

How to Mitigate the Effects of Data Gravity



SPONSORED BY



DIGITAL REALTY



Introduction




The world has changed. We now live in a data-rich, digitally driven world where remote working is prevalent, and new connected technologies are creating fresh potential for businesses. But the information infrastructure that underpins Industry 4.0 brings with it novel challenges that forward-thinking IT leaders are now having to address. And the sheer volume of the world's data is proving especially challenging.

According to [Statista](#), the 'total amount of data created, captured, copied and consumed globally is forecast to increase rapidly, reaching 64.2 zettabytes in 2020.' By 2025, this is set to increase to more than 180 zettabytes – or 180 quintillion gigabytes of data. Although just a small percentage of this data is kept, with just two percent saved and retained between 2020 and 2021, this too is set to increase. For instance, [the IDC predicts](#) a compound annual growth rate of 26% between now and 2024.

While this growth in data represents a kind of gold rush for many companies, gathering, processing and storing this data has and will become problematic in its own right. In order to overcome issues of latency and to ensure companies remain efficient, IT leaders must bring their applications closer to their data. But as they do, and their data grows, it will become more difficult to move the data and grow the information infrastructure.

This is the unique challenge of data gravity. And it's becoming more challenging by the day.

In this report, we'll look at what data gravity is (to ensure every reader is up to speed – if you're familiar, skip ahead) before delving into how IT, cloud and data leaders can deal with data gravity. We'll both assess the problems and recommend steps to take to overcome them. This is the first step towards unlocking new value for your business – are you ready?



What Is Data Gravity?

Data Gravity as a term first arose over a decade ago when Dave McCrory used it in a [blog post](#) to describe the relationship and attraction between data and applications. He highlighted that as datasets were growing ever larger, they were getting more and more difficult to move. To keep applications and services operational, IT leaders were therefore having to move their apps closer to the data. Hence the term data gravity as an analogy: as in nature, objects of a larger mass (in this case, data) will attract other, smaller objects (apps).

The term also draws attention to two critical issues surrounding big data: those of latency and non-portability.

The Challenges of Data Gravity

Firstly, let's look at latency. To get the most out of your data, and to run optimal workloads and workflows at speed, your applications need to be relatively close to the data or you'll suffer latency issues. For instance, it's no good having an application running on one side of the globe but using a dataset stored on the other. It's clear it would be faster for you to run the application closer to the dataset so that you benefit from heightened workload performance.

But trying to limit the detrimental effects of latency would also mean limiting your operational ability – and your market. You would have to move your business's primary location closer to your main data centre, and you would find it difficult to expand into new markets because latent speeds would limit your ability to operate in more distant areas. In a similar vein, if your company's dataset outgrew your current infrastructure, you would also face issues of latency as systems struggle to deliver data for you.

This throws up the other critical issue: non-portability. If a company stores all of its data in one location, and that data continues to grow in both real terms and its importance for workloads, it will become increasingly difficult to move it to a new location. That might not seem like an issue right now, but forward-thinking data leaders must understand that eventually the company and its data will outgrow its existing data infrastructure. New solutions will be needed, and data will have to be migrated.

Unfortunately, that migration will likely come at great cost and expense. The temptation is therefore to stick with the existing data centre choice, bringing more and more applications to the data. In other words, bringing more of the company's

core operational processes to the data. In effect, this hampers the potential for the company to grow – meaning it becomes even more costly to move data.

Data gravity, as such, is a huge problem for enterprises and smaller organisations alike, both now and in the future. But there are options available for IT and data leaders looking to mitigate the damage of data gravity.



Emerging Tech Processing Requirements Also Fuel Data Gravity

Cloud technology has certainly powered a new era of data expansion, serving to increase data gravity. Compounding the issue is the rise of activity at the edge, with the Internet of Things (IoT) and technologies like Artificial Intelligence and Machine Learning (AI/ML) requiring more data processing than an on-prem architecture could ever handle. This edge activity is powered by easy access to data, meaning companies must move more of their apps into the cloud to deliver seamless processing.

But this extra processing creates even more data gravity. Traditionally, architectures would centre on data backhaul and bring the data to the app for processing. Now the tendency is to take the app to the data, as the apps require so much data that to do it any other way would be too costly and inefficient to consider.

As emerging technologies increasingly operate at the edge, this idea of bringing processing to the data will only gain more of a foothold – and data gravity will increase as a result. IT leaders need to offset these issues by planning ahead now.

How Do We Deal with Data Gravity?

There are multiple ways to mitigate the effects of data gravity on your organisation. Here are the top things to think about.

Plot Not Just Where Your Data Is, But Where It Will Be

Data gravity affects every business to some degree or another. Where some people fall down is underestimating how much data their architecture will have to handle, both now and in the future. It isn't enough to plot and strategise for where your data is; you must plan ahead so that your architecture, and therefore your company, has the capacity to scale and adapt to changing circumstances and the likelihood that you'll have more data as time goes by.

Failing to plan for growth will mean you have to migrate your data sooner rather than later, adding to the cost of your data not just in terms of architecture, but in lost potential revenue and business value. However, plotting ahead isn't as simple as thinking that all you need is more capacity. You also need to consider how to offset issues of latency, thereby ensuring you can maintain operational efficiency and have the potential to expand into new markets.

The Data Gravity Index

In 2020, Digital Realty released the exclusive Data Gravity Index report, where they measured the intensity of data gravity and its effect on the Global 2000 Enterprises. The forecasts demonstrate how much of a problem data gravity will become through to 2024, and the sectors that will see the most intensity.

[Download your copy now](#) to understand how data gravity will impact your country and sector.

Replicate Your Data Across Multiple Data Centres

That's where replication could help. Instead of having a single location for your data, replicating it across multiple data centres could reduce your latency issues and keep your business efficient regardless of geography.

To reduce unnecessary replication, it may be worth assessing what data each territory needs in order to operate. If your US operation, for instance, requires access to more datasets than your APAC operation, only replicate the relevant data the APAC operation requires in order to fulfil its function.

Not only does replication help to reduce latency issues, but it also helps with portability. By replicating your datasets, you'll have increased your company's capacity to scale and therefore won't have to worry as much about porting your data in the future.

Utilise a Hybrid Cloud Approach

With a hybrid cloud approach, you can further mitigate the negative effects of data gravity.

As we've discussed data gravity throughout this report, we've made it clear that large datasets have a tendency to pull apps closer to the data to avoid latency issues. But you should ask yourself: should I bring my data to where I'm doing the processing or bring my processing to the data? In other words, do you give in to the data gravity and bring your apps to the data centre or not?

A hybrid cloud structure could help. For instance, if you work in a regulated sector and handle sensitive data, you could host some of your data on-prem while the bulk of your data is in the data centre. In this example, you wouldn't suffer problems of latency (and potential non-portability) on some key workloads because your data is in-house.

However, this may not be feasible in a world where remote working is an everyday norm. Your workforce may need access to data from anywhere. It's up to you to identify where and how you should store your data so your company can succeed. Speaking to a data strategist from companies like Digital Realty could help.

It's also worth noting that edge computing is transforming the way the world does business. Data processing is increasingly being done closer to IoT devices, with some simple processing being done on the devices themselves and heavier processing offloaded to a local/edge data centre. This allows for lower latency than a centralised model, but isn't restricted to on-device processing. Nonetheless, this aggregated use of data and transformation is only increasing the amount of data further. Heightened connectivity is required to reduce latency and the effects of data gravity. Later, we'll talk about a connected community approach that might prove the right way forwards.

Identify Bottlenecks in Your Data Centre Strategy

The continued success of your business relies on your ability to utilise data effectively and seamlessly. Therefore, you need to identify where current limitations in your data centre strategy exist and address them.

These bottlenecks could be to do with having an unsuitable location for your data, making it harder to access data quickly and increasing latency. Likewise, it could be because of a hard limit on connectivity, so you find it difficult to operate even though your data centre is in a relatively close proximity to your company and apps.

But a bottleneck could also arise if you have a skills shortage. In February 2021, [ZDNet reported](#) that APAC would see severe shortages in data, cloud and cybersecurity skills by 2025 unless countermeasures are introduced. Combatting data gravity's negative effects means investing in the necessary expertise to overcome them. It's important to remember that not every problem with data gravity arises because of infrastructure – people impact how effectively your data is performing, too.

Finally, if your business is currently trying to meet sustainability goals then your data centre strategy must factor into your efforts. Data centres consume a lot of electricity, and therefore contribute to your company's carbon footprint. There's a growing movement within the data

centre industry for greater sustainability, with renewable energy now powering many data centres, so you have options here. Our previous white paper, [The Rise of Sustainable Data Centres](#), has more on this. You can also hear about it in a recent episode of the Data Centre podcast, [The Importance of the Sustainability Imperative for Data Centres](#).



A Connected Community Approach

Another way in which data leaders can offset the impact of data gravity is to champion a connected community approach between enterprises, connectivity, cloud and content providers.

Key to this approach is the idea that core, cloud and edge centres of data exchange become integrated, drawing on a hybrid IT architecture that is data-centric and secure.

Given that edge products will power the next wave of digital transformation, there's a greater need for a hybrid IT approach.

Some resources will have to remain in-house, but many apps will have to migrate to the cloud in order to ensure data and processing can be delivered and occur seamlessly for end-users. This also helps to reduce cost while fuelling expansion, as no company alone can expect to build the necessary data infrastructure required for life in the 21st century entirely in-house.

Think: how can you partner with your cloud and data centre providers to help bring your services to end-users and improve workload performance? Taking a connected community approach is the only way.



The Rise of Hybrid IT

All of the above points indicate a shift towards one thing: hybrid IT. Although it can present its own challenges, in terms of combatting data gravity, hybrid IT could have more benefits than drawbacks. It would allow enterprises to forego issues of latency on critical business functions while having the ability to harness public cloud processing power as and when needed.

Likewise, a hybrid IT approach is secure by design and reduces cloud migration issues. Nevertheless, the approach could sometimes slow agility, but if the approach means delivering a service that has faster workload performance and negates issues of portability, it must be considered.

If you're worried about data gravity and how it will impact your company, [head to the Digital Centre](#) now. Through our partnership with Digital Realty, we're raising awareness of data gravity and discussing solutions for mitigating its negative effects. You can even speak to a Digital Realty consultant today to restructure your data strategy and address the issue. [Get in touch](#) with us to arrange your consultation!





Unlock More APAC Insights with the Digital Centre

As data centres continue to evolve to meet growing demand,
we want to enable you and support you to continue delivering
your business as usual.

**You can book your complimentary digital centre consultation
today to find out more.**

**BOOK YOUR
FREE CONSULTATION**

**GET IN TOUCH WITH
DIGITAL REALTY**