

INFINIDAT

STORING THE FUTURE

# A More Sustainable Approach to Enterprise Data Storage: Reducing Power Consumption with InfiniBox™

Executive Brief

# Table of Contents

---

- Abstract..... 3**
- Introduction..... 4**
- Storage Power Consumption ..... 5**
- Designed for Sustainability: INFINIDAT Storage Architecture..... 6**
  - Hardware..... 6
  - Software..... 7
    - InfiniRAID..... 7
    - What Zeros?..... 7
    - Provisioning..... 7
    - Compression..... 8
- InfiniBox Power Consumption by the Numbers..... 8**
  - By Comparison..... 9
- Conclusion..... 10**

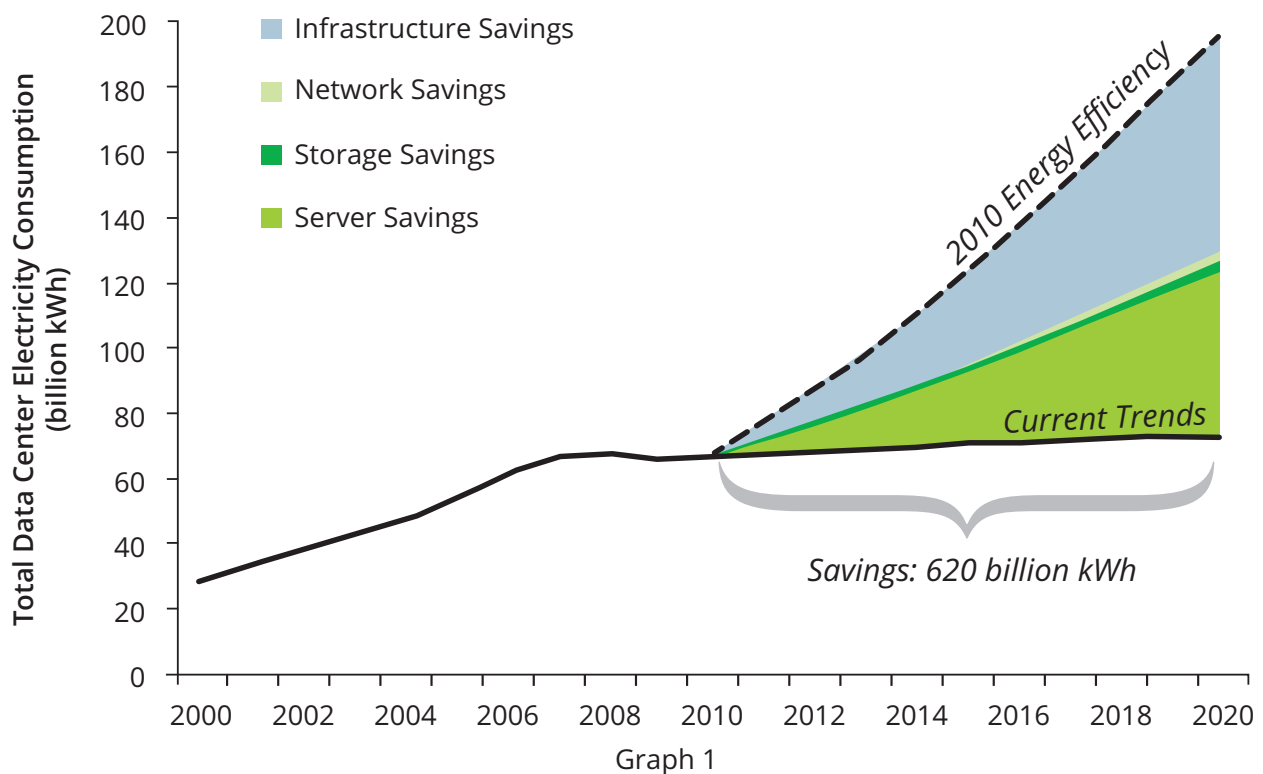
# Abstract

---

Data continues to grow at an astounding pace. As a result, data center space is becoming more scarce, as more arrays are acquired to store all of this data. Along with this data taking up space, it is also utilizing a great deal of power and cooling. In fact, the average data center in the U.S. uses approximately 34,000 kW of electricity each year, costing \$180,000 in annual energy costs. As INFINIDAT set out to revolutionize the storage industry, one of our goals was to help consumers of storage build a more sustainable infrastructure that would be not only better for the environment, but also help them to save money as well. All of our patents come together to form InfiniBox™, a storage solution that does just this.

# Introduction

Data has become the lifeblood of the economy. Big Data, analytics, cyber security and machines themselves are churning out vast amounts of data like never before. A byproduct of this explosion of data has been the unfortunate uptick in power consumption / energy use by data centers which house the storage and compute infrastructures. According to the most recent study, U.S. data centers consumed as much as 2% of the country's total energy consumption, or about 70 billion kW hours.<sup>1</sup> Data center power is predicted to increase to 140 billion kilowatt-hours annually by 2020 in the U.S. alone, which is the equivalent annual output of 50 power plants.<sup>2</sup> While these are only U.S. numbers, this phenomenon is a worldwide problem, even for developing countries. Every data center around the world is feeling the “power crunch,” either through growing costs or the physical, inability to obtain more power.



*This chart shows past and projected growth rate of total U.S. data center energy use from 2000 until 2020. It also illustrates how much faster data center energy use would grow if the industry, hypothetically, did not make any further efficiency improvements after 2010. (Source: U.S. Department of Energy, Lawrence Berkeley National Laboratory)*

<sup>1</sup>[https://betterbuildingsinitiative.energy.gov/sites/default/files/attachments/Better\\_Buildings\\_Data\\_Center\\_Overview-FAQ\\_0.pdf](https://betterbuildingsinitiative.energy.gov/sites/default/files/attachments/Better_Buildings_Data_Center_Overview-FAQ_0.pdf)  
<sup>2</sup><http://www.2datacenterknowledge.com/archives/2016/06/27/heres-how-much-energy-all-us-data-centers-consume/>

In 2016, the U.S. Department of Energy issued a Better Buildings Accelerator challenge<sup>3</sup> specifically to data centers with the goal of reducing power consumption by 25% over five years. If all U.S. data centers reduced their power requirements by just 20%, it would save over \$2B per year. With this goal in mind, what kind of role can storage play in helping to reduce power consumption in the data center?

## Storage Power Consumption

The average data center consumes 34,000 kW of electricity per year, which equates to approximately \$180,000 in annual energy costs. It is estimated that as much as 40% of this power is consumed by storage.<sup>4</sup> Storage power consumption is increasing due, in large part, to two factors. First, data is doubling in size every two years,<sup>5</sup> requiring more capacity and more power-hungry arrays. Second, the increase in server virtualization has influenced a 3% drop in annual server shipment growth since 2010, driving up storage needs and power consumption along with it. Researchers expect this trend to continue through at least 2020.<sup>6</sup>

In addition, cooling all this storage capacity is playing a major role in the overall data center energy costs for storage. As arrays work harder, they run hotter, and they need to be kept cool in order to maintain the proper working of all the components in the array. If a small error with a disk or network function were to occur due to overheating, and the wrong data is served to an application, it can take down your entire operation.

In order to lower power consumption in the data center, storage arrays must achieve higher energy efficiency. According to Energy Star, which published 12 ways to save energy in the data center, better management of data storage<sup>7</sup> can play a significant role. Best practices for increasing storage energy efficiency include:

- Automated Storage Provisioning
- Data Compression
- Deduplication
- Snapshots
- Thin-provisioning
- RAID Levels
- Storage Tiering

<sup>3</sup><https://betterbuildingsinitiative.energy.gov/accelerators/data-centers>

<sup>4</sup>[https://www.storageio.com/DownloadItems/CMG/MSP\\_CMG\\_May03\\_2007.pdf](https://www.storageio.com/DownloadItems/CMG/MSP_CMG_May03_2007.pdf)

<sup>5</sup><http://www.emc.com/leadership/digital-universe/2014iview/executive-summary.htm>

<sup>6</sup><http://www.datacenterknowledge.com/archives/2016/06/27/heres-how-much-energy-all-us-data-centers-consume/>

<sup>7</sup>[https://www.energystar.gov/products/low\\_carbon\\_it\\_campaign/12\\_ways\\_save\\_energy\\_data\\_center/better\\_management\\_data\\_storage](https://www.energystar.gov/products/low_carbon_it_campaign/12_ways_save_energy_data_center/better_management_data_storage)



The storage hardware that is chosen also plays a large part in power consumption. Using higher density, lower speed drives, SSDs, and an array that can spin drives down, all play a part in lowering storage and data center power consumption.

## **Designed for Sustainability: INFINIDAT Storage Architecture**

INFINIDAT understands that sustainable technologies are essential — not just for achieving costs savings, but for preserving the environment in which we live. INFINIDAT offers a much greener and therefore more sustainable solution for storing data with the highest levels of performance, availability and scalability that enterprises require. This is made possible by the revolutionary architecture of the INFINIDAT InfiniBox storage array.

Building high-performance storage arrays has been made easier by the advent of NAND devices. Unlike many competing solutions, INFINIDAT was able to innovate beyond expensive media band-aids and created a new, more sustainable approach. This approach makes use of off-the-shelf components and increases the ratio of low power components to high power components (e.g., disk:controller ratios). The result is a solution that is far better for the environment, and is able to deliver the performance of an All-Flash array, with a lower overall TCO (CapEx and OpEx).

InfiniBox boasts an incredibly low power consumption for our flagship F6000 storage array platform. InfiniBox F6000 can store multiple petabytes of data in a single 42U rack while utilizing only 8kW of power. This is over 12 times more efficient than other storage vendors storing the same amount. In fact, INFINIDAT's design incorporates many of the hardware and software best practices outlined by Energy Star resulting in incredibly low power consumption for a large amount of capacity. Aiding this is the utilization of high capacity, low speed, near-line SAS drives, as well as SSDs. In addition to following many of the Energy Star best practices, INFINIDAT's storage software architecture incorporates an extensive patent portfolio that has enabled the InfiniBox to be incredibly efficient.

## **Hardware**

The INFINIDAT storage software architecture is designed and built to be hardware agnostic and is pre-integrated with off-the-shelf hardware to deliver the most reliable storage array on the market. There is no element about the design of the software that ties it to any particular hardware platform. InfiniBox uses multi-terabyte NL-SAS drives to store a copy of the data. The most active data is read and written to DRAM or SSDs; both utilize low power and provide exceptional performance.

Additionally, the ratio of servers (controllers) to disk (3:480) means less power is consumed while managing more drives. Compare this to a number of new “hyperconverged” systems where each time more capacity is required, a new server (controller) with as few as 14 drives (1:14) needs to be deployed, further increasing power consumption.

Finally, InfiniBox does not power each drive with a battery backup unit (BBU), instead it only powers the servers with BBU's and saves power conversion losses.

## Software

There are a number of the key storage services that contribute to power efficiency that have been redesigned and patented by INFINIDAT:

### InfiniRAID™

The reliability of a storage system is largely a function of the RAID it uses. The issue is, RAID solutions can utilize as much as 50% of an array's overall capacity for data protection. Other RAID levels may use less capacity, but have an impact on system performance. The question is, how to provide high reliability without impacting performance and increasing the amount of data that can be put in a smaller footprint. One of INFINIDAT's patents is InfiniRAID which enables InfiniBox to have 72% storage utilization while providing dual-parity like reliability and RAID 1-like performance.<sup>8</sup> INFINIDAT accomplishes this without the need for hot spare disks. InfiniRAID provides 12% higher storage efficiency than most other arrays on the market.

### What Zeros?

Another INFINIDAT patent is the ability, when storing data to not store the zeros that are a part of the data. There are a few arrays on the market today that go back and spend an extra cycle to clean up the zeros, but again this takes power to accomplish this task. InfiniBox does this as data is written, so no extra cleanup or processing is required. Also, by not writing zeros, this lends itself to InfiniBox having an even higher areal density which allow users to store more data in a smaller footprint.

### Provisioning

InfiniBox runs with all disks spinning, as with all arrays, however, due to our provisioning and intelligent data layout, InfiniBox removes the power peaks and valleys in reading data as each drive is prepared to serve up data quickly and efficiently. The data layout is also sequential such that it enables us to prefetch data more intelligently and move it to SSD or DRAM for reading, rather than always going to spinning disk, again reducing power consumption.

<sup>8</sup>Industry standard RAID 6 is considered to provide high reliability, handling a 2 drive failure while RAID 1 is considered to provide the best performance.

## Compression

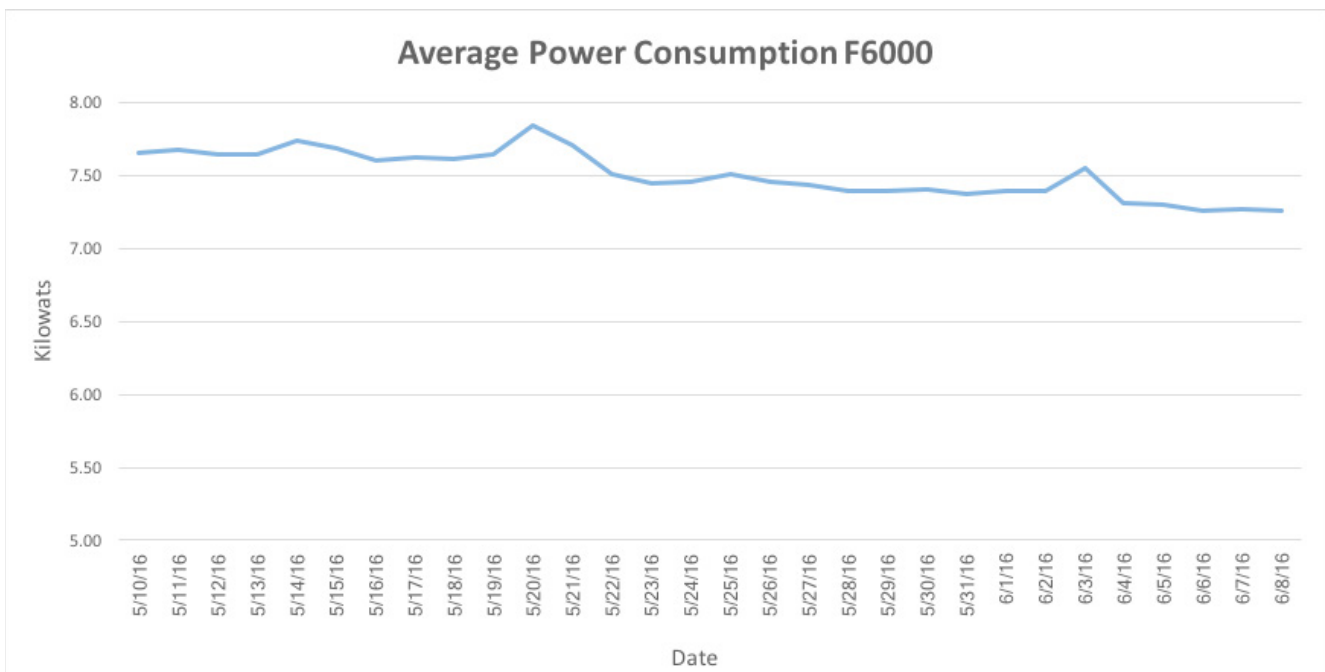
INFINIDAT has also implemented an intelligent compression platform in InfiniBox. While most systems apply one type of compression, InfiniBox being a data-driven platform, can identify different data types and pick from a library of compression algorithms that will compress the different data types better. A higher compression ratio means that more capacity is freed up allowing users to store more data in the already dense InfiniBox.

It is the combination of these technologies that enable the InfiniBox system to run much more efficiently and not require as much power. In addition, these technologies allow the InfiniBox to run at high performance, without building up a lot of heat, therefore needing less cooling than that of traditional arrays. InfiniBox's unique architecture allows clients to save on both power and cooling costs.

## InfiniBox Power Consumption by the Numbers

Through a combination of these software features, InfiniBox boasts an incredibly low, 3W per terabyte, of power consumption. Saying this is one thing, so we monitored our "smart" PDUs (power distribution units) on 120 of our flagship F6000 systems in the field running real customer workloads and collected the daily average power consumption.

The average F6000 power consumption of all systems monitored over the course of the 30 days is about 6.9kW. The graph below shows the power consumption of one of the systems monitored over the course of the month.



Graph 2: Reflects Average Power Consumption of One F6000 System



The systems monitored were F6000 systems of varying configurations and capacities. We looked deeper into our data to see if any configuration in particular changed the power consumption and found that the configuration was negligible. The workload had more to do with the power consumption than the configuration.

## By Comparison

INFINIDAT has also worked a real-world consolidation / replacement exercise against two well-known storage brands where they require seven 23" cabinets to provide 800TB of useable capacity, which is the equivalent of nine floor tiles in a climate controlled data center. These seven racks have two engine bays and five drive bays totalling 9.3kWh. Using \$.07/kWh for both power and cooling, the total energy costs to power 800TB of capacity comes out to \$42,066 per year. INFINIDAT storage is much more efficient. In one 19" cabinet on one floor tile, the InfiniBox can provide 2PB (or 2000TB) of capacity for just under 8kWh. Assuming the same \$.07/kWh figure, the annual power costs for INFINIDAT are \$9,811 per year and INFINIDAT is providing 150% more capacity, while consolidating these five units into one.<sup>9</sup>

Vendor Name	Storage Tier	Bin File Changes/ Month	Engine Bay Power	# Engine Power	Drive Bay Power	# Drive Bays	# of Frames	Total KW Power	Total Power (Cost/Yr)	Carbon Footprint (lbs/Yr)	Floor Tiles Used	DC Real Estate Cost/ 1 Sqft Tile	EOL Migration Per Host	Average # Host Per Frame	Disk Drive Retention Per Disk
EMC	VMAX 1	\$1,000	3.9	2	5.3	5	1	34.30	\$42,066	367171.90	3.00	\$4,000	\$750	300	\$2,500
EMC	VMAX 2	\$1,000	3.9	2	5.3	4	1	29.00	\$35,566	310436.88	3.00	\$4,000	\$750	300	\$2,500
HDS	VSP (001)	\$0	26.7	1	0.0	1	1	26.72	\$32,769	286030.12	7.47	\$1,200	\$750	250	\$2,500
HDS	VMAX (002)	\$0	26.7	1	0.0	1	1	26.72	\$32,769	286030.12	7.47	\$1,200	\$750	250	\$2,500
INFINIDAT	F6140	0	8.0	0	0.0	0	2	16.00	\$19,622	171275.52	2.00	\$4,000	\$0	200	\$450

<sup>9</sup>[http://content.infinidat.com/power\\_consumption](http://content.infinidat.com/power_consumption)

<sup>10</sup>These costs do not include cooling

## Conclusion

INFINIDAT has several innovative patents that not only help InfiniBox to provide high performance, superior reliability, and high density, these innovations also work together to help reduce power consumption. Power consumption is not only low per system, InfiniBox's density enables clients to consolidate, in one case, from 22 systems down to two InfiniBox arrays. This helped the client save not only data center footprint, but also on power consumption for the arrays and cooling systems as well. Utilizing an average cost of \$.0603 per kilowatt hour, a competitive system would cost as much as \$1.70 / hour to run versus the InfiniBox which would cost \$.48 / hour to run. At approximately 3PB, a competitive system would cost \$14,892 / year to run, where as the InfiniBox would cost \$4226.<sup>10</sup> The InfiniBox is on average 3.5 times more efficient than competitive systems providing similar capacity, and when consolidating even more arrays we have found InfiniBox to be 10 to 12 times more efficient.