

Software-Defined Storage

SNIA Emerald Metric Analysis Working Group
ENERGY STAR® Discussion January 19, 2017

Agenda

- Definitions
- Operations Fundamentals
- Implications for ENERGY STAR

SDS - Definitions

- Software-defined storage (SDS) is an approach to data storage in which the programming that controls storage-related tasks is decoupled from the physical storage. *Computer Weekly*
- Abstraction of logical storage services and capabilities from the underlying physical storage systems, and in some cases pooling across multiple different implementations. Since data movement is relatively expensive and slow compared to computation and services (the "data gravity" problem in [infonomics](#)), pooling approaches sometimes suggest leaving it in place and creating a mapping layer to it that spans arrays. *Wikipedia*
- The aspect of SDS that differentiates it is apparent from how some of the products are deployed. It is that Data Services can be executed either in servers, storage, or both spanning the historical boundaries of where they execute. This has potential impacts on security and reliability, and may be an interesting revival for Direct Attached Storage (DAS) in some cases. *SNIA*

In other words:

- SDS products are **applications** designed to be agnostic regarding
 - the brand of hardware platform
 - Potentially to the operating system running on those servers.
- They rely on the same sorts of installation guidelines as any other application when it comes to a suitable run-time environment,
 - listing requirements such as instruction set, operating system, memory required, disk space needed for installation
- The customer is free to install this, along with the corresponding disk drives that will be used by the SDS Software, on any server, including ones running other applications.
- The primary point of using SDS products is to
 - Avoid the purchase of a dedicated appliance
 - Allow the storage application to run on the same virtualized server as the applications that will be using the SDS.

Implications – Part 1

- SDS breaks the link between sophisticated data center storage functionality and a particular brand or configuration of hardware.
- Data center storage energy efficiency is determined by two independent variables, the behavior of the software and the behavior of the hardware.
 - Each specific storage appliance combines a specific controller design, dedicated i/o infrastructure, and a specific storage application. This fixed combination provides repeatable results.
 - SDS product performance is a function of the storage application, the specific server configuration (including the memory, i/o controllers, etc. needed for all applications) and the other applications in use on the server. Even if the apps are quiesced, the hardware configuration will still reflect these additional needs. Running the same SDS on two different "storage servers" will yield different results.
- This non-appliance approach does not meet the expectations of ENERGY STAR testing approaches

Implications – Part 2

- Servers running SDS applications will require unique configurations.
- Minimum focus areas are
 - Additional I/O channels for added disk traffic
 - Appropriate quantities of high-performance I/O devices for data transfer to/from the server
 - Appropriate quantities and types of storage media
- Software-defined storage products are outside the scope of ENERGY STAR, as the unique product being sold is a piece of software rather than an appliance.