DeiC Conference 2021

“What is the optimal setup for HPC Type 1-4 in a changing world?”

Microsoft

Agenda

1. Our commitments to research
2. Support and accelerate HPC
3. “What is the optimal setup for HPC Type 1-4 in a changing world?”

Jan Cordtz, Microsoft Denmark

Snr. Cloud Solution Architect

Our commitment to Denmark and Danish research

DEIC
Fourth year as participant

Close to educational institutes
Large exposure and usage of products and services

Danish Datacenter
Building a local datacenter

Committed to learn and help
"Don't be a know-it-all, be a learn-it-all"

Our sustainability commitments

Carbon negative
by 2030

Water positive
by 2030

Zero waste
by 2030

Build the
Planetary
Computer

Carbon negative by 2030; erase footprint by 2050

Ground our work in the best available science and the most accurate math

Take responsibility by achieving carbon negative by 2030

Fund $1B for carbon reduction and removal

Support and empower suppliers and all customers with new tools, products and partnerships

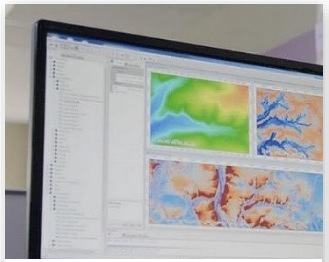

Work to advance transparency for reporting on emissions and removals

Use our voice on carbon-related public policy issues

Enlisting our employees to enable them to contribute to our efforts

We're taking action ourselves

We're empowering our customers

We're investing in broader innovation

We're supporting government action

Building a Planetary Computer

Our commitments to protecting global ecosystems

Putting data and digital technology to work with a Planetary Computer

Empowering customers

Using our voice on ecosystem-related public policy issues

Taking responsibility for our land footprint

Zero waste by 2030

Our plans

Achieve zero waste company-wide

Our campuses and datacenters will be zero waste certified and we're creating Surface device and product packaging with 100% recyclable materials. We'll eliminate the use of single use plastics from our packaging by 2025.

Establish Circular Centers

We're locating new Microsoft Circular Centers on every new major datacenter campus and existing datacenters to increase our reuse of servers and components up to 90% by 2025.

Invest in digitizing waste accounting

We'll invest in identifying opportunities to improve waste data collection and potential digital solutions for our operations.

Climate Innovation Fund investment

We're investing $30 million in Closed Loop Partners' funds to help accelerate innovation to build a more circular economy at scale.

Empower our customers

Microsoft technologies – from Power BI to AI – are being used today in several ways that will help reduce waste, from material development to disposal.

Enlist our employees

We're inviting our employees to participate in our waste reduction efforts by giving them visibility into the impact of their actions.

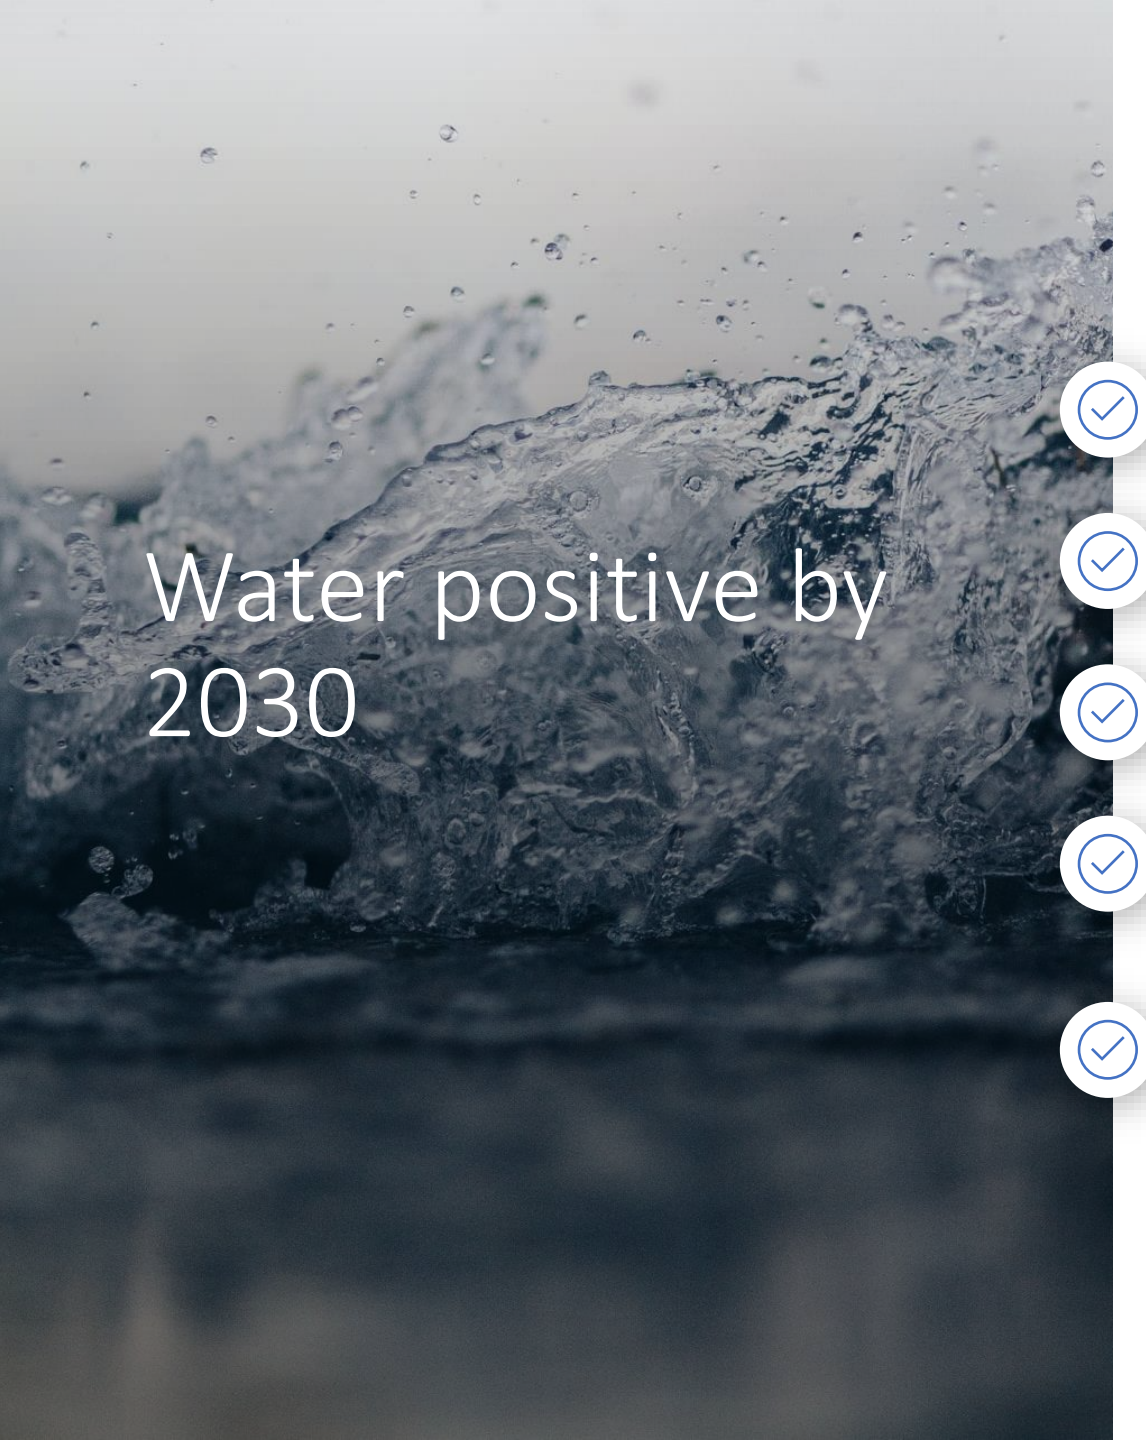

Water positive by 2030

- Digitize water data through AI for Earth and our partners
- Climate Innovation Fund investment
- Partner with Water.org
- Join Water Resilience Coalition, WaterEurope, and Center for the Fourth Industrial Revolution Norway
- Enlist our employees through volunteer opportunities

Be part of an evolution towards more research – rather than wait for an revolution.

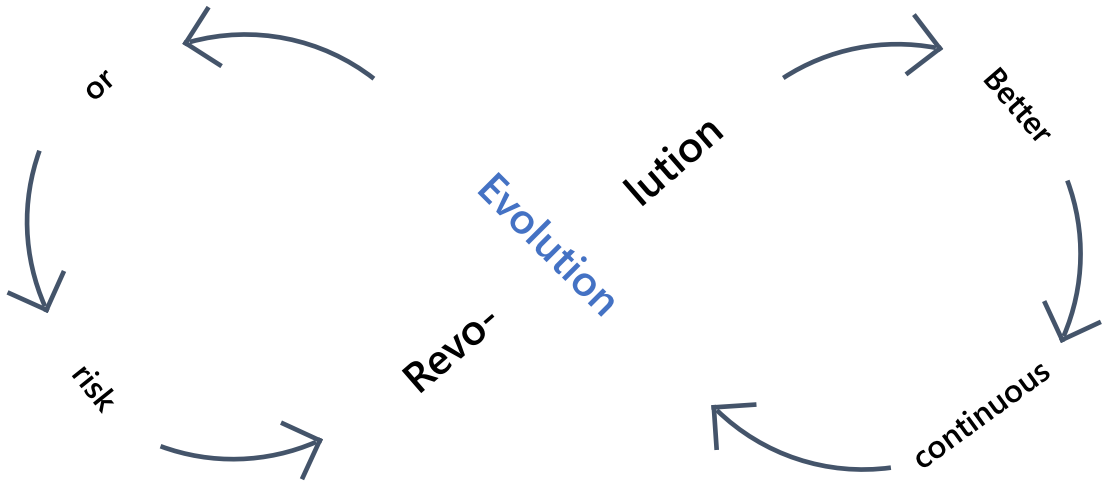

Microsoft Cloud - Azure

Largest geographical footprint of any cloud provider with more than 60 Azure regions

>3,500 full-time security professionals

6.5 trillion global signals daily

$1 billion annual cybersecurity investment

Compliant

91 Compliance offerings

GLOBAL

- ISO 27001:2013
- ISO 27017:2015
- ISO 27018:2014
- ISO 22301:2012
- ISO 9001:2015
- ISO 20000-1:2011
- SOC 1 Type 2
- SOC 2 Type 2
- SOC 3
- CIS Benchmark
- CSA STAR Certification
- CSA STAR Attestation
- CSA STAR Self-Assessment
- WCAG 2.0 (ISO 40500:2012)

U.S. GOVT

- FedRAMP High
- FedRAMP Moderate
- EAR
- ITAR
- DoD DISA SRG Level 5
- DoD DISA SRG Level 4
- DoD DISA SRG Level 2
- DFARS
- DoE 10 CFR Part 810
- NIST SP 800-171
- NIST CSF
- Section 508 VPATs
- FIPS 140-2
- CJIS
- IRS 1075
- CNSSI 1253

INDUSTRY

- PCI DSS Level 1
- GLBA (US)
- FFIEC (US)
- Shared Assessments (US)
- SEC 17a-4 (US)
- CFTC 1.31 (US)
- FINRA 4511 (US)
- SOX (US)
- 23 NYCRR 500 (US)
- OSFI (Canada)
- FCA + PRA (UK)
- APRA (Australia)
- FINMA (Switzerland)
- FSA (Denmark)
- RBI + IRDAI (India)
- MAS + ABS (Singapore)
- NBB + FSMA (Belgium)
- AFM + DNB (Netherlands)
- AMF + ACPR (France)
- KNF (Poland)
- European Banking Authority (EBA)
- FISC (Japan)
- HIPAA BAA (US)
- HITRUST Certification
- GxP (FDA 21 CFR Part 11)
- MARS-E (US)
- NHS IG Toolkit (UK)
- NEN 7510:2011 (Netherlands)
- FERPA (US)
- CDSA
- MPAA (US)
- FACT (UK)
- DPP (UK)

REGIONAL

- Argentina PDPA
- Australia IRAP Unclassified
- Australia IRAP PROTECTED
- Canada Privacy Laws
- China GB 18030:2005
- China DJCP (MLPS) Level 3
- China TRUCS / CCCPPF
- EU EN 301 549
- EU ENISA IAF
- EU Model Clauses
- EU – US Privacy Shield
- GDPR
- Germany C5
- Germany IT-Grundschutz workbook
- India MeitY
- Japan CS Mark Gold
- Japan My Number Act
- Netherlands BIR 2012
- New Zealand Gov CIO Framework
- Singapore MTCS Level 3
- Spain ENS High
- Spain DPA
- UK Cyber Essentials Plus
- UK G-Cloud
- UK PASF

The Azure Toolbox

”Azure should not be seen as the goal - but as an enabler and accelerator of advanced IT services”

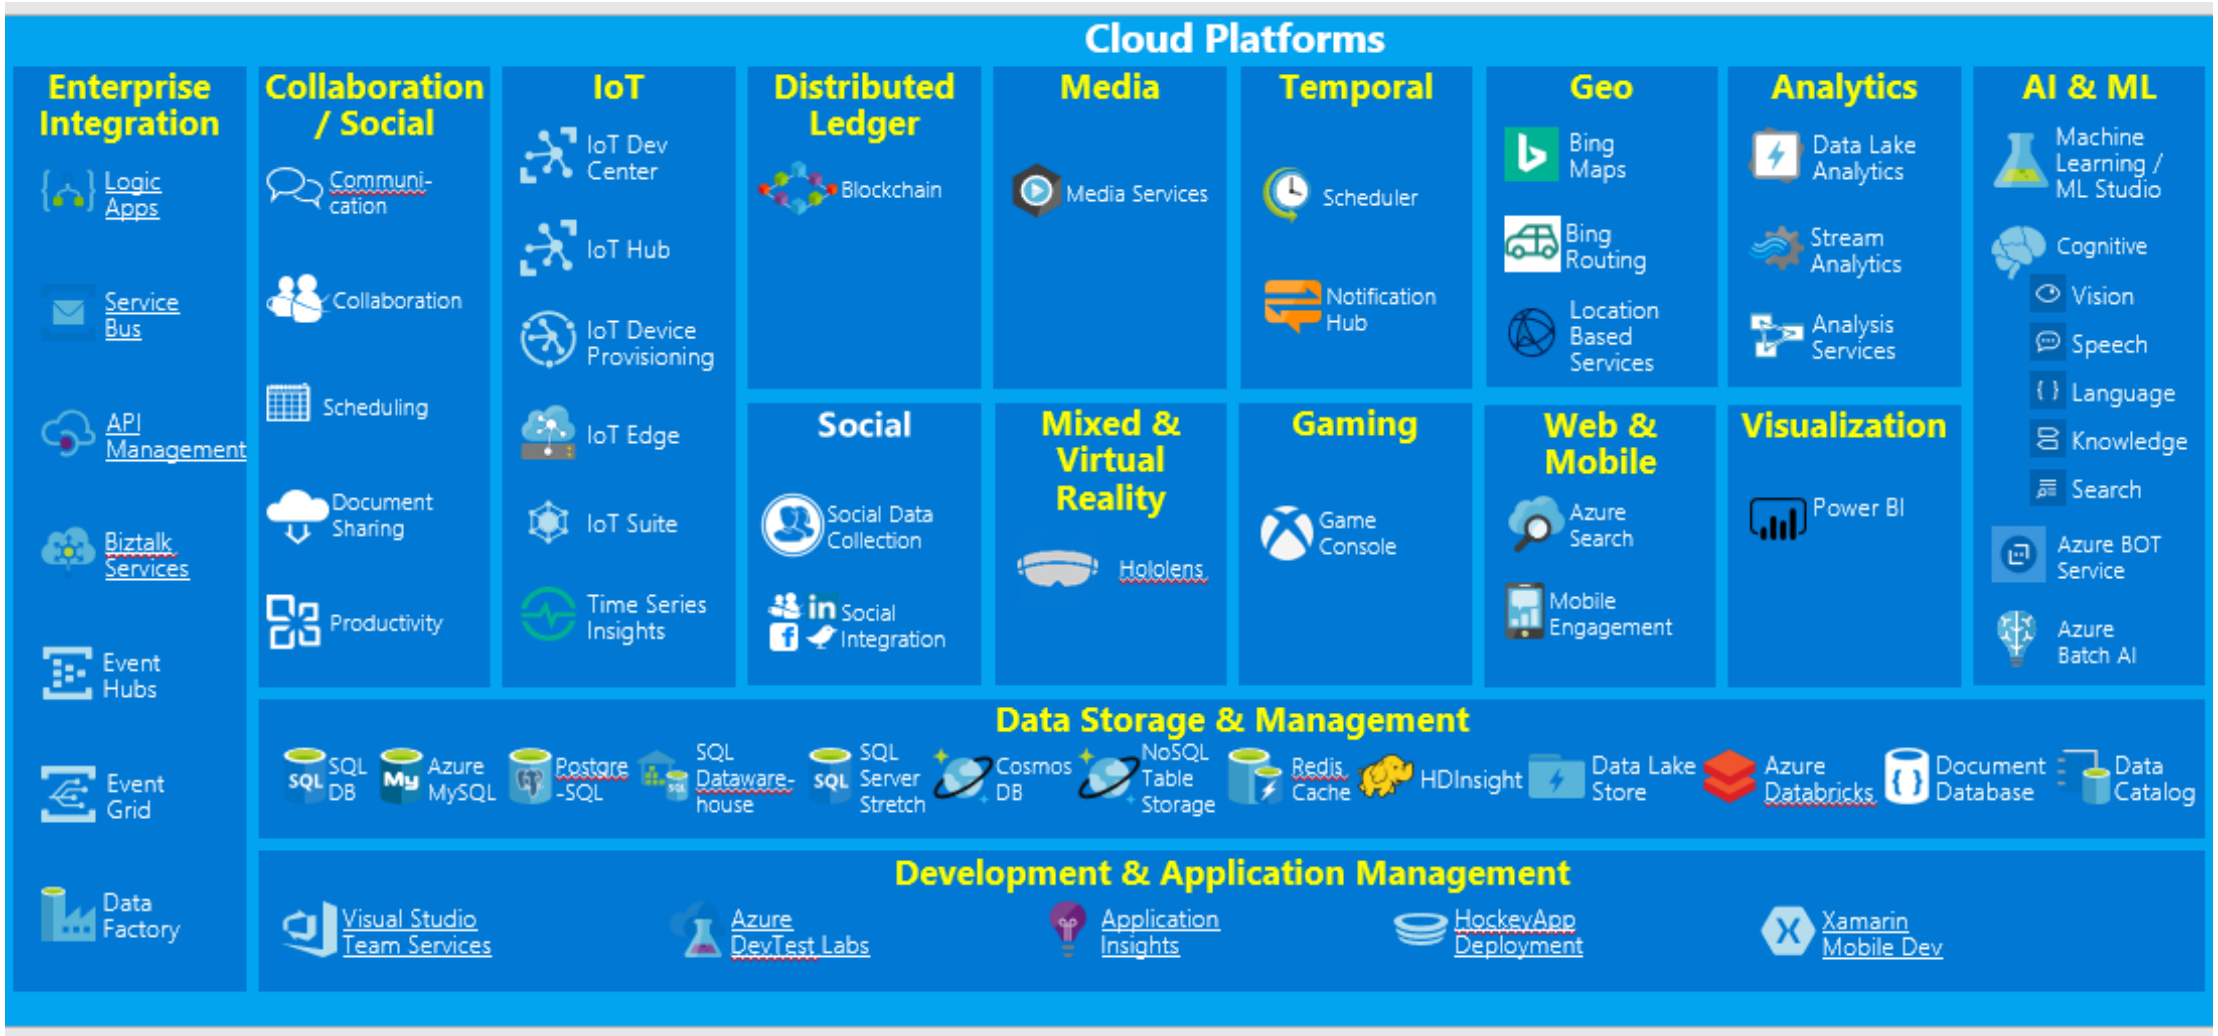

Azure is an open cloud – works with what you use

The hybrid (multi-)cloud opportunity for you?

Keeping up with demands in research

Maximizing your investment

Maximizing collaboration

Variable demand

Fixed demand

Regulations AI IoT Data analytics

Continuously, random, unpredictable spikes in demand for HPC can come from any new or existing application workflow

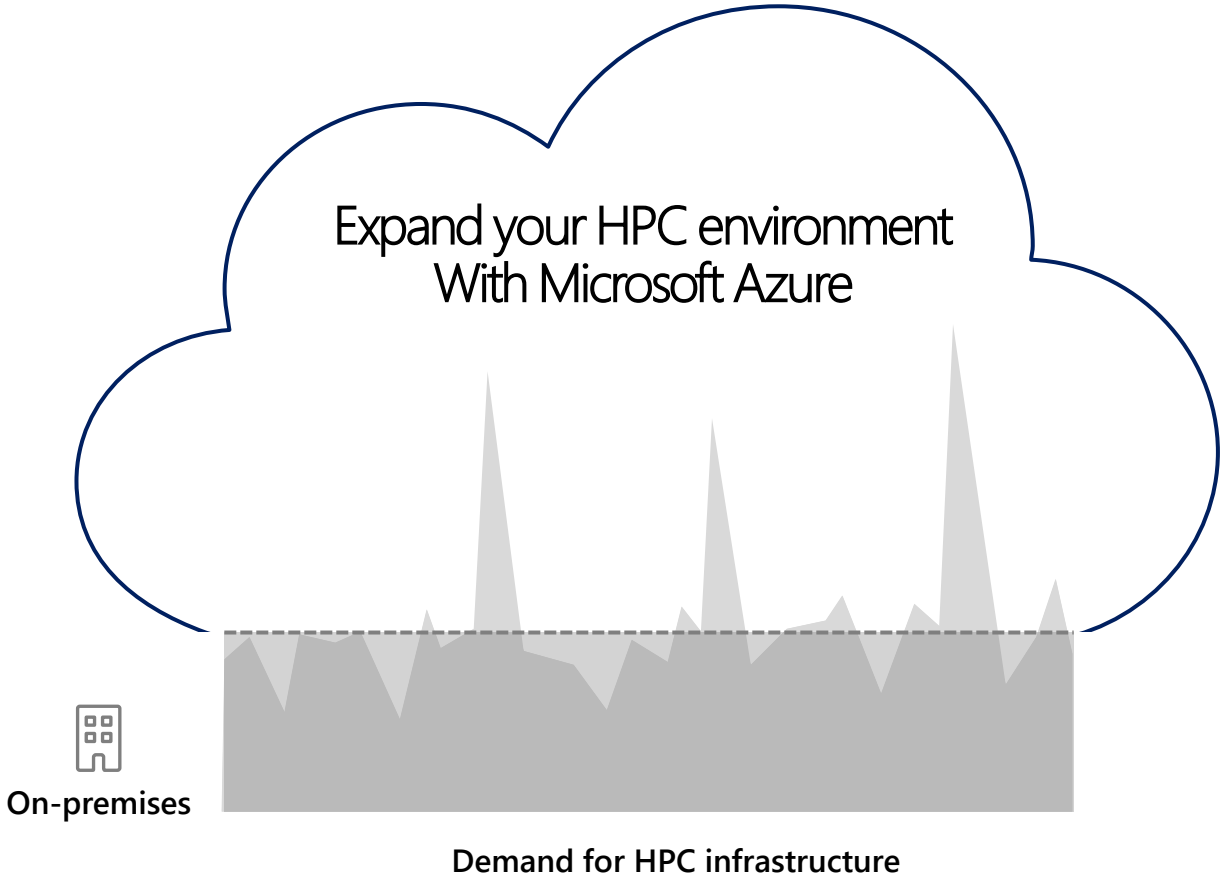

HPC in Azure: VMs with RDMA, GPU, FPGA and Cray

Entry Level VMs

Dev/Test Workolads

General Purpose VMs

Common Applications, Web servers etc

Compute Optimized VMs

Gaming, Analytics

Large Memory VMs

Large Databases

>80,000 IOPs Premium Storage

Low latency, high throughput apps

Storage optimized VMs

No SQL Databases (Cassandra, MongoDB), Data warehousing

Infiniband

High Performance VMs

Batch processing, fluid dynamics, Monte Carlo simulation

GPU-enabled VMs

NV – Graphic based applications

NC2 – Advanced Sim

ND1 – AI Inferencing

ND2 – AI Training

FPGA

Virtual Machines – HPC

FPGA Microservices – AI/Edge

Cray in Azure

Aries Connected CPU/GPU/Storage available in cloud

Mix and Choose

Choose OS

redhat

SUSE

ubuntu

CentOS

Windows Server

BYO Scheduler

Altair | PBS Works™

slurm workload manager

grid engine

UNIVA

TIBCO DataSynapse GridServer Engine

IBM Spectrum Symphony

Bright Computing

openlava Open Source Workload Management

Microsoft Azure Batch™ HPC Pack™

Adaptive COMPUTING

BYO Parallel FS

lustre

BeeGFS® developed by Fraunhofer

GlusterFS

IBM Spectrum Scale

BYO Platform

Uber Cloud

rescale

CYCLE COMPUTING

PBS Cloud Manager

pbscloud.io

DEIC – the different types

Titel	Type 1 Interactive HPC	Type 2 Through HPC	Type 3	Type 4	Type 5
Beskrivelse	Fokus på interaktive beregningsressourcer og let tilgang for nye brugere. Desuden vil denne type kunne udgøre en fælles platform for uddannelse. Denne type vil forventeligt øge andelen af nye brugere og bruges til generel prototyping for mere generelle/erfarne HPC-brugere.	Denne anlæ har typisk e antal kerne kan være e mellem kost tive og bereg effektive e med stort th put kapacite ofte høj fok sikkerhed. Id mange små o emstore job r.			
Primær(e) beregnings-enhed(er)	CPU, GPU; Tynde og/eller tykke noder	CPU			
Intercon-nect	Ikke nødvendigt	Low latency			
Access	Interaktiv	Jobbaseret			

Azure Services

- Data Science Lab VM
- Azure Machine Learning Studio
- Azure DataBricks
- Azure Virtual Desktop
- Azure Synapse Studio
- Azure Data Lake

Azure ”Raw”

- Azure Marketplace
- VM based on Linux/Windows
- Azure Powershell/CLI
- Azure BluePrints

Arealdata.dk – an open data platform

Thank you


Microsoft
Appendix slides

Economic benefit: Well-optimized cloud usage can free-up excess capacity and allows customers to optimize spend

C Unused Capacity – Extra usage which is rarely if ever used

Immediate Savings

B Fluctuations on top of base workloads – seasonal patterns or occasional bursts

Hourly pricing for the hours or days needed

A Base workloads – steady state, typically covers all day everyday use.

Use Reserved Instances

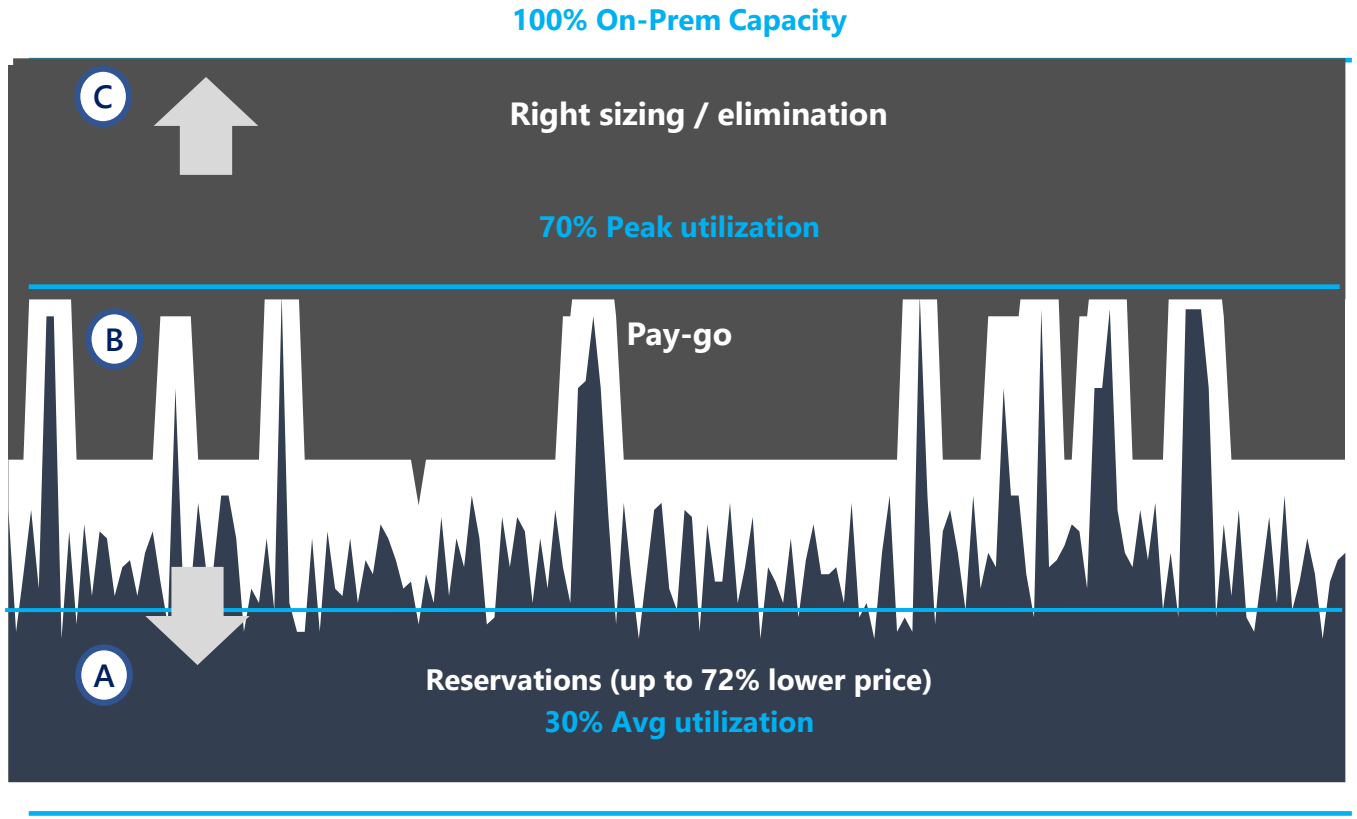

Extremely Large Data Sets

Large Data Sets

Medium & Small Data Sets

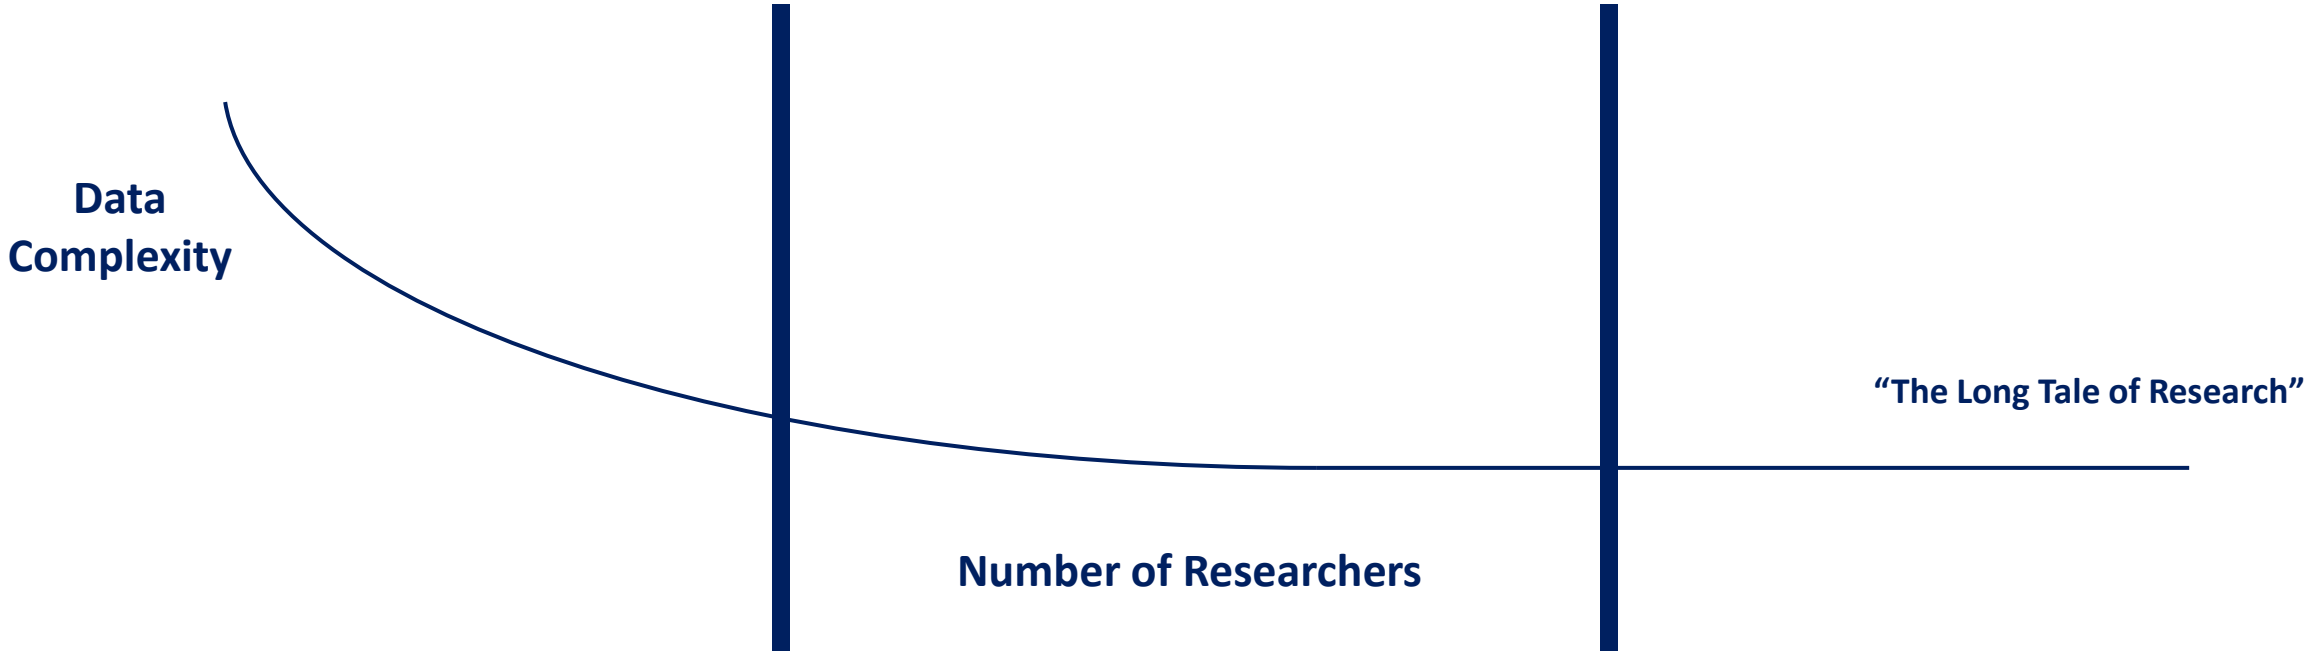

- Expensive to move
- Domain standards
- High computational needs
- Supercomputers, HPC, Grids

e.g. High Energy Physics, Astronomy

- Some Standards within Domains
- Shared Datacenters & Clusters
- Research Collaborations
- e.g. Environmental Science, Clinical Data & Healthcare Informatics

- Widely diverse data; Few standards
- Local Servers & PCs
- Flat Files, Excel
- e.g. Social Sciences, Humanities, Economics

DEIC – the different types

Titel	Type 1 Interactive HPC	Type 2 Throughput HPC	Type 3 Large Memory HPC	Type 4 Accelerated HPC	Type 5 Capability HPC
Beskrivelse	Fokus på interaktive beregningsressourcer og let tilgang for nye brugere. Desuden vil denne type kunne udgøre en fælles platform for uddannelse. Denne type vil forventeligt øge andelen af nye brugere og bruges til generel prototyping for mere generelle/erfarne HPC-brugere.	Denne anlægstype har typisk et stort antal kerner som kan være et mix mellem kost-effektive og beregningseffektive enheder med stort throughput kapacitet med ofte høj fokus på sikkerhed. Ideel til mange små og mellemstore jobs, der bruger store data/filer.	Fokus på applikationer som ikke nemt eller effektivt kan distribueres mellem mange computer-noder. Krav om stort fladt memory-space som ved store matrixproblemer eller anvendelser med stor mængde memory og forholdsvist lille antal effektive kerner.	En anlægstype hvis primære beregningskapacitet kommer fra acceleratorer af forskellig slags. Denne muliggør, at danske forskere fremadrettet klædes på til næste generations supercomputere. Det kan være spirende teknologier som FPGA, in-memory og in-storage computing.	En anlægstype med særlig fokus på at løse problemer som kræver et mellemstort til stort antal computer-noder samtidigt. Få og store problemer samtidigt som løses hurtigst muligt. Dette kan fx være EuroHPC LUMI sub-exa-scale anlægget.
Primær(e) beregnings-enhed(er)	CPU, GPU; Tynde og/eller tykke noder	CPU	CPU	CPU, FPGA, "Next generation GPUs", computational storage og evt. andre relevante enheder	CPU, GPU
Interconnect	Ikke nødvendigt	Low latency	Shared: Ikke nødvendigt Cluster: Low latency	Low latency	Low latency
Access	Interaktiv	Jobbaseret	Jobbaseret	Jobbaseret	Jobbaseret

DEIC – the different types

Titel	Type 1 Interactive HPC	Type 2 Through HPC	Type 3	Type 4	Type 5
Beskrivelse	Fokus på interaktive beregningsressourcer og let tilgang for nye brugere. Desuden vil denne type kunne udgøre en fælles platform for uddannelse. Denne type vil forventeligt øge andelen af nye brugere og bruges til generel prototyping for mere generelle/erfarne HPC-brugere.	Denne anlæ har typisk e antal kerne kan være e mellem kost tive og bereg effektive e med stort th put kapacite ofte høj fok sikkerhed. Ic mange små o emstore job r.			
Primær(e) beregnings-enhed(er)	CPU, GPU; Tynde og/eller tykke noder	CPU			
Intercon-nect	Ikke nødvendigt	Low latency			
Access	Interaktiv	Jobbaseret			

Azure Services

- Data Science Lab VM
- Azure Machine Learning
- Azure DataBricks
- Azure VDI
- Azure Synapse Studio
- Azure Data Lake

Azure ”Raw”

- Azure Marketplace
- VM based on Linux/Windows
- Azure Powershell/CLI
- Azure BluePrints

DEIC – the different types

Titel	Type 1 Interactive HPC	Type 2 Throughput HPC	L
Beskrivelse	Fokus på interaktive beregningsressourcer og let tilgang for nye brugere. Desuden vil denne type kunne udgøre en fælles platform for uddannelse. Denne type vil forventeligt øge andelen af nye brugere og bruges til generel prototyping for mere generelle/erfarne HPC-brugere.	Denne anlægstype har typisk et stort antal kerner som kan være et mix mellem kost-effektive og beregningseffektive enheder med stort throughput kapacitet med ofte høj fokus på sikkerhed. I el til mange sm lemstore bruger store ata/filer.	Fo sc fe re co or m st el st og
Primær(e) beregningsenhed(er)	CPU, GPU; Tynde og/eller tykke noder	CPU	Cl
Interconnect	Ikke nødvendigt	Low latency	Sh di Cl
Access	Interaktiv	Jobbaseret	Jo

Azure Services

- HPC VM based on Linux/Windows
- Azure Container Services
- Azure Batch
- Azure Cray
- CPU and GPU based workloads

Azure ”Raw”

- Azure Marketplace
- VM based on Linux/Windows
- Azure ML API
- Azure Powershell/CLI
- Azure BluePrints
- Azure ARM

DEIC – the different types

Azure Services

- HPC VM based on Linux/Windows
- Azure Container Services
- Azure Batch
- Azure Cray
- CPU and GPU based workloads

Azure ”Raw”

- Azure Marketplace
- VM based on Linux/Windows
- Azure ML API
- Azure Powershell/CLI
- Azure BluePrints
- Azure ARM

	Type 3 Large Memory HPC	Type 4 Accelerated HPC	Type 5 Capability HPC
ut	type stort som mix ffek- ngs- eder ugh- med på l til a/fi-	Fokus på applikationer som ikke nemt eller effektivt kan distribueres mellem mange computer-noder. Krav om stort fladt memory-space som ved store matrixproblemer eller anvendelser med stor mængde memory og forl dsvist lille an- kerner.	En anlægstype hvis primære beregningskapacitet kommer fra acceleratorer af forskellig slags. Denne muliggør, at danske forskere fremadrettet klædes på til næste generations supercomputere. Det kan være spirende teknologier som FPGA, in-memory og in-storage computing.	En anlægstype med særlig fokus på at løse problemer som kræver et mellemstort til stort antal computer-noder samtidigt. Få og store problemer samtidigt som løses hurtigst muligt. Dette kan fx være EuroHPC LUMI sub-exa-scale anlægget.
	CPU	CPU, FPGA, ”Next generation GPUs”, computational storage og evt. andre relevante enheder	CPU, GPU
	Shared: Ikke nødvendigt Cluster: Low latency	Low latency	Low latency
	Jobbaseret	Jobbaseret	Jobbaseret

DEIC – the different types

Azure Services

- HPC VM based on Linux/Windows
- Azure Batch
- Azure Cray
- CPU and GPU based workloads
- FPGA based VM workloads

Azure ”Raw”

- Azure Marketplace
- VM based on Linux/Windows
- Azure ML API
- Azure Powershell/CLI
- Azure BluePrints
- Azure ARM

Type 3 rge Memory HPC	Type 4 Accelerated HPC	Type 5 Capability HPC
s på applikationer ikke nemt eller ef- vt kan distribue- mellem mange puter-noder. Krav stort fladt me- -space som ved matrixproblemer anvendelser med mængde memory rholdsvist lille an-	En anlægstype hvis primære beregningskapacitet kommer fra acceleratorer af forskellig slags. Denne muliggør, at danske forskere fremadrettet klædes på til næste generations supercomputere. Det kan e spirende teknologier som FPGA, in-memory og in-storage computing.	En anlægstype med særlig fokus på at løse problemer som kræver et mellemstort til stort antal computer-noder samtidigt. Få og store problemer samtidigt som løses hurtigst muligt. Dette kan fx være EuroHPC LUMI sub-exa-scale anlægget.
	CPU, FPGA, ”Next generation GPUs”, computational storage og evt. andre relevante enheder	CPU, GPU
ed: Ikke nødven- er: Low latency	Low latency	Low latency
aseret	Jobbaseret	Jobbaseret

DEIC – the different types

Azure Services

- HPC VM based on Linux/Windows
- Azure Batch
- CPU and GPU based workloads

Azure ”Raw”

- Azure Marketplace
- VM based on Linux/Windows
- Azure ML API
- Azure Powershell/CLI
- Azure BluePrints
- Azure ARM

Type 4 ccelerated HPC	Type 5 Capability HPC
anlægstype hvis nære bereg- skapacitet kom- fra accelerato- af forskellig s. Denne mulig- at danske for- e fremadrettet des på til næste erations super- putere. Det kan	En anlægstype med særlig fokus på at løse problemer som kræver et mellem-stort til stort antal computer-noder samtidigt. Få og store problemer samtidigt som løses hurtigst muligt. Dette kan fx være HPC LUMI sub-scale anlægget.
nemory og in- age computing.	
, FPGA, ”Next eration GPUs”, putational age og evt. an- relevante enhe-	CPU, GPU
latency	Low latency
baseret	Jobbaseret

Sustainability – In the news

Ørsted indgår elkøbsaftale med Microsoft Corporation

10.08.2021 09:02

Ørsted og Microsoft Corporation har indgået en aftale om, at Microsoft skal købe energi fra Old 300 Solar Center i Fort Bend County i Texas.

[Ørsted indgår elkøbsaftale med Microsoft Corporation]

Successful carbon removal depends on these 3 conditions

Reducing the market and technology risks of carbon removal solutions

[Removing CO2 depends on these 3 conditions | World Economic]