


# The Future of Storage:

Four reasons why open source storage should be part of your strategy





**The comfortable world of storage appliance vendors of ten years ago is under assault from a wave of technical and market changes that are rewriting how storage and data protection processes are managed and governed. Enterprises need to plan now to deal with those changes — or face a future where data is locked into vendors’ platforms and products, hard or expensive to release, and business agility is threatened by a lack of data mobility.**

The competition between storage vendors – in the cloud, on premise and across both is fierce, and technology is changing radically: in this white paper, SUSE argues that enterprises need to maintain internal skill sets to take advantage of market opportunity, develop cost savings, and avoid the threat of vendor lock down – with the accompanying loss of agility and cost disadvantage that this can entail.

## Exponential data growth: the challenge that just won’t go away

Data growth is not like economic growth or growth on savings in a bank – a sluggish few percentage points grudgingly added over long periods – it’s incredibly rapid, long into double figures, and the interest rate is compound. In their attempts to describe this growth, media outlets and analysts alike pile on the superlatives and adjectives, with many favouring the term ‘exponential’ a word previously confined to advanced maths students and lecturers, and even ‘explosive’ – bringing to the mind the rate at which a fireball expands when a bomb is detonated – with the added connotations of damage.

Alongside the charts and the forecasts and the long technical words are the anecdotes – the stories used to illustrate growth – here’s a few choice examples:

- Decoding the human genome was one of the biggest science stories of a generation. It took ten years: on today’s tech it could be done in two weeks.
- In 2010, the economist published a story about Walmart’s systems handling a staggering 75 million transactions an hour, feeding into databases estimated to hold a total of 2.5 petabytes. This year in 2017, Forbes cites them as processing 2.5 petabytes an hour.
- When the Sloan Digital Sky Survey began in 2000, more data was collected in the first few weeks than had been previously been gathered in the entire history of astronomy. When the large synoptic Survey Telescope goes online in Mexico 2022, it will start with 16 petabytes of storage—scaling immediately—and the volume of data collected will be so large as to be impossible for humans to analyse it – this is AI or bust on a system with 100 terraflops of computing power.

Benjamin Franklin is widely credited with saying ‘nothing in this world can be said to be certain, with the exception of death and taxes’. The typical storage pro added data growth to that list a long time ago. And we haven’t even mentioned the IoT.

Storage costs might have come down a lot over the last few years, but those savings are more than negated by the volume growth rate, creating a major issues for many, as reduced capital spending has become the norm – as Fortune reports, CapEx spending is on the way down – and not just for IT. ‘Do more with less’ says the CFO to the CIO, and the CIO to everyone else. We’ve reached the limits of ‘if it ain’t broke don’t fix it’. Everyone knows things have to change. But how?

## How the hyperscalers broke the mould

Given a free hand, it's tempting to think that most appliance manufacturers would have continued exactly as they always had done: continuing to sell expensive proprietary hardware operated by expensive proprietary software, with a built-in limitation on interoperability – frustrating the movement of data from one platform to another by design. Coupled of course, with a pricing model that encouraged the customer to 'bet the farm': the vendor offers customers who give them the majority of the estate aggressive discount – locking them down, and the competition out with the same strategy. Hence the expression of a Dell or HP or IBM 'shop'. Good business for big boys, not so good for the customer, who can end up with a limited choice of product features, and may find long term they have backed the wrong horse.

It was all good till the hyperscalers broke the mould: the economics of massive cloud provision doesn't allow for the margins of appliance manufacturers. Wired reported that Google's software consists of something approaching 2bn lines of code: if they were paying Microsoft for the server operating system licenses and HP for the storage space we'd be seeing a very large difference on those three firms' comparative profitability, and we'd all sell our Google shares.

Faced with this unsustainable bill for proprietary hardware and software, Google (grudgingly!) teamed up with the likes of Facebook in the Open Compute Project, and outsourced all server provision to unheard of companies in Asia on its own designs: thus the hyperscalers drove unprecedented use of the 'white box' and made it the norm. How else could Cloud providers compete with highly virtualised on-premise infrastructure? The cost saving have to come from somewhere. The scale of data centres got so large it became possible to rent them out – giving us the Cloud.

Today we have the world of hybrid cloud, where pretty much all major enterprises have a mixed environment: some hardware on premise, some in the cloud, some software run completely by others as SaaS, increasingly complex APIs connecting it all up; speed, and scale at the touch of a button, but, difficult to control and architect – and here we go again as the competing vendors try to get us all to 'bet the farm' as our tech Goliaths swing their weight around – with the occasional David thrown into the mix.



## Software defined—the place where the storage vendors staked their futures

So, if massive data volume growth and the example of the hyperscale reaction to it has caused a rethink in enterprises' approach to storage, where now lies the future, and how are vendors reacting? SUSE would argue, as would practically every analyst in the world, that the answer lies in software defined.

The simple truth is that storage at scale on proprietary hardware is unaffordable. The reaction of storage vendors is to confine hardware sales to the high performance end of the market – massive throughput, low latency storage for workloads literally coupled to the servers running the compute in the same stack – frequently hyperconverged; all else is about the software. 'We are now software companies', say the storage hardware vendors, 'work with us, and you will need only one software layer for your entire infrastructure – you know that change that took over compute with the dawn of virtualisation? well we are bringing that power to your storage'. Pooled resources, increased utilisation, cost savings. This is inevitably going to cause consolidation: there will be winners and losers in this environment.

And it doesn't take a genius to work out there's a problem with this approach from the enterprise's perspective: it only solves half the cost problem because you are still paying for expensive proprietary software. Moving to software defined storage with any vendor will save you some money, but not anything like as much as you could with open source.

## How enterprises hedge their bets in an uncertain world

Lock down is a perennial fear for enterprises: nobody wants to find having bet the farm that they are at the mercy of a vendor who knows the customer cannot operate without them. So how do you get your storage strategy right? Storage strategy has to follow the main enterprise strategy – and for the following reasons, open source should be part of your strategy.

### 1.

#### **'Betting the farm' is dangerous. Maintain multiple storage vendors**

It can be tempting to take hold of a vendor's short and medium term pricing strategy by placing all of your business with them, generating operational simplicity in the process – one set of storage tools and processes does make things easier. It may look particularly attractive in the Cloud. But if you go 'all in' you are playing poker with your storage budget and gambling that your vendor partners will not punish you with price hikes later.

### 2.

#### **Pay attention to the Cloud war—it hasn't been won**

Amazon unquestionably currently has the lead in adoption over Microsoft Azure and Google Compute. Nevertheless, everyone knows that Amazon is playing a 'long game' of profit tomorrow, not today. Hence, many have a foot in Azure or Google Compute, even when they have a leg in AWS, because there must be an exit plan. But this comes with a price – and that price is operational complexity – a price that can be particularly high in the world of storage, where the new pricing models can be about how much data you move down the wire rather than how much you own.

### 3.

#### **Maintain and expand your skill sets to avoid lock down**

It's tempting to reduce complexity by standardising on a small set of suppliers. The upside is that you get simplicity – one approach to storage means it's easy to train staff, some are no longer necessary in a cloud scenario, and arguably, you can get on with that 'core business' proprietary vendors love to tell you about of serving customers. However, if you don't know how to exit AWS to move to Azure without crippling operations, if you don't know how much it costs to repatriate data, and if you've got nowhere to put it when you do, you are locked down. See point #1.

### 4.

#### **Use open source software defined storage—or pay more**

If you use only cloud or only proprietary software, your software and hardware costs will always be greater than they need to be. This is a simple fact – open source means costs savings from moving to commodity hardware, and the total elimination of proprietary software costs. Proprietary storage vendors will tell you—rightly—that cost can reappear as skilled headcount, consultancy and support. But then, if you don't have skilled headcount, how are you going to maintain your capability to switch cloud providers, and how are you going to assess which vendors to use? Well, given that the obvious answer is hire expensive consultants, we'd argue that the proprietary sales argument is somewhat duplicitous.



**For more information,  
contact your local SUSE  
Solutions Provider, visit us  
online or call SUSE at:**

1-800-796-3700 (U.S. and Canada)

1-801-861-4500 (Worldwide)

SUSE

Maxfeldstrasse 5

90409 Nuremberg

Germany

**[www.suse.com](http://www.suse.com)**