HPC Breakthroughs 2027

Unleashing scientific discovery with the power of computing



Including

- Does scientific computing need FP64?
- Europe's shrinking supercomputer budgets
- The water footprint of supercomputers
- Science is generating more data than it can store
- How NASA balances resources

Brought to you by

SCIENTIFIC COMPUTING WORLD


Google Cloud
GLOBAL COMMUNITY OF PRACTICE
Connect.
Collaborate.
Accelerate Discovery.
Join researchers, engineers, and HPC leaders advancing the frontiers of AI, quantum, and high performance computing on Google Cloud.
• WEEKLY TUESDAY SESSIONS (9:00 AM PT) • IN-PERSON WORKSHOPS (SC26 AND ISC) • PEER EXCHANGE
Join the community at
sites.google.com/view/
advancedcomputingcommunity
SCAN TO JOIN COMMUNITY

CONTENTS

4 Editor's Welcome
Making the most of advanced computing infrastructure

6 High-performance cuts
European HPC directors explain to us how soaring hardware costs are altering system designs.

10 The storage squeeze
The growth of scientific data has collided with a constrained storage market. Research computing leaders explain what has to change.

14 Precision requirements
Pekka Manninen from CSC talks to us about whether scientific computing still needs native double precision.

18 Water usage in AI
HPC practitioners examine how cooling architecture drives supercomputer water consumption across 11 global sites.

22 NASA's compute strategy
Andrew Mitchell from NASA examines how matching computing environments to scientific workloads supports mission success.

24 Imperial's HEX platforms
Andrew Richards from Imperial College London explains how replacing a single monolithic cluster with six purpose-built platforms matches modern research requirements.

26 UCL ARC strategy
Ben Thomas from UCL ARC examines how private cloud systems and Ethernet architecture enhance accessibility for diverse research communities.

28 Software-defined HPC
Pekka Manninen explains how the LUMI AI Factory is adapting hardware and software infrastructure to serve simulation and inference.

30 Architecting the future
Mark Papermaster from AMD highlights chiplet innovation, 2nm manufacturing and hardware for AI for science.

32 Quantum hardware scaling
Mark Webber from Universal Quantum examines the architectural shift needed to build fault-tolerant, million-qubit systems for practical quantum simulation.

34 Post-exascale roadmap
Jean-Yves Berthou talks about Europe's post-exascale strategy and HPC convergence

36 CSC's supercomputing evolution
Kimmo Koski traces CSC's evolution from early vector systems to GPU-heavy computing with LUMI and Roihu

38 Vascular digital twins
Amanda Randles on using HPC to move healthcare from reactive to proactive

40 Uncertainty in AI
Peter Coveney from UCL explains how exascale physics models and uncertainty quantification to improve AI-driven drug discovery.

42 Quantum reactor simulation
Andrew Childs from the University of Maryland explains his research on hybrid quantum algorithms for neutron transport.

44 AI in UK science
Tom Griffin from STFC talks about integrating AI into scientific workflows and managing growing data demands.

Thanks to our commercial partners

CoolIT An Ecolab Company · Google Cloud · HAMMERSPACE · SEAGATE · STREAM High Performance Computing

CONTENT AND MARKETING
Tel: +44 (0)1223 221030

Editor
Robert Roe
robert.roe@europascience.com

Senior designer
Zoe Wade
zoe.wade@europascience.com

Head of content
Finbarr O'Reilly
fin.oreilly@europascience.com

Head of product
Laura Schyberg
laura.schyberg@europascience.com

Audience development manager
Andrew Knight
andrew.knight@europascience.com

ADVERTISING & PRODUCTION

Account Director
Tim Richters
tim.richters@europascience.com
+44 (0)1223 221039

Client Success Manager
Nick Clark
nick.clark@europascience.com

CORPORATE TEAM

Chief operating officer
Mark Elliott
mark.elliott@europascience.com

Chief executive officer
Warren Clark
warren.clark@europascience.com

Scientific Computing World is published by Europa Science Ltd, St John's Innovation Centre | Cowley Road | Cambridge | CB4 OWS | ● ISSN 1744-8026
Tel: +44 (0) 1223 211170 ● Fax: +44 (0) 1223 213385
Web: www.scientific-computing.com

Receive our weekly newsletter: Go to www.scientific-computing.com/register

While every care has been taken in the compilation of this magazine, errors or omissions are not the responsibility of the publishers or of the editorial staff. Opinions expressed are not necessarily those of the publishers or editorial staff. All rights reserved. Unless specifically stated, goods or services mentioned are not formally endorsed by Europa Science Ltd, which does not guarantee or endorse or accept any liability for any goods and/or services featured in this publication.

©2026 Europa Science Ltd.

Navigating the convergence of HPC, AI and economics



Welcome to the latest edition of HPC Breakthroughs. As we take our annual look at the state of high-performance computing, one theme emerges above all others – users of high performance computing and the people who run supercomputing resources on their behalf are navigating an era of profound convergence and unprecedented constraint.

According to this year's crop of expert interviewees, the conventional model of building monolithic supercomputers to serve every possible scientific discipline is fracturing. In its place, research computing directors are designing bespoke, hybrid architectures that blend traditional simulation with the exploding demands of artificial intelligence.

However, this evolution is colliding with brutal market realities. Across Europe's leading HPC centres, research budgets remain fixed while hardware costs soar. Memory prices have surged five- to six-fold, and GPUs remain in short supply. As Kimmo Koski of Finland's CSC bluntly notes, system architects are no longer solely asking what the science needs, but what the market will allow. These economic pressures are forcing quiet compromises, such as delayed GPU tiers and hybrid storage solutions replacing pure flash.

The data deluge compounds this infrastructure squeeze. Scientific instruments – from larger telescopes to faster detectors – are producing information at staggering rates. As experts from NASA and STFC highlight, the growth of scientific data has collided with a constrained storage market, exacerbated by AI's massive appetite for hard drives. The uncomfortable reality is that scientists may no longer be able to afford to store all raw data, necessitating a cultural shift toward AI-enabled, real-time data triage and on-the-fly processing.

Thank you to our partners

HPC Breakthroughs would not have been possible without the support of its sponsors. Scientific Computing World is proud to partner with these inspiring organisations and help showcase their work in this project

CoolIT™
An Ecolab Company

CoolIT specialises in scalable liquid cooling solutions for the world's most demanding high-performance computing environments. Through its broad portfolio of modular direct liquid cooling (DLC) products, CoolIT enables dramatic increases in rack density, component performance and power efficiency across any high-density computing scenario. CoolIT collaborates with the top semiconductor manufacturers and cloud service providers to develop cooling technologies that enable ever greater computing performance and efficiency. Through its Liquid Lab™ R&D centers in Calgary and Taipei, global manufacturing operations in Canada, China, and Vietnam, and on-site support in more than 80 countries, CoolIT serves customers at scale worldwide. Together, CoolIT and its partners are leading the way for widespread adoption of efficient high-density, high-performance computing.

Google Cloud

Google Cloud is the new way to the cloud, providing AI, infrastructure, developer, data, security, and collaboration tools built for today and tomorrow. Google Cloud offers a powerful, fully integrated and optimised AI stack with its own planet-scale infrastructure, custom-built chips, generative AI models and development platform, as well as AI-powered applications, to help organisations transform. Customers in more than 200 countries and territories turn to Google Cloud as their trusted technology partner.

Amid these hardware constraints, the very architecture of supercomputing is diverging. The AI boom is driving vendor roadmaps toward lower-precision formats such as FP8 and FP4. Yet, traditional scientific simulations – such as climate modelling and fluid dynamics – still demand native IEEE 754 double precision (FP64). While NVIDIA pivots heavily toward AI workloads, AMD has strategically positioned its Instinct MI430X to support legacy scientific workloads with high FP64 performance.

We are also seeing a deliberate shift away from monolithic systems toward flexible, right-sized infrastructure. Imperial College London recently replaced its 80,000-core monolithic cluster with six purpose-built HEX platforms, balancing traditional HPC, interactive AI training, and teaching on shared infrastructure. Similarly, UCL's Centre for Advanced Research Computing has embraced flexible, Ethernet-based private clouds to democratise HPC access for non-traditional academic communities. On an even larger scale, of course, NASA's Science Mission Directorate manages 150 petabytes of data through a hybrid strategy, carefully balancing on-premises high-end computing for tightly coupled simulations with cloud infrastructure for interactive science. Furthermore, Europe's flagship supercomputer, LUMI, is evolving into a software-defined AI platform to dynamically serve simulation, training, and inference.

At STFC, Tom Griffin emphasises that AI is touching every part of the facility workflow, from experiment design to accelerator fault diagnosis. However, this interconnectivity brings heightened cybersecurity risks, with the shrinking gap between vulnerability discovery and exploitation demanding robust, defense-in-depth network security.

Sustainability has also stepped into the spotlight. A landmark study of 11 global supercomputing sites reveals that cooling architecture, rather than geographical water stress, is the primary driver of a facility's water consumption. As power densities climb into the tens of megawatts, HPC operators must carefully balance energy demands with critical water conservation.

Despite these infrastructural hurdles, the scientific breakthroughs powered by HPC have never been more inspiring. At Duke University, Amanda Randles is leveraging supercomputing and wearable sensor data to build patient-specific vascular digital twins, moving healthcare from reactive treatments to proactive disease detection. Similarly, Peter Coveney at UCL demonstrates how exascale physics modelling and uncertainty quantification are essential to disciplining AI in drug discovery, ensuring that digital twins remain actionable and trustworthy.

Looking further ahead, we explore the horizons of quantum computing. Universal Quantum's Mark Webber details how trapped-ion shuttling could bypass photonic bottlenecks to scale fault-tolerant systems to the millions of qubits required for practical drug discovery. Meanwhile, Andrew Childs from the University of Maryland outlines the ongoing challenge of proving hybrid quantum speedups against increasingly sophisticated classical algorithms when modelling complex nuclear reactor dynamics.

Europe is already preparing for the next paradigm. Through the Horizon Europe-funded SPE project, leaders such as Jean-Yves Berthou are defining the post-exascale roadmap, where future performance will be measured in watts and tokens rather than just flops. It is a future defined by the deep convergence of HPC, AI, and quantum computing.

In this supplement, we bring you the voices of the practitioners, architects and researchers navigating these complex shifts. We hope their insights inspire your own work as we collectively build the next generation of scientific discovery.

Robert Roe,
Editor,
Scientific Computing World

HAMMERSPACE

Hammerspace is the high-performance data platform built to simplify and optimise AI infrastructure at scale. It makes all your data immediately accessible – anywhere across on-premises and cloud environments – without copying or migrating data into new silos. By integrating with existing storage, networking, and applications, Hammerspace creates a unified, high-speed data backbone for AI, enabling organizations to accelerate every stage of the AI pipeline while eliminating data silos.

SEAGATE

As data becomes the foundation of modern scientific discovery, research organisations are under increasing pressure to manage, access and preserve ever-growing volumes of data. Seagate Technology plays a critical role in enabling this transformation, delivering the infrastructure needed to support data-intensive research at scale. From genomics and high-throughput sequencing to computational modelling and AI-driven analysis, Seagate helps institutions overcome the challenges of managing petabyte-scale datasets, ensuring data remains accessible, reliable and ready to drive insight. With a focus on long-term data integrity, efficiency and sustainability, Seagate supports researchers in accelerating innovation while building the data foundations needed for the next generation of breakthroughs.

STREAM
High Performance Computing

Ever used ChatGPT? Then you've used GPU-libraries from StreamHPC. Stream HPC is a 40-person leading specialist in GPU computing, helping research teams turn compute-intensive scientific software into faster, more capable tools. The company develops and optimises CPU- and GPU-based software for simulation, data analysis and AI, from algorithms and architecture to production code. StreamHPC's software performance engineers remove bottlenecks and accelerate scientific discovery.

Budget crises force academic HPC cuts

As component costs surge six-fold, European supercomputing directors reveal to **Robert Roe** how supply constraints are quietly reshaping research architecture

Research computing budgets haven't shrunk, but what they buy has. Across Europe's leading HPC centres, the same story is playing out: memory prices up five- to six-fold, GPU allocations deferred, storage tiers diluted from flash to a flash-disk hybrid. The system architects building the next generation of scientific infrastructure are no longer asking what the science needs – they're asking what the market will let them have.

On 31 August, EuroHPC signed off a €387.8m contract for LUMI-AI, Europe's largest AI Factory procurement to date. But the headline figure obscures a harder story underneath. Kimmo Koski, director of Finland's CSC, puts its scale plainly. "Market conditions impact LUMI-AI just as they impact every major hardware deployment worldwide. It certainly would have been more cost-effective to procure components a year earlier. Memory prices, for instance, have increased five- to six-fold."

How rising HPC costs are changing hardware purchases

At UCL's Centre for Advanced Research Computing, head of research computing Ben Thomas had a clear plan: consolidate two ageing systems, one built for high-throughput batch work and one for traditional multi-node MPI, into a single machine that could serve both communities. High core counts and a fast interconnect made the merge credible, and consolidating two geographically separated systems solved a data centre problem at the same time. The plan also included GPUs. It no longer does.

"The intention, before the world turned upside down, was to have both CPU and GPU capability within that system," Thomas says. "Because of how things have changed over the last few months, we decided to lay the groundwork and buy CPU-only, then use subsequent funding to build out GPU capability later. So that's where we decided to take the hit, not buying both."

Even the CPU purchase came in short. "We ended up buying significantly less than we would have hoped. We just about covered our existing capabilities, because although we've got fewer nodes and fewer sockets, the core counts are higher, and the cores can get through significantly more work than the old machines."

The systems being replaced date from 2018 and 2019/2020, so generational gains cover the gap. "So we're okay in that respect, but it's far less than we'd have hoped. Funding for the GPU tier has since been secured for the coming year, but at a smaller scale than originally envisaged. The price changes were brutal."

Two further compromises followed. The single-system design relied on high-speed memory for memory-bound researchers, which is another area hit by the AI bubble. "We wanted high-speed memory, MRDIMMs, for the memory-bound researchers, and you can't get those for love nor money at the moment," Thomas says. "Even if you've got the money, you can't get your hands on them."

Then the storage. "I'd also kind of forgotten about the storage; we were originally going to go 100% flash on the scratch tier, and that just wasn't feasible. So instead we've gone for a hybrid mixture of flash and disk. We'll 'suck it and see' whether we need to bulk out the flash, or whether we can get away with disk on the back end."

Why timing is everything in HPC hardware procurement

The difference between a comfortable procurement and a painful one is now measured in months. Dr Andrew Richards, director of research computing services at Imperial College London, knows exactly which side of the line his most recent systems fell on.

"With our HX2 and HX3 platforms, we were really lucky; we placed the purchase orders last summer, before the world went completely crazy on pricing," he says. "If we'd done that a year later, for our storage system we'd probably be getting a quarter of what we got, just from the pricing changes we've seen over the last year."

Richards is in the fortunate position of sitting between service refreshes, which buys time. "We're personally lucky right now, in a window between service refreshes where we don't need to panic this year, so we can ride this out and hope things are better in 12 months." The work that cannot wait, though, still gets hit. "Things we have to do this year, like the secure platform work, where we're seeding a secure storage platform and some compute, we just have to accept we'll get less for our money. We have a fixed budget; we need to seed the system and get it going, and we'll end up with less storage than if we'd done this a year ago."

Even internal processes have become expensive. Richards describes a new professorship where Imperial declined to commit spend until the appointment was fully signed off. "By the time that concluded, pricing on what we were trying to buy had gone up by about £100k."

Koski's experience at CSC points the same way. Roihu, Finland's next national supercomputer, "was contracted before the sharpest price increases took effect". For LUMI-AI, the response was to fight harder at the table: "We negotiated aggressively to mitigate price surges."

How costs are changing HPC system architecture

What emerges is that system design is increasingly influenced by market forces rather than the demands of scientific workloads. As Koski puts it: "Market conditions force trade-offs in system architecture, whether that means adjusting

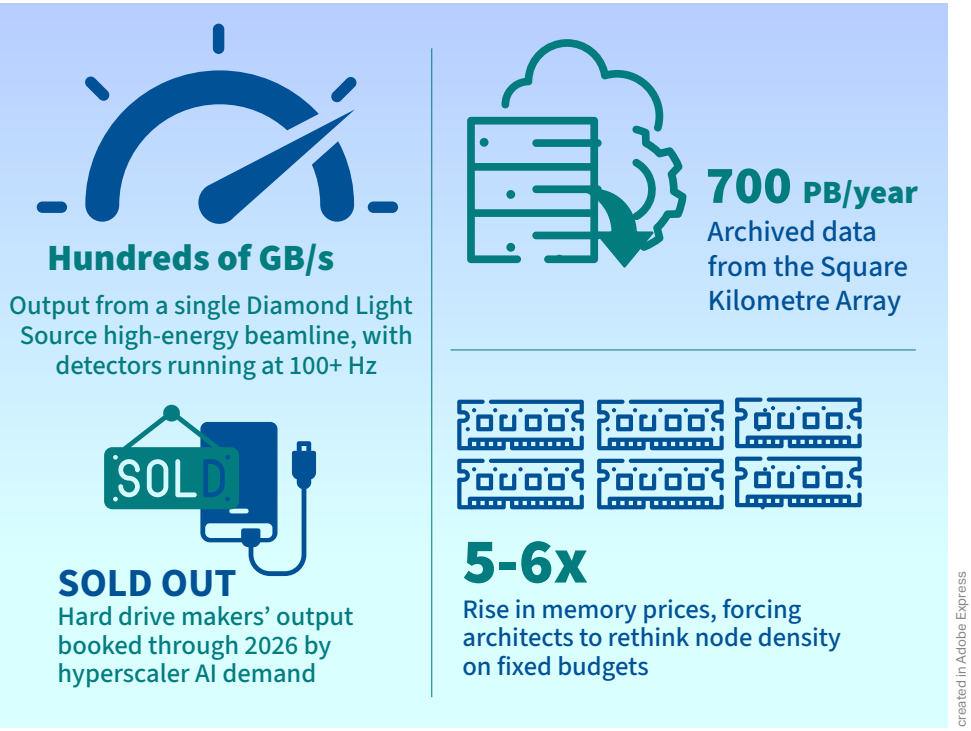

"Market conditions force trade-offs in system architecture, whether that means adjusting node density, fine-tuning memory configurations per node, or re-evaluating total node counts relative to the budget"

node density, fine-tuning memory configurations per node, or re-evaluating total node counts relative to the budget."

"Geopolitical factors, global supply turbulence, and manufacturing concentrations in regions like Taiwan mean you cannot wait indefinitely for market conditions to improve," he says. "You have to optimise within your constraints and move forward."

Richards describes the same collapse of choice on the accelerator side, where the technical evaluation has largely been replaced by an availability check. "Five years ago, you could have a rational conversation about which GPU to buy and look at the specs," he says. By the time Imperial procured HX1 around four years ago, the question had become "what's available in the timeframe we need? That narrows it to one or two options. It has to be delivered by a certain date, and the alternative is on a longer lead time and four times more expensive, so the only choice is, say, the H100 at that point in time."

The pattern repeated with the HEX AI platform. Community preference kept the procurement on NVIDIA, Richards says, because researchers "wanted a simple on-ramp to the new system and didn't want to explore different technologies". Within that constraint, "H200s were the only thing deliverable in our timeframe on a sensible lead time." B200s are now arriving as part of two separate systems that together make up the HEX AI platform.

And the sizing? "It's largely been driven by what we can afford and how many GPUs we get. When we built the HEX AI platform, we made a case to the university, got the management board to sign off a capital amount, and that determined how many GPUs we got. It wasn't driven by any more complex maths than that."

That appetite for a single vendor may be softening, and not only at the institutional scale. Europe's largest AI Factory procurement has just gone the same way: LUMI-AI will run on AMD Instinct MI430X accelerators rather than NVIDIA silicon. "Although that's changing now, we're seeing growing interest in AMD and other technologies," Richards says. His team runs a deliberately unstructured test bed to encourage it: no queuing system, no production constraints, just hardware handed to individuals or groups who can reimagine it as they like. One group treated an MI300 badly enough that it was simply wiped and restarted. "There's no attempt to run this in any kind of production sense or lock it down; it's use it, tell us what you think, let's learn from it."

Why rising energy prices mean CPUs trump GPUs

For Tom Griffin, director of STFC's Scientific Computing Department, the cost pressure arrives on a second front. Hardware is expensive to buy, and in the UK it is expensive to run.

"The other challenge we specifically have in the UK, to be blunt, is energy costs," he says. "Compute efficiency is more important than raw performance in many ways. In many cases, we are looking for the most power-efficient way to deliver what we need, and in some cases, CPU is still the answer." A significant amount of STFC's work runs on ARM rather than GPUs for the same reason. "It's not always about raw performance for us; it's about power

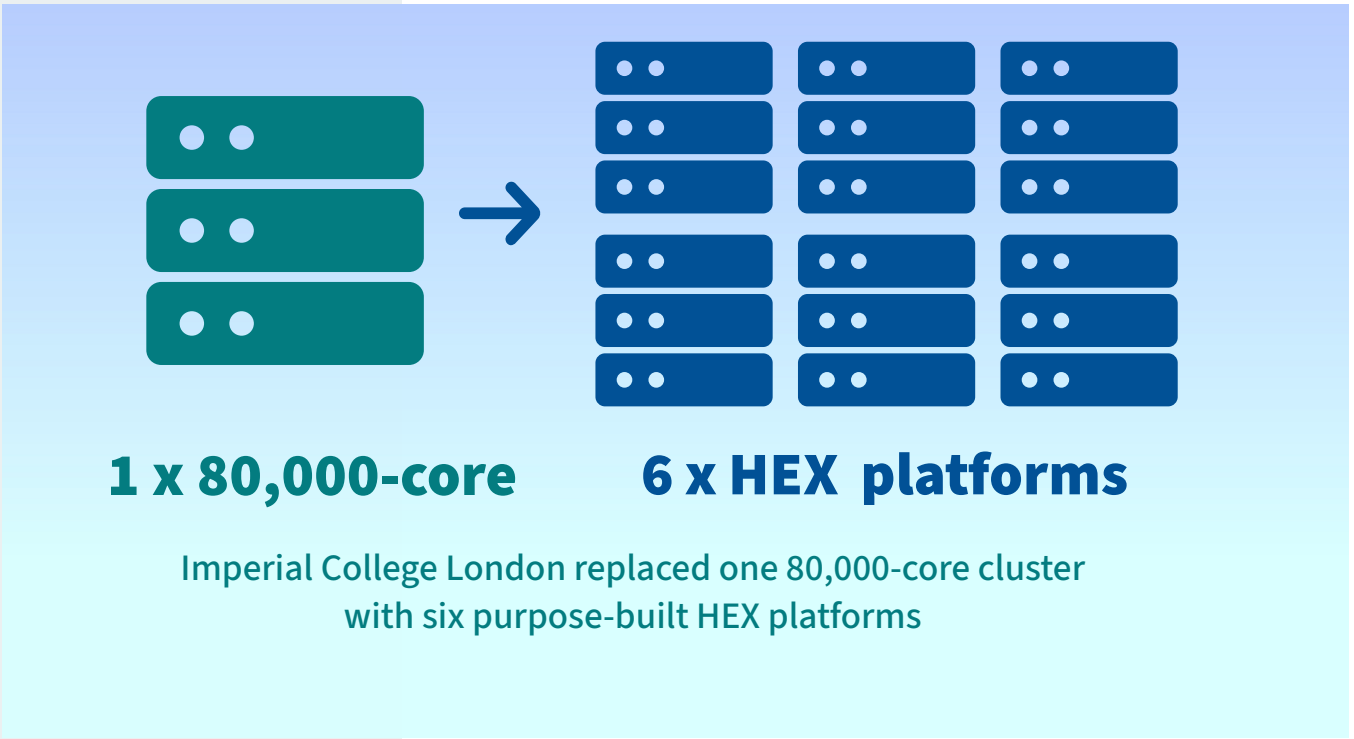

efficiency, cost efficiency, and total cost of ownership. Often that means not buying the newest and sexiest chips."

The scale-out side of the estate is governed by the same logic. "A lot of what we do for high-data-intensity compute is still scale-out, large Ceph installations," Griffin says. "For the really large-scale data stuff, we simply don't have the resources to do that; we have to use lower-cost technologies to get the scale-out benefit."

This makes lock-in a financial risk as much as a technical one. "We are trying to avoid device locking and vendor locking as much as possible. That is a definite concern for science, where codebases often live for very long periods, and we don't want to be locked into anything. We are seeing such a rapid rate of change compared to the last two decades that it is a new challenge."

Griffin also flags a budget line that rarely makes the headlines and is growing, regardless of what happens to memory prices. "The other thing that is often not mentioned is networking. If I look at the fraction of our budget spent on networking, it has increased significantly due to higher data volume, and the network is absolutely critical to operations. Allied to that is cybersecurity and increasing pressure there, so more money is going toward security networks and firewalls."

"Five years ago, you could have a rational conversation about which GPU to buy and look at the specs"

What rising HPC costs mean for future systems

The LUMI-AI announcement trailed by Koski arrived on 31 August, when the EuroHPC Joint Undertaking signed a €387.8m procurement contract with Bull covering acquisition, delivery, installation and maintenance. The system will be the computing backbone of the LUMI AI Factory, deployed at CSC's new Kajaani data centre in the second half of 2027, and is expected to deliver a tenfold increase in AI capacity while nearly doubling the HPC capability of the current LUMI. Funding is split evenly between EuroHPC and the six-country LUMI AI Factory consortium of Finland, Czechia, Denmark, Estonia, Norway and Poland.

The specification is worth reading against Koski's comments about optimising within constraints. The liquid-cooled machine is built on Bull's BullSequana XH3500 architecture, featuring AMD Instinct MI430X GPUs and 6th Gen AMD EPYC 256-core processors, with IBM Storage Scale providing storage, Nokia contributing data centre networking, and Bull's BXI interconnect. Node density, memory configuration and node count were, by Koski's own account, all live variables while the market moved underneath the procurement.

Elsewhere, UCL has funding for GPUs in the coming year. Imperial expects to keep growing its platforms once conditions ease.

But the shape of what gets built has changed. GPU tiers deferred, flash tiers diluted, memory specifications abandoned because the parts cannot be bought at any price, and system sizes set by a capital figure rather than a scientific requirement. Everyone is hoping for the same thing: that the window reopens before the next refresh comes due. As Richards puts it, the plan is to ride this out and hope things are better in 12 months.

The alternative, as Koski notes, is waiting indefinitely. Nobody has that long. ■

HAMMERSPACE

Your **Cluster** Is Fast. Is Your Data?

Stop letting fragmented storage and slow I/O hold back scientific computing.

Simulation. Checkpointing. AI training. Visualization. Collaboration. Archive.

Hammerspace automatically places active data closer to CPUs and GPUs while providing unified access across storage infrastructure, so applications and researchers get fast access to the data they need without disruptive migrations or workflow changes.

ONE PLATFORM. FROM TIER 0 TO ARCHIVE.

- Accelerate I/O-intensive workloads
- Unify existing storage without replacing infrastructure
- Put active data closer to compute
- Automate data placement across storage, sites, and clouds

4.4× Higher Throughput

56% Lower Latency

66% Faster Checkpointing

Accelerate Your HPC Infrastructure

Any Storage. Any Cloud. Anywhere.

SC26 CHICAGO NOV 15–20

Visit us at SC26 - Booth #4238

See Hammerspace for HPC

More data than we can keep

Scientific instruments are producing data faster than storage systems can absorb it and the AI boom has devoured cheap capacity. Experts from NASA and STFC tell **Robert Roe** what happens when science generates more data than anyone can afford to store

Growth in scientific data is gathering pace, making it increasingly difficult to manage the data being generated. Scientific instruments are improving faster than the infrastructure built to catch what they produce. Detector pixel counts, sequencer throughput, sensor array density and simulation resolution have all risen sharply over the past decade, while the cost per byte of storage has fallen more slowly and the bandwidth available to move data off an instrument has barely kept pace. The result is a widening gap between what can be measured and what can realistically be stored and effectively managed.

The conventional model of scientific computing assumed that the gap could be closed by buying more disk space – record everything the apparatus emits, archive it, and let researchers work through it. That model is now failing at several points at once. Archives are a recurring cost rather than a one-off purchase, since data must be migrated, replicated and curated for as long as it remains useful. Even where capacity exists, the pipes between the instrument and the data centre often do not: a detector generating terabits per second cannot simply ship its output across a wide-area network. And the sheer volume of an undifferentiated archive imposes its own tax, because finding the interesting fraction later can cost more than the storage itself.

Real-time analysis offers a potential solution, but this also creates challenges with keeping data generation and compute elements tied closely together. Instead of storing first and deciding later, the decision about what matters is made while the data is still in flight, close to the instrument, and only the selected fraction is committed to permanent storage.

Particle physics has worked this way for decades and offers a clear illustration. The LHC produces bunch crossings (the rate at which particle bunches cross each other at the centre of a detector every second) at around 40 MHz, a raw rate far beyond anything that could be recorded. The trigger and data acquisition (DAQ) system uses a hardware trigger built from FPGAs or ASICs to make a keep-or-discard decision within microseconds, cutting the rate by several hundredfold, and a software farm then reduces it further before anything reaches disk. The overwhelming majority of collisions are never stored and cannot be recovered.

Tom Griffin, Director of STFC's Scientific Computing Department, commented on the growth of data coming from scientific instruments in the UK, such as the Diamond Light Source. He noted that some beamlines are growing at a rate far beyond 10x. "I wish [data growth] was only 10 times as much. Detectors are getting bigger, pixels are getting smaller," and time resolution is compressing as well, "with detectors operating at 100Hz or higher. The source's brightness is increasing, so you have fewer zeros and a lot more ones and noise. Add all those together, and a 10-times increase in data generation is a dream; the reality is even worse."

At NASA, the totals are already large and moving. Andrew Mitchell, Deputy Chief Science Data Officer within NASA's Science Mission Directorate, puts the current figure at "over 150 petabytes of scientific data representing 50,000 unique science datasets" across Earth science, astrophysics, biological and physical sciences, heliophysics and planetary sciences.

"NASA missions are collecting more data at higher resolution over longer periods of

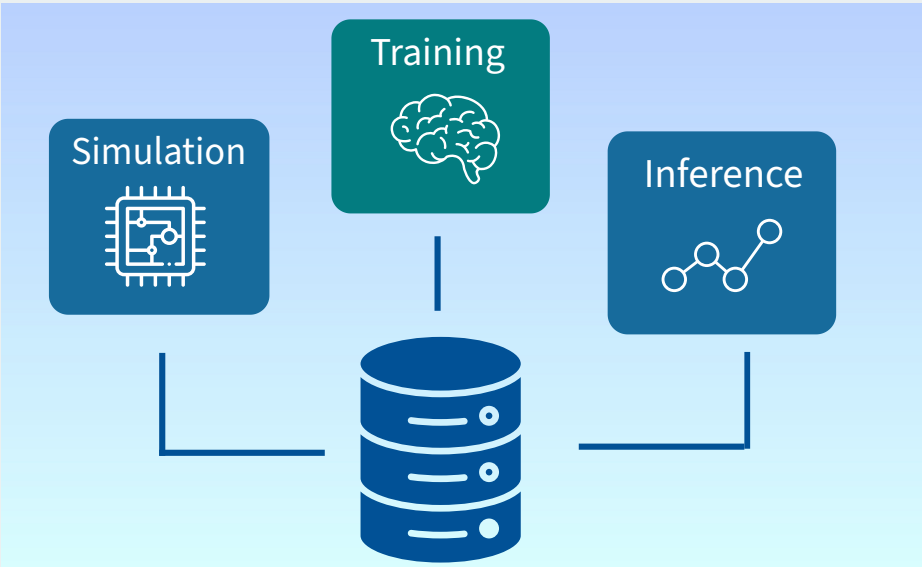

LUMI Software-defined supercomputing

CSC is turning LUMI into an AI Factory that runs simulation, training, and inference on shared hardware

created in Adobe Express

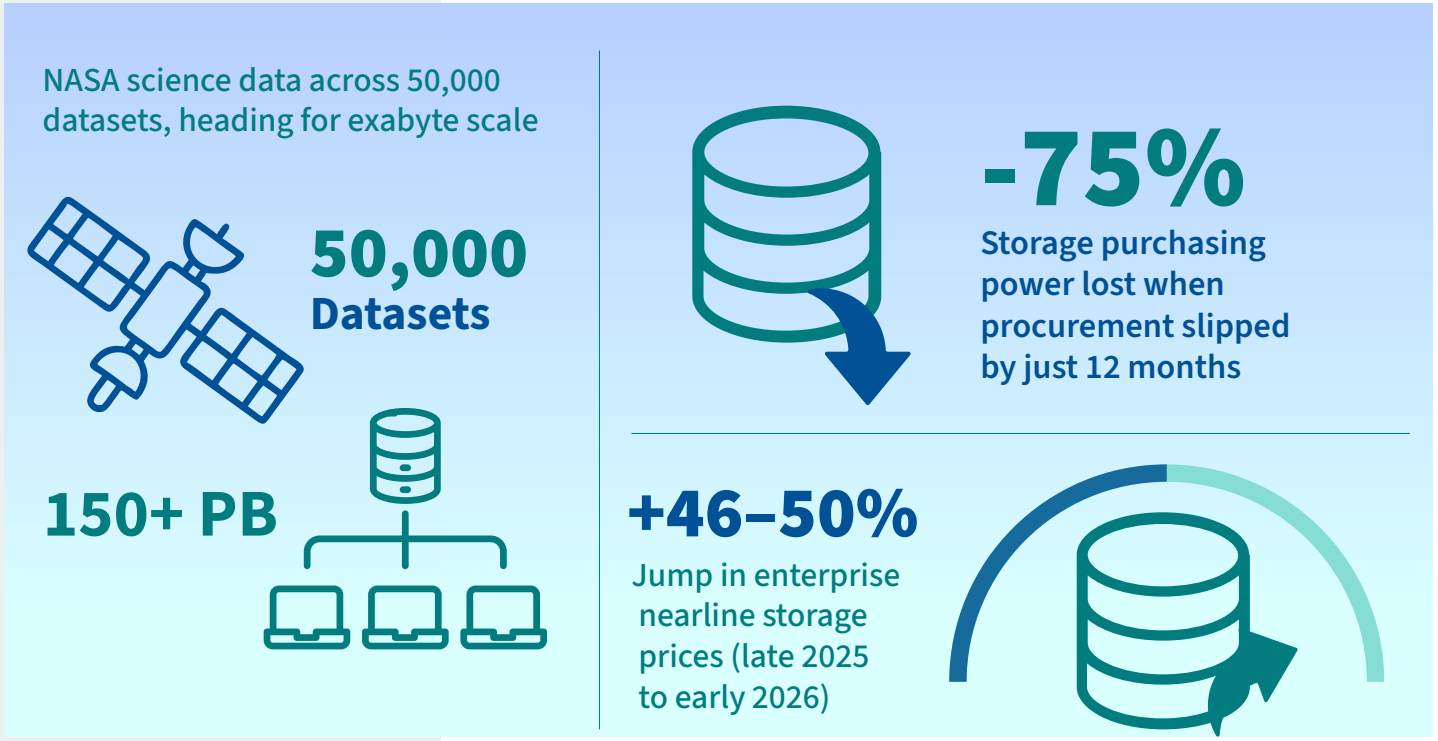

time than ever before," he says. The growth is not only in volume. "Our scientists expect to work with that information more interactively than they did previously when things were on-premises."

Scientific storage costs are rising

The second half of the problem arrived from an unrelated direction. Bulk capacity, the one part of the stack research computing could historically rely on getting cheaper, has become scarce.

Western Digital told investors it was effectively sold out for calendar year 2026, with firm orders from its largest customers and long-term agreements extending into 2027 and 2028. Reported pricing moved roughly 46-50% between September 2025 and early 2026. The reason? Hyperscaler demand for nearline capacity to hold training data, model checkpoints and logs. Enterprise flash moved further and faster because the same wafer capacity was being redirected toward high-bandwidth memory. Analysts at Omdia downgraded their outlook for the general-purpose enterprise storage market even as they raised overall data centre capital expenditure forecasts.

The practical consequence for a science facility is that the traditional assumption that next year's capacity will cost less per terabyte than you paid this year can no longer be relied on.

Griffin's response is to buy the most cost-effective technology. High-data-intensity compute at STFC is "still scale-out", he says, built on "large Ceph installations", while "for other AI workloads, there are more turnkey solutions for niche use cases". The reasoning is budgetary rather than architectural. "For the really large-scale data stuff, we simply don't have the resources; we have to use lower-cost technologies to get the scale-out benefit."

There is a related discipline in how those systems are procured. "We are trying to avoid device locking and vendor locking as much as possible," Griffin says. "That is a definite concern for science, where codebases often live for very long periods, and we don't want to be locked into anything. We are seeing such a rapid rate of change compared to the last two decades that it is a new challenge. AI tools like Claude can help with that, but it's not a panacea."

All scientific data can no longer be stored

If the data is growing faster than the budget, something has to give, and increasingly it is the assumption that raw data is kept by default.

Griffin is direct about where this leads. "The reality is we are looking at different ways of operating. It is no longer the case that we can store all data," he says. "In high-data areas, we need to adopt approaches used in particle physics, where AI-enabled or AI-delivered rapid data analysis effectively discards the first-pass data and retains only the most relevant or scientifically meaningful data. Obviously, that's a cultural change for communities – that requires us taking people with us.

"To be clear: for some beamlines, this isn't likely to ever be a problem," he says. Those teams will carry on analysing everything they collect. "But for high-data-rate beamlines, we need to use a triggering-like approach to discard data that is definitely not interesting," he adds, giving the example of "a frame with no crystal in it".

"We also have to look at keeping a fraction of all data so you don't throw everything

"Add all those together, and a 10-times increase in data generation is a dream; the reality is even worse"



away and can go back over it, not least for future training of AI models to increase efficiency."

NASA is facing the same pressure from the processing side as from the acquisition side. "Some of the other things we are looking at include on-the-fly or on-demand processing, so we do not have to store all levels of our datasets at the exact same time," Mitchell says. "That hybrid approach spans not just the infrastructure itself, but also how we manage that data and make it available. We have also been working with our international partners on those strategies."

HPC storage is moving between on-premises and cloud

Neither organisation is solving this with a single system, and both describe an architecture that spans on-premises and cloud rather than choosing between them.

"On the cloud side, we heavily use services like S3, and we use deep storage options like Glacier," Mitchell says. "For our on-premises high-end computing, we utilise scratch storage. It is a hybrid approach to managing that 150 petabytes, which will soon grow to exabyte scale."

He is candid that stitching those environments together is where the difficulty actually lives. Integrating HPC and cloud, he says, is "neither easy nor automatic, presenting real challenges. First, managing data movement costs across environments at petabyte and exabyte scales. Second, aligning security accreditation processes built for different infrastructure models, with the industry still catching up. Third, building a workforce with expertise across both niche environments."

Scientific data storage needs to support AI and real-time analysis

Underneath all of this is a change in what users expect from a storage system. Archival capacity is no longer sufficient on its own, because researchers want to work against large holdings in near real time.

Mitchell describes the target as convergence rather than coexistence. NASA is "working toward an integrated computing environment where compute, storage, AI capabilities, and network services are not just designed together, but managed together rather than as separate systems".

AI is a substantial part of why. NASA has released foundation models for Earth observation and solar physics, and Mitchell frames the shift in terms of infrastructure: "We are no longer designing infrastructure solely to run simulations; we design it to support the complete scientific workflow combining observation, simulation, AI, and visualisation."

That workflow places different demands on storage than a simulation campaign does. Training and inference want throughput and random access against large corpora rather than sequential streaming from tape. "Where traditional scientific workflows focused on CPUs, there is now increasing demand for GPU acceleration, training systems, massive storage, and high-speed networks," Mitchell says. He adds a caveat that applies squarely to any data reduction strategy: "We must ensure those models are verified, validated, and applied to support sound science rather than simply speed."

Griffin, meanwhile, is sceptical that this capability is genuinely available to everyone who wants it. "Democratisation of HPC was a term used a few years ago, but I don't think it ever really became true," he says. "If you think about the UK's investment in national HPC, the intention is to serve very large use cases, not thousands of users doing small jobs. There is a trade-off between the very large leading science problems you want to solve versus availability for everyone."

The same stratification, he argues, now applies to AI: "Lots of people are using AI, but not very many are using AI at massive scale. The level of expertise, skill, and need is quite different."

How scientific storage strategies are changing

There is one piece of good news in Mitchell's account: the argument for investment has become easier to make.

"If you had asked me that a couple of years ago, I would have said yes," he says of whether it is hard to convey the challenge upward. "But management, industry, and the world are coming around to the importance and need for data and data storage. This is not a NASA-specific problem."

That is largely an AI dividend. Executives who have spent two years being told that data is the strategic asset are more receptive to a request for storage than they were when it was a line item on a science budget.

What the money buys, though, looks different from what it bought five years ago. Less of it goes on keeping everything and more on deciding what to keep. More of it goes on network and security than the storage label implies. Some of it goes on cheaper, less impressive technology chosen for scaling out affordably rather than for benchmarking well. And a growing share goes on the software and the cultural work required to persuade research communities that discarding a frame with no crystal in it is not a loss of science.

The instruments are not going to slow down. The storage market is not going to relax on the timescale that facilities plan on. What is left is the decision about which data is worth the money, and that is a conversation with scientists rather than with vendors. ■

SEAGATE | nexstor | SCIENTIFIC COMPUTING WORLD

We asked the scientific computing community. Here's what they told us.

Scientific data is growing fast. AI and simulation are reshaping workloads. And the infrastructure underneath it all is under pressure.

What the survey found:

77%

say data movement is challenging.

47%

cite storage performance as a bottleneck.

60%

say simulation is putting pressure on storage.

83%

lack confidence that their infrastructure is ready for what's next.



What are organisations doing about it?

Find out in *The State of Scientific Data Storage 2026.*

Silicon roadmaps diverge as simulation demands precision

Can low-precision AI hardware truly replace native double precision in scientific simulation, or are procurement teams right to demand FP64? **Robert Roe** reports

As the GPU market is dominated by two players, their architectural choices have a dramatic impact on technology choices made for scientific computing. As NVIDIA moves increasingly towards mixed precision performance, are traditional simulation workloads that require double precision arithmetic (FP64) a thing of the past?

On 31 August, the EuroHPC Joint Undertaking signed a €387.8m contract with Bull for LUMI-AI, the successor to Europe's most prominent AMD-based system. The machine will land at CSC's new data centre in Kajaani in the second half of 2027, built on BullSequana XH3500 racks with AMD's Instinct MI430X GPUs and 6th Gen EPYC 256-core processors, and is billed as delivering roughly 10 times LUMI's AI capacity without quite doubling its high-performance computing capability.

CSC's own announcement acknowledges it directly: new GPU architectures have advanced far faster on AI-driven tasks than on the highly precise floating-point arithmetic that scientific simulation depends on. Ten times the AI. Not quite twice the science.

Which raises the question that procurement teams across Europe and the US have tried to answer for the last two years. If the silicon roadmap is sprinting toward FP8 and FP4, is IEEE 754 double precision still a requirement?

Why AMD and NVIDIA are taking different approaches to FP64

The numbers are not subtle. NVIDIA's B200 retained approximately 40 TFLOPS of dense FP64 tensor throughput. Blackwell Ultra, the B300, rebalanced its tensor cores for NVFP4 and dropped to roughly 1.2 TFLOPS dense FP64 per GPU, according to the published datasheet – a reduction of more than an order of magnitude in a single half-step of the same architecture. At SC25, while presenting the TOP500 list, Jack Dongarra pointed out that Blackwell's 64-bit floating-point capability had not improved over that of Hopper at all.

AMD has gone the other way. The Instinct MI430X, due in 2027, is specified to deliver up to 288 TFLOPS of hardware-based peak FP64 performance, with 432GB of HBM4 and 23.3TB/s of memory bandwidth. AMD's own comparison puts that at 8.7 times the FP64 of NVIDIA's Vera Rubin, which it pegs at 33 TFLOPS vector. Alongside LUMI-AI, the part is slated for the US Department of Energy's Discovery system at Oak Ridge, France's Alice Recoque, and HLRS's Herder.

Why scientific computing still needs native FP64

Pekka Manninen, Director of Science and Technology at CSC, was blunt about where precision sat in the LUMI-AI design process.

"That was one of the big questions during the design phase," he says. "HPC centres and the scientific computing user base use legacy software; there's a lot of software containing decades of built-up wisdom, and the user communities trust it. It's very much about trusting the results and knowing how they work, so even though software development is very easy today with AI agents and so on, we're still not at the point where we could always build a new simulation suite from scratch, because of all the validation and verification needed; it has to match scientific tradition."

The word doing the work there is trust. A climate model or a lattice QCD code is not simply a program; it is an accumulated argument about physics that has been checked against observations, published, and replicated. Rewriting it to tolerate lower precision requires rerunning the entire validation chain.

"There's a lot of need to support legacy simulation software, and many algorithms are sensitive to precision," Manninen continues. "Native FP64 precision is still important in this generation of systems for everything I know of in scientific computing, though there may be different opinions in the field, and whether that's still true five years from now, I don't know."

He is more direct about the procurement consequence: "Looking at my user base today, if we bought a system with no FP64 support, I don't think it would be well received."

How FP64 is changing the HPC GPU market

The interesting second-order result is what this is doing to vendor relationships. For a decade, accelerated HPC in most universities has meant CUDA [NVIDIA's development environment for creating GPU-accelerated applications], and the path of least resistance has been NVIDIA.

Andrew Richards, Director of Research Computing Services at Imperial College London, describes an estate now actively looking around. Imperial announced a strategic collaboration with AMD in June, covering AI-enabled scientific discovery, sovereign AI infrastructure, and next-generation HPC, spanning multiphysics simulation, materials discovery, climate and earth system modelling, genomics, and more.

"There's a direction of travel where, whilst the NVIDIA ecosystem is great and easy to adopt (CUDA's been around a long time, even though it changes a lot under the covers, and NVIDIA have done a very good sales job keeping people engaged), Intel and AMD have a lot to learn in that regard," Richards says. "But there's a real competitive push now to understand where other technologies can fit. With NVIDIA's direction of travel away from FP64, there's growing interest in the community in what else is out there, because we still need FP64; we're not all doing AI workloads where we can afford lower precision."

Note the sequencing. The precision roadmap is not merely annoying users; it is opening a door that had been effectively

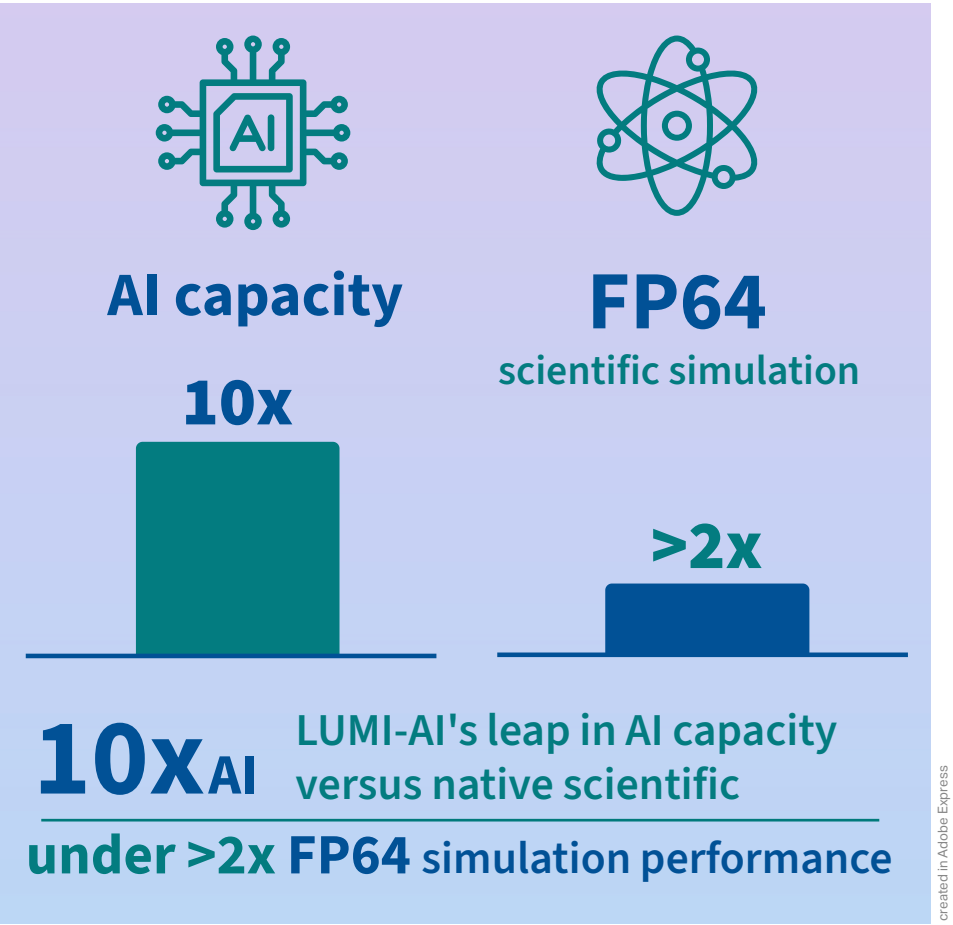

shut. “So there’s growing interest in seriously exploring other technologies and how we can bring them into our estate,” he adds. “That’s where AMD seems to be stepping in at the moment. ROCm [AMD’s open-source software stack for GPU programming] has improved significantly, and they’re having the right conversations with academics about what this means.”

Richards is markedly less warm about the third player. “In the past, we’ve seen Intel try to have these conversations, but they seem to keep losing conviction to stick in one direction long enough to make it work.”

He also points to something procurement narratives usually miss. This is happening as a community rather than institution by institution. “Places like Cambridge have a lot of AMD activity, and there’s intent (with the fusion community out at UKAEA and the links there) to build up use of AMD technology and make it more applicable in some of those domains. There’s a growing community in the UK interested in this, and if you look at some of the fusion work at UKAEA, there are links to Imperial as well. There’s a lot of commonality in the research base across multiple institutes. It’s something we’ll probably see develop further over the next year.”

“Looking at my user base today, if we bought a system with no FP64 support, I don’t think it would be well received”

Why AMD is doubling down on FP64 for HPC

From the vendor side, the obvious question is why a company whose revenue growth is overwhelmingly driven by AI would spend R&D time on a data type the AI market does not use. Mark Papermaster, CTO and SVP at AMD, frames it as a consequence of packaging rather than sentiment.

“It is indeed strategic for AMD,” he says. “We made a commitment to the high-performance computing and AI-for-science market when we launched the MI300A back in 2024 that we’d continue that support. The way we’ve been able to do that is by leveraging our modular approach to generate a version of our GPU that supports FP64 by swapping out a single chiplet, while keeping the rest of the infrastructure the same as in the much higher-volume AI market, which doesn’t require that highest precision.”

That is the commercial trick, and it explains how a specialist part like the MI430X can exist alongside the MI455X without a separate programme’s worth of investment. “The very underpinnings of how we architected our implementation gave us the flexibility to support the HPC and AI-for-science market with native, high-precision FP64 in a way our competitors couldn’t match without a much more significant investment,” Papermaster says.

He argues the broader trend is toward specialisation, not consolidation. “What you’re seeing now is unique configurations required for agentic AI throughput, unique configurations for building up AI training and very-large-language-model inferencing at a head node with a GPU or accelerator

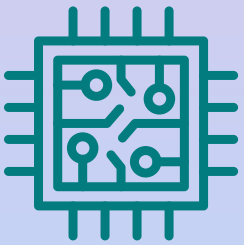

256
cores per socket

AMD EPYC Venice (2nm), with PCIe Gen 6 and CXL 3.1



up to 288

Peak FP64 on the Instinct MI430X, with 432GB HBM4 and 23.3 TB/s bandwidth, 8.7x the FP64 vector performance of Vera Rubin

created in Adobe Express

“When you’re accelerating discovery, trying to accelerate solving fusion, doing new molecular studies, or modelling the atmosphere and weather more accurately, that all requires high precision”

cluster, and unique configurations required by sovereign nations and HPC labs for AI for science.” The answer, he says, “is a chiplet approach, tailoring the engines, the I/O and memory configuration, and the very math formats you support, to more efficiently serve the algorithms at hand for a given market.”

On the science itself, he is unequivocal: “When you’re accelerating discovery, trying to crack how to accelerate solving fusion, doing new molecular studies, or modelling the atmosphere and weather more accurately, that all requires high precision, native FP64, the highest double-precision floating point. You can’t get there with AI approximations alone.”

The emulation argument – can low-precision GPUs replace native FP64?

NVIDIA, unsurprisingly, disputes the framing. The company has said publicly that it is not abandoning 64-bit computing, that FP64 remains core because simulation provides the ground truth against which AI surrogate models are trained and validated, and that future generations will improve. Its answer, in the meantime, is emulation: using the enormous low-precision tensor throughput to reconstruct double-precision results via CUDA-X library support and schemes derived from the Ozaki approach. NVIDIA has cited a 1.8x speedup for FP64 matrix multiplication through emulation in cuBLAS, and reported figures suggest emulated FP64 matrix performance on Rubin is in the region of 200 TFLOPS, compared with roughly 33 TFLOPS native.

There is real research behind this. Recent papers demonstrate error-corrected Ozaki variants that match DGEMM accuracy with modest overhead, and at least one argues bluntly that hardware FP64 is no longer the irreducible foundation of scientific computing. Others are sceptical, including researchers at AMD, who have suggested the techniques are worth investigating but not yet production-ready.

The practical objections are well rehearsed. Emulation accelerates dense matrix multiply, which is not where most HPC codes spend their time; a great deal of scientific computing is memory-bandwidth-bound with low arithmetic intensity, and a faster GEMM does not help a sparse solver or a stencil kernel. Edge-case semantics differ. And an application that produces a plausible but subtly different answer is a worse outcome for a scientific user than one that runs slowly.

Will HPC still need FP64 in five years?

None of this means FP64 is permanent. Manninen, who has more reason than most to consider what the machine after LUMI-AI will look like, is open to the idea that the requirement may fade.

“There may be a paradigm shift as algorithms become more tolerant of lower-precision arithmetic, which would be a good thing if it happens,” he says. “We can also generate software much faster than ever before, so we may increasingly move toward rewriting even the most iconic software packages from scratch, which opens possibilities for algorithmic changes and for leveraging ultra-fast, low-precision compute.”

That is a genuinely interesting scenario, and it is one in which AI eliminates the FP64 requirement on the software side rather than the hardware side. The barrier to mixed-precision science has never really been that researchers preferred 64 bits; it has been that rewriting and revalidating a 40-year-old code base was economically impossible. If that cost collapses, the argument changes.

Papermaster’s framing points in the same direction from a different angle, describing HPC and AI for science as increasingly a single market driven by a shared demand for efficiency. AMD’s own positioning for the MI430X is convergence: FP64 for the simulation that generates trustworthy training data, and modern low-precision formats for the models trained on it, in the same package.

For now, though, the answer to the question is straightforwardly yes. Discovery, Alice Recoque, Herder, and LUMI-AI have all been specified with native double precision and will run through the early 2030s. The users asked for it, the codes require it, and one vendor found a way to supply it without having to rebuild its roadmap. The five-year question stays open. The procurement question, for this generation, is closed. ■

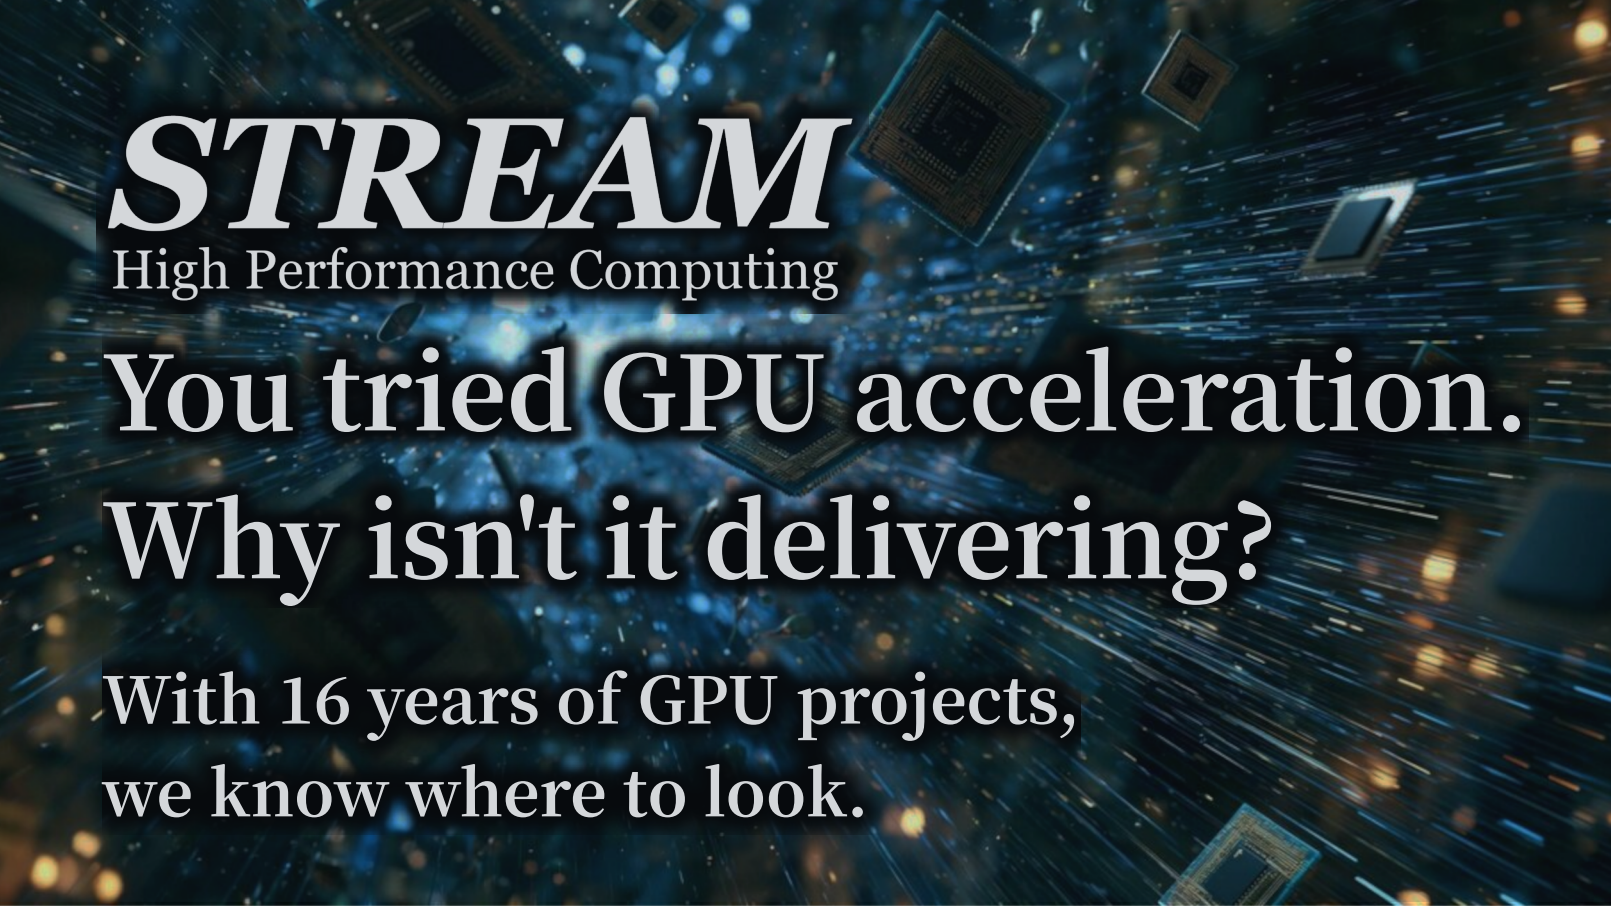

Our **GPU Performance Baseline & Bottleneck Audit** gives you the answers, backed by numbers. "What's the potential?" "What would it take?" "Is our approach right?"

Fixed price. Fixed scope. Typically two to four weeks.

Book a one-hour intake.
Tell us your case.

intake@streamhpc.com or scan the QR-code

We help R&D-intense companies, research institutes and semiconductor companies. Those who need fast software to keep ahead of the market. We do this via GPU software engineering, performance audits, trainings and co-research.

How can supercomputers cut their water footprint?

Data shared with the Energy Efficient HPC Working Group by 11 top HPC facilities shows why cooling architecture choice is crucial to balancing energy demands with water conservation. **Robert Roe** reports

As artificial intelligence and high-performance computing (HPC) workloads grow denser and hotter, the conversation about their environmental footprint is expanding beyond electricity.

Supercomputing sites feel the squeeze earlier than most because they have always run at densities that the commercial AI build-out is only now reaching. El Capitan draws close to 30MW and Aurora nearly 39MW. Preparing Livermore's computing centre for exascale output meant taking its electrical capacity from 45MW to 85MW and installing half a dozen 3,000-ton cooling towers. At that scale, small differences in cooling philosophy translate into millions of litres a year, or into megawatts.

A new study, compiled by the Energy Efficient HPC Working Group with contributions from practitioners at some of the world's most powerful supercomputing sites, turns the spotlight on water, a resource that has historically received far less scrutiny than energy in data centre operations.

The study, presented as a scientific poster titled "AI and HPC Water Usage: Findings from 11 Leading Supercomputing Centers," draws on a detailed survey of operational facilities including LLNL's El Capitan, ORNL's Frontier, RIKEN's Fugaku, CSCS's Alps, and eight other flagship systems. Together, these machines rank among the highest-ranked supercomputers in the world for both raw performance (TOP500) and energy efficiency (GREEN500). The findings suggest that while the HPC community has made real progress in tracking energy metrics such as Power Usage Effectiveness (PUE), water management remains comparatively immature, is measured inconsistently, if at all, and is often poorly understood even by the facilities themselves.

Cooling has always been central to supercomputing. As chips and accelerators pack more transistors into smaller spaces, they generate more heat per square metre, and someone has to remove it. For decades, the go-to solution has been evaporative cooling towers, which are energy-efficient, but consume significant volumes of water through evaporation. As AI training clusters push power densities even higher, the water demands of cooling infrastructure are becoming a serious consideration for facility siting, permitting, and community relations, particularly in regions already under water stress.

Despite this, water use in data centres has lacked the kind of standardised tracking that energy has enjoyed for years. This gap prompted a collaborative team of HPC practitioners, facility managers, industry vendors, and academic researchers to develop a 35-question survey covering cooling architecture, annual water consumption, water usage effectiveness (WUE), blowdown, water chemistry, sourcing and the barriers facilities face in reducing their water footprint.

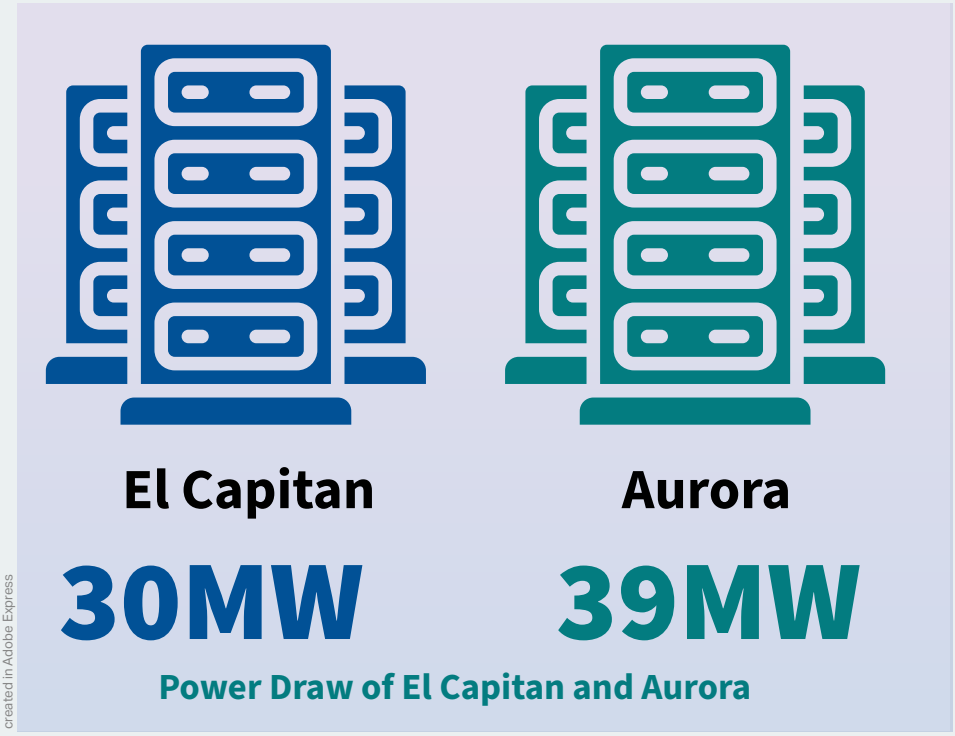

Power Draw of El Capitan and Aurora

created in Adobe Express

How 11 supercomputing sites use water for cooling

The survey was distributed to 23 major supercomputing sites worldwide. Twelve responded, and one was excluded because its facility was still under construction, leaving 11 operational sites for analysis. The respondents span a broad range of geographies, water-scarcity conditions, and cooling philosophies: Lawrence Livermore National Laboratory (El Capitan), Oak Ridge National Laboratory



(Frontier), RIKEN (Fugaku), the Swiss National Supercomputing Centre (Alps), KAUST (Shaheen III), Los Alamos National Laboratory (Venado), TU Dresden (Capella), Idaho National Laboratory (Teton), the Leibniz Supercomputing Centre (SuperMUC-NG), the National Laboratory of the Rockies (Kestrel), and the European Centre for Medium-Range Weather Forecasts (HPCF).

All data were self-reported by named site representatives, and the authors note no systematic bias was identified, although they acknowledge scope for refining future surveys.

Why cooling architecture matters more than location for water efficiency

Perhaps the study's most striking finding is that a facility's cooling architecture is more closely linked to its water efficiency than the underlying water stress of its location, as classified by the World Resources Institute's Aqueduct 4.0 dataset. In other words, where a data centre is built matters less to its water footprint than how it is cooled.

The two facilities using hybrid or adiabatic cooling reported by far the lowest WUE values; TU Dresden's Capella system recorded just 0.147 litres per kilowatt-hour, and ECMWF's HPCF reported 0.471 L/kWh. Lower WUE indicates better water efficiency, so both figures represent strong performance.

By contrast, sites relying on evaporative cooling generally reported considerably higher water use. The National Laboratory of the Rockies reported 2.09 L/kWh, while RIKEN's Fugaku, one of the most powerful machines in the world, reported between 4.2 and 4.7 L/kWh. LANL, LRZ, and ORNL each reported WUE values above 1.5 L/kWh.

CSCS stands apart from both categories. Its Alps system draws roughly eight million cubic metres of water annually from Lake Lugano in a once-through cooling arrangement, returning the water directly to the lake rather than losing it to evaporation. This unusual configuration makes direct WUE comparisons difficult, and CSCS is one of three sites, alongside LLNL and KAUST, that does not currently measure WUE.

How blowdown affects cooling tower water use

Cooling towers lose water in two main ways: evaporation, which actually cools the system, and blowdown, the deliberate discharge of water to prevent mineral buildup as dissolved solids concentrate during repeated evaporation cycles. The relationship between the two is governed by CoC [cycles of concentration] – essentially, how many times water can be recirculated through the tower before it must be flushed and replaced.

Across the surveyed evaporative sites, blowdown accounted for roughly 15-33% of total facility make-up water, a substantial share of overall consumption. The National Laboratory of the Rockies' Kestrel system offers a useful illustration. Operating at a CoC of approximately seven, the facility used 32,986 cubic metres of makeup water annually, with blowdown accounting for 4,744 cubic metres, or about 14.4% of the total. The poster's authors calculate that if CoC were reduced to three, the blowdown fraction would nearly double, to roughly 33%. Put simply: pushing >

"Where a data centre is built matters less to its water footprint than how it is cooled"

created in Adobe Express

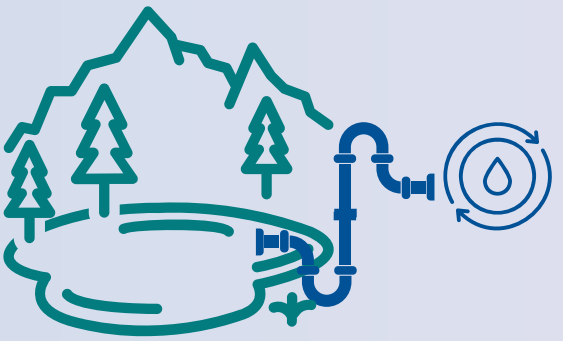

8million m³

Water drawn yearly from Lake Lugano to cool CSCS Alps, returned with no evaporative loss

Share of makeup water lost to blowdown in evaporative cooling

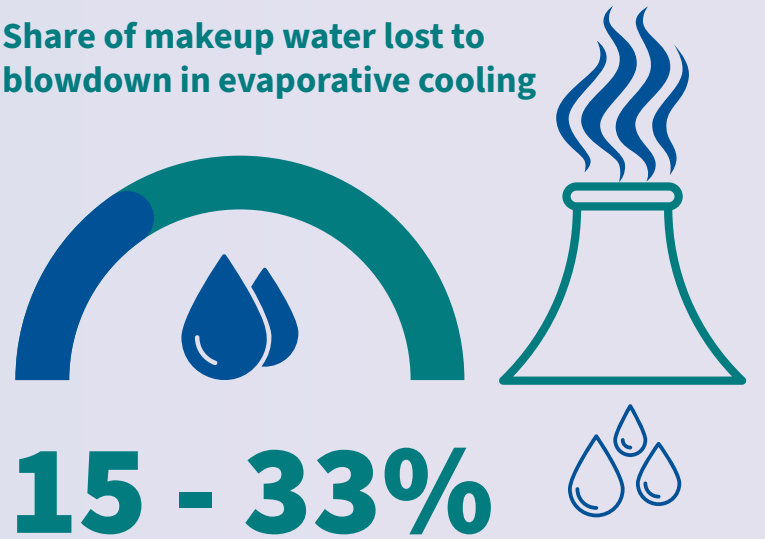

15 - 33%

> CoC higher squeezes considerably more value out of every litre drawn into the system.

However, higher CoC isn't simply a dial to be turned up at will. Water chemistry imposes a ceiling. The survey found corrosion issues at seven of the 11 sites, scaling at five, and microbiological growth at four, all of which can degrade heat-transfer efficiency and equipment reliability and tend to worsen as CoC increases. Chemical treatment is therefore essential to sustaining higher cycles of concentration without compromising system performance.

Why water-efficient cooling can increase energy use

The report also highlights a trade-off that complicates any simple push towards lower water use: some water-saving strategies, such as dry cooling or reverse osmosis treatment, increase energy consumption. Because weather conditions change throughout the year, several facilities adopt seasonal or hybrid strategies, relying on dry cooling for much of the year, and switching to water-assisted cooling only during periods of high outdoor temperature. The clear implication is that water and energy efficiency cannot sensibly be optimised in isolation from one another; facility operators need to weigh both together.

This point extends to a broader observation about liquid cooling within IT racks. Closed-loop systems that circulate coolant directly to chips are often assumed to reduce or eliminate water use, but the survey found this isn't necessarily true. Such systems still need to reject heat somewhere outside the rack, and if that outdoor heat rejection relies on evaporative cooling, on-site water consumption persists regardless of how efficient the in-rack loop itself might be.

The cost and practical challenges of water-efficient cooling

Even where facilities recognise opportunities to conserve water, practical constraints often stand in the way. Physical space and cost were each cited by seven of the 11 sites as barriers to further conservation efforts, a reminder that retrofitting treatment, recovery, or alternative cooling technologies into existing infrastructure isn't always straightforward, even when the engineering case is sound. Reported water costs also varied considerably among the five facilities that provided pricing data, reflecting the different ways water is priced and regulated across regions.

What the study reveals about HPC water efficiency

The authors are careful to frame this as an operational baseline rather than a definitive judgement on any single facility or cooling approach. The data clearly show that water efficiency in AI and HPC facilities is shaped by a web of interacting factors: cooling architecture, local water chemistry, energy trade-offs, measurement practices, and physical or financial constraints. Hybrid and adiabatic systems indicate what's achievable, but evaporative cooling remains widespread, and improving its water efficiency depends heavily on better tracking of blowdown and CoC, as well as more sophisticated water treatment.

As AI infrastructure continues its rapid expansion, this kind of cross-institutional benchmarking may prove just as valuable as the energy-efficiency work that preceded it. Establishing consistent measurement, starting with the basics of tracking WUE and blowdown at every site, appears to be the essential first step towards treating water, like energy, as a resource to be engineered and managed rather than simply consumed. ■

This article is based on a poster produced by the Energy Efficient HPC Working Group's Water Management Team, drawing on survey contributions from representatives at LANL, CSCS, RIKEN, TU Dresden, LLNL, KAUST, ECMWF, NLR, ORNL, LRZ, and INL.

"Water and energy efficiency cannot sensibly be optimised in isolation from one another; facility operators need to weigh both together"

Over a decade delivering fanless designs.

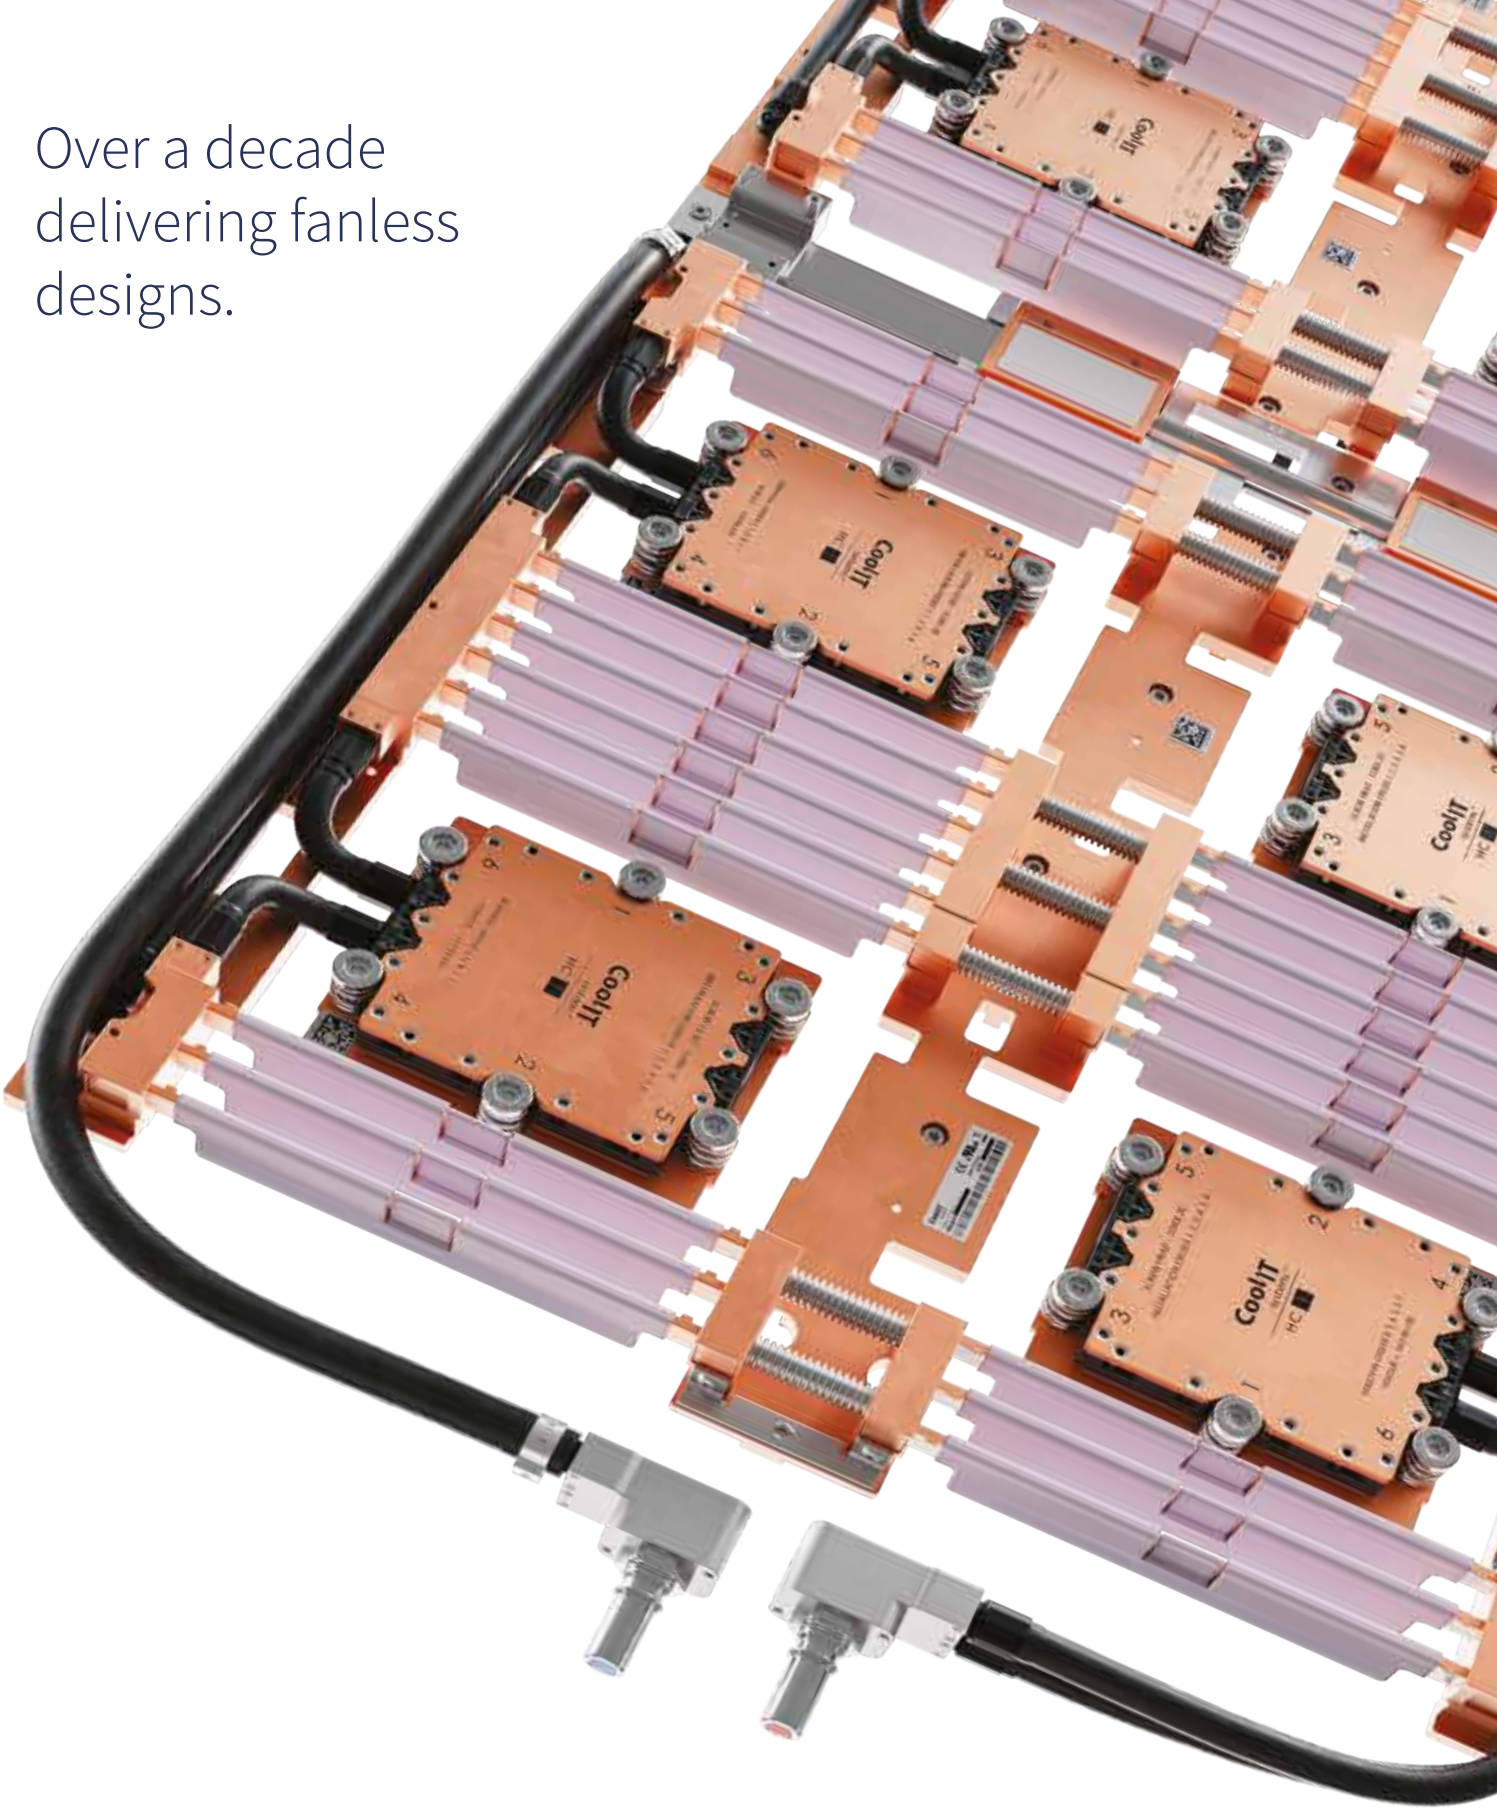

6TH GENERATION
100% HEAT CAPTURE
SUPERCOMPUTER BLADE

CoolIT™
An Ecolab Company

Right-sizing computing environments for NASA science



Andrew Mitchell explains how NASA balances supercomputing with cloud infrastructure to process 150 petabytes of scientific data across five research divisions

NASA's Science Mission Directorate now manages more than 150 petabytes of data spanning Earth science, astrophysics, heliophysics, planetary science, and biological and physical sciences. At that scale, matching the right computing environment to the right scientific problem is not just a technical detail – it is mission-critical. Get it wrong, and simulations get bottlenecked or collaborative science grinds to a halt; get it right, and missions such as Artemis II run on infrastructure built for exactly what they need.

Andrew Mitchell, Deputy Chief Science Data Officer within NASA's Science Mission Directorate (SMD), is a program manager and technical leader in scientific data systems and computing services, overseeing initiatives that enhance data management and computing strategies within NASA's Office of the Chief Science Data Officer. Here, he tells SCW how the agency's hybrid strategy pairs high-end computing for tightly coupled simulations – supporting work such as Artemis II and wind tunnel testing at Ames – with cloud infrastructure for interactive, collaborative science.

What does NASA's Chief Science Data Officer do?

Andrew Mitchell: Within the NASA SMD, we have five science-focused divisions: Earth science, astrophysics, biological and physical sciences, heliophysics, and planetary sciences. Currently, we manage access, distribution, and analysis for around 150 petabytes of scientific data. Within the Office of the Chief Science Data Officer, we coordinate that work, providing guidance to the science divisions on how to manage those data systems.

How do you design a computing strategy that can support NASA users across the scientific areas that you oversee?

AM: We use a flexible hybrid computing strategy. It combines our high-end computing, our managed cloud environments, such as our science cloud, along with our integrated data services and advanced networking to support mission needs across the entire agency. A good example is with our engineers preparing for the recent Artemis II mission. They relied heavily on our high-end computing to run tightly coupled computational fluid dynamics (CFD) simulations. Meanwhile, our Earth scientists studying wildfire damage might use our science cloud to analyse satellite observations or collaborate through cloud-native research environments. Rather than forcing every workload into the same infrastructure, we match the computing environment to the scientific problem we are trying to solve. The goal is to give researchers an environment that best supports the work they are trying to accomplish, whether that is an astrophysicist processing telescope data, a planetary scientist, or an engineer running structural simulations.

How is data growth changing NASA's computing and storage needs?

AM: NASA missions are collecting more data at higher resolution over longer periods of time than ever before. Within the Science Mission Directorate, we manage over 150 petabytes of scientific data representing 50,000 unique science datasets. Our scientists expect to work with that information more interactively than they did previously when things were on-premises. Because of that, we are working toward an integrated computing environment where compute, storage, AI capabilities, and network services are not just designed together, but managed together rather than as separate systems. Another big game-changer for us is the use of AI, machine learning, and advanced analytics in our scientific workflows. Recently, we released two foundation models, one for Earth observation and one for solar physics. These are helping researchers identify patterns, prioritise observations and accelerate scientific advancements.

How does NASA make petabyte-scale data interactive for researchers?

AM: That goes back to the hybrid infrastructure we are building. On the cloud side, we heavily use services like S3 (Amazon's Cloud Object Storage), and we use deep storage options like Glacier (a low-cost cloud storage service from Amazon Web Services (AWS) built for long-term data backup, digital preservation, and archiving). For our on-premises high-end computing, we utilise scratch storage. It is a hybrid approach to managing that 150 petabytes, which will soon grow to exabyte scale.

Is it difficult to keep pace with data growth? How do you convey this challenge to NASA management?

AM: If you had asked me that a couple of years ago, I would have said yes. But management, industry and the world are coming around to the importance and need for data and data storage. This is not a NASA-specific problem. Some of the other things we are looking at include on-the-fly or on-demand processing, so we do not have to store all levels of our datasets at the exact same time. That hybrid approach spans not just the infrastructure itself but also how we manage that data and make it available. We have also been working with our international partners on those strategies.

How does NASA decide which infrastructure fits each workload?

AM: Our decisions are driven by intended outcomes. We evaluate how researchers and engineers will use the data, considering multiple factors: whether they need low latency, large memory, GPU acceleration, AI capabilities, how they will scale in the cloud, and proximity to the data. Some users prefer not to move data. Our measure of success is how effectively users accomplish their mission, not necessarily the technology they use.

For example, in our wind tunnels at the Ames Research Center, engineers performing computational fluid dynamics simulations needed an environment closely tied to the wind tunnels, rather than just the fastest available computer. By connecting our high-end computing systems directly to the wind tunnel workflows, researchers could compare computational results in near real time. This used to take weeks or months, but now they accomplish it within hours. We prefer a right-sizing approach.

How do you determine which workloads go to HPC versus the cloud?

AM: First, we begin by understanding what types of computing the mission requires. High-end computing remains the best choice for large-scale, tightly coupled simulations that demand extremely high performance, low latency, and predictable execution, such as CFD, spacecraft design, and climate modelling.

Cloud environments serve researchers who need flexibility, rapid scaling, or collaboration. The cloud market also offers shared data services and modern AI technologies. Some projects use both: a team might prepare data in the cloud, run a large simulation on high-end computing resources, and return results to the cloud for analysis, visualisation, and sharing.

We have a systems engineering-led team that works directly with end users to determine fit-for-purpose based on mission requirements. That can change over time as the research progresses, so it is not stagnant.

Do you see HPC and cloud as complementary parts of a single research ecosystem?

AM: Definitely. Different missions require different computing environments, and some leverage both. Integrating HPC and cloud environments is neither easy nor automatic, presenting real challenges. First, managing data movement costs across environments at petabyte and exabyte scales. Second, aligning security accreditation processes built for different infrastructure models, with the industry still catching up. Third, building a workforce with expertise across both niche environments. We manage this through our hybrid computing strategy, optimising for performance, scalability, and security.

On workforce development, does NASA develop those skills and training, or do you search for external talent?

AM: Both. NASA currently has a major push to in-source capabilities germane to our core mission. On the flip side, where we can leverage industry expertise, we rely heavily on partners. We look to partner with industry rather than compete, including addressing workforce demand.

How does NASA plan long-term infrastructure amid fast-changing science?

AM: We make infrastructure investments that remain viable over time while recognising that scientific priorities and technologies continue to evolve. A good example is our newest supercomputer, Athena, launched at Ames Research Center.

As NASA's largest system, with over 20 petaflops of peak performance, it significantly reduces the agency's supercomputing utility costs. We managed this by housing Athena in a modular supercomputer facility designed for flexibility, accommodating future needs while keeping operating costs low. Under our hybrid computing strategy, we rely on cloud resources for quick access to emerging technologies or rapid scaling to meet changing scientific needs, while balancing the multi-year procurement cycles of high-end computing.

Does that agility allow you to test new technologies in the cloud before bringing them in-house, or do you use dedicated testbed systems?

AM: Both. In addition to testing in the cloud and on our HPC systems, we utilise industry partnerships. We test on emerging infrastructure provided by industry partners, domestic agencies, and international partners, working with organisations such as IBM and Microsoft to access cutting-edge technologies.

How is AI reshaping NASA's scientific computing infrastructure?

AM: One of the biggest changes is that we are no longer designing infrastructure solely to run simulations; we design it to support the complete scientific workflow combining observation, simulation, AI and visualisation. AI has become part of nearly every stage of the research lifecycle, helping prepare data, identify patterns and discover relationships.

We are doing significant work in agentic AI across the agency, which changes infrastructure demands at scale. Where traditional scientific workflows focused on CPUs, there is now increasing demand for GPU acceleration, training systems, massive storage and high-speed networks. AI has become central to our scientific process, but we must ensure those models are verified, validated and applied to support sound science rather than simply speed. ■

"We are no longer designing infrastructure solely to run simulations; we design it to support the complete scientific workflow combining observation, simulation, AI, and visualisation"

Andrew Mitchell is the Deputy Chief Science Data Officer within NASA's Science Mission Directorate Office of the Chief Science Data Officer.

How Imperial College is replacing monolithic HPC with HEX


Imperial College London

Andrew Richards discusses replacing a single cluster with six purpose-built platforms to match modern research requirements

Five years ago, Imperial College London ran its research computing on a single monolithic HPC system: one 80,000-core platform trying to serve everybody at once, from capability-class simulation to high-throughput batch jobs to the first stirrings of AI workloads. Today that system has been split into six purpose-built platforms, known as HEX, named after the magical ant-powered computer in Terry Pratchett's Discworld novels.

Andrew Richards, Director of Research Computing Services at Imperial, tells SCW how splitting resources into distinct environments within a single framework – alongside GPU virtualisation to meet growing demand for interactive AI workloads – has let the university balance research and teaching on shared infrastructure without either side compromising the other.

What does Imperial's Director of Research Computing Services oversee?

Andrew Richards: I sit within the central ICT department. Within ICT, we have a digital plan that supports our work across research, education, and business; that's the core ICT role. While there are many departmental and faculty-level activities supporting research, we are the core provider of research IT support for the university. Specifically in Research Computing, I have a team of about 32-33 people. We support any IT-related research, including HPC and associated large-scale data storage platforms.

How is Imperial's research computing team structured?

AR: On the HPC side, we've now defined what we call our HEX platforms, a suite of HPC platforms supporting AI, simulation, and modelling. We have an associated Research Data Facility, which provides storage for active projects and the ability to archive data long-term, supporting the HPC platforms in doing their best work.

> **"The ask from the AI community across departments was, in a nutshell, that they want something interactive; really, they want a big workstation under their desk..."**

We also have several other teams. One focuses on research software engineering [RSE], a team of software developers engaged with projects across the whole breadth of the university, essentially software developers for hire, working on different things at any point in time.

We have a Research Management Applications team that supports products that underpin research activity, from grant applications through to publications and dataset publishing. Whether you're involved in the initial costing activity or you want to publish, we have a public-facing service called Spiral and a new data repository called Helix where people can publish data linked to their publications.

The final team in my area focuses on Secure Research Services. We have a team that's the front door for researchers handling sensitive data, whether patient-sensitive or commercially sensitive, for projects with more sensitive needs. We then have a team dedicated to developing Trusted Research Environments and Secure Data Environments, working with colleagues across the hospital trust area to underpin research activities that require access to data more sensitive than that used in day-to-day HPC activities. So the breadth of Research Computing covers not just HPC but anything that strays into the research IT space, supported centrally by ICT.

Are Imperial's HPC resources centralised or departmental?

AR: There's a mixture. A lot of it has been consolidating, and we have more collaborations taking place now. For example, with the AI stuff, we've been working with the Department of Computing, and that's been a very joined-up activity, bringing that work into the creation of our new AI platform. But there are still some departments, like Physics and Electrical Engineering, running local clusters at their level. In Physics, for example, they're still heavily involved in things like the GridPP activity in the UK and the particle physics community. Whilst they get support from ICT for the data centre, hosting, and networking, as a system they are still running it at a departmental level, and that hasn't been folded into the central HPC provision.

Why did Imperial move from one HPC system to six?

AR: Of all the places I've worked, Imperial is one of those universities where we have users from every single department; there isn't a single department without users. I've worked at Oxford before, and there are many divisions and departments that almost didn't even know you existed. Imperial is quite rare in that sense. We're a STEM university with a business school. We have three main faculties: the Faculty of Medicine, the Faculty of Natural Sciences, and the Faculty of Engineering. Across all of those faculties and their associated departments, we have users. We have a lot of simulation and modelling, AI permeating across all those areas, and growing use across the business school as well.

When I came to Imperial five years ago, we had a very monolithic HPC service. It was a very big system, about 80,000 cores back then, and it's still around 80,000 cores overall. But it was one big monolithic system trying to do everything for everybody. The true HPC users doing capability workloads were competing with the high-throughput workloads, and the nascent AI workloads were in there somewhere too. It was just trying to be fair to everybody on one platform.

Over the last five years, we created what we call the HEX platforms. There are six distinct platforms within that. HX1 is our true high-performance compute system for people doing capability workloads in the classic simulation and modelling areas. HX2 is our high-throughput compute machine, which handles the bulk of our workload: people using single servers or a few cores at a time, the classic batch processing that serves many of our workloads. That's allowed us to separate out the people who want to do true HPC and put them on a distinct platform; they have to get extra permission to access it, so there's a bit of gatekeeping – a light-touch approach, but it's there.

Similarly, for AI workloads we've created HX3, the HEX AI platform. That's something new; we're not using the traditional job schedulers on it. On our other systems, we use a mixture of PBS Pro and Slurm; on HX3 we're using Run:ai. The ask from the AI community across departments, in a nutshell, was that they want something interactive; really, they want a big workstation under their desks that's exclusively for them.

We've separated the true AI workload from the rest, though some people are still running AI on other systems because they're handling offline workloads that don't require interaction; it just runs as a classic job. But people who want to be interactive in their training, manipulating the model as it goes, are now steered towards HX3.

With HX3, there was a specific request from the university to support both research and teaching on that platform. That's another reason we adopted Run:AI – it gives us the ability to virtualise the GPUs, so we can allocate whole GPUs for research workloads, with people accessing as many GPUs as they need. For teaching workloads, we can make a single GPU appear to be many tens of smaller GPUs. From a teaching point of view, where you're teaching the concept rather than worrying about the final result, students get what appears to be a GPU, a very small fraction of one, but to all intents and purposes it looks like the real thing and supports the teaching principles.

That's also changed the operating model around the AI platform. Teaching departments now have to nominate a resource in their department, what we call "RCS [Research Computing Services Associates", who work closely with our operations team to support that local teaching workload. So the challenge in addressing different workload types across the university isn't just about separating the infrastructure into distinct physical platforms for different needs, but also about changing the staffing model to get departments more directly involved in support, particularly in teaching.

To complete the HEX story: HX4 is our secure work platform, handling the trusted research environment stuff and secure data processing. HX5 is something we're just about to launch: visualisation. One of the asks was how to better visualise large datasets without having to keep moving data around the ecosystem, so the idea is to put a visualisation capability next to the data and the HPC. HX6 is just our R&D platform, where we put things to the test on a single server to see what something is actually like before we adopt it and scale it up into production.

How is AI changing the way researchers use HPC?

AR: I think it has opened things up a lot more for people. What we're seeing, particularly within the RSE team (which ties into how we support people on these platforms), is that many more people are now writing their own software than before, because it's become much easier to experiment.

Where projects would previously have had a problem and come to the RSE team, saying "We need help, where do we start?", we'd do some of that initial uplift to get something going. Now people arrive saying: "Here is something I've hacked together with Claude. It doesn't quite work as I expected; where do we go from here?" It's changed the entry point into those discussions, because it's a lot easier to get started. ■

Andrew Richards is the Director of Research Computing Services at Imperial College London.

"Now people are arriving saying, 'here's something I've hacked together with Claude; it doesn't quite work as I expected; where do we go from here?"

Solving HPC accessibility barriers with flexible cloud



Can Ethernet-based private cloud systems bridge the gap for non-traditional academic communities? **Ben Thomas** discusses modernising university compute architecture

Research computing at UCL has quietly reinvented itself. Where dense, InfiniBand-connected HPC clusters once set the template, the university's Centre for Advanced Research Computing (ARC) has built its first Ethernet-based system, blending bare metal and virtualisation to give researchers – from physicists to arts and humanities scholars who've historically been left behind by traditional HPC – far more flexibility in how they work.

Leading that effort is Ben Thomas, Head of Research Computing at ARC, where he oversees a 31-strong team spanning HPC, research data storage, cloud services, and a newly formed Unified AI team focused on model experimentation rather than pure inference. Thomas's route into the role has been anything but linear: trained as a computer scientist, he joined UCL in 2007 as a PhD student, completed a doctorate in medical imaging, held postdoctoral posts at UCL and A*STAR in Singapore, and later led infrastructure projects – including the Data Safe Haven Refresh – before rejoining ARC permanently in 2022.

Here, he talks to Scientific Computing World about designing systems for wildly different research communities, why cloud computing is opening HPC up to disciplines that scheduler-based systems have historically shut out, and the thinking behind ARC's move to a hybrid Ethernet-OmniPath architecture.

Tell us about UCL's research computing operation

Ben Thomas: Sure. I'm the Head of Research Computing, and the Research Computing section sits within ARC. I run a team of about 31 people, delivering services for researchers predominantly at UCL, but also national services, in the areas of traditional high-performance computing.

We manage the institutional research data storage and provide public and private cloud services for the university. We have a relatively new team called Unified AI, mainly focused on delivering an AI platform for early ideation – not really inference, not just the consumption of LLMs, but where people actually want to play and mess with the models. Then there's another team, Bespoke Solutions, which picks up everything that doesn't necessarily fit into the previous categories, a bit more consultancy-like. That's my role in essence; I have five direct reports who make all of that happen. My background is in computer science; I did my PhD in medical imaging, but these days I am more on the IT infrastructure and application side.

> **"The thing about traditional HPC is that there's a very steep learning curve, and for some subject areas that barrier is basically too high"**

How do you design research computing for such diverse requirements?

BT: That's a good question, and also a tricky one. When it comes to design, we have our bread-and-butter users – maths, physics, chemistry, that sort of thing – and I think we acknowledge that we don't support the humanities anywhere near as well as we could or should. Also, because of a tradition of close links with hospitals, the medical and biosciences side comes with it.

When we're designing these systems, we're still biased towards our large communities, but the advent of cloud computing has given us a bit more flexibility, and that's where I think we can pick up a lot, or certainly make life better for arts and humanities researchers. The thing about traditional HPC is that there's a very steep learning curve, and for some subject areas that barrier is basically too high. Cloud computing gives us much more flexibility in the interfaces people use to access the systems and in how we let them work, rather than forcing them to use a scheduler, which is completely alien to many of those researchers, who are domain experts, not IT experts.

So being able to provide web interfaces, notebooks, and that sort of thing, which are very easy to provide in a cloud environment, and not that easy in an HPC environment, is where we're starting to see fruit in terms of serving those communities. We've put quite a lot of effort in the past few years into building a private cloud capability to address some of those use cases.

What cloud resources do you have, and how are they managed?

BT: What we did was take a typical hardware footprint you might use for an HPC system, dense CPU-based nodes, plus some GPU capability, and build our first ever Ethernet-based system. Up until that point we'd always used InfiniBand or OmniPath, so the cloud system was our first purely Ethernet-based system, and we run it all on-prem. It's nice because it gives us a lot of options. We can run on bare metal, run things virtualised in VMs depending on what we need, or have a mixture. So if people need a couple of VMs to run one thing plus the grunt of a bare-metal node, we can slice and dice it as we need. That kind of mixed configuration is very hard to pull off in a traditional HPC system.

The other thing worth pointing out is that,

because we've gone with Ethernet, it lets us use the typical approaches to network isolation you'd use in an enterprise IT system. The beauty of that is it lets us jail researchers, or a group of researchers, into their own little bubble within the network, so they can only harm themselves, not other people. The great thing is we can then be confident enough to give them sudo rights [elevated privileges that allow regular users to execute commands] on the machines.

In much the same way as plugging your credit card into your favourite public cloud provider gives you free rein to do whatever you want, because of the network isolation we can give researchers, within limits, the freedom to do as they please.

So that feeling of being restricted by central services, which some researchers have, is somewhat less in a cloud-like environment, because we can put those controls around the perimeter instead.

The other added bonus is that as a section we're an enterprise Linux shop through and through; everything we do is Red Hat or AlmaLinux. But because we can put researchers into this safe bubble, we can say to them, "you can run pretty much any operating system you like; if it's not enterprise Linux, we're probably not going to give you a lot of support, but you have the freedom to do it if you want".

Of course, most researchers don't actually reach for enterprise Linux; they tend to go for Ubuntu, but we can facilitate that on the cloud, and some researchers really like having that capability.

Is that kind of set-up particularly good for teaching, where you can have safe but limited guardrails?

BT: We do some teaching; our remit, strictly speaking, is research, and we're a service provider back into the university for research purposes, but there's obviously an overlap with teaching.

We do run some teaching services that were originally isolated bare-metal systems. We're moving them to the cloud, where we can do exactly the same thing.

The other thing we've done is use our infrastructure and that same multi-tenancy concept to let students, as part of a teaching module run or partly run by ARC staff, build infrastructure as part of their coursework. That way, they learn not just the concept, but actually get to use infrastructure-as-code tools and stand things up themselves. We ran the first year of that with a public cloud vendor, but we've now brought it on-prem as well.

The nice thing about doing it on-prem is that people are less scared of cost overruns, because, worst case, they've just used a bit more CPU or storage, we're not getting billed from an OpEx perspective. That's gone really well; we've done it on-prem for two years now, and we've had great feedback from students because they're not always able to be really hands-on with this stuff. We can let them spin things up, break them, and play with virtualisation and containers.

What drove the decision to go for your first Ethernet-based system?

BT: Ethernet is finally good for high-performance computing. Sure, in terms of maxing out line speed, the overhead is still higher than InfiniBand or OmniPath, but it's now so fast that, with the movement towards Ultra Ethernet, we're far less concerned about it. Having said that, the new HPC system we've just procured for UCL has both; we've made the maybe-exotic choice to go for OmniPath 400-gig as well. So we've got both, and the intent is that if we want the absolute fastest point-to-point connection, we'll use OmniPath, but we also have the option to slice and dice with Ethernet. I've never really badged it as a hybrid HPC-cloud system, but in reality, that's kind of the intent behind it. ■

Ben Thomas is the Head of Research Computing at the Centre for Advanced Research Computing, UCL.

"Ethernet is finally good for high-performance computing… with the movement towards Ultra Ethernet, we're far less concerned about it"

How do supercomputing hubs balance converging HPC workloads?



Can traditional exascale infrastructure survive the shift towards interactive inference and agentic workflows? **Pekka Manninen** tells us how CSC is evolving LUMI into a software-defined framework equipped with real-time digital assistants and confidential hardware

For decades, supercomputers worked the same way: a researcher compiled a binary, submitted a job and emailed support if something went wrong. Pekka Manninen says that model no longer holds. With AI now woven into every stage of research, demand is shifting towards inference, orchestrated workflows and AI agents, and the machines serving them have to change too.

Manninen is Director of Science and Technology at CSC – IT Center for Science, where he also leads the LUMI AI Factory and the current LUMI supercomputer. He tells us how the AI Factory was launched on existing LUMI infrastructure to respond quickly to political pressure from the EU, why the new LUMI AI system will combine a versatile, general-purpose GPU with a software-defined platform running separate tenants, how memory and SSD price rises cost the project nodes, and how AI assistants could put a specialist alongside every user.

Can you tell me about your roles at CSC?

Pekka Manninen: I have the title of Director of Science and Technology. I lead a unit that is in charge of our scientific computing, basically research services that support our activities in AI and high-performance computing, but also technology development like quantum computing, digital twin technologies, chip technologies, and so on. I also lead the LUMI AI Factory and the current LUMI supercomputer.

The AI Factory is a service centre, but it also comes with its own infrastructure. We started off the AI Factory using the current LUMI supercomputer as the infrastructure. But we are meanwhile also building new infrastructure; there will be a LUMI AI Factory supercomputer, a kind of LUMI 2 or LUMI AI, and a connected quantum computer, LUMI-Q. These are part of the AI Factory initiative. But we've been building the services on top of the existing LUMI supercomputer.

> **"Supercomputers used to work the same way for decades, and now there's a complete paradigm shift in how people want to use them"**

Does that give you a chance to better understand the users?

PM: That's one aspect. It gives us valuable feedback on what's important and what isn't for building the new infrastructure. But at the same time, using the old HPC infrastructure as the platform for the AI Factory at the outset enabled us to launch it as quickly as possible. There was a lot of political pressure in the European Union to deploy these hubs for AI and industry AI adoption, and that was our way of responding. We already had this supercomputer, so why not start with that? No matter how fast you want to be, it always takes time to build computing infrastructure. It's impossible to deploy that the following week, no matter how much money you have. So that's why we decided to start delivering as quickly as possible and let the new infrastructure come up behind it.

How do you go about provisioning an AI supercomputer today? Do you have to make compromises on the infrastructure you want to provide?

PM: We just finalised the LUMI AI contract, and of course we're not immune; we did get hit by the memory and SSD price surges. We got the system we wanted in terms of functionality, architecture, and vision, but we ended up with fewer nodes than we originally planned when the budget was set. So yes, it's there, and HPC procurements are by no means immune to it.

How are the computational demands of your user group changing over time?

PM: I don't think there's a certain base of workloads and demand that's always there. If we want to be a versatile innovation platform where different fields of science can come together, and academia can meet industry, we need a highly versatile infrastructure and a wide range of capabilities. I try to identify where different needs of the modelling and simulation community and of artificial intelligence can be met on the same platform. That's the platform's basic remit; the workloads demanding these capabilities remain. We need capabilities for simulation science and exascale computing. There is still a big need for AI training, but the emerging thing that now needs to be addressed is inference. There's an increasing scientific inference load, and we should have a good answer for how we cater for it on an HPC platform, rather than just deploying it on some cloud platform. From an R&D perspective, it's becoming part of the development cycle of how researchers work; there will be more and more inference loads, more automation, and scientific or AI agents. AI has penetrated R&D completely. Those are the main differences from building a supercomputer two, four, or five years ago.

Does that change the technology or how you build your infrastructure?

PM: There are multiple takes on this across the HPC provider community, and we'll see which one is the right, most optimal one. Our take is to choose a cheap, "take-all-comers," jack-of-all-trades-style GPU so we can support simulation, training, and inference on the same GPU. It's not necessarily optimal for any one of them, but it can do them all.

The main thing on the software and user environment side is: how do we cater for a batch-queue supercomputer and interactive inference on the same platform? Our answer is a software-defined platform that lets us run different tenants on the same HPC system, dynamically hosting inference capabilities, the batch supercomputer, or even software-defined clusters for users on the same platform.

The main difference I see is a shift away from a monolithic, highly homogeneous system toward a more platform-oriented way of thinking, converging and bringing ideas from the cloud world into HPC system design. We have different hardware modalities in the platform, but on a software level we can carve out tenants, where one tenant can consist of multiple hardware capabilities, and at that tenant level we design what they're meant for: whether it's a batch-job supercomputer that looks and feels like an HPC system, or something where the user group defines their own stack, or a place to run inference workloads.

With this approach, I don't need to predict each user's needs in advance; I just have different hardware modalities, and with software definitions, I can mix and match these capabilities to create a valid use case.

What lessons have you learned from designing and operating LUMI?

PM: There are a few things in particular that the AI community, and industry or enterprise users, need us to serve better since the AI Factories have a mandate to do that. We've understood the importance of elevated security and confidential computing support, built into the system from the ground up. The world has changed over the past few years, and IT security is even more important than it used to be. Then there are the different ways of working with very complex, orchestrated workflows, and how to support them.

It's no longer just one researcher submitting jobs and analysing results; that's going to be a minority use case in reality, even within the modelling community. Ways of working have become very complex, and agentic AI is increasingly important in research, too. More people are orchestrating complex sets of different codes, different analyses, and incoming dynamic data.

There's a lot more going on in workloads today than there used to be, and building a platform for that reality is something we needed to do. Also, moving on from compiling a binary and running it, people are increasingly using frameworks and high-level libraries, which places different demands on the computing environment and requires different kinds of API access, so you can execute things without just the use of SSH.

Supercomputers used to work the same way for decades, and now there's a complete paradigm shift in how people want to use them, so new generations of systems need to respond to that to stay relevant and fulfil our mandate of catalysing new discoveries.

Confidential computing is one of those things supporting new usage patterns The value of data, large input datasets, is another thing; it's not only about output data, as it used to be in HPC; having good, well-curated, well-maintained datasets is a third big change in what we're doing differently now. And then there's user support.

For decades it worked like this: if there was a problem, the user logged in, worked, and if something didn't work and they couldn't figure it out themselves, they'd send an email, someone would read it and start investigating, and after some exchange the problem was solved and the researcher could continue.

We want to build a new system that can help the user right away by running different AI assistants for productivity, troubleshooting, programming, orchestration, and even generating ideas for the research problem itself. Having a specialist sitting next to the user all the time, but as an AI. Building and leveraging AI capabilities for user support is a major differentiator between the LUMI generation and LUMI AI. ■

Pekka Manninen is the Director of Science and Technology at CSC and the Director of the LUMI AI Factory.

"Our answer is a software-defined platform that lets us run different tenants on the same HPC system, dynamically hosting inference capabilities"

How AMD will tackle exascale limits with ‘Zen’ chiplets



Can modular 2nm silicon and high-precision GPUs maintain AMD’s position at the summit of global research clusters? CTO **Mark Papermaster** tells us how strategic hardware design is redefining AI for science

Four of the top 10 supercomputers on the latest TOP500 list run on AMD silicon – a remarkable turnaround for a company that, little more than a decade ago, was fighting to stay relevant in the CPU market. Mark Papermaster, AMD’s Chief Technology Officer and Executive Vice-President, has spent the past 15 years driving the strategy behind that resurgence.

In this interview, Papermaster traces the architectural bets that redefined AMD’s trajectory: the ground-up rebuild of its CPU line with the ‘Zen’ family, the parallel push for leadership GPUs, and the modular Infinity architecture that ties them together. That early commitment to chiplet technology – now a high-volume approach across the industry – has let AMD tailor hardware to increasingly specialised workloads, from agentic AI to the highest-FP64-precision computing science demands. He tells Scientific Computing World about the move to two-nanometre manufacturing with EPYC Venice, the memory-bandwidth pressures of inferencing and agentic AI, and how AMD is positioning its hardware and software roadmap to shape scientific discovery and AI over the next decade.

Can you tell me about your role at AMD?

Mark Papermaster: For most of the last 15 years, I’ve had two roles, CTO and head of what’s called Technology and Engineering, driving the bulk of R&D across the company. That was important in the formative years of remaking AMD into such a force in HPC. It was important to have one leader setting the technology direction while driving a very rigorous execution and culture of innovation and collaboration across the company, and that served us well. For the past year or so, I’ve been focused primarily on the CTO role.

That role continues to drive the company’s technical direction; we do so very rigorously, and I’ve been doing it for 15 years. We look at where technology and applications are going. Fifteen years ago was the dawn of AI, and really the dawn of heterogeneous computing, CPUs and GPUs, and how you bring those together to move the state of the art of computation forward. Back then, it was really just the domain of huge labs and sovereign labs across various organisations, but as AI and AI for science have gained momentum, you now see massive, heterogeneous compute clusters – CPUs alongside GPUs and other bespoke accelerators – commonplace across the world. That’s a great example of the CTO’s role: anticipating these trends and positioning our product roadmap and portfolio to be front and centre, because that’s what AMD needs to be: a provider of preeminent computing technology for the industry at large.

Which architectural decision has had the greatest impact on that resurgence, particularly in HPC and the data centre?

MP: There was a twofold decision we made years ago when Lisa Su and I were recruited to AMD – to focus on providing the best high performance, to differentiate AMD on that basis. That led to some very important decisions, starting with the CPU – scrapping the prior CPU efforts in place and creating a new CPU family targeted to be leadership from its very inception, followed by investment to have a leadership GPU too.

That upfront decision to invest in heterogeneous computing drove a number of phenomenally important architectural decisions. Alongside starting the new compute engine efforts, we simultaneously made the decision to create what we call Infinity Fabric, which let us invest in a modular approach that enabled not just heterogeneous compute cores, but also flexibility in how we partitioned into what became a chiplet approach. AMD led the industry in investing in a high-volume chiplet approach to computing, and that flexibility was extremely instrumental in winning bids for the largest supercomputers launched over the last five or six years.

It all traces back to that decision in three ways: One, invest in leadership compute engines, CPU and GPU. Two, enable that with a modular approach to how they’re stitched together. Three, invest in chiplet technology to leverage that modularity, so you can mix and match and create different configurations for different compute needs. That’s what we saw coming – there wouldn’t be a vanilla need for computing, but workloads would need specialised configurations to operate most efficiently. That’s exactly how the industry has played out, and we called that over a decade ago.

What about EPYC Venice and the move to 2nm? What new capabilities does this enable for scientific computing?

MP: The 6th-gen EPYC is really a statement of our continued commitment to leadership. The Zen 6 core, the new generation of the core, brings significant improvements in instructions per clock. We’d already invested in a native 512-bit-wide vector engine, AVX-512–native, and we’re making it more robust now by bringing even more memory bandwidth to those cores. We bump up to PCIe Gen 6, which provides more I/O bandwidth, and we also bump up to CXL 3.1, giving system providers yet more options for bringing peripheral devices on board and extending memory.

Two-nanometre manufacturing gives us two things. Firstly, more performance per watt in every CPU core. For our high-performance versions, we didn’t sacrifice the ability to ramp up to high frequency. We are a leader with what we call 3D V-Cache, which

lets you stack SRAM; that combination lets us deliver high frequency, high performance, and much larger local caching.

With the two-nanometre node, you also get density. For agentic workloads, we have another version with a higher density of cores, up to 256 CPU cores on a single EPYC socket, which is ideal for agentic workflows needing massive throughput, because the CPU's role isn't just as a head node running a high-performance Venice CPU; it's also the orchestrator, the brain running all the analysis of context windows across broad AI operations.

How does the rise in agentic AI impact your architectural decision-making? Is memory bandwidth going to remain a bottleneck?

MP: Inferencing is indeed memory-bandwidth-dependent, and that very much influenced our decisions on Venice. One was more memory options – you can run DDR or MRDIMM, which can run at very high bandwidth with the registered DIMM approach. We also have a variant called Verano with LPDDR. What you're seeing in the current era is not just the need for heterogeneous engines, different CPU and GPU configurations, different CPU core counts, but also different memory needs. With the flexibility and very high volume we have, our market share on EPYC is now rapidly approaching 50%.

We have scale, and with that scale we can really give customers options. Leveraging our chiplet approach, we've added more variants to address unique needs across HPC, head-node applications for large AI clusters, and agentic clusters, which require super-high-density throughput and very high memory bandwidth.

Interconnect is always extremely important for us because we're focused on high performance. If you're providing the leading high-performance solutions in the industry, you have to, as I've always put it, 'feed the beast'. You can't grow your compute engines without equally growing your connectivity to memory and I/O, and the ability to give customers configurability.

You saw that historically with AMD; look at what we did in late 2023 through 2024 with the MI300 family. We offered the MI300X, a beast for AI inference, and the MI300A, the HPC option, which integrated Zen CPUs with the Instinct GPU.

That's the processor in the Lawrence Livermore supercomputer today – a top-five supercomputer; I think it's number two, having been number one for a number of years – as well as a number of installations across the world, including LUMI, HLRS in Germany, and other UK and European installations. It's been a huge stalwart of HPC implementations across Europe. We have our MI455 in the Helios rack we've announced, shipping next quarter in volume, followed by another variant called the MI430 that brings in a chiplet change optimised for HPC. On all of these, interconnect plays a huge role, because we require every variant we build to be optimised for the memory and I/O of its targeted market.

How will the chiplet design continue to evolve? Do you see demand for more specialised chiplets?

MP: I do, because the need for incredible efficiency in HPC and AI is driving what used to be a specialised need of HPC more broadly. There's now mass demand for HPC, and AI for science requires ever greater efficiency.

One way to achieve that efficiency is to tailor hardware computation to the algorithms required for a specific set of applications. So what you're seeing now is unique configurations required for agentic AI throughput, unique configurations for building up AI training and very-large-language-model inferencing at a head node with a GPU or accelerator cluster, and unique configurations required by sovereign nations and HPC labs for AI for science.

Is it a strategic decision to focus on FP64 and double precision in your GPUs?

MP: It is indeed strategic for AMD. We made a commitment to the HPC and AI-for-science market when we launched the MI300A back in 2024. The way we've been able to [continue that support] is by leveraging our modular approach to generate a version of our GPU that supports FP64 by swapping out a single chiplet, while keeping the rest of the infrastructure the same as in the much higher-volume AI market, which doesn't require that highest precision.

The very underpinnings of how we architected our implementation gave us the flexibility to support the HPC and AI-for-science market with native, high-precision FP64 in a way our competitors couldn't match without a much more significant investment.

When you're accelerating discovery, trying to solve fusion, doing new molecular studies, or modelling the atmosphere and weather more accurately, that all requires high precision, native FP64, the highest double-precision floating point. You can't get there with AI approximations alone.

The most efficient way to handle all those unique configurations for bespoke applications is a chiplet approach, tailoring the engines, the I/O and memory configuration, and the very math formats you support, to more efficiently serve the algorithms at hand for a given market, and we're incredibly well poised to do that with unique chiplets. What's also moving very rapidly is the technology to stitch those chiplets together even more efficiently. ■

"If you're providing the leading high-performance solutions in the industry, you have to, as I've always put it, feed the beast"

Mark Papermaster is the Chief Technology Officer and EVP at AMD.

How do we overcome quantum hardware scaling barriers?


Universal Quantum

Mark Webber tells us why current qubit limitations hinder drug discovery algorithms and how a novel trapped-ion shuttling architecture bypasses the performance bottlenecks of photonic interconnects

Twelve months ago, Universal Quantum's Mark Webber spoke to us about the Open Quantum Institute (OQI) – an initiative backed by CERN and GESDA [Geneva Science and Diplomacy Anticipator] – applying quantum simulation to endometriosis, a chronic condition affecting roughly 10% of women worldwide. We caught up with Webber again to find out what's changed since, and to talk through where the field stands on the road to fault-tolerant, error-corrected quantum computing.

Bridging the gap between physical qubits and real-world applications

Mark Webber: As we continue our work on OQI, we've been working with industry players who aren't necessarily quantum experts to understand where their bottlenecks are today and where things can change, and mapping that to the quantum algorithms we expect to provide an advantage, mostly in the fault-tolerant regime. The qubit resources required to run these applications are quite large relative to current hardware.

Hardware players have on the order of hundreds of qubits, but to run these algorithms in a way that provides an advantage, we need the fault-tolerant, error-corrected regime, where the qubit quality required is unlikely to be reached by the underlying physical hardware. That's why we need error correction, which, in short, uses many physical qubits working together to represent what we call a logical qubit. A logical qubit isn't necessarily perfect, but it has much superior error characteristics compared to the underlying physical qubits. It comes with a big overhead, though: you need many physical qubits to represent one logical qubit, and there's also the efficiency of logical operations, because now you want to run operations on the logical qubits rather than the physical ones, which introduces a big overhead too. So to run the applications we're excited about in drug discovery, we're still talking hundreds of thousands, or maybe millions, of qubits – orders of magnitude more than what anyone has today.

That ties in nicely with the Universal Quantum story in general, where we've always focused on the scaling challenge: how do we build these machines with hundreds of thousands or millions of qubits as quickly and cheaply as possible? That's a very different design challenge from getting the highest-quality 100-qubit machine; you optimise for very different things between those two cases.

Moving beyond single-chip hardware limitations

MW: Internally, we have very ambitious projects to move to next-generation hardware, where our research aims are to validate the underlying scaling criteria that correspond to our design-modular chip manufacturing and to determine how we connect these modules. That's a unique feature of Universal Quantum, and we're moving towards further demonstrations of it using different techniques.

For trapped ions, the more commonly considered approach across the industry is a photonic interconnect. Modularity is key for scaling because it's much more feasible to make one high-quality module with hundreds to thousands of qubits than to fabricate a single million-qubit chip directly; device yield and other factors make that infeasible.

So we need modularity, and if you have modules, you need a way to connect them, where the fidelity and rate of those connections are very important. The more commonly considered approach for trapped ions is the photonic interconnect. Each module sits in its own vacuum system, and at the large-scale, million-qubit regime you'd suddenly have a thousand of these modules, each in its own vacuum system, connected together with photonic interconnects – think of them like a fibre optic cable.

The challenge is that the quality and rate of photonic interconnects demonstrated by industry and research groups are orders of magnitude worse than what those same groups can achieve within a single module performing local ion-ion-mediated entanglement. We see photonic interconnects as a major scaling challenge and an application performance challenge, so we've gone in a different direction.

For us, the large-scale machine operates within a single vacuum system, and our individual modules are physically aligned, so connectivity between modules is enabled by the same technology that enables ion connectivity within a single module. Ions are exciting in general, compared to superconducting systems, because of their high-fidelity, long-range connectivity, achieved through shuttling, in which we physically move qubits around the device. For us, the scaling story and modularity come from precisely aligning modules within the same vacuum system: we can physically shuttle an ion from one module to another with the same fidelity and rate as moving it within a module.

There's a lot of exciting progress here, as we move beyond existing demonstrations between two modules, which are very promising, towards scaling with a tiled array of many modules, along with the technology and manufacturability behind it.

What are the benefits or limitations of trapped-ion shuttling as opposed to the photonic interconnect?

MW: We see shuttling, in isolation, as a huge advantage for modularity and for how we connect modules. Ion shuttling provides very high-fidelity, high-rate connectivity relative to our other operations, and isn't available on many other hardware platforms; superconducting systems, for example, are famous for only nearest-neighbour connectivity. Whether you're running a NISQ [noisy intermediate-scale quantum] algorithm or trying to do quantum error correction more efficiently, we focus on high-fidelity, long-range connectivity, which is a huge advantage

“To run the applications we’re excited about in drug discovery, we’re still talking hundreds of thousands, or maybe millions, of qubits orders of magnitude more than what anyone has today”

over the more commonly considered error-correction protocols.

The best-studied one is the surface code, which relies on the assumption that you have only nearest-neighbour connectivity. By broadening that horizon and exploring the possible error-correction codes we can invent, given our high-fidelity, long-range connectivity, we can select codes with much more favourable ratios of physical qubits needed to reach a target logical error rate – a big advantage.

In the photonic interconnect picture, not all connectivity is enabled via photonic interconnect within a module; it generally assumes some physical shuttling to enable that connectivity. It’s only between modules that the photonic interconnect enters the story, but the scaling concerns there are rather dire. There are a few points – today’s highest-quality demonstrations rely on relatively large lenses that capture photons from the ion and send them through the cable, and these lenses are large relative to the module beneath them, which limits how many photonic interconnects can come out of one module. This, in turn, limits data bandwidth. If you have a thousand ions in one module but can only fit a small handful of these photonic interconnects, your data bandwidth is very limited.

They are investigating alternatives, such as integrated optics, but it’s a research challenge that’s arguably still unsolved today. The other challenge: imagine a million-qubit device with a thousand 1,000-qubit modules. How do you get all-to-all connectivity between those modules? You can’t connect every module to every other module with a single cable; that’s a combinatorial problem that doesn’t work. So the suggestion is to route these photonic interconnects into a quantum switch, similar to telecoms, with high N x N switches that have many cables coming in and out, directing signals. But because this is a quantum signal – a single photon – you can’t use the same telecoms infrastructure and solutions; you need to design them differently.

The state of the art there, in terms of achievable N x N and loss within the switch, is a major research challenge; you occasionally lose photons, which adds another significant error component to the overall budget. Some of these same challenges around high N x N switches also apply to players like PsiQuantum, which is going the pure-photonics route and faces a similar switching-design challenge.

As for using ion shuttling to connect modules, it’s true the modules need to be precisely aligned. But through our investigations into what that tolerance actually is, we’ve identified manufacturable solutions that are well within the required tolerances, so we don’t see that as a concerning challenge.

Do you need different algorithms for NISQ versus fault-tolerant, error-corrected systems?

MW: The types of algorithms differ significantly between designing for a NISQ machine versus a fault-tolerant one. For NISQ, you’re working in a very focused way to design low-depth circuits; the depth needs to be very short because you’re working with a very tight error budget. In chemistry, variational quantum eigensolvers are a well-known NISQ approach. For a fault-tolerant algorithm, depth is still important, but it’s not a constraining factor in the same way for quantum chemistry; phase estimation and qubitisation run at very high depth, perhaps hundreds of millions of gates, achievable because error correction lets you reach the needed logical error rates. The algorithms work in a very different way.

That’s not to say there’s no value in NISQ proof-of-principle demonstrations, but it’s similar to machine learning, where it’s hard to know in advance that a given approach will solve your problem; it’s more a matter of trying it and seeing what happens. NISQ is generally more like that, which means we’re limited by the hardware available and are waiting for next-generation hardware with lower error rates and more qubits to test what application performance we’ll actually get. ■

Mark Webber is the Head of Algorithms at Universal Quantum.

Planning Europe's transition to post-exascale supercomputing



How can Europe maintain HPC sovereignty as supercomputing converges with AI and quantum architectures? **Jean-Yves Berthou** discusses leading the SPE project to shape future European policy

Europe is trying to work out what comes after exascale computing – and Jean-Yves Berthou is coordinating the project tasked with answering that question. As coordinator of SPE (Strategy for Post Exascale), a €2.5m Horizon Europe project bringing together 23 organisations including CERN, CEA and the Barcelona Supercomputing Centre, Berthou is building the roadmap that will shape how the European Commission funds HPC, AI and quantum computing convergence in its 2028-2034 work programme.

It's the latest chapter in a career spent at the intersection of high-performance computing and policy. Berthou has been Inria's Deputy CEO for Territorial Strategy since October 2024, overseeing the institute's regional site policy through its network of University Centres; before that, he spent five years directing Inria's Saclay centre, growing it from 28 to 42 research teams. He also co-directs NumPEx, France's national exascale computing programme, and has advised the European Commission on HPC and research strategy for two decades.

Here, he talks to *Scientific Computing World* about why "flops" are giving way to watts and tokens as the metrics that matter, how Europe should balance sovereignty against reliance on international partners, and what two decades moving between industry and academia have taught him about building infrastructure that lasts.

How do you help to shape the strategy for HPC/AI in France?

Jean-Yves Berthou: The challenge has changed: we now have to understand the future of scientific computing, and more broadly the future of science, in light of how AI is evolving. The next point is the evolution that quantum computing will bring, which is much longer-term, but it's still something we want to address.

I'm coordinating a European project launched in April this year to develop a post-exascale roadmap. NumPEx is a purely French programme dedicated to software development for exascale and post-exascale computing, but the roadmapping itself is being done through this European project, which I coordinate. It's called SPE, for post-exascale strategy. *[See panel, right.]*

What is the role of SPE?

J-YB: This is a project launched by DG CONNECT, and its role is to give recommendations for the next work programme, 2028-2034. We plan to produce our first recommendation by the end of next month, which is quite fast. We also expect to regularly provide policy papers to advise the Commission on which directions should be funded in this area.

What do you mean by post-exascale?

It's the phase after exascale, but clearly post-exascale means convergence between AI, HPC, and quantum; that's exactly how it's phrased, by the way, in the context published by the Commission, which wanted a vision, a roadmap, and recommendations on what should be promoted in this area in Europe, for the next seven years, in the next work programme.

"Another measure is tokens, so in future we'll talk more in terms of tokens and watts than flops; that's my belief"

The next phase requires different KPIs and ways of measuring performance. Flops mean less and less now. We're not just talking about flops anymore, but about what data centres consume: power, measured in watts. Another measure is tokens, so in future we'll talk more in terms of tokens and watts than flops; that's my belief. We still have the TOP500, of course, but it's becoming less and less suited to a rapidly changing world, again because of how AI is evolving.

What we really want to understand is which tools scientists will use in the next few years in this fast-moving environment. We're following very closely what our colleagues are doing elsewhere, particularly in the US and Japan; it's harder to know what China is doing, but they're following the same path. We see new infrastructures emerging from time to time; one appeared in the last TOP500, for example.

We also see new models; we're following developments in the US and Japan, generative AI programmes being one example, and this is clearly close to what we believe the future of post-exascale will look like. That said, it's still a bit early to define clearly what post-exascale actually is. I'd say AI factories, and then gigafactories, were a first step in preparing the shift from HPC towards a convergence of HPC and AI.

How do you predict what HPC infrastructure you'll need for the next five years?

J-YB: We try to understand the evolution of technology and science. Of course, there are geopolitical and industrial issues related to that too, so it's a bit difficult for a project like ours to address all these factors at once. What we try to do is recognise that sovereignty is becoming increasingly important for Europe, because we don't know what our allies will do in the coming years. But in the meantime, we're not going to reinvent the

What is SPE and what does it mean for European HPC?

Strategy for Post-Exascale (SPE) is a Horizon Europe-funded project, coordinated by Inria, that brings together 23 leading European organisations, including CERN, CEA, BSC, Fraunhofer, GENCI, CINECA, and the universities of Bordeaux, Edinburgh and Turin, to build a European roadmap for the convergence of HPC, AI and Quantum Computing Systems (QCS).

Launched in April 2026 and running until 2029 with €2.5m in EU funding, SPE aims to help Europe anticipate, rather than react to, the "post-exascale" transition by identifying research gaps, emerging technologies, and new players across the full computing stack. Its work is built on three pillars: a continuously updated roadmap paired with a "Policy Paper Factory" producing rapid briefs for policymakers; a Community of Practice connecting academia, industry and research infrastructures; and international coordination through the InPEx initiative. The project's recommendations are intended to feed directly into the European Commission's next Horizon Europe work programme, covering 2028-2034.

"What we'll do inside SPE is understand where it's worth Europe investing, so the results are really measurable; we won't try to close every gap"

entire hardware industry in Europe over the next two years, so we have to deal with that reality. It will be a mix of sovereignty and partnerships, chosen according to the best path for Europe.

What we'll do inside SPE is understand where it's worth Europe investing, so the results are genuinely measurable; we won't try to close every gap. There are some areas where we're too far from the state of the art, but others where we're very close to the best level in the world. So our job is to identify where to invest our efforts to deliver the maximum value and sovereignty for Europe."

How has working in both academia and industry shaped your view of what scientific computing needs to deliver?

J-YB: On the industry side, I worked at EDF, which is a somewhat particular case, because in France it's mainly driven by nuclear power, and nuclear needs a roadmap that lasts 20 to 30 years, so it's actually quite well aligned with the timescales of public research. That's not the case for industries that deliver smartphones or anything that evolves every six months. When you work in an industry like that, you have time, and you can invest in things like quantum, for example, because you know it's something you need to prepare for. It might not be useful for another 20 years, but you have to start working on it now to be ready.

What I'm saying is, it's really important, and this is something we believed at EDF, and believe at Inria too, to have joint teams between industry and academia. That's something we're doing: sharing a strategic agenda at the highest level of the company and of Inria, with the CEOs aligned on that agenda, and then putting in place joint teams with permanent people from both industry and academia, not just PhD students working on an industry topic, but a real, ten-year agenda. We believe this works, and delivers genuine added value and real transfer of what public research can offer, turning results into useful outcomes for industry, which is something only industry can actually do; turning results into added value isn't a job for the public sector. Having joint teams gives you the full spectrum, from fundamental science right through to the delivery of new products and services. ■

Jean-Yves Berthou is Deputy CEO of Inria.

Scaling supercomputing to power European research ambitions



How does a national centre adapt when scientific computing shifts from niche vector processing to GPU-heavy AI workloads? **Kimmo Koski** discusses phased deployments, European collaboration and the balancing of national-tier systems with extreme-scale infrastructure

Kimmo Koski has led CSC – IT Center for Science through nearly the entire arc of Finnish supercomputing history, joining in 1989 alongside the country's first Cray X-MP and becoming CEO in 2004.

In this interview, Koski traces CSC's evolution from a niche technical provider serving a handful of specialist users into a 780-strong organisation with a national mandate spanning HPC, AI, quantum computing, research networking, and university digitalisation.

He discusses the lessons learned and how they still shape procurement strategy today, the design philosophy behind Finland's new national system, Roihu, and its relationship to the EuroHPC-funded LUMI, and the Finnish government's landmark €250m commitment to LUMI-AI, five times its original LUMI investment, as the country positions itself at the centre of Europe's AI computing ambitions.

Can you tell me about your role at CSC?

Kimmo Koski: I have been the CEO for 22 years now, since 2004. However, my history with CSC started back in 1989, when the first Finnish supercomputer, a Cray X-MP with four processors, 256 megabytes of memory, and a price tag of 40 million Finnish marks, was acquired. That system arrived in 1989, and I joined alongside it, not in the same box, of course, but to maintain one of the front-end machines.

Throughout the 1990s, I worked at CSC in various roles: purchasing supercomputers, running the technical group, and so forth. I then spent one year at CERN and a bit more than four years at Nokia before returning as CEO in 2004.

It is quite an interesting story. In 1989, we had about 35 people. I remember when I left for Nokia in 2000, we were around 80. The famous words of the previous CEO were always that CSC would never grow beyond 100 people. Today, we have 780 staff members. The main reason for this growth is that in Finland, everything is centralised under one roof. While you are interested in HPC, AI, and quantum computing, we also run the national research network, handle the digitalisation of universities, and manage extensive data services, not only for research, but also for culture, such as museums, archives, and scientific libraries, as well as data for education. It is a very wide mandate in Finland at the moment, probably wider than anywhere else.

> **"One of the main cultural shifts we have driven, and a key reason why we have grown, is prioritising customer satisfaction"**

How has the role of CSC changed over the years?

KK: The role has developed significantly over the years. In the early 1990s, we were practically running the national research network, Finnish University and Research Network (Funet), alongside supercomputing services for a relatively limited number of users and a narrower scope. Eventually, more data services emerged, first for research data, and later through library collaboration. AI came along more recently, and the newest addition has been the digitalisation of universities.

To give you an idea of our structure today, out of our 780 employees, roughly 200-plus work in scientific computing, 200-plus in data, and about 100 in digitalisation, alongside our international collaboration teams.

Another interesting dynamic: when I returned to CSC as CEO in 2004, we were involved in just two EU projects. Our roles were quite minor, mostly communications or secondary tasks outside our core business. The board, since we are a non-profit limited company with a board like any corporation, gave me the task of making CSC more international. Today, we have more than 70 EU projects underway, with LUMI among the largest. The organisation has evolved continuously, sometimes by targeted strategy and sometimes almost by accident.

What shifts are you seeing in the computational demands of the user community?

KK: Looking back over a 20-year timeframe, or even going back to when the Cray X-MP arrived, that machine was a vector processor. At the time, only a specific set of applications could use it. The way CSC operated back then was also quite different: we were technocrats, and a customer was almost seen as a disturbance to technical work.

One of the main cultural shifts we have

driven, and a key reason why we have grown, is prioritising customer satisfaction. Even though many of our staff have technical backgrounds, we focus on customer needs rather than our own technical preferences.

The second major shift is recognising our core identity: CSC does not conduct its own research; we support other people's research. About 15% to 17% of our staff hold doctorates, including myself. Our task is strictly to support external research. Thirty years ago, that distinction was not as clearly understood as it is today. Staff can conduct research if they take leave to attend a university and return later, but our primary mission remains support.

In terms of hardware and workload areas: we moved from vector processing to RISC processors and clusters. In the mid-to-late 1990s, we faced justified criticism for providing only certain types of high-end resources. We shifted our approach to offer a broader spectrum of resources, combining high-end machines with medium-range clusters.

At that time, systems were entirely CPU-based. Recently, GPUs and AI computing have entered the picture in a big way. When we competed for EuroHPC and established the LUMI collaboration, performance benchmarking was conducted a few years before its 2022 deployment. We had to ask ourselves whether we could build such a GPU-dominant system, given that over 95% of LUMI's compute cycles come from GPUs, and less than 5% from CPUs. Even that 5% is still considerable given the system's massive scale.

It turned out to be the right decision. Today, workloads have converted heavily toward GPUs, though CPUs still retain a role. Quantum computing remains an unknown for the future. I do not believe there is a proven quantum advantage anywhere yet, but many are investing effort into it. As you might have noticed, in the second week of July, we announced the €40m IQM deal linked to EuroHPC.

While things are moving strongly toward AI, we must remember that it is not only about AI. Amid the AI hype, people sometimes forget that traditional computing persists. However, GPU- and AI-type workloads can now be applied across far more disciplines than was possible a few years ago.

Has the introduction of AI changed how the user community wants to consume these resources?

KK: The difference is very clear, and it actually predates the current AI boom. In the 1990s, when vector processors were dominated by climate scientists running heavy models, bioinformaticians were beginning to emerge. Bioinformaticians did not need massive vector operations; they wanted a large number of smaller processors to run vast amounts of simultaneous calculations.

The same applied to certain subfields of physics: they needed high volumes of small jobs. In the mid-to-late 1990s, we illustrated this using a resource pyramid. Extreme-scale, high-end computing was necessary at the peak, but many disciplines required middle-tier cluster computing instead.

Historically, physicists, fluid dynamics researchers, and meteorologists were traditional CPU users. Bioinformaticians, economists, and others adopted alternative, less traditional methods earlier on. Looking at LUMI's current usage, about 50% of the system is dedicated to AI-type workloads, while the remaining 50% is used for other GPU-accelerated computing.

Furthermore, younger generations graduating today have not necessarily lived in the world of command-line text editors like Vi or Emacs. They enter with entirely different toolsets. We are seeing a major user transformation driven not just by academic discipline but also by generational shifts.

What lessons did you take from LUMI that shaped the design of your next national supercomputer, Roihu?

KK: LUMI is a European system, formally owned by EuroHPC and hosted by our consortium. However, we maintain our national tier because it serves a different function: it is under our direct control and can therefore be tailored to national priorities.

Roihu was not designed directly as a derivative of LUMI, but rather to complement it. Because Finland owns a quarter of LUMI's capacity, in addition to what our researchers can apply for via the general EU quota, we designed Roihu to be more CPU-oriented and less GPU-heavy. Roihu went through an open European procurement process, which was won by Bull (Atos) using Nvidia and AMD components, whereas LUMI was built by HPE using AMD GPUs and CPUs.

The biggest single historical lesson came from that first Cray X-MP purchase. It was a massive investment at the time and we bought the entire system upfront on day one. By year three, technology had advanced so much that cheaper systems were matching its performance in certain applications, prompting criticism. From that experience, we learned two core principles that still guide us today.

The first is phased deployments – we do not acquire a full five-year system all at once. We secure an initial environment to start immediately, followed by a main system or expansion phase roughly two years later. You are effectively purchasing capacity over time.

The second is the ecosystem – we do not rely on a single massive machine to handle every workload. Even within a system like LUMI, there are distinct tiers of memory, storage, CPUs, and GPUs. The goal is to build an interconnected compute ecosystem rather than just run an isolated machine. ■

Kimmo Koski is the CEO of CSC.

"Roihu was not designed directly as a derivative of LUMI, but rather to complement it. Finland owns a quarter of LUMI's capacity… we designed Roihu to be more CPU-oriented and less GPU-heavy"

Modelling a life, not just a heartbeat



Amanda Randles of Duke University on why proactive, patient-specific digital twins need HPC to move medicine beyond the clinical snapshot, and what it takes to get there

Heart failure, stroke and cancer rarely announce themselves. They develop quietly, sometimes for years, before a patient ever feels unwell enough to see a doctor – by which point the disease has often already progressed. Amanda Randles wants to change that timeline entirely, using patient-specific digital twins that combine high-performance computing, machine learning and multiscale biophysical simulation to catch disease before symptoms do.

Randles is the Alfred Winborne Mordecai and Victoria Stover Mordecai Associate Professor of Biomedical Sciences and Biomedical Engineering at Duke University, where she directs the Duke Center for Computational and Digital Health Innovation. Her work has earned the ACM Prize in Computing, the NIH Pioneer Award, the NSF CAREER Award, the ACM Grace Hopper Award, the Jack Dongarra Early Career Award, and the Sony and Nature Women in Technology Award. She holds a PhD in Applied Physics and an MS in Computer Science from Harvard, and a BA in Computer Science and Physics from Duke, and began her career as a software engineer on IBM's Blue Gene supercomputing team.

SCW spoke to Randles about what it takes to build a vascular digital twin, how HPC has pushed the scale of these simulations from single heartbeats to whole-body models, and how the technology is already changing decisions in the clinic.

Why is this type of simulation important?

Amanda Randles: Heart failure, stroke and cancer don't really appear overnight – they come on slowly and silently, developing long before you see symptoms, before whatever it is that sends you to hospital happens. By the time symptoms appear, the disease has often progressed to an advanced stage.

What we want to know is: what would happen if healthcare didn't just react to these crises, but could identify them long before you had chest pain, long before you had a reason to see your doctor, and doctors could act proactively and stop it happening in the first place?

> **"We did what every great HPC team does and threw as many GPUs as we could at it: we used all of Oak Ridge's Summit supercomputer, about 24,000 GPUs, and modelled about 550 million red blood cells"**

Imagine you're a patient living with heart failure. Long before you start to feel worse, your pulmonary artery pressure is already changing, and you're having subtle changes in blood flow that indicate something is going wrong. Today, there are implantable sensors that sit in the pulmonary artery and can detect these subtle shifts before you start having issues. They're already reducing hospitalisations. But the problem is that these sensors are invasive and typically take only one measurement a day while the patient is lying down. That's where we see digital twins playing a role: getting continuous data lets us identify these subtle shifts much earlier, before symptoms arise.

Most medical decisions are still based on clinical snapshots, the discrete points where you're actually interacting with the healthcare system. But health isn't only changing when we collect these measurements. It's changing slowly, all the time. It's not the single snapshots we really care about. It's the change in your specific trajectory that helps us identify issues. That matters because diseases that develop slowly like this – heart failure, stroke, neurodegenerative disease – represent the largest healthcare burden today: over $480bn in the US alone.

What do you actually need to build a vascular digital twin?

AR: Digital twins aren't new. They've already been used throughout engineering, for jet engines, for bridges. The key point is these models aren't static: it's not a model of an idealised, representative jet engine – it's a digital twin of a specific one, moving along with real sensor data.

Healthcare is just now starting to adopt these principles. We're modelling a patient instead, and the goal is very similar: we want to give every patient their own 'check engine' light, an early warning system the doctor can identify and check in with, earlier and remotely.

When we're building this digital twin, we start with medical imaging. It needs to be patient-specific, tuned to that particular person. We need CT or MRI scans to get the 3D geometry of your specific vascular system. Then we feed that into a high-resolution simulation using HARVEY, our large-scale fluid dynamics software, to build large-scale, high-fidelity computational fluid dynamics models over time, using HPC.

We then connect that to wearable sensors, obtaining continuous streaming data, and use it to create longitudinal haemodynamic maps, moving from single snapshots to understanding cumulative exposure and cumulative stress over weeks or months.

Building the digital twin is really just the beginning. To move from a digital twin to predictive, proactive health, we need to

address three challenges: scale – capturing from cells to organs to larger parts of the body; time – moving from single heartbeats to much longer periods; and usefulness – letting doctors make clinical decisions and forecast outcomes. 'If I gave this patient a statin, how would that help them? If I gave them stent A versus stent B, what would work best for this specific patient?'

How has HPC pushed the scale of these health simulations?

AR: HPC has played a huge role in moving both the scale we could capture and the resolution we could get. In 2010, we had a Gordon Bell finalist looking at blood flow in the coronary arteries at cellular resolution, modelling a whole heartbeat at the cellular scale.

In 2015, using the Vulcan supercomputer at Lawrence Livermore, we captured the entire vasculature, all the arteries in the body, one centimetre in diameter and above.

For cancer, we wanted to look at the flow of a cancer cell over long distances, ideally the entire body, with all the cells moving. We did what every great HPC team does and threw as many GPUs as we could at it: we used all of Oak Ridge's Summit supercomputer, about 24,000 GPUs, and modelled about 550 million red blood cells. That got us a tiny cube, about 3.1mm on each side. That showed brute force wasn't going to be enough; we needed new algorithms.

So we developed what we call the adaptive physics refinement, or APR, model, inspired by adaptive mesh refinement. We have a high-resolution window around a specific cancer cell, where every red blood cell is explicitly modelled using a full fluid-structure-interaction physics model, and outside that window we model a bulk fluid to capture overall movement.

That lets us reach much longer distances: about 6mm of vessel took only 96 node-hours in the cloud to simulate, compared to using all of Summit, 24,000 GPUs, just to get a 3mm section. Suddenly, we can use cloud resources to reach much longer distances while maintaining the same fidelity, and we've since pushed the resolution even further, capturing individual ligand receptors on the nanometre scale to understand why cancer cells adhere in some locations and not others.

How are these digital twins being used in clinical care today?

AR: The clearest example today is fractional flow reserve, helping doctors decide whether a patient needs a stent. Conventionally, that means threading a guide wire into the artery and measuring pressure before and after a narrowing. We've shown that digital twin models, using tools like HARVEY, can match invasive measurements to within less than 5% error, and there are already FDA-approved tools doing this commercially, such as Heartflow and CathWorks. It's one of the first places we're seeing digital twins have a real impact in clinical care.

We've also shown that a non-invasive digital twin can match the pulmonary artery pressure measurements captured by implantable haemodynamic monitors in patients with heart failure. And we've built a tool, written in Unity, so doctors can interact with a patient's vascular geometry in virtual reality, even on a Vision Pro, and test different treatment options before they ever go into the operating room: if I put stent A here versus stent B, if I take this vessel from the main branch versus the side branch, what's the better option for this specific patient? Combined with machine learning, that gives immediate, real-time feedback. We've validated the predictions for peripheral vasculature in the leg against Doppler ultrasound measurements taken after treatment, and the model's pre-treatment predictions matched what actually happened. ■

"It's not the single snapshots we really care about. It's the change in your specific trajectory that helps us identify issue"

Amanda Randles is Associate Professor of Biomedical Sciences and Biomedical Engineering at Duke University.

Overcoming AI's unpredictability in drug discovery



University College London

Can exascale physics modelling fix AI's reliability problem in therapeutic development? **Peter Coveney** explains how uncertainty quantification makes molecular digital twins actionable.

Digital twin technology promises a new kind of medicine: detailed, personalised models of the human body or the molecules used to treat it that let researchers test and predict outcomes before committing to real-world experiments or clinical trials. But that promise only holds if the models clear two separate bars: they must be trustworthy enough to act on, and fast enough to inform a decision while it still matters.

Peter Coveney, Director of the Centre for Computational Science, Professor of Physical Chemistry, Honorary Professor of Computer Science and Research Fellow at the Advanced Research Computing Centre within University College London, sets out where the field stands today, from the wide uncertainty bands hidden inside a well-known COVID-19 model to the current, more sobering reality of AI-driven drug discovery.

How can you make AI more trustworthy?

Peter Coveney: A lot of issues in fields like this are about prediction, seeing the future. That's the apotheosis of the scientific method: it starts from experiment, the only source of truth in the game, and then you try to explain experiments. But as your theories get better, you expect to be able to predict things before they occur. That's pretty tricky in general, and the methodologies we have available are theory, modelling, and simulation, with artificial intelligence recently joining the mix.

But if you're making decisions that are actionable in areas people care about, you need to be sure your predictions are reliable and trustworthy, and that's always a question levelled at AI: how can we be sure about its predictions? I want to talk about verification, validation and uncertainty quantification, and an example in drug discovery that, at first glance, might not look like a digital twin, but you'll see that it is, given the definition of a digital twin.

What is a digital twin?

PC: The digital twin vision, for those uncertain about it, is in danger of being a buzzword, just a new way to raise funds. But what we're really trying to do is develop a virtuous circle between our theories and modelling and the world we're observing: an analogue entity, in the human case, say myself, and a digital version of that, with prediction and observation tuning both the model and the predictions of the real-world system.

The medical field is particularly challenging because we're dealing with an immensely complex, multi-scale domain. But the prize potentially available is a digital twin of the human, not just an average human, but a personalised one from the get-go, which would achieve far more than current medicine can.

Why is uncertainty quantification important?

PC: We always have a mathematical model of some sort in play, and because those are hard to analyse, we put them on computers, which immediately creates issues around validation from observation and verification: whether your implementation matches the model you're working with. It's worth stressing the term "analogue": the digital is not a guaranteed, verified version of the real world. Then there's uncertainty quantification, which arises because the models have many parameters; parameters are the bugbear of artificial intelligence, and aleatoric uncertainty plays a role.

For some simple systems you can get away with off-the-shelf methods, but for real-world systems, described by partial differential equations or operating in the molecular domain, digital representations have very large numbers of parameters to cope with, and you need new methods for handling them, or the curse of dimensionality prevents progress.

What did you learn from the COVID-19 example?

PC: Climate and weather are best studied in the context of uncertainty quantification, because people have long cared about the accuracy of these predictions. Some work we did in the COVID context plotted the cumulative number of deaths in the UK as a function of when COVID began, using an individual- or agent-based model built by Neil Ferguson at Imperial College London, who made these predictions for the then prime minister and a committee. What they did when advising the Prime Minister was run a single simulation and report the results. Our study of the uncertainties in what's now called the COVIDSim model, which was inaccessible to the general scientific community at the time, shows the distribution is, in fact, very wide: the number of deaths compatible with the model could range from a thousand to tens of thousands. How do you take action based on that level of uncertainty?

How do you build tools for uncertainty quantification?

PC: With some funding from the European Union, we built a VVUQ (verification, validation and uncertainty quantification) toolkit to help people across general domains, open-source and available to anyone. Its core, EasyVVUQ, is an orchestrator that defines your sampling strategy and campaigns, implementing a range of off-the-shelf techniques. It doesn't launch jobs itself – that's handled by separate tools such as FabSim3. A later UK project extended these capabilities, with US partners including DOE labs, and we now work particularly with the Oak Ridge Leadership Computing Facility, applying uncertainty quantification on the exascale machine Frontier and others to come, through DOE INCITE [the US Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment] awards.

The challenge of molecular dynamics

PC: Turning to the physics of molecular dynamics, heavily relevant to drug discovery: one of its problems is that it parted company with modern dynamical systems theory almost as soon as it was developed. People who work in molecular dynamics, to this day, tend to want to run only a single simulation, but how can you do that? The dynamics is chaotic; for rather fundamental reasons, chaos is a requirement to reach a thermodynamic equilibrium state, which is what we're trying to measure. Quite often we're calculating a thermodynamic property like the Gibbs free energy of binding, so a one-off simulation is self-evidently not reproducible; the next person who runs it will get a different answer, and no single run is representative of an equilibrium state. The codes are highly dependent on the random seed, but they also have thousands of parameters: the interatomic potentials that govern interactions between atoms.

This has been a very open problem in the field: how do you know which parameters you need to be more careful about than others? We have answers for that now published in a book on molecular dynamics that's really unlike any other and worth referring to by its full title[1] about probability and uncertainty, and how reliable any molecular simulation actually is.

Combining AI with physics-based methods

PC: What we're after is to do a whole lot better than the pharma industry, which takes 10 years, $3bn, and succeeds only about 4% of the time. The work that conveys the message combines AI with physics-based methods and experiments, with the experiments aimed at making real analogues of the molecules and comparing them with the predictions.

At the centre of this is a generative AI tool developed by AstraZeneca, combined with physics-based methods that discipline the AI. Because we know the uncertainty in its predictions, we can drive the molecules it creates towards lower free energies. What we're doing on [Argonne's] Aurora and Frontier combines the entire workflow from the start with docking calculations. This isn't a funnel; it's more like a blunderbuss: you multiply many options, injecting new molecules through different stages of generative prediction and correction by the physics-based methods.

We've run this once with a "dark" protein, not well known, called WDR91. In the first run, making only about half a dozen molecules, some of them chiral, we found four to be nanomolar-potent by integrating system-level, physics-based methods and experiments to guide us in making the molecules.

From molecules to the clinic

PC: The last part of this goes to the higher level: if you've done the molecular end of the game, you're only at the start. You need the physiological, clinical-trial end too, bringing the higher-level testing of these molecules into the real world, testing on humans, or digital twins of them, where there's less chance of things going wrong. We need models that operate between the molecular level and organ-based systems like the heart, enabling cardiotoxicity evaluations based on molecules that affect simulations of real hearts, in silico trials, and feeding back into the workflow described earlier. You can identify specific arrhythmias that develop during these tests when run on virtual patients.

Professor Peter Coveney is the Director of the Centre for Computational Science at University College London.

[1] [*Molecular Dynamics: Probability and Uncertainty,* by Peter V. Coveney and Shunzhou Wan]

"The digital twin vision, for those uncertain about it, is in danger of being a buzzword, just a new way to raise funds"

Simulating nuclear criticality with quantum algorithms



Joint Center for Quantum Information and Computer Science (QuICS)

Can hybrid quantum algorithms outpace classical methods when modelling complex reactor dynamics? **Andrew Childs** talks about his latest work on neutron diffusion equations and the realities of demonstrating speedup

Nuclear reactors run on a delicate balance: too few neutrons and the chain reaction dies out, too many and it runs away. Calculating that balance – the neutron diffusion k-eigenvalue problem – is fundamental to reactor safety, fuel management and radiation transport modelling, but realistic models with complex geometries and heterogeneous materials can produce billions or even trillions of unknowns. Andrew Childs and his colleagues have proposed a hybrid classical-quantum algorithm that takes this problem on, drawing on quantum linear-system algorithms, quantum preconditioning and Hamiltonian simulation.

Childs is a Professor of Computer Science at the University of Maryland, Interim Director of the University of Maryland Institute for Advanced Computer Studies (UMIACS), Director of the NSF Quantum Leap Challenge Institute for Robust Quantum Simulation (RQS), and a Fellow at the Joint Center for Quantum Information and Computer Science (QuICS). His research centres on quantum algorithms – particularly quantum simulation and methods for reducing the computational resources needed to simulate physical systems.

SCW spoke to Childs about the state of quantum algorithms research, why comparing classical and quantum performance is harder than it looks, and what the neutron transport work reveals about where real quantum speedups might – and might not – be found.

Tell us about the key developments in quantum algorithms research over the past year

Andrew Childs: My area is really quantum algorithms, which includes trying to better understand applications of algorithms that have already been developed, for example, how to use quantum simulation algorithms to simulate particular kinds of physical systems.

Another class of things I'm interested in is using quantum algorithms for scientific computing, giving algorithms for differential equations on quantum computers, and understanding how much speedup they can offer. I'm also pretty interested in fundamental questions about the power of quantum computers and in relating their performance to that of classical computers in abstract models, where you have a clean formulation of a problem and try to make a precise mathematical statement about how they relate.

The closer you get to applications, the harder it is to make precise statements, because the performance of algorithms can be difficult to understand. When you want to compare the performance of classical and quantum algorithms, many of the classical algorithms people use are pretty sophisticated, and people don't necessarily worry so much about proving theorems about how they perform; they just have algorithms, implement them, and they work in some computing system.

But in the quantum computing world, we still don't have really good devices – the actual computers are very limited in what they can do. The progress is amazing; there's a lot happening, but there are still no large-scale, utility-scale devices where you can really run useful simulations. So, in practice, comparing classical and quantum algorithms is largely limited to understanding classical performance and then comparing it to what we can prove about quantum performance, which makes it pretty hard to draw meaningful comparisons.

How do you measure quantum algorithm performance against classical algorithms?

AC: There are cases where you can give some analysis of the performance of a quantum algorithm; you can say the running time will scale in a certain way with some parameters of the problem. Those kinds of statements are generally possible to make, depending on the details, for some kinds of algorithms, and that gives you some indication of how the algorithm performs.

That doesn't necessarily mean it couldn't actually perform better; you do some analysis, make some approximations, and something might be loose somewhere, so the actual performance could be better. But it establishes a baseline for the algorithm's performance. The trickier side is comparing that to the performance of classical algorithms, because generally, putting an upper bound on the running time of an algorithm is doable, but lower-bounding how long an algorithm will take is harder, especially if you want to think about the time it takes on practical, real-world instances, which is hard to characterise with a theoretical framework.

How does quantum computing apply to nuclear reactor simulation?

AC: A concrete example of this is a paper we recently wrote on understanding neutron transport. These kinds of problems can model the behaviour of nuclear reactors: you have neutrons moving around in the reactor, and a differential equation describing that neutron transport problem, and you try to understand properties of it with an algorithm. This is obviously a practical thing people actually want to do, because they want to understand how these things perform.

We developed a quantum algorithm for this, using ideas from quantum algorithms for differential equations, and we have some analysis of its performance. To get it to perform well, we have to do things in the right way; there's essentially a linear-system-solving problem involved, and we care about the condition number of that linear system, so we have to do some preconditioning to make the condition number good, and show how to implement that preconditioned linear system efficiently on a quantum computer. There are things to work out, but once we do, we can bound how long it would take the quantum computer to run the algorithm and produce information about the solution. But the hard part of this problem is understanding how it compares to the best classical algorithms.

In the paper we wrote, the quantum algorithm essentially discretises the system

of differential equations on a uniform grid, and we compare its performance against an analogous classical algorithm that discretises everything on a uniform grid in the same way.

There, it looks like there's a polynomial quantum speedup between the two approaches; the speedup is somewhat limited, but in that context, there does appear to be some quantum speedup.

But then, of course, you can consider other classical algorithms. This isn't really my area, but as I understand it, people generally approach these differential equations using finite element analysis, and just discretising uniformly isn't necessarily the best approach. If there's some part of the system where something interesting is happening, you should discretise more finely there, and more coarsely where things are smoother and better behaved. If you do that, the speedup could disappear as the classical algorithm improves.

How can quantum algorithms keep their speedup over improving classical methods?

AC: One way is to just wait until we have big devices and implement these things. If they're useful, they're useful, and it doesn't really matter what you can prove about them; that's definitely an approach. But it's also possible to do what we can with the tools we have.

For example, for particular model problems you're studying, you can try implementing alternatives to the uniform discretisation, a reasonable choice of non-uniform discretisation, and see how it performs. Maybe empirically the speedup goes away, and that's an indication that it's not a place to look for speedup, which is still a useful result, because it can tell you what you should do differently to find a problem where you'll get more advantage from quantum algorithms.

Exploring particular classical approaches is definitely something one can do, especially since people who dedicate their lives to developing methods for solving differential equations on classical computers have a lot of intuition about the right way to do it. So this is probably a place for collaboration between people thinking about quantum algorithms and people with expertise in classical methods for solving these problems.

Are there trade-offs from having so many different underlying technologies, topologies and architectures?

AC: It makes things more complicated, but that also means more choice, more things to explore. I don't feel like it's a huge burden though. If you want to focus on the big picture of what quantum computers can do, you just work in the basic quantum circuit model, where you can do gates between any pair of qubits; that's what the vast majority of my papers consider, and I think it gives a good idea of what's going on. It adds some complexity in some cases, but I don't think it's the main barrier to understanding the power of quantum computers; it's manageable.

What is the biggest barrier to practical quantum computing?

AC: I think the most important barrier is coming up with quantum algorithms for new problems and understanding their performance. And, of course, the challenge of actually implementing quantum computers, which will really change the landscape, because if you can just try things out and be much more empirical, it's a completely different ball game. ■

Andrew Childs is a Professor of Computer Science at the University of Maryland. Childs is the Interim Director of the University of Maryland Institute for Advanced Computer Studies (UMIACS), Director of the NSF Quantum Leap Challenge Institute for Robust Quantum Simulation (RQS) and a Fellow at the Joint Center for Quantum Information and Computer Science (QuICS).

Preprint research paper: Quantum Algorithms for Heterogeneous PDEs: The Neutron Diffusion Eigenvalue Problem

"I think the most important barrier is the difficulty of coming up with quantum algorithms for new problems and understanding their performance"

How AI is transforming UK scientific research



Tom Griffin explains how Britain's Science and Technology Facilities Council is integrating AI across research workflows, managing petabytes of scientific data and adapting e-infrastructure to meet growing compute demands

AI is touching almost every step of the modern scientific workflow, but keeping pace with its exponential growth is no easy feat. With detector data outputs soaring into hundreds of gigabits per second and hardware prices rising, traditional computing architectures are being pushed to their limits.

In an exclusive interview with SCW, Tom Griffin, Director of Scientific Computing at the Science and Technology Facilities Council (STFC), breaks down how the UK is reshaping its national research infrastructure for an AI-driven future.

Tell us about your role at STFC

Tom Griffin: My main job is Director of Scientific Computing. That is the department that delivers all of the scientific computing infrastructure that underpins the national labs, particle accelerators, the astronomy programme, particle physics programme, material science – all of that – and that's everything from computational science through to large-scale HPC platforms and 100 petabytes of data storage and access for all of those. I'm seconded at the moment into a role as Digital Transformation Director, which is related to all of STFC's transformation and prioritisation, which I'm sure you've heard all about.

What does digital transformation look like at STFC?

TG: It potentially means different things to different people. In this context, it's very specifically about the STFC prioritisation exercise and reshaping the organisation to be as efficient and effective as it can be. Part of that involves examining how we use AI and other tools to make the organisation more efficient and effective. It's also looking at how we deliver all of our scientific computing to meet our science's needs while being effective and efficient.

> **"We can already see AI touching pretty much every part of that facility's workflow"**

How much of an impact will AI have on workloads?

TG: That's the million-dollar question. If I could answer that, I'd solve all the problems. Within the facilities – and this goes back to questions we've heard before – we see AI. By facilities, I mean the Diamond Light Source, the ISIS neutron source, and the Central Laser Facility. We can already see AI touching pretty much every part of that facility's workflow. Facilities are big, football field-sized accelerators used to look at materials, viruses, and genomes, and AI is used at every step of that workflow to a greater or lesser extent: from experiment design, the pre-modelling and simulation that we do, running the experiment, determining when to stop the experiment, driving the experiment and data analysis, to running all underlying support services like cooling systems, cryostats, particle accelerator operations and fault diagnosis.

In the design phase, it's surrogate modelling. In particle accelerators, it's modelling for design, but also automated AI fault-finding, feeding the diagnostics of particle accelerators into an AI model (a digital twin, if you like) to speed up fault diagnosis and, even better, map blind faults before they occur, spotting subtle patterns in accelerator diagnostics so we can detect faults before they impact stability and operations. It really is everywhere. Across the wider organisation, it's much more about using turnkey AI tools, like Claude for coding, right through to Copilot to automate business processes within the organisation.

Where do you see STFC's role evolving within the wider research computing ecosystem?

TG: The key thing to remember about STFC

“Some use cases require that very high-end NVMe-type approach, and we have to buy what we can afford”

is that we are a dual body: we are both the funder and the doer of research. As a funder, we have long supported data-intensive doctoral training centres (CDTs) and AI CDTs, investing in postdoc and PhD schemes to train and deploy AI in high-energy physics. Within the labs, as doers, we have a role in developing those tools and putting them in the hands of users of our experimental facilities. AI is many things to many people. In many ways, it will give us better tools, and users will still run those experiments; it’s just that the control system will be AI-enhanced rather than manually programmed. STFC has published an AI strategy, and that clearly shows we have that dual role.

The other significant role we play is as a generator of high-quality data. Our facilities are data-generating machines, and if you look at examples like DeepMind and specifically AlphaFold, that model was developed because of data deposited over many years into repositories like the Protein Data Bank. A huge fraction of that data comes from synchrotron sources like the Diamond Light Source, which I’m looking at out of my window right now. So our role is to provide data, crucially, AI-ready data, and to ensure it is appropriately annotated and open, with a data policy that allows it to be reused by AI. STFC touches loads of different parts of the AI and data landscape.

Datasets need to be appropriately stored and managed for later use with AI. Are there best practices or data standards that you recommend for the wider community?

TG: It varies a lot by domain. If you look at particle physics, there are long-established open formats into which data in that domain is collected. Whether it’s machine learning, AI, or statistics, these approaches have been used for a long time in those communities, and they are adapting to adopt more modern AI approaches.

Looking at multidisciplinary facilities, such as materials science or biology, the degree of variation across scientific domains is large – probably greater than it is across facilities. Areas like bio in general and protein crystallography specifically have well-established data formats that are widely used and contain all the metadata. Others are behind the curve on that. We have established facility standards for raw data, but metadata quality and reusability vary widely across domains. Interestingly, we’re seeing a trend of users becoming aware that they need to provide higher-quality metadata to enable their own future reuse.

The other interesting side is that AI is driving how we need to store data. Beyond having the right metadata, we are seeing demand for different data storage styles. Because we store hundreds of petabytes of >



> scientific data, we do a lot of that on tape because tape is extremely good value – it is a cheap and very energy-efficient way of storing large volumes of data. However, it is not compatible with most AI use cases where you want to scan over huge volumes of data at speed. So we are having to make new investments in different performance tiers of data storage to facilitate AI use cases, and that comes at a high cost, especially given the memory crisis.

How do you strategise around the huge price increases?

TG: In some cases, the use case requires that very high-end NVMe-type approach, and we have to buy what we can afford. In others, we benefit from scale-out features. We operate well over 50 petabytes of spinning disk storage, and the scale-out nature of platforms like Ceph offers massive parallelisability, which fits some use cases incredibly well.

How has demand for these types of resources changed?

TG: It has grown significantly, and in many cases is outpacing Moore's law, certainly outpacing our budgets. A lot of what we support within the national labs and multidisciplinary facilities is for close-to-real-time use: in-experiment HPC. We are seeing much more demand for during-experiment support, modelling, and data simulation. That doesn't fit traditional HPC batch models and is pushing us toward using more cloud-like technologies. We have a very large OpenStack instance and see different usage patterns on that.

That's a challenge for efficiency, because batch systems are inherently efficient; they operate at over 90% efficiency in most cases. With cloud workloads, we have to balance the cost-benefit of having resources sitting idle versus direct support for science.

The most obvious change we're seeing is the rapid, massive increase in data volume, driven by increased source brightness. The Diamond Synchrotron is upgrading to a new accelerator, which will lead to 10 to 100 times the brightness in many cases. With detector increases and bigger, faster detectors, single detectors will output data on the order of hundreds of gigabits per second in the next few years. A facility like Diamond has around 30 beamlines, so you can multiply these numbers up several times.

Are there negative impacts of AI?

TG: The AI bubble has led to price rises in memory and GPUs, which is a problem against relatively flat budgets. On the positive side, DSIT [the UK's Department for Science, Innovation and Technology] has made significant UK investments in initiatives such as the AI Research Resource (AIRR) for the Isambard-AI machine and the AIRR-2 heterogeneous computing resource. One of our challenges is how best to make use of those. As I mentioned earlier, there is a tension between our need for in-experiment use of AI versus national services that are based around traditional batch-based processes.

Another change in facilities, such as Diamond and ISIS, is opening them up to a wider range of users. Twenty or thirty years ago, users were specialists in neutron or synchrotron facilities who knew how to do everything. Now we see a much broader range of science done on these facilities, which means providing different levels of user services, data analysis support, instrument scientists, and simpler interfaces and tools. People often come to use these facilities for just a couple of days a year or throughout their entire PhD, as one tool among many. We see much more diverse use, which changes how we need to provide data and systems to make them easier to use.

So you now have a more diverse user base?

TG: Very definitely. Scientists are pushing forward really exciting and innovative ways of using AI within their workflows; it has become a much easier tool to use. It is challenging for us when users arrive with a local LLM on their Mac, wanting data fed directly into it, which requires us to modify our systems to allow that outside the data centre, along with all the security and safety considerations that come with it.

There seems to be a lot of worry about the potential impact of AI on cybersecurity. How much of this is real fear about the here and now, and how much is based on the future of AI development?

TG: It is a tricky question to answer, but it does look pretty scary. The shrinking gap between when [vulnerabilities] are discovered and when they are exploited makes the timeframe [to patch] much shorter. The classic answer is defence in depth; we cannot rely on any single method of protection. It requires network segmentation, protection, patching, and threat analysis coming together. The rate of change, or the time from discovery to announcement to exploit, is terrifyingly short.

The availability of easy-to-use hacking tools that have the latest vulnerabilities available to exploit is very noticeable.■

Tom Griffin is the Director of STFC's Scientific Computing Department.

"Batch systems are inherently efficient; they operate at over 90% efficiency in most cases. With cloud workloads, we have to balance the cost-benefit of having resources sitting idle versus direct support for science"

Laboratory Informatics Guide

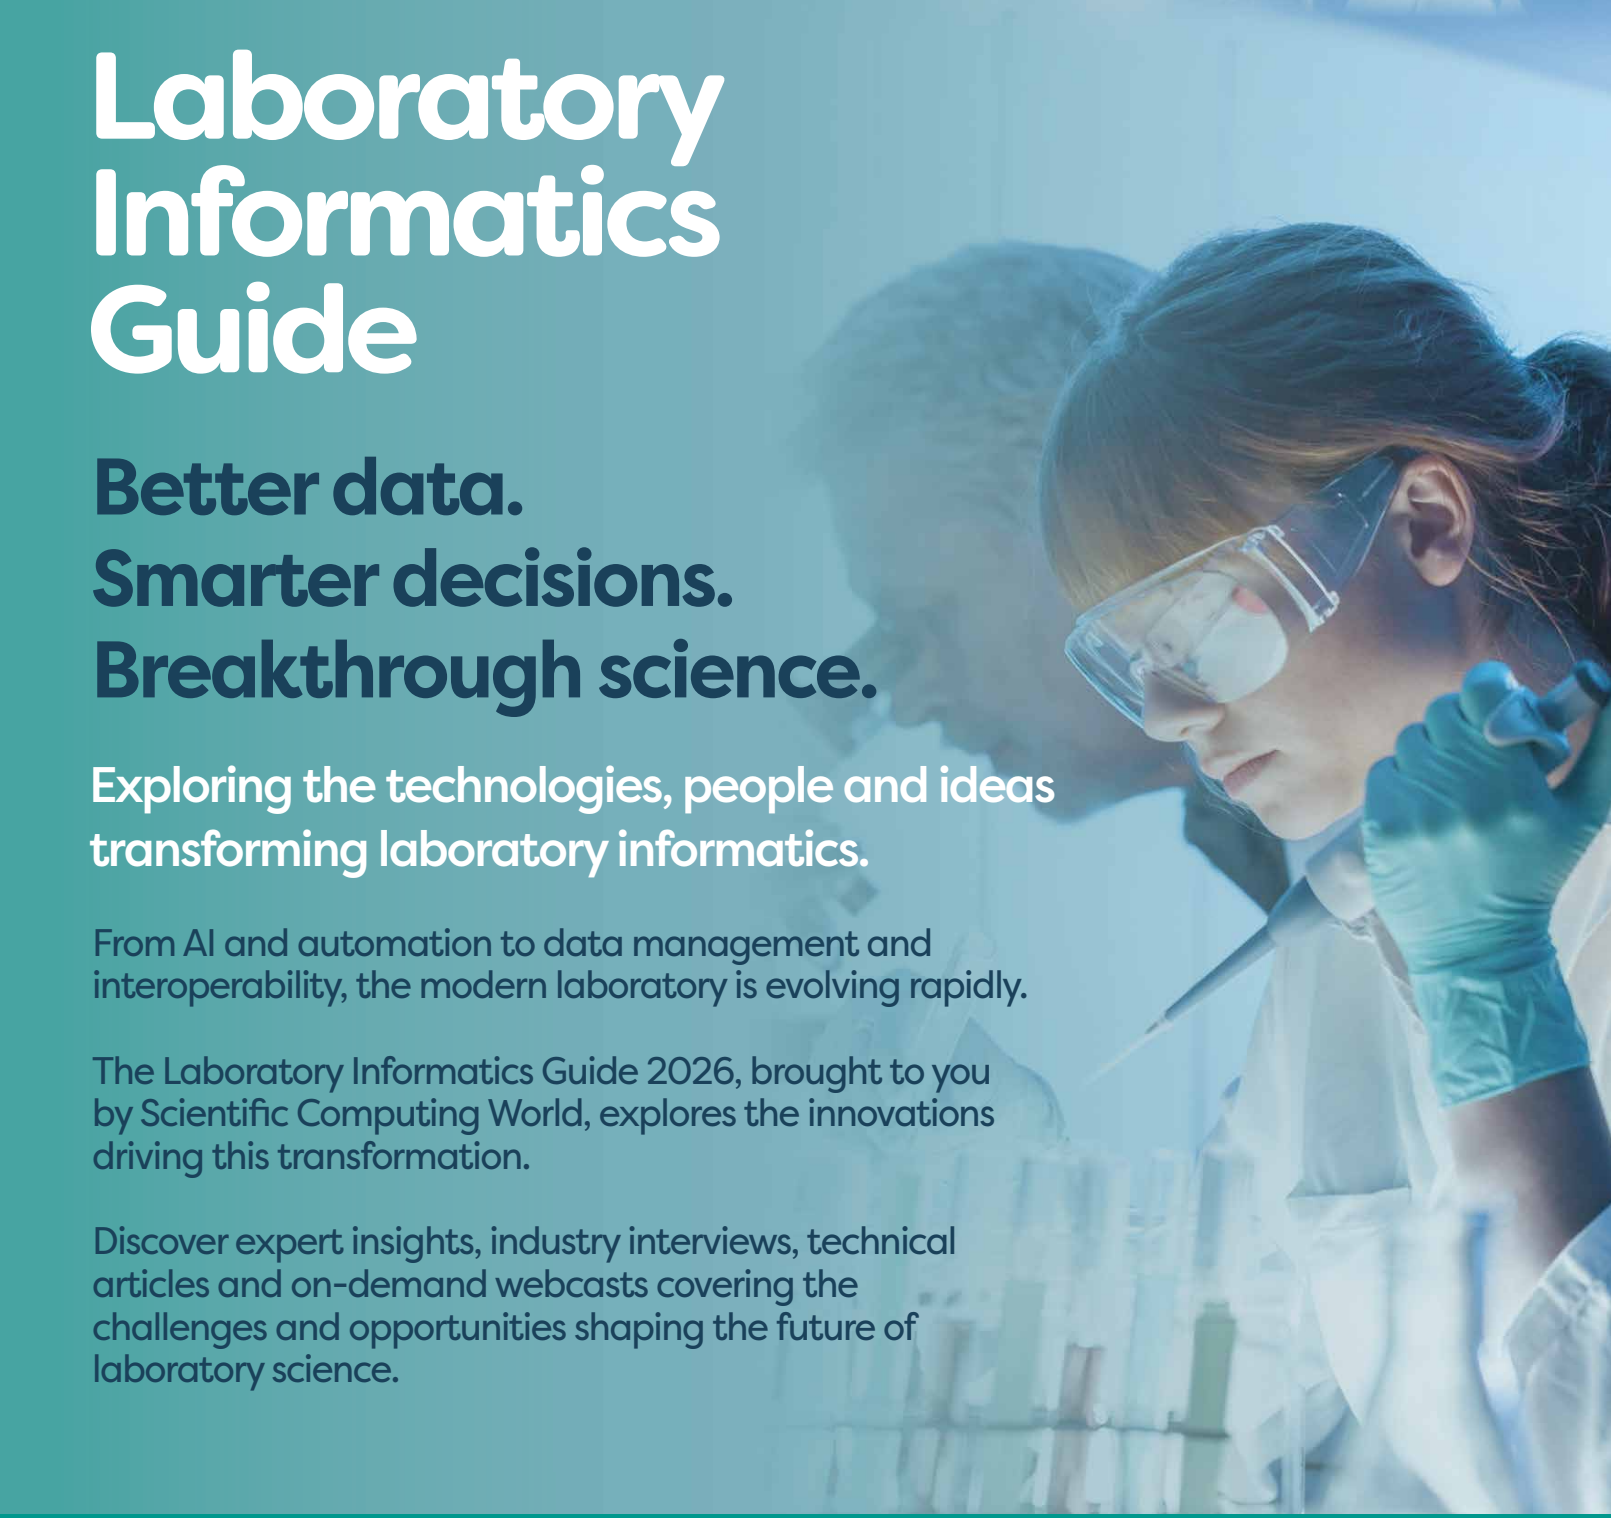

Better data. Smarter decisions. Breakthrough science.

Exploring the technologies, people and ideas transforming laboratory informatics.

From AI and automation to data management and interoperability, the modern laboratory is evolving rapidly.

The Laboratory Informatics Guide 2026, brought to you by Scientific Computing World, explores the innovations driving this transformation.

Discover expert insights, industry interviews, technical articles and on-demand webcasts covering the challenges and opportunities shaping the future of laboratory science.

Explore the guide. Discover what's next.

SCAN TO EXPLORE THE GUIDE

POWERED BY

SCIENTIFIC COMPUTING WORLD

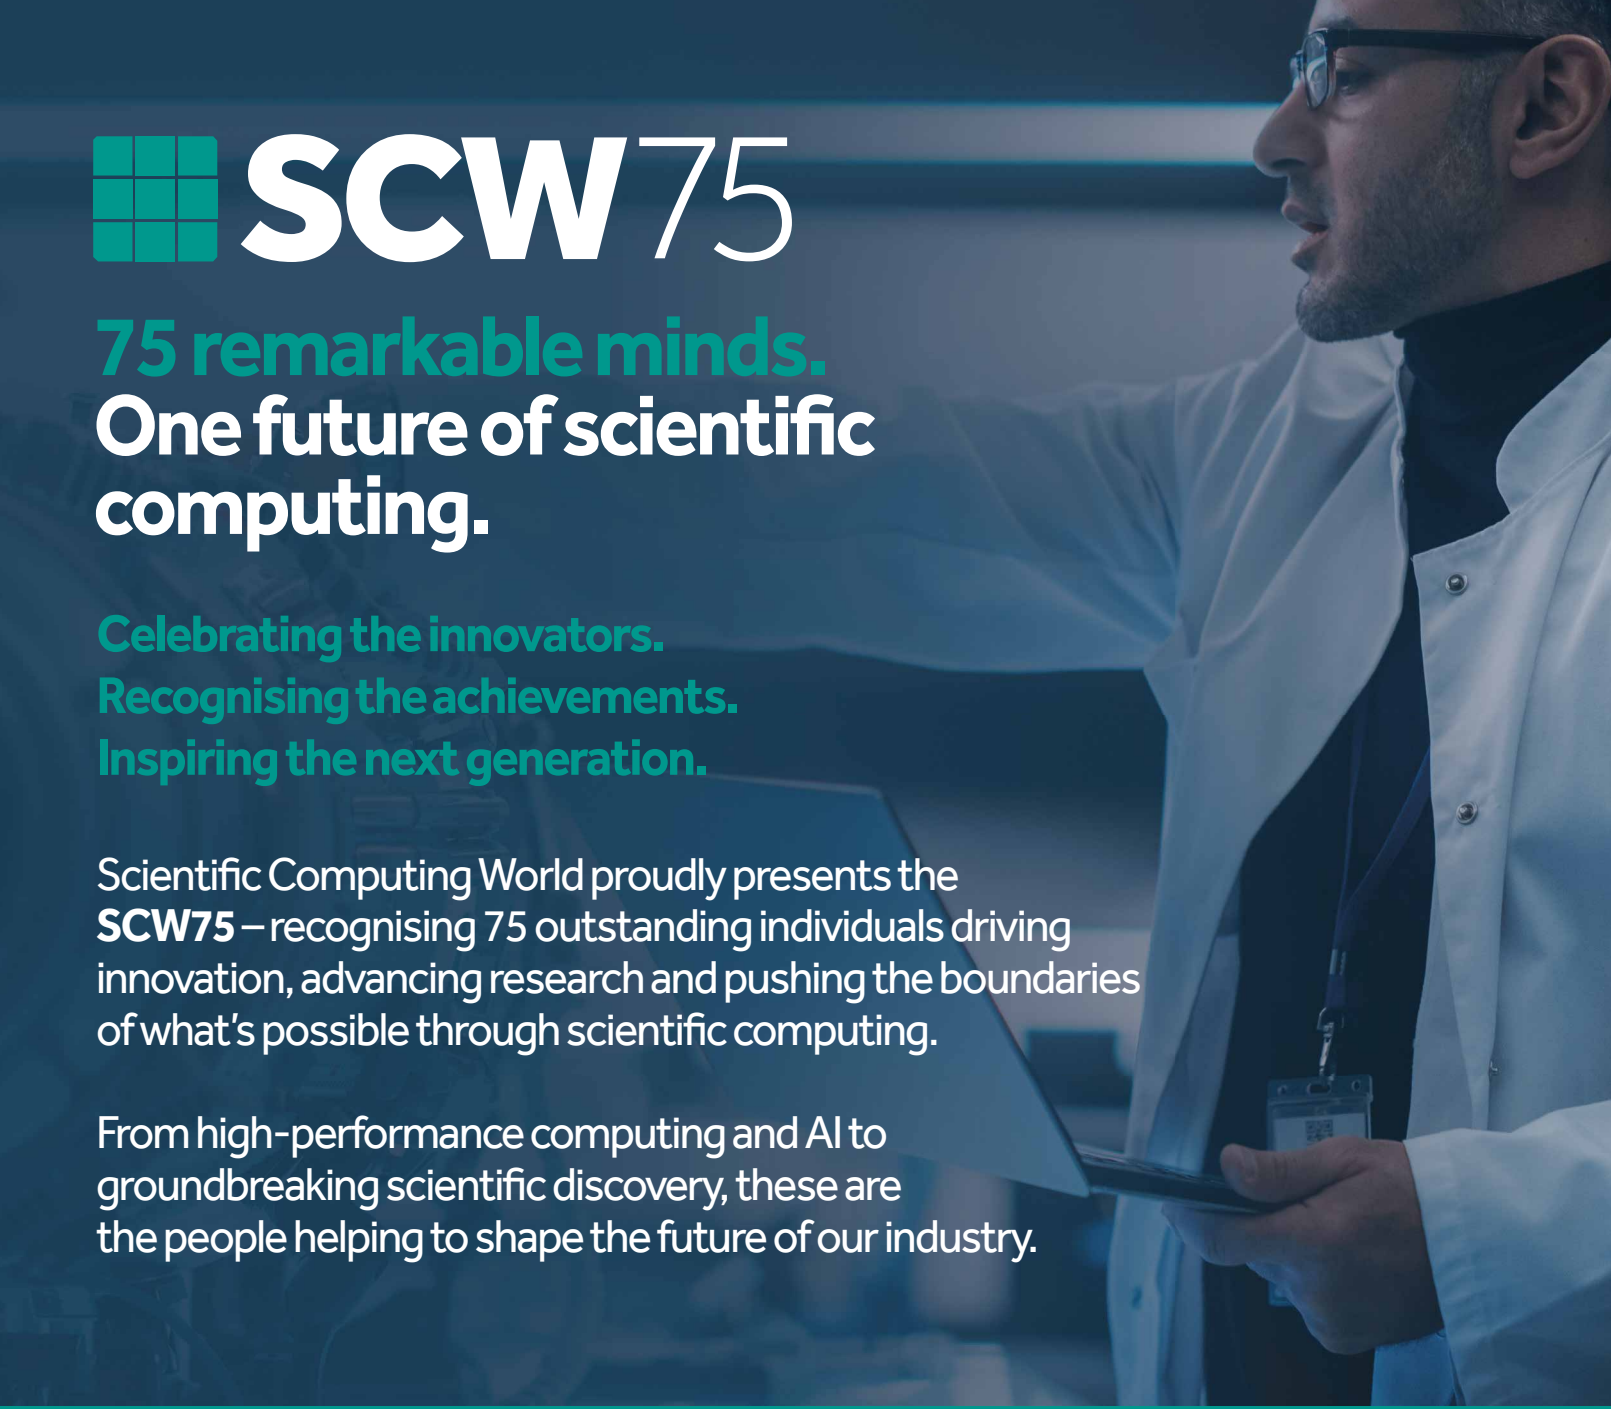

SCW75

75 remarkable minds.
One future of scientific computing.

Celebrating the innovators.
Recognising the achievements.
Inspiring the next generation.

Scientific Computing World proudly presents the **SCW75** – recognising 75 outstanding individuals driving innovation, advancing research and pushing the boundaries of what's possible through scientific computing.

From high-performance computing and AI to groundbreaking scientific discovery, these are the people helping to shape the future of our industry.

Discover the SCW75. Meet the innovators. Explore their stories.

SCAN TO EXPLORE THE SCW75

POWERED BY
SCIENTIFIC COMPUTING WORLD