




Inside Track Research Note

In association with



Hyper-Scale Data Management

An open source-based approach
to Software Defined Storage

January 2015

About this Inside Track

The insights presented in this document are derived from independent research conducted by Freeform Dynamics. Inputs include in-depth discussions on the latest technology developments with IT vendors and service providers, along with intelligence gathered from mainstream enterprises during broader market studies.

Coming up with a single definition of SDS is hard. It's therefore better to focus on the shapes of solution that seem to be emerging.

In a nutshell

The term 'Software Defined Storage' is used to describe a wide range of ideas and offerings. This makes it hard to pinpoint specific opportunities. Fully supported enterprise class solutions based on open source technology, however, are driving practical and economic benefits in the context of today's storage challenges.

Software defined everything

If you work in IT, the term 'software defined' will not have escaped you as one of the latest ideas to be promoted in the drive for greater efficiency and flexibility. It has surfaced across the industry in a number of different ways.

Server virtualisation led to private cloud architecture, which was in turn extended to the 'Software Defined Datacentre'. Some of the latest ideas in decoupling key elements of communications technology were then branded 'Software Defined Networking', and not to be left out, the data management community is now talking about 'Software Defined Storage', or 'SDS' for short.

The common thread that runs through all of this is the idea of moving control and management functionality from the hardware tier into an independent software layer. The way in which the infrastructure is configured and run is then no longer reliant on the many and varied proprietary embedded capabilities that exist in most data centres today. A more consistent set of software tools is used instead.

From theory to practice

In theory, everything from policy definition, through resource provisioning and configuration, to ongoing optimisation, monitoring and administration, can be done centrally in a more joined up and flexible manner once your infrastructure is software defined. The challenge is that pretty much every vendor selling storage tooling or middleware has taken this as a licence to reposition their offerings as SDS. Put this together with the term also being used in the context of scale-out systems to handle huge data volumes, and the result is a lot of confusion over what SDS actually is.

Against this background, coming up with a single definition of SDS is hard. It's therefore better to focus on the shapes of solution that seem to be emerging:

Mainstream SDS: This is about abstracting functionality such as thin-provisioning, compression, deduplication, replication, snapshotting and backup/recovery into generic software that can operate in a heterogeneous hardware environment.

Storage Virtualisation: Some regard software that pools and virtualises capacity across devices as SDS, so we have included it here. However, this approach is best thought of as a complement to SDS. It is possible to implement Mainstream SDS with or without virtualising.

Hyper-scale SDS: The focus here is on high-performance distributed processing software that can be used to create a massively scalable, automated and resilient environment. The aim is to meet escalating storage needs in a flexible, unconstrained and cost-effective manner.