

# Why top CTOs are building unique data fabrics

There is so much insight in your data. Knowing what your users are after will help you unlock the best of modern data storage solutions.

All organisations have the same kind of problem – how to take their infrastructure investments and get as much insight from their data in an economical way. From scientific research to the enterprise, storage topology is changing to derive greater business value through integration of on-premises and cloud.

CTOs across industries are having to navigate this changing topology. So, what's really at stake with modern data storage and the capability versus cost equation of cloud economics?

Organisations are adopting everything from flash-based storage on-premises to the cloud - and increasingly a hybrid of both - to meet exponentially growing volumes of data and demand to access it quickly to make sense of it and identify trends.

### INSIGHTS ON CLOUD ECONOMICS

In a poll conducted during the IBM-iTnews webinar, 63% of attendees said cloud alone is an “overrated” or “overhyped” solution to today's IT problems.

However, data storage alone can't solve all these challenges either.

University of Queensland Research Computing Centre CTO Jake Carroll says, “Information infrastructure alone is not good enough - it's infrastructure plus an understanding of the pressures the infrastructure will be under. Otherwise, all you've got is metal.”

This understanding is critical to optimising where to run workloads to get the most out of modern data storage.

Andy Walls, IBM Fellow, CTO and Chief Architect of Storage is unsurprised but cautious in interpreting the poll results. “I think we need to be careful equating overhype with not important,” he says. “What really is settling now is what we call the hybrid cloud where you have a combination of cloud and on-premises,

and you can shift data between them, because that gives you an economic advantage to do what's best for a particular application - or even for that application to use both.”

### STAYING CLOSE TO USER REQUIREMENTS

Carroll - in his own words - is “not a distant CTO” and is closely connected to what these applications need to deliver.

“I'm a person that sits with our researchers on a regular basis. I would prefer to be with them in the lab, in the data centre or in front of their scientific research infrastructure. It's incredibly important to me to interact with our users at the coalface because only then do I actually understand really well what people are doing,” Carroll tells the IBM-iTnews webinar.

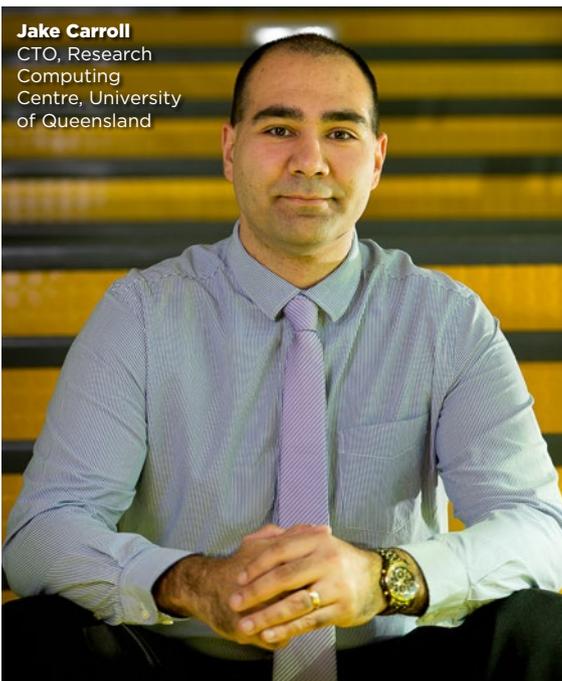
“When you've got too much distance, you don't necessarily fully appreciate the difficulties people have, the trials they go through or what they really need. By being close to users, we can translate their needs with technology partners to build infrastructure for our people.”

For partners like IBM, one-to-one customer interaction and feedback is key to ensuring they are improving the performance of their storage and that their solutions are fit-for-purpose and meeting users' needs.

It's clear CTOs like Carroll have a deep understanding of their user's requirements so that they can build the right data fabrics and make data more accessible to their users to enable them to quickly gain new insights from their data.

“It's extremely important to me to take feedback, both from our own internal research as well as research partners like Jake because it's critical to put our products in [there], see how they use them and then have that feedback come back to IBM,” Walls says. “Partnership is critical because

**Jake Carroll**  
CTO, Research  
Computing  
Centre, University  
of Queensland



“Information infrastructure alone is not good enough - it’s infrastructure plus an understanding of the pressures the infrastructure will be under.”

**Jake Carroll**

CTO, Research Computing Centre,  
University of Queensland

[UQ is] using some of our products right now and it will be invaluable to see how they use them and what we can do better.”

## MANIPULATING DATA SETS AT SPEED

Some of UQ’s most important research right now that requires the support of strong information infrastructure and quick access to data relates to the COVID-19 virus. UQ is currently in pre-clinical testing of a vaccine.

“Research is always exciting, but lately things have become really vibrant and interesting. What’s driving me is making sure that we are providing solutions and value to our researchers who are tackling really serious problems at the moment ... like the COVID-19 virus. We are spending a lot of time and infrastructural effort, and we’re paying a lot of attention to efforts to empower those users, because it is a global fight and it is extremely serious,” explains Carroll.

Both Carroll and Walls’ clients share many of the same challenges when it comes to data and storage: the need

for low latency and high performance to be able to manipulate data sets at speed.

Outside of COVID-19, the University has Institutes and research focused on ageing and dementia, along with neuroscience, material science - and the list goes on. “Everything you can imagine, we’ve probably got one of, and we’re heavily involved somehow in accelerating [answers to] their problems.”

“To get data everywhere and into people’s hands ... the challenges are enormous,” Carroll says. “You have to have amazing networks and amazing storage, a pretty keen vision, and a willingness to walk on the journey to do this. It’s hard stuff, but it’s how research has to operate now - unencumbered - if you want success.”

## BALANCING CAPABILITY AGAINST COST

The challenges faced by Carroll and UQ are also common across enterprise environments, says Walls.

“In some ways, all clients have the same kind of problem as Jake, and what they need to be able to do is take their infrastructure and get as much performance as they possibly can from it in an economical way, so they can get data and make sense out of it, and not just store it,” Walls explains. “All clients are really facing the same challenges, even if they’re not doing scientific research.”

By understanding the needs of users, organisations can choose appropriate information infrastructure. Carroll says that solution should strike an appropriate



**Andy Walls**  
IBM Fellow,  
CTO & Chief Architect,  
IBM Storage

“economic and functional balance”. Walls similarly says it’s about balancing capability against cost.

While organisations may have sought to do that in the past by going all-in on cloud, Walls says some are now walking workloads back onto on-premises hardware and private cloud infrastructure, or a mix of on-prem and cloud. “You’re never going to get completely rid of your on-premise storage but the requirement is how to make that on-prem storage work together with the cloud” Walls says.

## PREPARING FOR THE CONSUMPTION OF CHANGE

The IBM-iTnews poll also found managing complexity due to data storage silos is a specific challenge in data storage environments, alongside inadequate performance, scalability and reliability.

“The problem is the amount of

data has grown exponentially but what hasn't grown is the budget or the number of people administering all that data or the storage behind it," says Walls.

UQ is solving the challenge by stitching together different data fabrics such that it can present data in multiple ways to users - bridging devices, operating systems and workspaces. "What's important here is the accessibility and the multiple ways people can actually get at the data," Carroll says. "Data consumption is the key."

IBM is addressing manageability directly in its hardware and cloud, including an ability to use the "same GUI, CLI and commands" on flash storage or the cloud. "The storage looks the same," Walls says. IBM is also cramming more storage and performance into smaller enclosures; it's FlashSystem 9100 and 9200, for

example, offer 24 FlashCore modules in a 2U product. "In 2U, you can fit 1.7 petabytes of storage. You can fit that much storage in three-and-a-half inches of one rack."

In the future, IBM storage may be able to manage itself, further reducing administrative burden, if current research comes to fruition. "We're looking to use artificial intelligence to make sense of all of the data coming back from storage products to detect anomalies, to help the administrators and the storage personnel figure out what's wrong, and to do it very simply," says Walls.

### SECURITY AT THE CORE

Data security concerns must also be addressed in architectural decisions around information infrastructure. "Security keeps me awake at night," Carroll says. "It is one of my top concerns along with data integrity

and making sure we have continuity of service and sustainability".

It also remains an inhibitor to cloud use, with the IBM-iTnews poll finding security and compliance concerns would stop 38% of webinar attendees moving a workload into the cloud (followed by cost, at 32%).

IBM's products and services are designed to be compliant with the likes of the EU's General Data Protection Regulation (GDPR) and similar schemes worldwide. "We take security very seriously at IBM," Walls says. "There's encryption to prevent people from getting in, but in case they do, there's cyber resiliency to rebound from it."

Walls also notes the importance of making security and compliance cost-effective, reflecting the poll research. "Security and compliance are extremely important, but they have to be affordable," he adds.



# Why top CTOs are building unique data fabrics

---

To watch the on-demand replay of the IBM-iTnews  
webinar, [click here](#) or for more information, visit  
<https://www.ibm.com/au-en/it-infrastructure/storage>

The logo for iTnews, featuring the letters 'it' in a lowercase, sans-serif font with a small red square above the 'i', followed by 'news' in a lowercase, sans-serif font.

The IBM logo, consisting of the letters 'IBM' in a bold, uppercase, sans-serif font, with a registered trademark symbol (®) to the right.