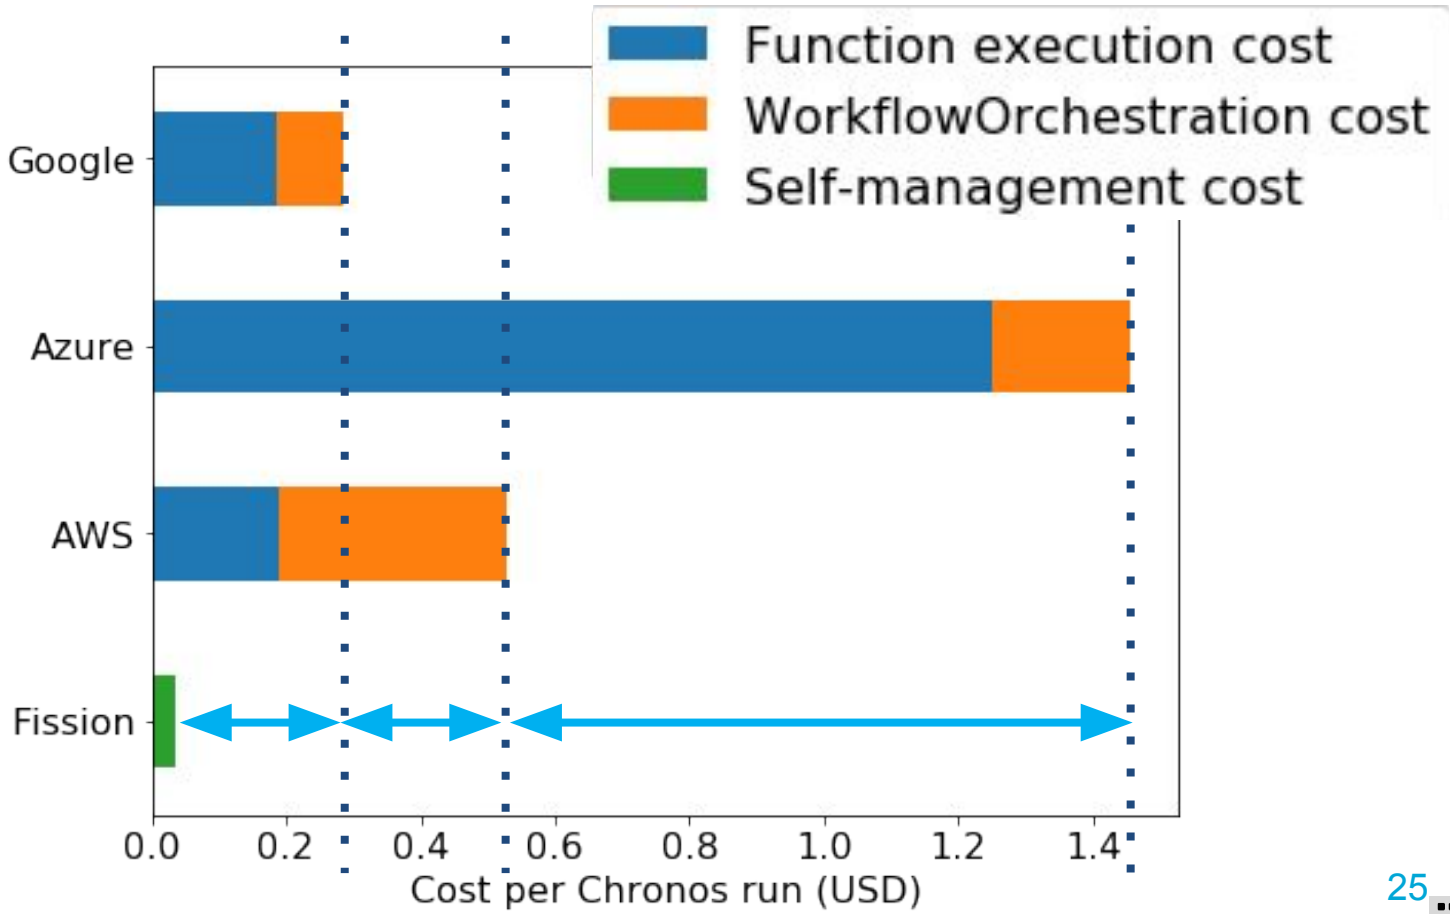

4c

Do Operators Tell the Truth and Nothing But the Truth? Characterization of User Reports

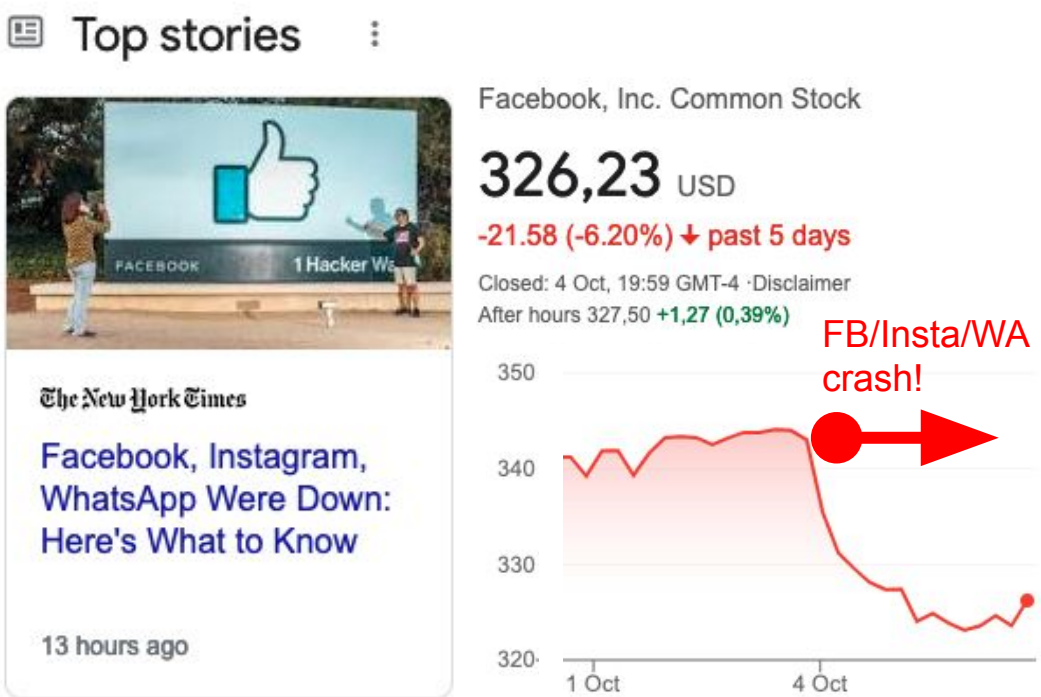

Source: NYTimes, Google market tracker, Oct 5, 2021

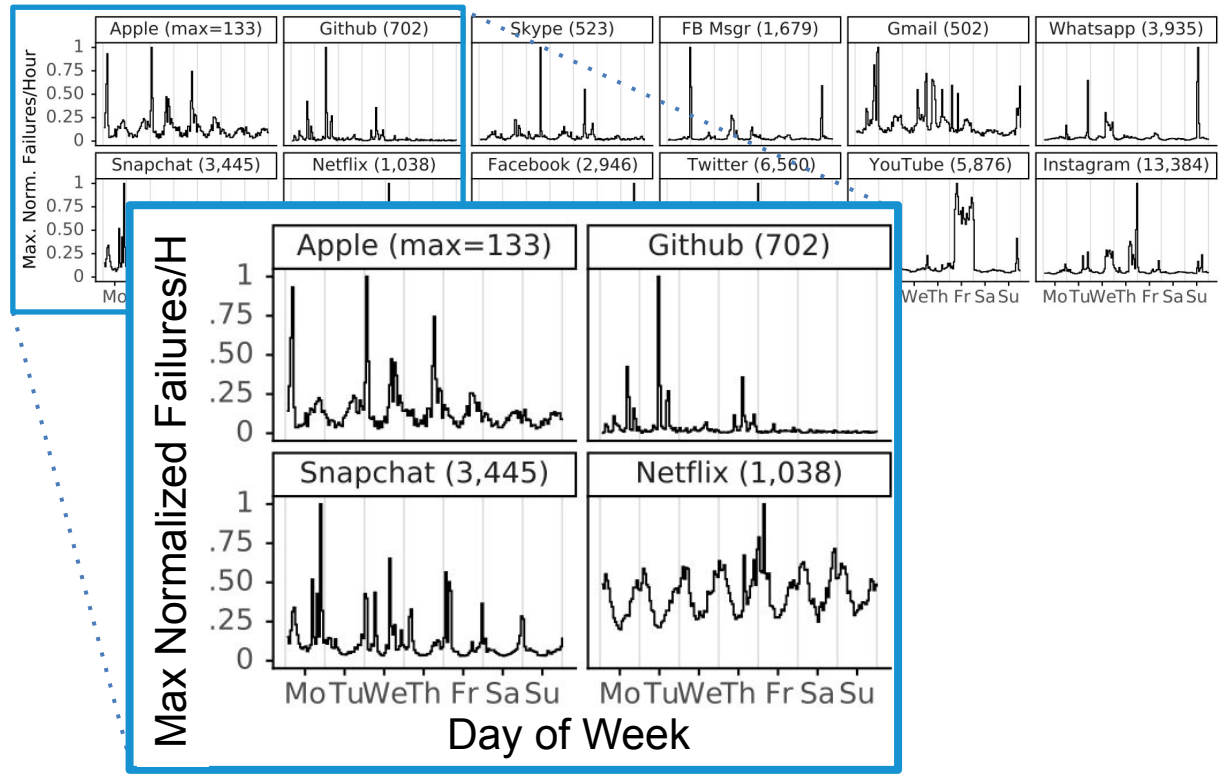

Talluri, Overweel, Versluis, Trivedi, Iosup (2021) Empirical Characterization of User Reports about Cloud Failures. ACSOS

SAMPLE VIEW OF THE DATA

4c

DO OPERATOR REPORTS SAY THIS?

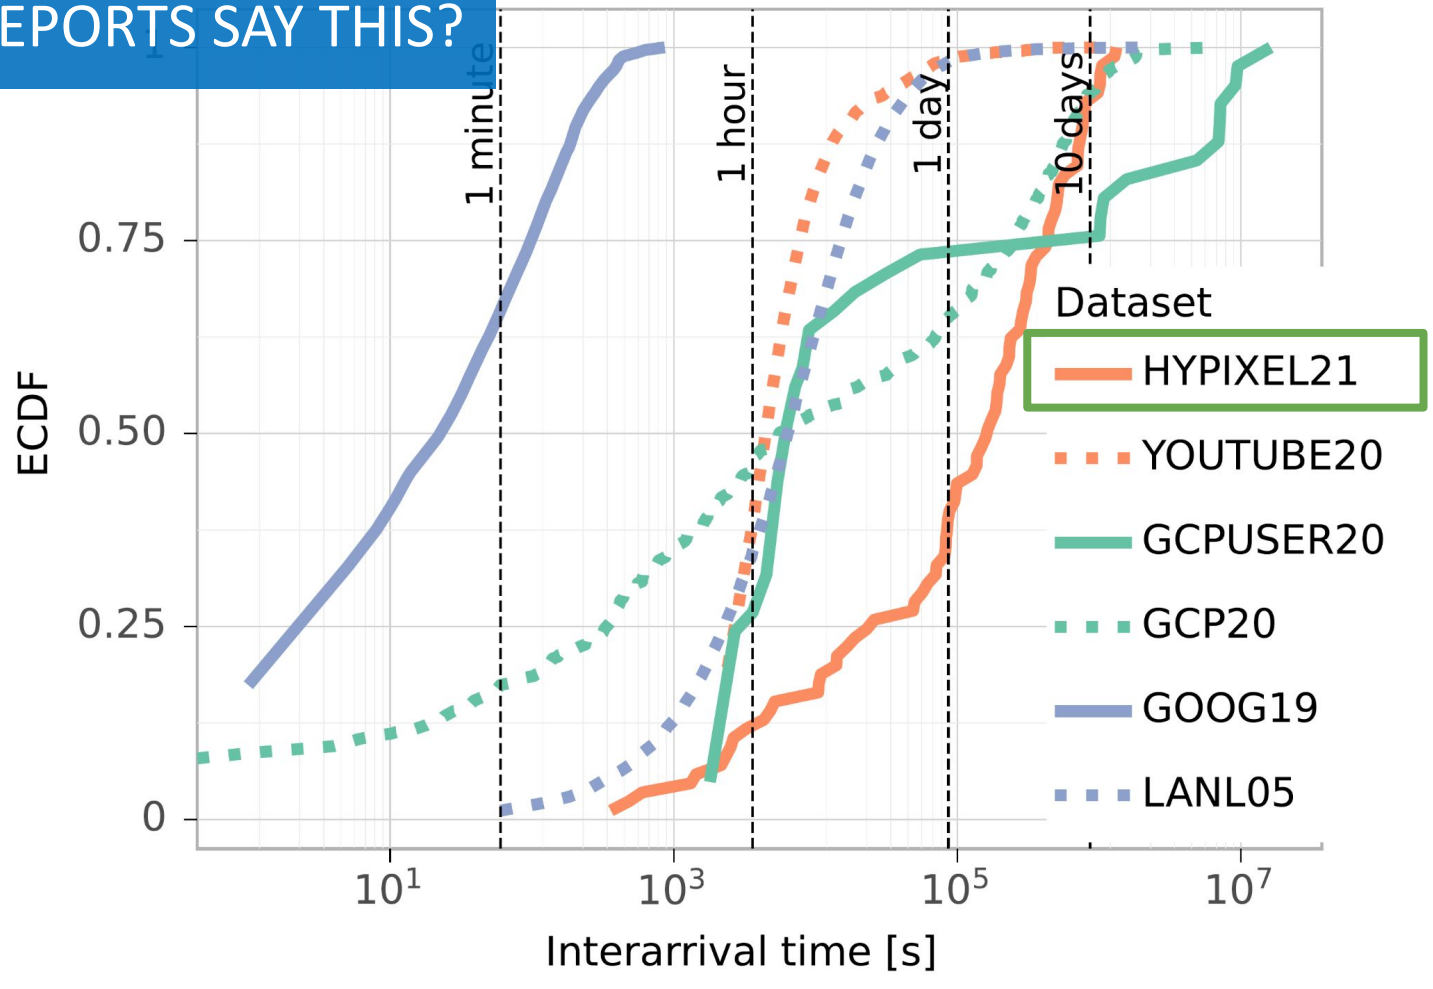

Games, Minecraft

SAMPLE VIEW OF THE DATA

4c

DO OPERATOR REPORTS SAY THIS?

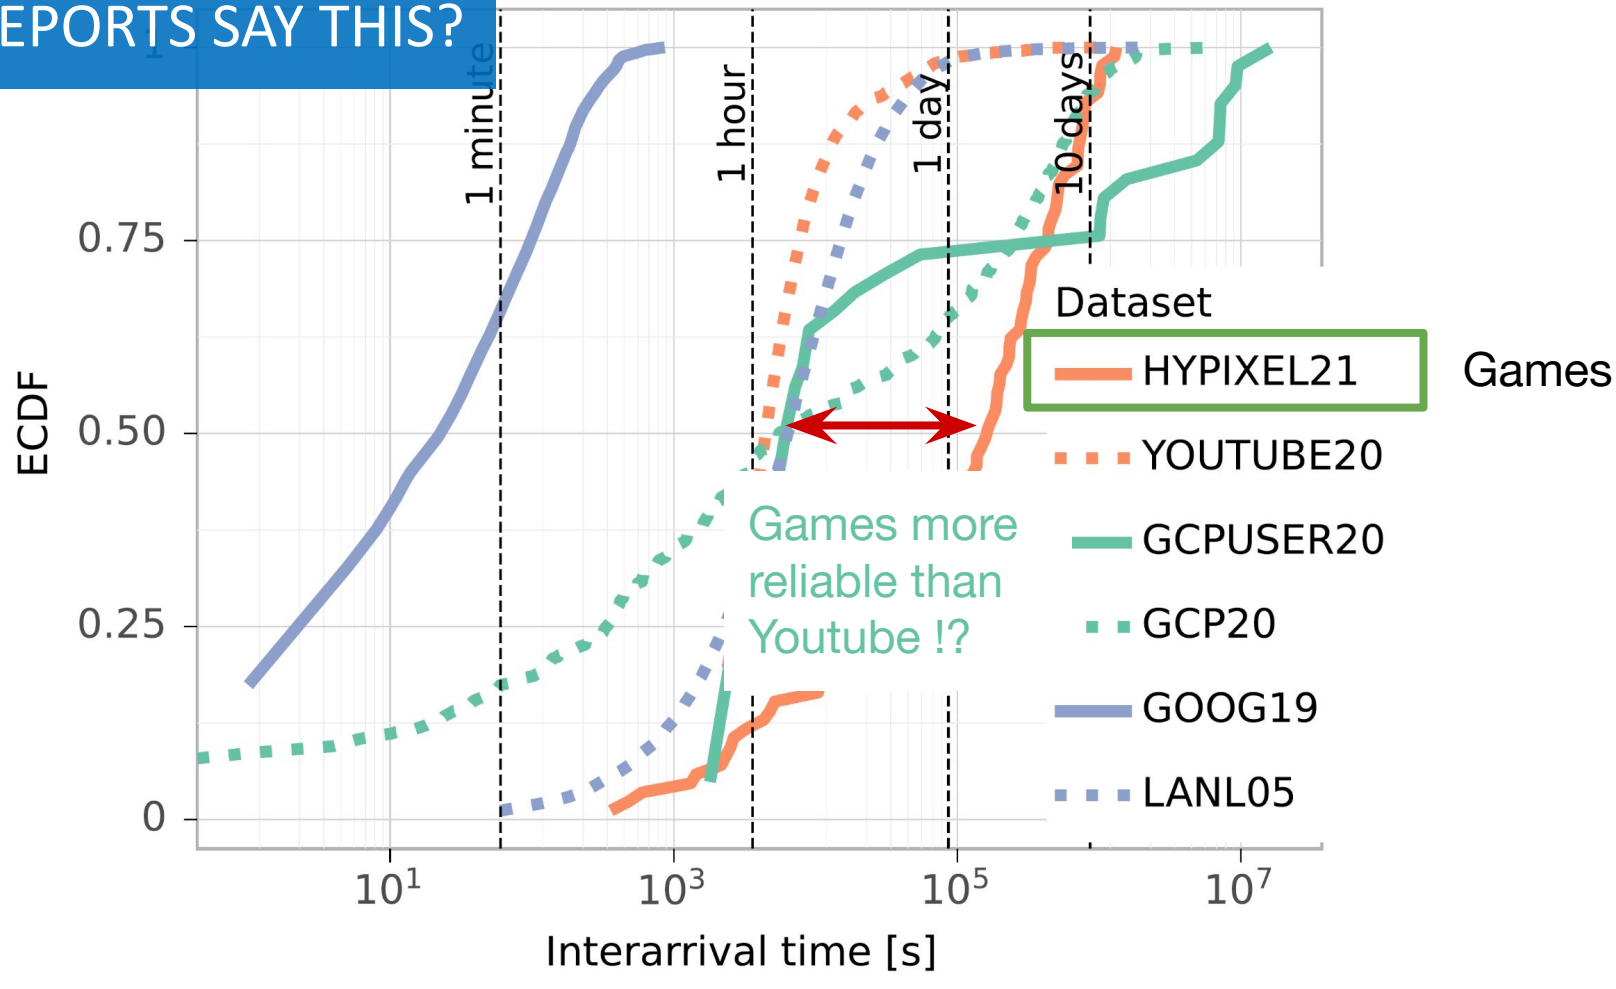

SAMPLE VIEW OF THE DATA

4c

DO OPERATOR REPORTS SAY THIS?

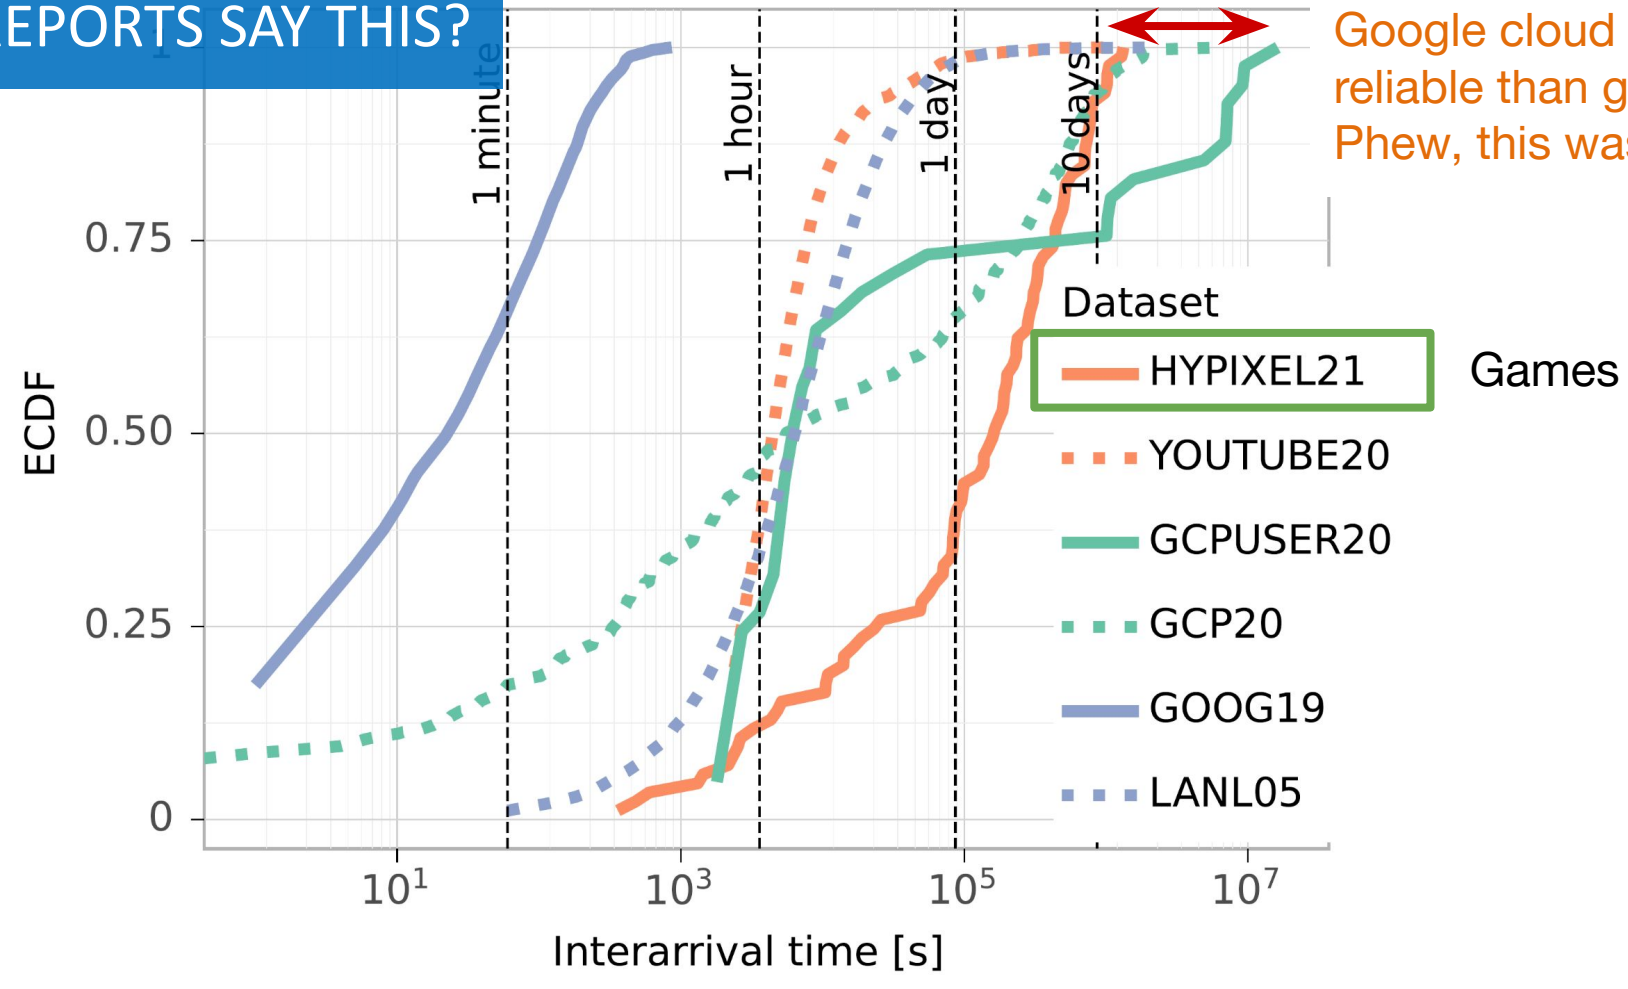

Google cloud more reliable than games. Phew, this was close!

Serverless should be aligned with funding and community concerns

5

ALIGNED WITH COMMUNITY CONCERNS…

The Manifesto on Computer Systems and Networking Research

Clear vision for the field in the NL, 2021-2035

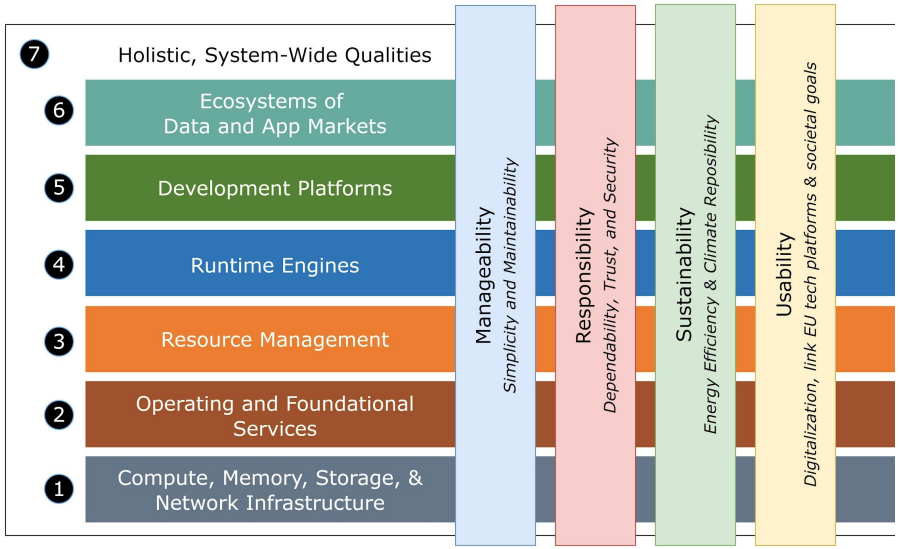

Signed

50+ PIs / Leads
7 universities
5 relevant societal stakeholders

Available

Full version (40+ pages) https://arxiv.org/pdf/2206.03259
Who's Who in CompSysNL? https://bit.ly/CompSysNLWhosWho

ALIGNED WITH FUNDING OPPORTUNITIES…

Big Graph Processing: Used in AI/ML, Sci/Pharma, Learning/Edu, Energy Mgmt., etc.

Vision: Massivizing computer systems approaches are key to enable big graph ecosystems

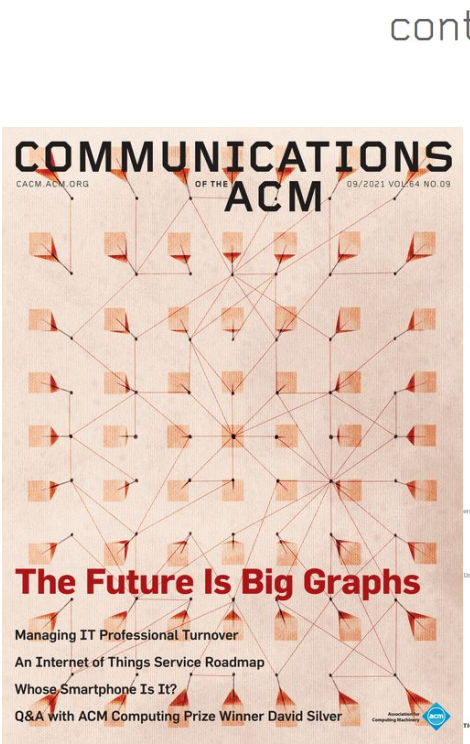

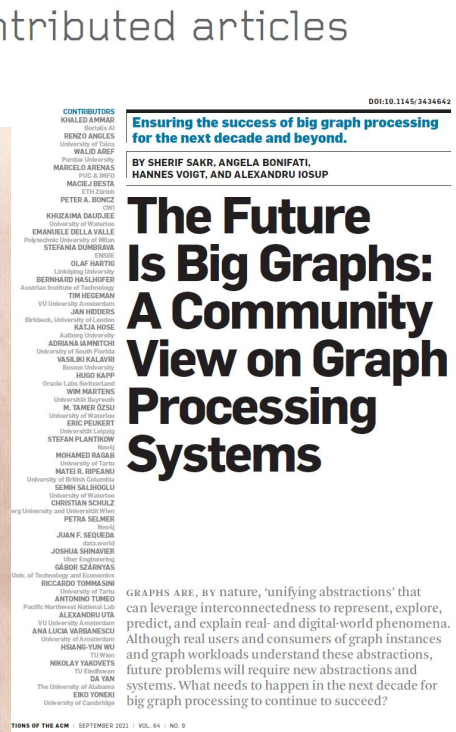

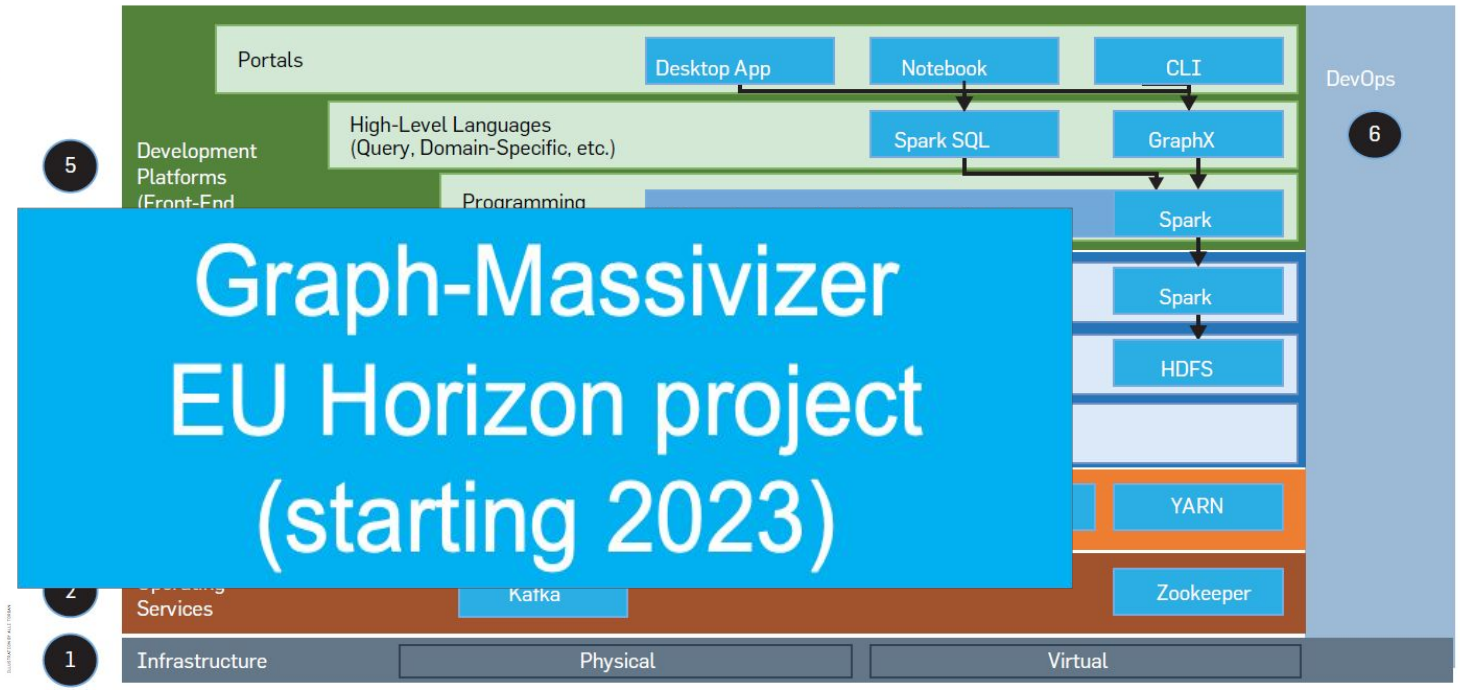

CACM Cover/Featured article, Sep 2021

Sakr, Bonifati, Voigt, Iosup, et al. (2021) The Future Is Big Graphs! CACM.

TAKE-HOME MESSAGE

Serverless = Extreme automation + fine-grained reporting + utilization-based billing

The serverless ecosystem: many apps, many platforms, may goals, many approaches

Many modern, open challenges: scheduling, telemetry, recovery, privacy/GDPR, etc.

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING

https://atlarge-research.com/publications.html

1. Iosup et al. Massivizing Computer Systems. ICDCS 2018 ← start here
2. Andreadis et al. A Reference Architecture for Datacenter Scheduling, SC18
3. Van Eyk et al. Serverless is More: From PaaS to Present Cloud Computing, IEEE IC Sep/Oct 2018
4. Uta et al. Exploring HPC and Big Data Convergence: A Graph Processing Study on Intel Knights Landing, IEEE Cluster 2018

5. Talluri et al. Big Data Storage Workload in the Cloud. ACM/SPEC ICPE 2019.
6. Toader et al. Graphless. IEEE ISPDC'19.
7. Jiang et al. Mirror. CCPE 2018.
8. Ilyushkin et al. Autoscalers. TOMPECS 2018.
9. Versluis et al. Autoscaling Workflows. CCGRID'18.
10. Uta et al. Elasticity in Graph Analytics? IEEE Cluster 2018.

11. Herbst et al. Ready for rain? TOMPECS 2018.
12. Guo et al. Streaming Graph-partitioning. JPDC'18.
13. Iosup et al. The OpenDC Vision. ISPDC'17.
14. Iosup et al. Self-Aware Computing Systems book.
15. Iosup et al. LDBC Graphalytics. PVLDB 2016.

Etc.

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING

https://atlarge-research.com/publications.html

1. Iosup et al. The AtLarge Vision on the Design of Distributed Systems and Ecosystems. ICDCS 2019 ← Start here
2. Uta et al. Is big data performance reproducible in modern cloud networks? NSDI 2020
3. Van Eyk et al. The SPEC-RG Reference Architecture for FaaS: From Microservices and Containers to Serverless Platforms, IEEE IC 2019
4. Papadopoulos et al. Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. TSE 2019 and (journal-first) ICSE 2020

5. van Beek et al. Portfolio Scheduling for Managing Operational and Disaster-Recovery Risks in Virtualized Datacenters Hosting Business-Critical Workloads. ISPDC 2019
6. van Beek et al. A CPU Contention Predictor for Business-Critical Workloads in Cloud Datacenters. HotCloudPerf19

\+ Iyushkin et al. Performance-Feedback Autoscaling with Budget Constraints for Cloud-based Workloads of Workflows. Under submission

Etc.

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING

https://atlarge-research.com/publications.html

1. Sakr, Bonifati, Voigt, Iosup, et al. (2021) The Future Is Big Graphs! CACM
2. Andreadis et al. (2021) Capelin: Data-Driven Capacity Procurement for Cloud Datacenters using Portfolios of Scenarios. TPDS, under review.
3. Versluis et al. The Workflow Trace Archive: Open-Access Data From Public and Private Computing Infrastructures. TPDS 2020.
4. Eismann et al. Serverless Applications: Why, When, and How? IEEE Softw. 38(1): 32-39 (2021)
5. Uta et al. (2020) Beneath the SURFace: An MRI-like View into the Life of a 21st-Century Datacenter. login USENIX
6. Iosup, Hegeman, et al. (2020) The LDBC Graphalytics Benchmark. CoRR. https://arxiv.org/abs/2011.15028
7. Hegeman et al. (2021) GradeML. HotCloudPerf.
8. Abad, Iosup, et al. An Analysis of Distributed Systems Syllabi With a Focus on Performance-Related Topics. WEPPE 2021. https://arxiv.org/abs/2103.01858

Etc.

EXTRAS

THIS IS THE GOLDEN AGE OF MASSIVE COMPUTING ECOSYSTEMS

Do you recognize this App?

Here is how it operates…

A TYPICAL ECOSYSTEM: SERVICE, DATACENTER, SCHEDULER

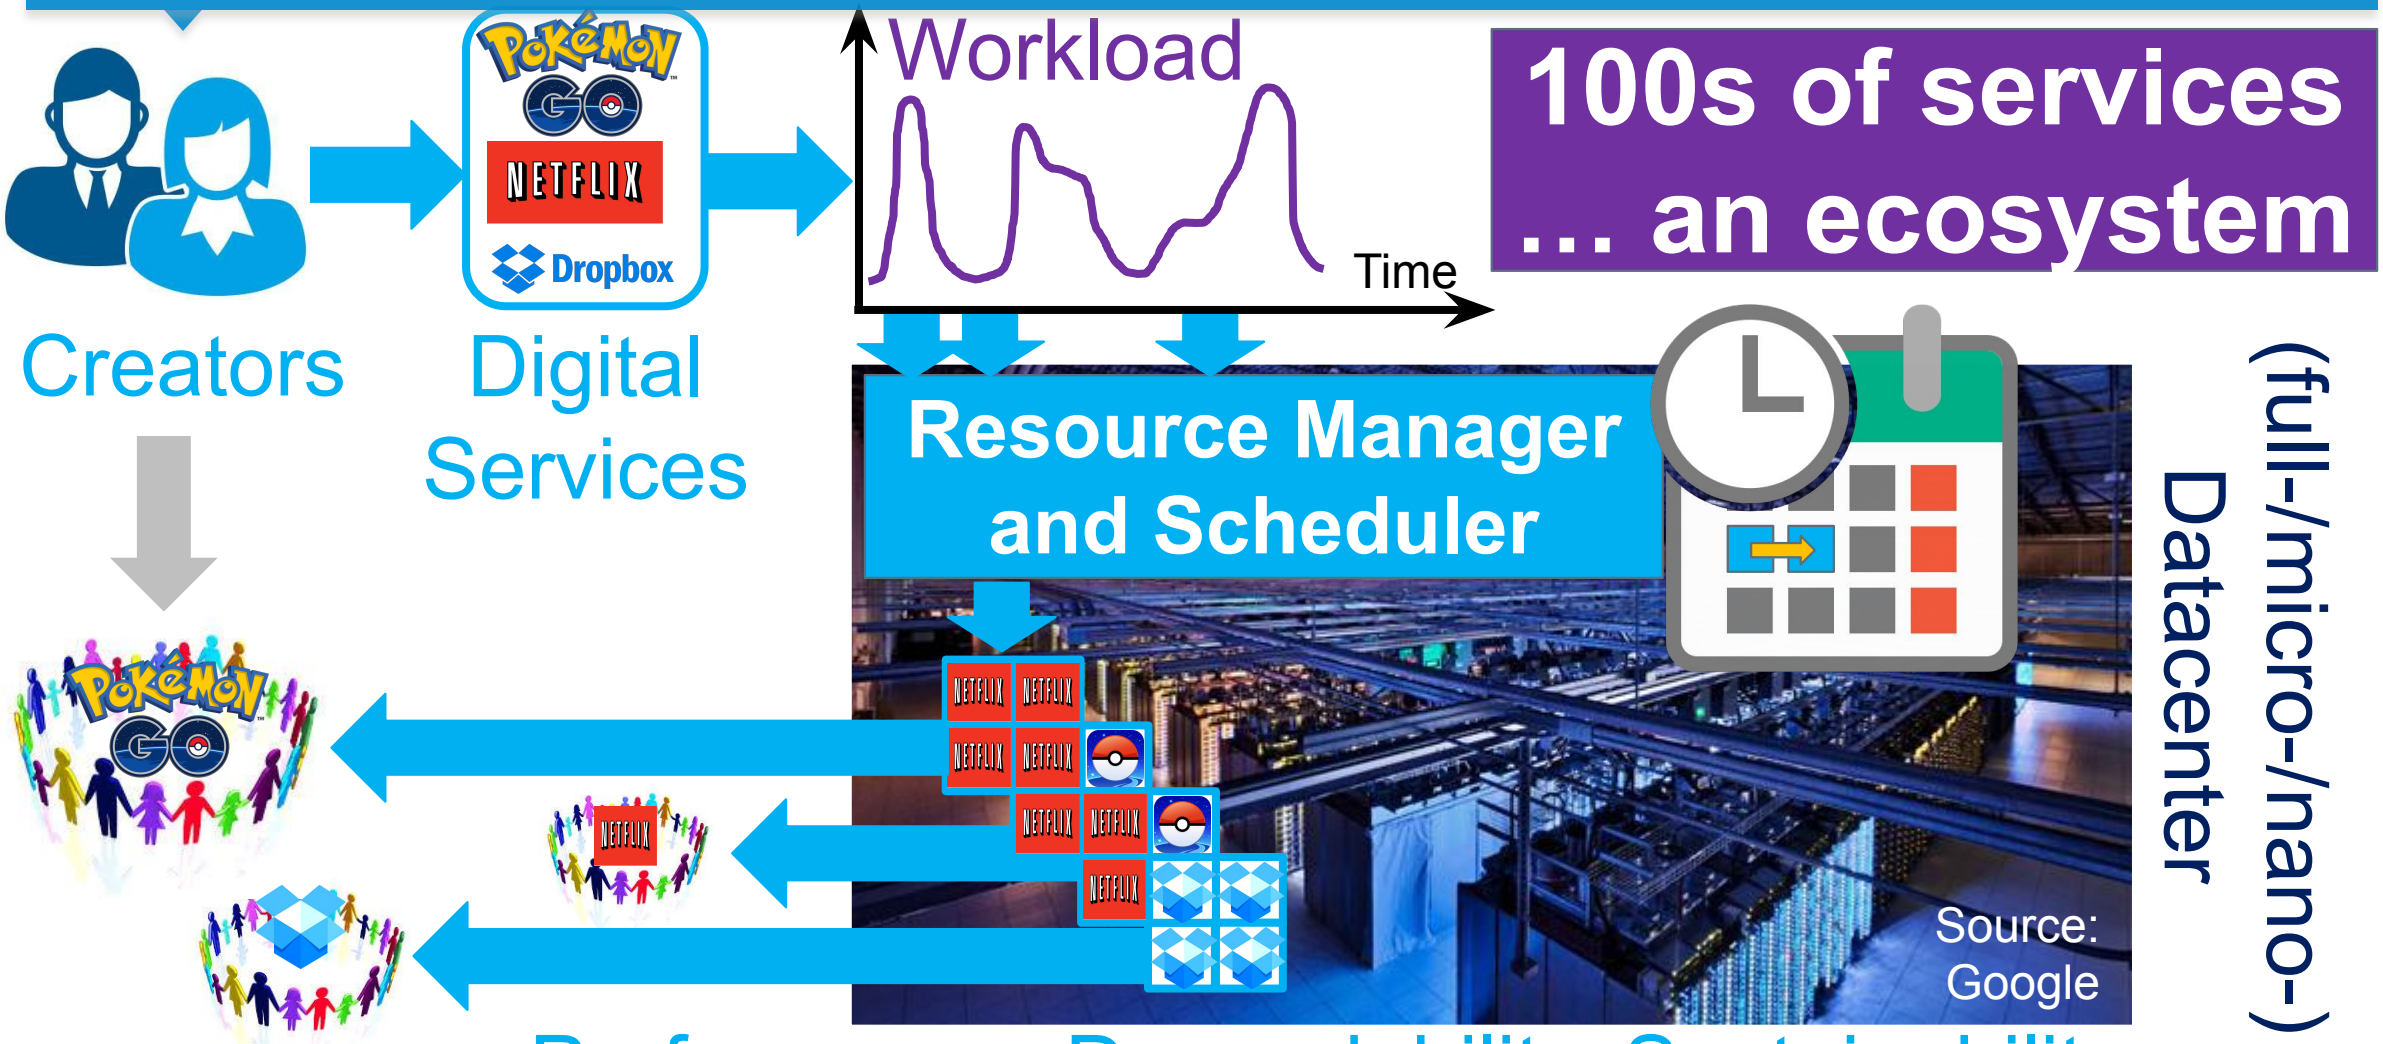

DISTRIBUTED ECOSYSTEMS, OUR DEFINITION

1. Set of 2+ constituents, often heterogeneous
2. Each constituent is a system or an ecosystem (recursively)
3. Constituents are autonomous, cooperative or in competition
4. Ecosystem structure and organization ensure responsibility
    1. Completing functions and providing services
    2. Providing desirable non-functional properties
    3. Fulfill agreements with both operators and clients, clients in the loop
5. Long and short-term dynamics occur in the ecosystem

Iosup et al., Lecture Notes in Distributed

Iosup et al., Modeling Computer Systems, ICDCS

MASSIVIZING COMPUTER SYSTEMS

A LARGER VISION OF HOW COMPUTING WILL HELP OUR SOCIETY

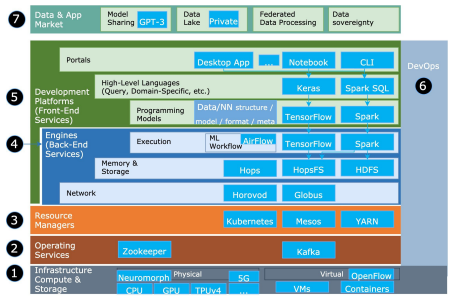

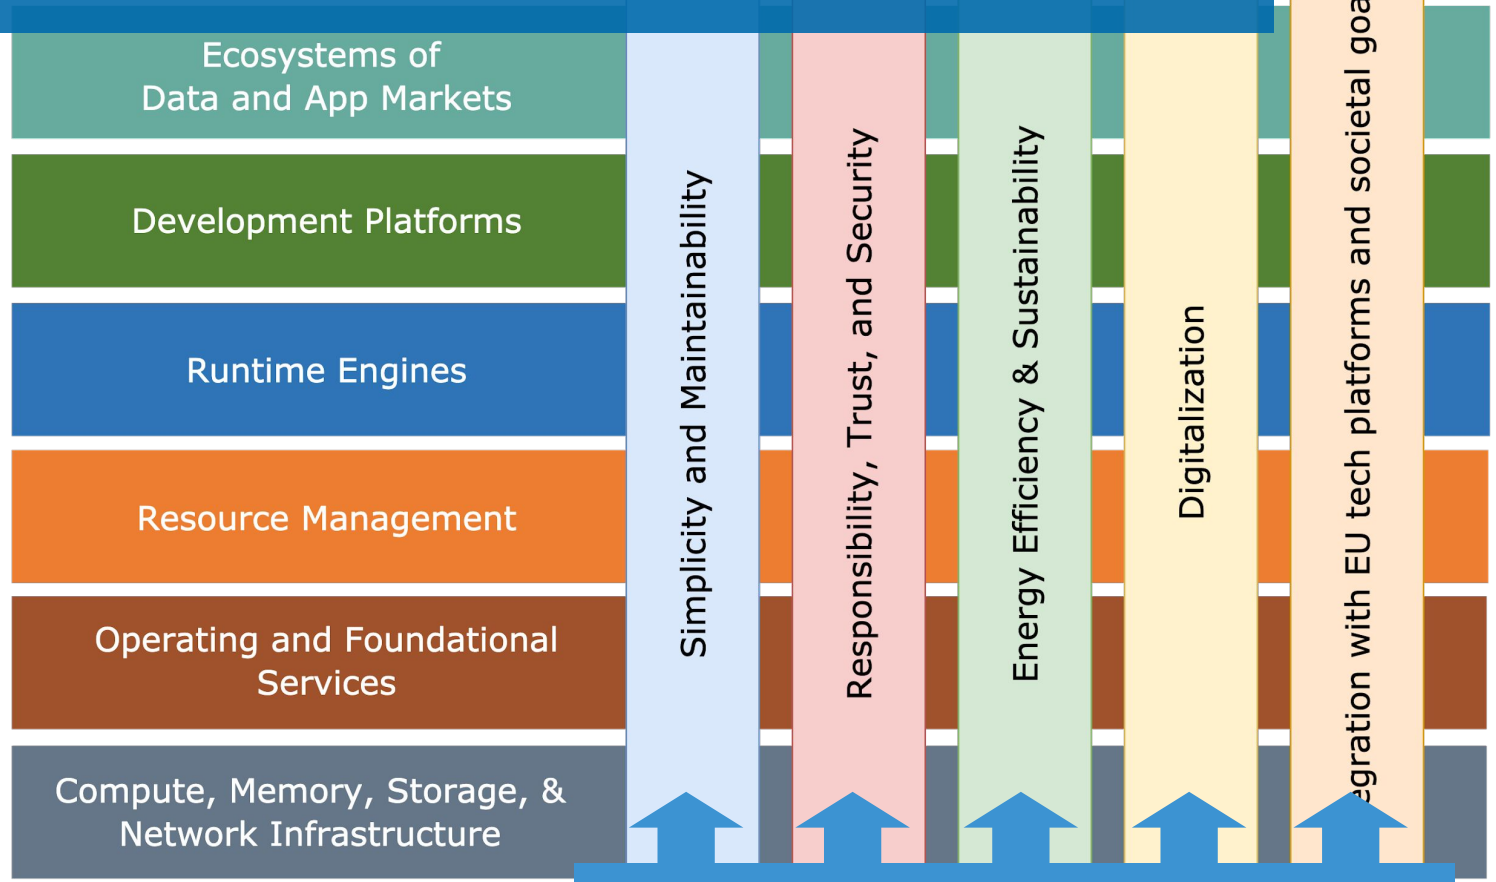

A.Iosup@vu.nl
http://atlarge.science