How do you know if you are achieving your data storage energy efficiency targets?

The continued rapid rise of data creation will provide immeasurable opportunities for companies to gain advantage in the market, but they must take proactive actions in order to be a step ahead of others. Business leaders need to increase their focus on the computing trends driving data growth over the next several years and revisit policies to assess the value of data throughout its lifecycle from creation, collection, utilization to its long-term management.

By Wayne M. Adams, SNIA Board of Directors and Chair of the SNIA Green Storage Initiative

International Data Corporation (IDC)1 has predicted that data creation will continue to grow year over year to a total of 163 zettabytes (ZB) by 2025. IDC also states that the industry will transition from a trend of consumers being the largest creators of the world's data, to where enterprises will become the larger creator again, creating 60% of the world's data in 2025. Computing trends like IoT, machine learning, and other types of Artificial Intelligence (AI) based data analysis/decision making and many others are behind these data growth trends

Within IT though the era of greenwashing has long passed, there remains the need to be energy efficient and to further optimize limited resources as the ongoing top priority to have increased computational resources and pools of data are required to drive a business. IT users continue to look for effective approaches to select technologies and products. For data center storage, there is a collection of standard energy efficiency metrics that enable IT decision makers to objectively compare a range of possible solutions and to manage a solution once deployed.

The SNIA Emerald™ Program provides a standardized way of reporting vendor-performed test results that characterize the several aspects of storage system energy usage and efficiency. For procurement metrics, SNIA's Emerald™ program and specification defines energy usage metrics for Block IO and the recently released File IO metrics that provides an energy usage profile on how a storage system will work in configurations optimized for transaction performance capacity or streaming.

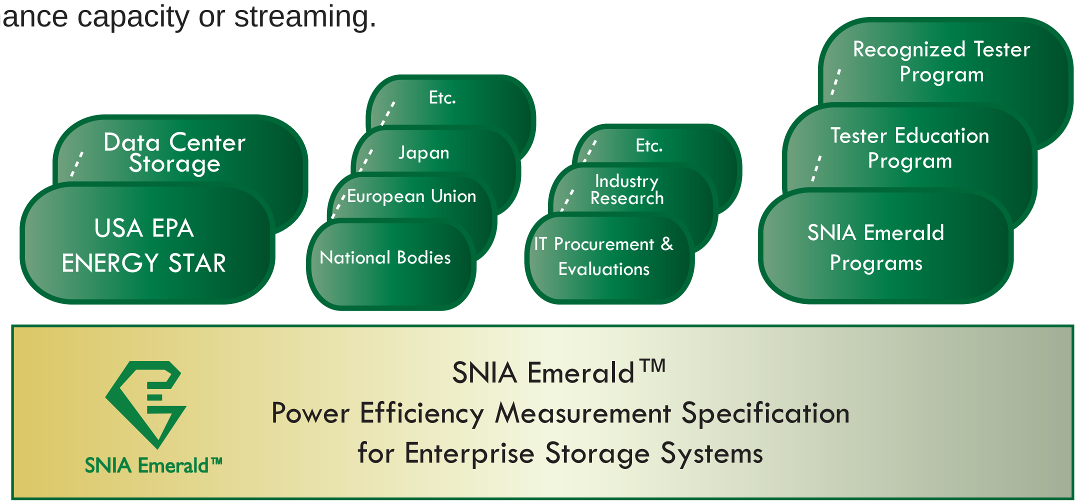

Procurement Metrics

The USA EPA Energy Star® Data Center Storage Program references the use of SNIA's Emerald™ Energy Efficiency Measurement Specification Metrics for Block IO storage system configurations. The EPA maintains a public repository of vendor tested products since 2014, where many many vendors are listed with a range of products.

The SNIA Emerald Block IO metrics are:

- Hot band (IO/second/Watt) as it relates to caching efficiency
- Random Read (IO/s/W); Random Write (IO/s/W)
- Sequential Read (Mebi-Bytes (MiB)/second/Watt)
- Sequential Write (MiB/s/W)
- Ready Idle (Giga-Bytes(GB)/Watt)

With the release of SNIA Emerald V3 specification in September 2017, the File IO metrics are based on the following application workloads:

- Video Data Acquisition (MiB/s/W)
- Database (MiB/s/W)
- Virtual Desktop Integration (MiB/s/W)
- Software Build (MiB/s/W)
- Ready Idle (GB/W)

RAS and Capacity Optimization

Additional considerations during the procurement phase to select a solution for your IT requirements includes a systems Reliability, Availability, Serviceability (RAS) features, capacity optimization features, and type of physical media being selected. All of these factor into a system's energy efficiency profile, so when contrasting solutions, keep these in mind.

The more RAS features will increase additional controller functionality and or systems to be running additional logic, which can add to the system energy usage. Capacity optimization technologies enable a system to store a data set size in a smaller physical storage size, which can reduce energy usage. Disk storage types from Hard Disk Drive (HDD) to Solid State Drives (SSD) have different energy usage profiles. Within HDD, there are rotational speeds and data placement considerations.

The SNIA Emerald specification recommends that these attributes be part of a system test report so the reader understands why there can be variations in metrics when looking at two systems configured with the same base hardware. RAS and capacity optimization technologies, each feature by themselves may be uneventful, but when combined there can be positive, additive improvements for reduced energy usage. Capacity optimization refers to a set of techniques which collectively reduce the amount of storage necessary to meet storage objectives. Reduced use of storage (or increased utilization of raw storage) will result in less energy usage for a given task or objective. Each of these techniques is known as a Capacity Optimizing Method (COM).

COMs are largely, though not completely, independent. They provide benefit in any combination, though their combined effect does not precisely equal the sum of their individual impacts. Nonetheless, since data sets vary greatly, a hybrid approach using as many techniques as possible is more likely to minimize the capacity requirements of any given data set, and therefore is also likely to achieve the best results over the universe of data sets. In addition, the space savings achievable through the different COMs are sufficiently close to one another that they are roughly equivalent in storage capacity impact.

Data Storage Space Consuming Common Practices

A commonl assumption is that certain space consuming practices are essential to the storage of data at a data center class service level.

- Disk-based redundancy: when one or more drives (or other storage devices) fail—the number depending on service level—no interruption in service or loss of data occurs
- Sufficient space: when an application shall have enough space provisioned for it, so that no downtime shall be incurred during normal operation
- Point-In-Time (PIT) copy availability: data center applications under test need access to PIT copies of data sets that they can manipulate without fear of interference with live data
- Disaster recovery: when data corruption or loss does occur, good copies of the data must be available to restore service

Capacity Optimization Methods (COMs) Characterized

In the SNIA Emerald™ Energy Measurement Specification, tests for the presence of the following COMs are defined as:

- Delta Snapshots: applicable to backup, PIT copy availability and disaster recovery. Both read-only and writeable delta snapshots are featured in shipping systems, but there are fundamental technical differences between them, and some systems implement only the read-only version;
- Thin Provisioning: a technology that allocates the physical capacity as applications write data, rather than pre-allocating all the physical capacity at the time of provisioning;
- Data Deduplication: addresses issues caused by multiple backups of the same data sets, and the tendency of large data sets, due to human usage patterns, to contain many copies of the same data
- Compression: takes advantage of the inherent compressibility of many data sets.
- The following COMs and additional COMs may exist, however no test for presence is defined in the Specification now.
- Parity RAID: addresses the need for disk-based redundancy.

Typical and common RAS Features in Enterprise Storage

- Dual Controller (No Single Point of Failure (SPOF) Controller)
- Mirroring (Local or remote, synchronous or asynchronous)
- RAID 1, 4/5, 6
- Snapshots (Full or Delta)
- Disk Scrubbing
- Multi-pathing
- Disk Sparing
- Drive-level Maintenance
- Dual Power Supply
- Variable-speed Fans

Operational Metrics

SNIA and the Green Grid organization collaborated on a whitepaper titled, "The Green Grid Data Center Storage Productivity Metrics: Application of Storage System Productivity Operational Metrics".

DCsP represents a set of operational phase metrics that observe storage system productivity while the data center runs normal or "real-world" workloads. These metrics are conceptually the same as those defined for the acquisition phase including aspects of capacity and performance. All are needed to completely characterize storage systems.

Although similar to the procurement metrics, these DCsP operational metrics differ in their measurement and usage aspects. The majority of "real-world" workloads represent actual data center information produced by at least one or more applications. Most of IT equipment, among other requirements, are required to run 24/7. This last-mentioned availability aspect is particularly important for storage systems as it makes the real time gathering of the operational metrics essential for good analysis.

The metrics to be calculated, based upon a storage systems operational information to be collected, polled, or stored in a Data Center Information Management (CIM) tool, based on the storage configuration for the applications it is supporting are as follows:

- Capacity - GB (or TB, PB…)/KWhr (KiloWatt Hour)
- Streaming - MB's/KWhr
- Transaction - IO's/KWhr

For more information on SNIA resources and programs for data storage energy efficiency, including the SNIA Emerald program, IT planning resources, and education materials, please visit https://www.snia.org/energy or email emerald@snia.org.

[1] *SDC reference: http://www.seagate.com/www-content/our-story/trends/files/Seagate-WP-DataAge2025-March-2017.pdf*