



Accelerate scientific breakthroughs

Microsoft Azure for research



Contents

- 4 **A science-led recovery**
- 6 **Help your research team achieve more using the cloud**
- 8 **Solve global problems at speed**
- 10 **Sharp minds need powerful technology**
- 12 **Advanced scientific tools in easy reach**
- 14 **Useful tools**
- 14 **Next steps**

A science-led recovery

As the world recovers from the COVID-19 pandemic, economic disruption is colliding with problems like climate change. To map the way forward in a post-carbon future, we need smart minds working together – and time is of the essence.



A new dawn for us

The world is at a historic inflection point. By investing in a science-led economic recovery, we can build a stronger and more resilient future for all. With the right strategy to guide us, we can transition faster to a knowledge nation built on big ideas and innovative technologies—one that's ready to tackle problems like climate change, food and energy security, global competitiveness, our aging population, and disease.

This strategy must focus on future-proofing our research capability to support excellent research at scale and must empower the entire innovation sector to work as one to conceive and commercialise the best ideas.

We have the expertise

We have the research talent to spearhead a science-led recovery. Universities across the world are researching fields like chemistry, earth sciences, and engineering; and researchers are pioneering everything from smart farms and agricultural robotics to solutions for space junk.

Many of their discoveries can literally be life-changing. For instance, surgeons can now see CT and MRI images in 3D during surgery. These crucial medical images expand their reality, enabling them to operate with greater precision. Or take a solar-based system that offers Kenyan farmers personalised recommendations and solutions through their mobile phones, to help them raise better quality highyield crops and make more money. Or consider the impact of a new generation of solar panels that can efficiently produce hydrogen gas out of thin air, of a novel process for cleaning up crude oil spills using magnets, or of a new vaccine that blocks the chronic pain caused by osteoarthritis.

It's ideas like these that will ensure our future. But to support their inception, we need to equip our researchers for how science is done now.

How science is done now

Sparking intense collaboration among researchers globally, the coronavirus outbreak accelerated a new research culture characterised by open datasets, rapid-fire exchange of knowledge on servers and an explosion in public papers. Scientists talk of an "extraordinary amount" of information.

To contribute to the myriad incremental increases in knowledge that nourish the greatest innovations, our researchers need to surface important findings, analyse massive datasets and share them so they can reproduce each other's experiments. To spearhead a science-led recovery, they need sustainable technology infrastructure that supports open, collaborative science – plus the skills to use it to discover and commercialise valuable ideas faster.



Extend your capabilities with ours

Microsoft Azure lets your research harness one of the world's most powerful public clouds for research.

- Infinitely scalable compute resources based globally
- 99.95% monthly availability with 24/7/365 support
- Secure, on-demand access to vast storage and supercomputing resources
- No additional investment in capital assets or IT infrastructure, freeing research funding
- Sustainable zero emissions technology

Diving into data

More and more, today's academic research teams are using high-performance computing (HPC), artificial intelligence (AI), and machine learning (ML) to delve into complex questions and power scientific enquiry across huge datasets. Scaling from desktops to supercomputers, they require large amounts of data storage, processing and analysis. Many are now using cloud computing to help meet these needs in a cost-effective way.

The public cloud

The public cloud provides modern research infrastructure to extend your own. Far more flexible than investing in your own technology, it allows researchers to dial up storage, compute when needed, and dial down when they're done. With appropriate security to protect intellectual property (IP), academic research teams can share immense datasets with research partners and industry, accelerating game-changing discoveries.



"Modern science is being called upon to solve global challenges at a rapidly increasing rate. The academic model that has driven research outputs in the past is struggling to keep up with the pace of change and a new focus on converting research into action ... The scientific community needs to embrace collaborative practises and harness the power of globally connected IT infrastructure to make solutions available when and where they are needed most."

Dr Justin Perry, Research Scientist and Project Leader, CSIRO

Help your research team achieve more using the cloud

Researchers around the world are using Microsoft Azure to accelerate their research. It provides an open, flexible, global platform that supports multiple programming languages, tools, and frameworks to enable faster results.



Faster discoveries

Work easily with huge datasets and use AI, HPC and Azure discipline-specific tools to arrive at new insight faster.



Virtual research teams

Share and manage code across research communities using GitHub repositories. Integrate IoT for remote sensing data and enable hybrid working for your team in a crisis.



Cost control

Pay for computing resources only when needed, plus use Azure tools to optimise cloud workload costs, reducing total cost of ownership.

Data management and governance

Make research discoverable, apply data policy and protect university IP with a best practice platform for automating governance.



Quality control

Students can view project data, upload results, and run analytics workloads with role-based access protecting data integrity.



Best practice code

Software engineers and researchers can use Azure's bestpractice coding toolsets to create custom research software to solve new scientific challenges.



Solve global problems at speed

To find answers to big challenges like climate change and pandemics, scientists in data-intensive fields like climatology and epidemiology are interrogating immense datasets to reveal the complex interplay of factors driving natural phenomena and disease. The Azure platform helps them turn evidence to conclusions more quickly.

Faster discoveries

Solve complex problems with big data

Make headway in data-intensive fields like astrophysics, geophysics, and genomics. In minutes, researchers can set up cost-effective Azure infrastructure to ingest, store, and share their biggest datasets. It's easy to add computing power to further investigate data using modelling, simulations, and analysis. Or to import Azure Open Datasets to train ML models using public data on weather, city safety, and more.

Build on what came before

With thousands of new findings coming out of the worldwide scientific community daily, Azure AI can help researchers digest the "extraordinary amount" of information and stay on top of their fields by surfacing relevant finds and data in scientific literature. And to help researchers reproduce findings and take the next step, Azure makes information more discoverable with best-practice governance and collaboration.

See things differently

Arrive at new insight with big data tools built for the Azure platform and available on demand, including HPC, AI, ML, and data analytics. Azure helps your team adapt code to the world of big data and the cloud, plus you can use our intelligent tools to reveal new patterns, relationships and trends in research data. There's no need for capital outlay—simply switch the tools you need on or off to suit research timelines.

CASE STUDY

IoT helps preserve ancient Roman ruins

A team of engineers at the Politecnico di Milano provided archeologists surveying a restricted underground site in Rome with the smart-sensor technology they needed to determine the impact of environmental factors on the sculptures and frescoes. The team used Azure to create a simplified IoT-based system that could deliver the collected data in customisable, easy-to-use graphs, with different tabs for different sensors and filters.

"When the archeologists were shown they could simply select the metric and a time window to see the data readings, it was fantastic for them. The Azure solution provided the data they needed in a way that made it valuable."

Luca Mottola: Associate Professor, Politecnico di Milano, Italy Massey





Virtual research teams

Collaboration

Azure lets everyone working on a research challenge access the latest data so you can collaborate across entire disciplines and industries. Data integrity is protected by access controls with multiple layers of security and governance, so you can co-create knowledge with global experts and citizen scientists alike. On BRUVnet, they're populating the world's largest training dataset of fish species this way. With Trove, you can pay people for the photos you need to train ML models.

Easily share ideas, methodology and research

Startling innovations can result when researchers connect data points and insights across existing knowledge and scrutinise and reproduce each other's work. Azure enables this with an open, flexible, fast, efficient infrastructure for sharing large datasets and findings. With research data at their fingertips, researchers can collaborate easily, plus replicate experiment results on other researchers' datasets to strengthen the foundation for future enquiry.

Resilient in a crisis

The COVID-19 pandemic has underscored the importance of resilient research infrastructure that supports flexible work styles, including remote working and hybrid working. With Azure, you can keep researchers connected and productive wherever they are. You can easily integrate Teams, SharePoint, and OneDrive so they can disseminate their results and communicate more effectively. Power BI helps them manage shifting priorities with up-to-the-minute reporting tools.

CASE STUDY

Using Azure to accelerate research at prominent Pakistani research center

Researchers at the International Center for Chemical and Biological Sciences in Pakistan were having trouble computing, storing, and managing the data from their DNA sequences. Azure created the perfect environment for accelerated research by providing scalable storage capacity and effective computing power, which users could access simultaneously. The platform was also used to create an interface for their learning website and store thousands of video-lectures.

"Azure has helped through massively scalable storage. The infrastructure on Azure streamlined our platform to match our needs, providing us with a single platform to monitor and analyse our data requirements and scale accordingly."

Dr. Ishtiaq Ahmad Khan, Assistant Professor, Jamil ur Rehman Center for Genome Research



Sharp minds need powerful technology

Now even the smallest research teams have the computing firepower to investigate science's biggest questions. On Azure, there's no capital outlay, no on-premise hardware to maintain, and data controls and governance are built in. Simply pay for what you use—and scale up and down as needed. Plus, spark new ways of seeing data with HPC, ML and other tools on demand.



Data management and governance

Manage rapid data growth easily

Scale your Azure research infrastructure easily to keep pace with scientific breakthroughs and collaboration opportunities. You can add on-demand storage and backup to suit any research requirement—all with best-practice security and governance. Data tags enable fine-grained life cycle management, curtailing data and budget bloat.

Make research findings discoverable

Data saved on hard drives and thumb drives doesn't fuel the future of collaborative science—and it's vulnerable to loss, theft, and corruption. Azure provides a universal platform to ingest and share every researcher's data. Policy-based tags impose a common logic, helping to surface data to the right scientific and industry collaborators.

Protect your IP

There's a misconception that the cloud isn't for sensitive data. In fact, it can help prevent catastrophic loss and data leaks resulting from compromised university servers. Security and sound governance are built into Azure and constantly updated. Fine-grained access controls protect data safety, integrity, and reproducibility—and your IP.


CASE STUDY

A massive high-performance computing experiment enabled by Azure

The University of New South Wales (UNSW) can respond faster to challenges like the COVID-19 pandemic with Azure providing a single source of truth for reporting. UNSW used Azure Databricks to create a second Azure data lake to store curated data. Workers can create their own reports from this data in record time using Power BI. UNSW is now exploring using AI and ML to analyse the data to identify contract cheating and more.

“Microsoft support was outstanding. Access to their software architects, specialising in open source, help in designing the experiment and providing 24/7 support.”

Assi Barak, Software Group Manager, Bar-Ilan University, Center for Research in Applied Cryptography and Cyber Security





Cost control

Scale out to a supercomputer

Empower researchers to take control of their technology budgets and eliminate waits for computing resources. Azure offers a cost-effective resource model with no capital outlay. You can start with what you need today and scale out to a supercomputer tomorrow. Switch on HPC, analytics, AI, and ML instantly, plus access discipline-specific tools like Cromwell or FarmBeats on demand to accelerate insight.

Reduce technical overhead

It makes no sense to duplicate or under-utilise expensive equipment. Why not pay only for the time you need? With Azure, you can make the most of scarce funding by reducing on-premise hardware and, with it, total cost of ownership. You can also maximise the value of your infrastructure investment by storing operational, compliance, production, and research data on the same Azure system.

Simplify your IT

Boost productivity with a cloud solution that's built to integrate with you. Azure automation makes it easy to ingest, manage, and analyse data, saving time. Alerts help control cloud workload costs so researchers stay on budget. And by outsourcing technology maintenance, upgrades, and IT resources, your IT team can deliver more value by working on your institution's strategic priorities instead.


CASE STUDY

Azure, asteroids and AI

John Hefele used Azure Batch to spin up 200 virtual machines in about 15 minutes. He ultimately scaled to 500 VMs. He was able to generate a million asteroid trajectories in just nine hours—a task requiring at least three weeks on the university's supercomputer, a shared resource that's hard to access. He used the data to train a neural network that identified 11 asteroids that could impact Earth and weren't yet flagged as a threat.

“[People] assume that cloud computing is going to be really expensive, but I was able to generate a million simulations using Azure Batch for just \$150.”

John Hefele, PhD candidate, Leiden University



Advanced scientific tools in easy reach

Now even the smallest research teams can use AI, HPC and discipline-specific tools created for the Azure platform. No huge overheads – simply pay for what you use.



Perform large-scale parallel and HPC batch jobs efficiently

Azure Batch and Batch Shipyard

Run large-scale applications efficiently in the cloud with cloud-scale job scheduling and compute management. Docker and Singularity containerised workflows don't require infrastructure setup. Add Azure Kubernetes Service (AKS) to help deploy and manage containerised applications.



Build, train, and deploy AI models collaboratively at scale

Azure Databricks

Work at scale with AI models using GPU-enabled clusters with this Apache Spark-based analytics service. Databricks Runtime comes preinstalled and preconfigured with deep learning frameworks and libraries e.g. TensorFlow, Keras, and XGBoost.



Integrate data from field devices or leverage AI in smart farming

Azure IoT, including FarmBeats

Securely connect, monitor, and manage billions of devices to develop Internet of Things (IoT) applications. Easily integrate meteorological, satellite, and sensor data plus other IoT feeds.



Build, train, and deploy ML algorithms and enable reproducible ML research

Azure Machine Learning

Whatever your level of AI expertise, get up and running quickly with everything from automated machine learning (AutoML) to low-code drag-and-drop tools (Azure ML designer), to full code with Jupyter Notebooks. Azure ML automates deployment and downscale of compute clusters.



Access the latest research

Microsoft Academic Graph (MAG)

Synchronise MAG to your Azure Storage accounts and see your searches in a heterogeneous knowledge graph—find scientific publication records and see citation relationships, authors, institutions, journals, conferences, and fields of study.



Manage expanding data efficiently with a scalable system

Azure Storage (Blob, Disks, Data Lake)

Azure Blob storage provides massively scalable, secure storage for your unstructured data. For critical workloads, disk storage is dynamically scalable. Extend your current storage systems easily by integrating an Azure Data Lake.



Analyse large genomic datasets

Cromwell on Azure

Orchestrate the computing tasks needed for genomic analysis with an open source scientific workflow management system that lets you run scripts on local machines or computing clusters, plus the cloud.



Manage HPC clusters on the cloud cost effectively

CycleCloud

Dynamically provision HPC Azure clusters and orchestrate data and jobs for hybrid and cloud workflows. Compatible with Slurm, Kafka, Zookeeper, PBS, Anaconda, and more.



Simplify data science in the cloud

Data Science Virtual Machines

Get started quickly with Azure Virtual Machine images (Linux/Windows), pre-installed, configured, and tested with tools widely used in data analytics, machine learning, and AI training.



Achieve real high performance computing (HPC) on a public cloud

InfiniBand-enabled HPC

Enable MPI and RDMA workloads to run rapidly at scale with compute and GPU resources connected over a low-latency, high-throughput network.



Share large datasets with research partners

Azure Data Share, Azure Industry Collaborative Service

Share research datasets with multiple collaborators with fine-grained controls. Enhance insights by easily combining data from third parties to enrich analytics and AI scenarios.



Manage low latency for realtime inference or model scoring requests

Field Programmable Gate Arrays (FPGAs)

Azure FPGAs enable ultra-low latency inferencing, even with a single batch size, using ResNet 50, ResNet 152, DenseNet-121, VGG-16.



Use off-the-shelf capabilities to train ML algorithms and keep AI costs down

Azure Cognitive Services

Tailor your AI capabilities to your specific needs using custom-developed models based on our broad range of AI services.



Make research reproducible

Azure Open Datasets

Improve the accuracy of your ML models and save time on data discovery and preparation by integrating ready-to-use public datasets straight from Azure services into your own ML workflows.

Useful tools

Researchers around the world are using Microsoft Azure to accelerate their research. It provides an open, flexible, global platform that supports multiple programming languages, tools, and frameworks to enable faster results.



Azure pricing

OCRE customers can find the current Azure public pricing by product/service at <https://azure.microsoft.com/en-us/pricing/#product-pricing>



Azure calculator

Customers wishing to confirm costs of deploying services to Azure can leverage the Azure pricing calculator to enter the type, quantity, and preferred purchase option i.e. PAYG or Azure Reservations by using the Azure Calculator at <https://azure.microsoft.com/en-us/pricing/calculator/>



Azure Cost Management

All OCRE members purchasing Azure under Volume Licensing will be entitled to use the Azure Cost Management solution free of charge. Azure Cost Management helps organisations visualise, manage, and optimise costs across Azure.

Next steps

Making Azure adoption easier

The OCRE framework, available to GÉANT members, enables over 10,000 universities and research institutions, served by 40 National Research and Education Network (NREN) organisations across Europe, to benefit from:

- Significant discounts across Microsoft Azure services
- Ready to use contracts designed to satisfy national and European public sector procurement regulations and GDPR requirements

Contact your [Microsoft OCRE Partner](#) who will work with you and support your Azure journey.

- Microsoft Azure – [Learn more](#)
- The Open Clouds for Research Environment (OCRE) Framework – [Learn more](#)
- GÉANT – [Learn more](#)





