


A look at data over the next decade

 By [Eran Brown](#)

24 Jan 2020

Over the next decade, we are going to see another wave of huge volumes of data created, driven by emerging technologies such as Artificial Intelligence (AI), the Internet of Things (IoT), Machine Learning and data lakes.



Eran Brown, CTO for EMEA at Infinidat

It's difficult to predict which of these will take the lead in terms of data proliferation but all these technologies rely heavily on data to drive business value. If we take IoT, for example, it produces massive data sets due to the ever-increasing number of IoT devices that create this data.

Gartner predicted that 14.2 billion connected things will be in use in 2019, and this will reach 25 billion by 2021, producing immense volumes of data. Moreover, we will use increasingly modern tools for pattern analysis but because we don't really know exactly what we are looking for, this data is not likely to ever get deleted.

Optimising the cost component

As we analyse more data and run more applications, performance becomes a key consideration, even more so than capacity.

Traditionally, we have grown storage media capacity much faster than its performance. Since data is not stored only for the sake of storage (which requires capacity) but rather needs to be analysed (which requires performance), that gap in their growth rate means customers keep paying more for performance, and the capacity problem is secondary.

Simply throwing the most expensive media at the problem is not a financially smart solution, because while we are optimising for performance, there is nothing that's offsetting the cost. The result? The overall cost of storing the data required to drive the business grows exponentially, which is what's been happening in the storage industry for years.

The more data we have, the greater the pressure to optimise the cost, which means that we need to drive higher performance requirements without inflating cost – this is definitely going to be the theme for the next five years. We must optimise the cost component, or the business model of these data-rich services starts to break down.

New consumption models

One of the biggest advancements that I expect to see in the storage industry is the adoption of new and more modern consumption models. Traditionally, customers would buy the amount of storage capacity they thought they needed for the next year or so, and slowly add as they grow. This was not only costly but procurement cycles would be between two and five months long causing delays in new application launches.

The hottest topic in my discussions with customers is how to bring cloud-like consumption models on-premises. This means vendors deliver on-premise agility, yet the enterprise is only required to pay for what they use. These “pay-as-you-grow” models will ensure that private clouds always have the capacity they need to onboard new applications in the private cloud, without having to pay for unused capacity in advance, resulting in higher business agility.

This is particularly important as tier 2 service providers need to streamline their operations in an increasingly competitive landscape with their larger cloud competitors. This new competition requires them to lower infrastructure costs (and storage prices are a big component of the overall cost) without sacrificing performance.

Moving from backup-centric to a recovery-centric approach

While backup speeds are still not what customers would have liked them to be, they are enough to reliably run their businesses. However, the combination of data growth and new types of data corruption such as Ransomware compel customers to refocus their requirements on the recovery of data and the time it takes them to resume their normal operations. This is also driven by the increasing cost of downtime to any business that conducts a lot of its transactions digitally. According to Gartner a minute of downtime currently costs over \$5k.

After giving up on tapes, customers relied on disk-based backup targets to achieve better speeds but those were always designed for backup speeds, as that's what customers would test in a Proof of Concept (POC). Frustrated with the low recovery speed, some customers started adopting Flash-based backup solutions.

Bear in mind that the backup data is already compressed and deduplicated, that means they are paying for the full cost of Flash without any data reduction leading to a surge in secondary storage costs.

With backup capacities often close to or surpassing primary storage capacity, this strategy is sure to consume the IT

budget and cripple any innovation. In 2020, we will see customers asking for disk-based solutions with better recovery speeds, to meet their Service Level Agreements (SLAs) to the business units

ABOUT ERAN BROWN

CTO for EMEA at Infinidat

- How agile data infrastructures can save the retail sector - 31 Jul 2020
- How can IT help beat Covid-19? - 4 May 2020
- A look at data over the next decade - 24 Jan 2020
- Elastic infrastructure is key to meeting current and future data storage demand - 12 Dec 2019
- Is AI the key to managing rampant data storage costs? - 23 May 2019

[View my profile and articles...](#)

For more, visit: <https://www.bizcommunity.com>