

EMC Symmetrix DMX-3 and DMX-4 Electrical Power Estimation and Configuration Planning

A Detailed Review

Abstract

This white paper provides a detailed description of the EMC[®] Symmetrix[®] DMX-3 and DMX-4 power systems design. It explains the various factors that contribute to power consumption and efficiency, with information on how to configure the storage array for the best efficiency.

June 2008

Copyright © 2006, 2007, 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H2399.2

Table of Contents

Executive summary	4
Introduction	4
Audience.....	4
Configuration planning	5
Configuring for power efficiency	5
Critical components in design	5
Disk drives	5
Drive density and configuration	6
Ultra-performing Flash drives	7
Tiered storage in large arrays	9
Virtual Provisioning	11
Configuration summary.....	12
Power density and cooling	13
Power Calculator tool: Getting the exact information	14
Conclusion	14
References	14

Executive summary

As the need to conserve energy becomes more important in the data center, having the information needed to acquire and configure efficient storage systems becomes important as well. The typical spec sheets provided for facilities planning assumes a maximum load with a fully configured system, and with a high margin for safety under all possible conditions. While this is appropriate for wiring and safety, it is not always helpful in predicting the cost of actual energy usage, nor in providing the proper capacity for cooling, data center layout, and other environmental factors.

EMC has taken great care to provide efficient storage systems, holding down the costs of operation. This white paper complements the facilities planning data by providing the information necessary to understand the various components in the array, how they contribute to power consumption, and how to configure Symmetrix[®] DMX-3 and DMX-4 systems for the highest efficiency. It also introduces tools available to get detailed information about actual usage.

Introduction

Energy consumption is rapidly growing in importance in IT, energy prices have risen over the past few years, and at the same time the amount of information to be stored continues to expand. Customers are demanding that the products that they buy tomorrow are designed with power efficiency in mind, and that they have access to the information that they need to make informed choices.

Energy management can be complex. Consumption varies with workload, availability requirements, time of year and other variables that make understanding the issues difficult. Power produces heat, which requires energy to cool, and all this must usually be managed for power failure as well.

Power usage and cost continue to grow.

A recent IDC survey¹ indicated that users expect power consumption to increase by 25 percent over the next two years. Since the amount of data to be stored is increasing by 56 percent per year, the challenge is to do more and store more while keeping power increases to much less than that.

There are a number of concerns related to power efficiency. The most obvious is cost. The cost of electricity is increasing steadily, and there is more at stake than just the energy used by the storage. For each dollar spent on electricity to power the storage, it typically takes another dollar or more to get rid of the heat produced, which means that the cost can double in some situations. In some cases, data center redesign is required to manage the excess heat.

Another problem can be even more severe. In some locations, additional power is unavailable at any price; the local utility is simply unable to provide additional capacity.

Of course, storage is not generally the main culprit. As servers become more dense, the proportion of energy that they use can overshadow that used by storage, and there are certainly ways to consolidate servers with VMware and other conservation methods, but clearly, anything that can be done to reduce power consumption is quite important. This white paper will provide valuable information on how to do that.

Audience

This document is intended for technical customers and EMC personnel who are configuring DMX-3 and DMX-4 storage arrays. It provides an understanding of the various factors that contribute to energy consumption, and how those apply to EMC DMX-3 and DMX-4 products. It assumes some familiarity with DMX-3 and DMX-4 storage and electrical concepts.

A useful prelude to this document can be found in the EMC white paper *Power Efficiency and Storage Arrays*, which can be found at EMC.com.

¹ IDC, End User Perspectives on Server Blade Adoption, Doc # TB20060525, May 2006

Configuration planning

Many customers are constrained by the amount of power and cooling available in a building. In addition, their operating budgets aren't growing as fast as storage requirements, so the rising cost of electricity can be problematic. If there is not enough power or cooling available for a storage array, some older equipment may need to be removed and consolidated with the newer equipment. EMC offers some configuration options and methodologies that can free up power and cooling levels and improve the customer's power efficiency in the data center.

The logical place to start is in the data center with the server population. Servers are faster, and consume more power, while blade server design packs more of them in smaller spaces, putting pressure on everything else in the data center, especially storage. Consolidating servers to reduce their numbers, perhaps using a product like VMware, can be quite beneficial.

The next step is to use the EMC ControlCenter[®] tools to measure utilization, particularly for multiple arrays, and StorageScope[™] to highlight opportunities for consolidation. With larger, scalable DMX-3 and DMX-4 arrays available, a significant power savings can be realized by consolidating many older arrays into one. Using the DMX-3 and DMX-4 and taking older equipment offline, can significantly improve power efficiency by reducing watts per terabyte. The removed equipment and its associated power consumption can be applied as a savings to the newer equipment.

The application of the two software packages, VMware and ControlCenter, allows the customer to realize immediate power savings and free up a portion of the power and cooling budget for the new storage array.

In cases where the availability of power is restricted, you will now know how much power and cooling are available for storage. As you will see, there are a number of tradeoffs available between capacity, performance, and efficiency, and it is important to configure the array to meet the intended applications. Since the DMX-3 and DMX-4 can start with a small configuration and be increased nondisruptively over time to very large configurations, systems need not be overconfigured today to support future growth, which provides for the best efficiency today and in the future.

Configuring for power efficiency

The best case for optimum power efficiency is an in-the-box tiered storage solution, which the DMX-3 and DMX-4 are specifically designed to do. It is important to be familiar with the *EMC Symmetrix DMX-3 Best Practices Technical Note* to fully understand the options and recommended configurations available. It is also necessary to understand the applications and environment prior to configuring the system. The EMC Sales teams have access to the EMC Power Calculator, a tool that calculate the front-end and back-end workload, performance requirements, power consumption, and more.

Critical components in design

It is necessary to understand the functioning of three critical components in the DMX-3 and DMX-4 to produce the most efficient design. These are:

- The drives themselves, speed and capacity
- The arrangement of the drives in DAEs and storage bays
- The advantages of using a minimum number of large arrays rather than more smaller arrays

We'll explore each of these topics in detail before returning for some final recommendations

Disk drives

Disk drives are the largest users of power in storage arrays, and a good place to start. As you will see, there are many choices to be made, resulting in tradeoffs between performance and efficiency. Finding the right balance is important and can be quite difficult, but EMC provides the expertise, information, and tools that are needed to make these critical choices. The unparalleled flexibility and scalability of the DMX-3 and DMX-4 mean that you can provide for multiple environments in a single array, and easily change or grow as the need arises.

EMC works closely with disk drive suppliers to ensure that Symmetrix arrays use drives that are designed for power efficiency. This ensures that the drives can be cooled properly, and that future upgrades can be done with the same power and cooling budget. Since the drives contribute the highest percentage to the overall power consumption, this translates into ongoing power savings.

The most important factor in drive power consumption is the actual speed of the drive. The faster that it spins, the more energy is required to keep it spinning. Thus, there is an immediate tradeoff available between power and performance, making it important to select the proper speed drive for the application.

Surprisingly, the capacity of the drive has very little effect on power consumption. For high speed drives, 73 GB and 146 GB drives use about the same power, and the same is true for 146 GB and 300 GB drives spinning at 10k rpm. This means that it is possible to immediately lower power consumption by specifying larger drives. But doubling the capacity reduces the number of spindles and may reduce the overall performance. Again, careful analysis is required to strike the right balance

The I/O rate acting on the drives also has some significance on power as shown in Figure 1. This is not surprising; the more activity that a drive is subjected to, the more power it draws. The effect can be quite noticeable at high I/O rates and should be allowed for in planning and budgeting.




Figure 1. Power consumption increases as a function of I/O rate

Drive density and configuration

A second factor in efficiency is the arrangement of the drives within a storage bay. A storage bay can contain as many as 240 drives, which are configured in 16 disk-array enclosures, or DAEs. Each DAE can contain up to 15 drives, so 15 x 16 will get you 240 total drives. The DAE provides power and connectivity to the individual drives, and the DAE chassis requires a certain amount of power itself, which does not increase with the number of drives, and is still required even when it contains no drives at all.

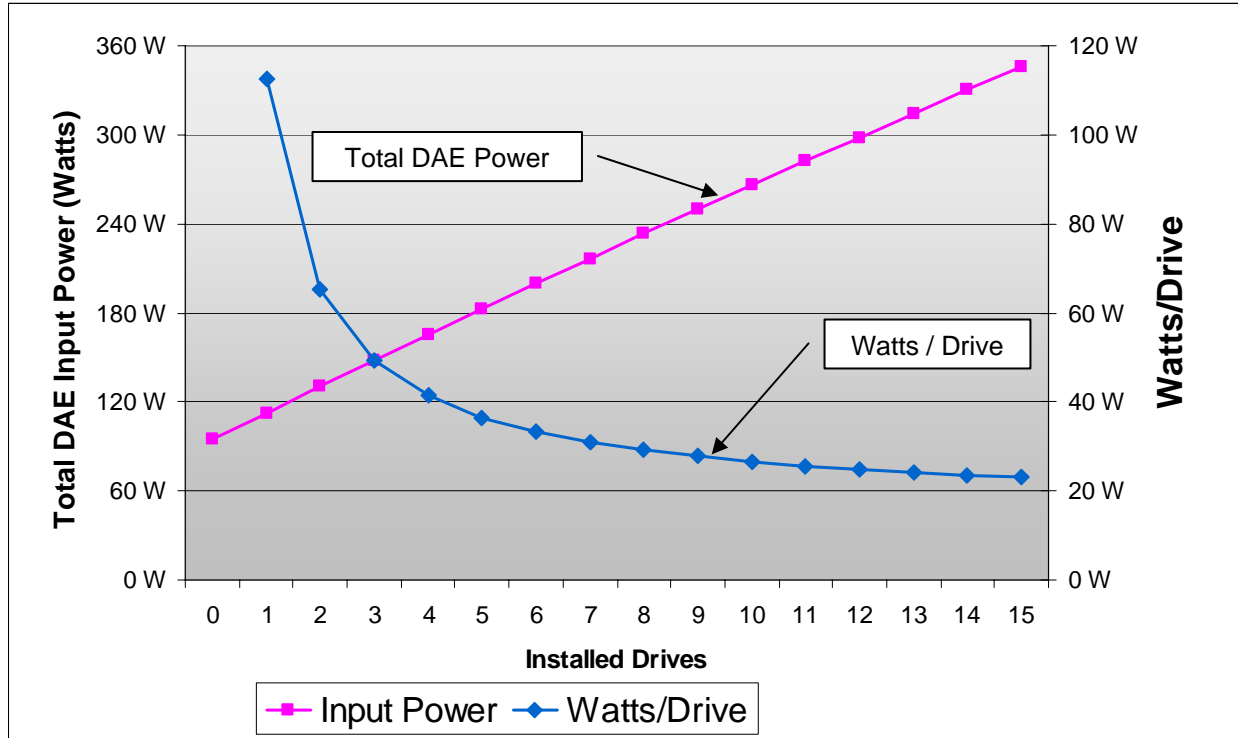


Figure 2. Storage bay input power versus number of installed drives using recommended drive distribution to optimize availability

Figure 2 illustrates this. The empty DAE with no drives is shown on the left axis, and you can see that it draws nearly 100W when empty. Since each individual drive adds about 17W, the average watts/drive is reduced significantly as the DAE fills up, dropping from 117W/drive to about 23W/drive with 15 drives included.

Ultra-performing Flash drives

Flash drives use less energy than rotating media, and weigh less as well. When compared to a 73 GB rotating drive, as can be seen in Table 1, the Flash drive uses 30 percent to 40 percent less energy. As the cost of these drives is reduced over time, this savings will become increasingly important.

The real energy savings, though, becomes evident when comparing the two drives based on comparable IOPS. Performance testing indicates that a single Flash drive can deliver the same IOPS as 30 of the 15k rpm drives. When this comparison is made, using Flash drives for performance can cut energy use by as much as 97.7 percent. If 120 rotating drives can be replaced by four of the Flash drives, the savings can be significant, as shown in Figure 3.

Existing DMX-4 and DMX-4 950 systems can be upgraded with Flash drives after upgrading to Engenuity 5773.

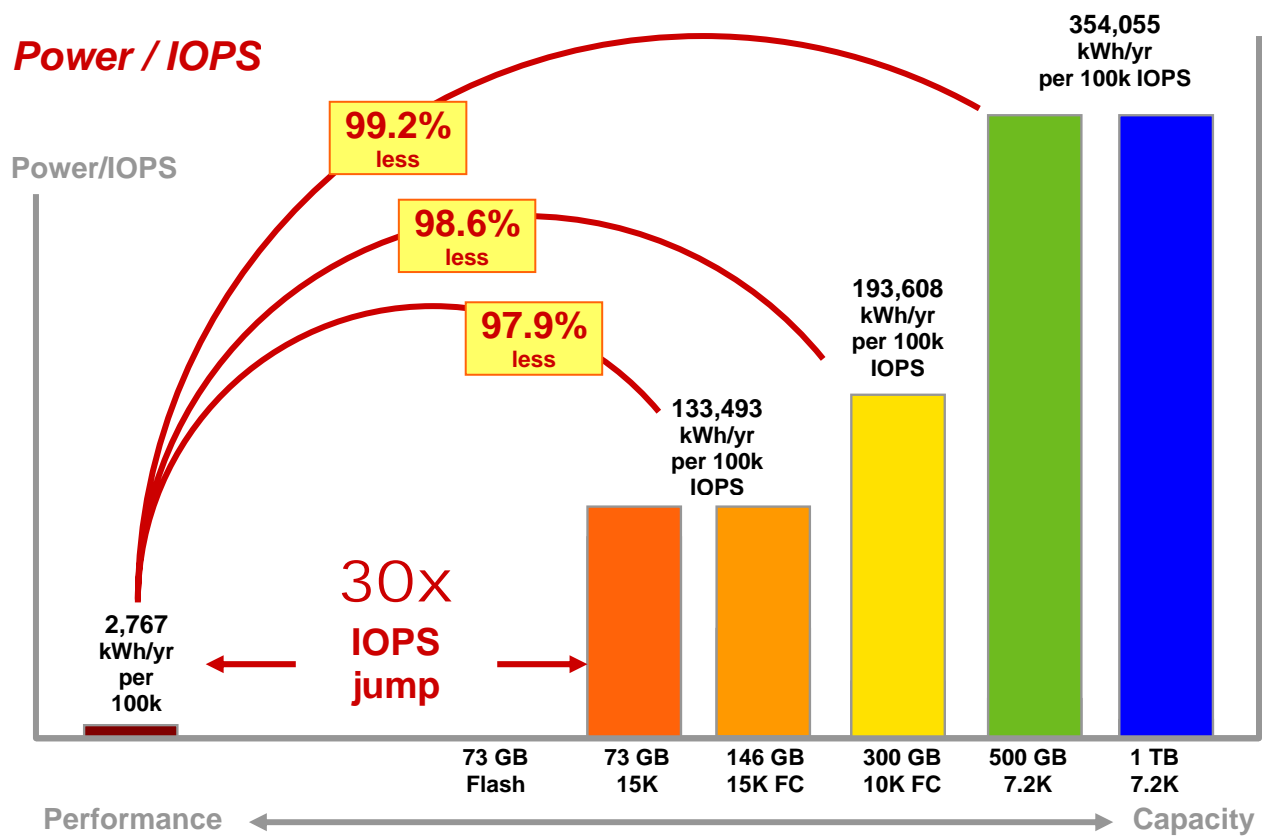


Figure 3. On a per-IOPS comparison, Flash drives require 98% less energy

And, of course, the savings in weight and floor space and even noise levels can be important as well. Table 1 compares individual power consumption and weight of Flash drives and HDDs.¹

Table 1. Power and weight comparison of 73 GB Flash drive and 73 GB 15k rpm HDD

Drive type	73 GB Flash drive	73 GB 15k rpm HDD
AC drive power in a system	15.80W	25.40W
Weight	0.781 lb	1.85 lb

¹ Power numbers in the table may vary depending on exact configuration. Consult your EMC representative for specific power usage derived from the EMC Power Calculator.

Tiered storage in large arrays

The third key to storage efficiency is consolidation. Instead of a larger number of smaller arrays, a smaller number of large, fully utilized arrays can mean a significant improvement in efficiency. Of course, this places a couple of critical requirements on storage.

Both the DMX-3 and DMX-4 allow you to grow capacity without suffering performance degradation. As storage increases, performance for critical applications must increase in a completely linear manner. One of the principal advantages of the DMX-3 and DMX-4 systems is the ability to configure multiple storage tiers in the same array. This can be done in a way that prevents the various tiers from affecting one another so that applications receive the performance that they need.

The in-the-box tiered storage solution essentially tailors the drive capacity to performance requirements for the intended applications. In this case, the drives are distributed so that certain DA pair processors have slower (high capacity) drives and others have faster (high performance) drives. There are four tiers frequently used to describe the performance requirements and criticality of the data, with Tier 0 being highest. Flash drives create a new “ultra performance” tier within a Symmetrix for applications such as algorithmic trading, Internet transactions, or foreign exchange. This “ultra” tier benefits from the full array of advanced capabilities that Symmetrix provides, including local and remote replication, cache partitioning, and priority controls.

The DMX design provides up to four pairs of disk adapters, or DAs, and each of these can support up to eight storage ports. A single (pair of) port(s) can manage drives in multiple bays cabled together, which are referred to as loops. Adding a large number of drives to a loop may affect performance.

Now, using Flash drives, Tier 0 applications can be closely coupled with other storage tiers within Symmetrix for consistency and efficiency, eliminating the need for time invested in manual data layout or end-of-day data transfers from separate RAM disk or specialized memory storage systems.

Tier 1 would include high performance applications, such as Exchange or Oracle, which are best suited for the 15-drive loops with the smaller, faster 73 or 146 GB 15k drives. Tier 2 applications include data warehousing and various noncritical database applications. These can be serviced with high capacity, medium speed, enterprise drives such as 300 GB 10k drives. Tier 3 is typically for data archiving and lower performance applications and uses the high capacity, low speed 500 or 1000 GB 7.2k drives with longer drive loops. Longer loops are accomplished via daisy chaining to subsequent storage bays. Configuring in this manner provides the highest performance with the best predictability for critical applications and a more cost-effective, lower-performance tier for less critical applications. Some considerations are:

- Small loops of high performance drives use the most power but offer the highest possible performance.
- Larger capacity drives conserve power, but fewer spindles may have performance implications in some situations.
- Extending loops permits more drives to share the DA pair, and thus spread its power over a larger number of drives, but may have performance implications.

Figure 4 illustrates the efficiency advantages of tiering quite clearly. It provides a direct comparison of power consumption per terabyte among drives of various sizes, and the power required at various I/O rates. Since the DMX-3 and the DMX-4 have the ability to mix multiple drive types in a single array, users have the ability to tailor the array to provide all the capacity and performance required with minimum power consumption.

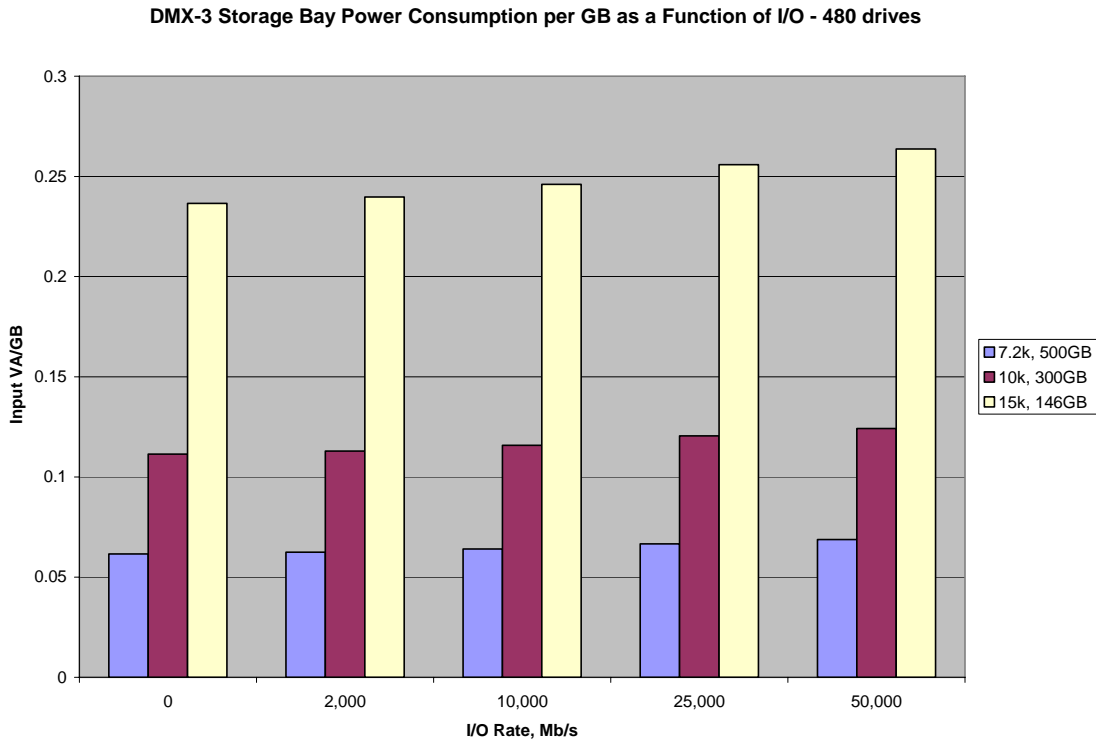


Figure 4. Power consumption per GB of hard disk drives as a function of I/O rate

As pointed out previously, for power efficiency it is also important that all slots in the DAEs be populated to fully utilize power overhead for cooling, logic, and power supply efficiency.

Figure 5 shows the scaling concept and use of in-the-box tiered storage. The color coding shows the different disk adapter (DA) pairs that “direct connect” to DAEs in adjacent storage bays and then they in turn “daisy chain” to DAEs in outer storage bays. Each DA pair supports 120 to 480 physical drives via eight Fibre Channel loops. Each loop contains from 15 to 60+ drives. Starting with the yellow DA pair shown in Figure 5, the adjacent cabinet can be direct connected with fibre cable. Subsequent storage bays can be daisy chain connected as shown in the lighter color. Depending on the applications and capacity and performance needed, the number of drives, type of drives, and number of DA pairs can be varied. Additional drives and performance can be realized by an additional DA pair, shown in green in Figure 5. It’s best to scale the drives one way and then add DA pairs to scale on the other side.




Figure 5. Scaling concept using different DA pairs

For instance, the high performance Tier 1 15k drives could be the direct connect dark green color code. The light green could be the lower performance 10k drives. If the loops are to remain small for high performance and avoid

contention, the light green drives might not be populated at all. For Tier 2 or 3 data, the 7.2k or 10k drives would be populated all the way to the outer bays.

For Tier 0, Flash drives are only supported in direct connect storage bays. It is recommended to dedicate one quadrant to Flash drives with a maximum of 32 Flash drives per quadrant. All members of a RAID 5 group should be configured on homogeneous Flash drives.

Virtual Provisioning

EMC Virtual Provisioning builds on the base “thin provisioning” functionality, which is the ability to have a large “thin” device (that is, volume) configured and presented to the host while consuming physical storage from a shared pool only as needed. Symmetrix Virtual Provisioning can improve storage capacity utilization and simplify storage management by presenting the application with sufficient capacity for an extended period of time, reducing the need to provision new storage frequently and avoiding costly allocated but unused storage. Energy is conserved using Virtual Provisioning by allocating less storage and therefore reducing hardware requirements.




Figure 6. Symmetrix Virtual Provisioning

Symmetrix Virtual Provisioning will operate on any new or existing Symmetrix DMX-3 or DMX-4 system and will integrate easily into organizations’ existing knowledge and management processes for the Symmetrix platform. Configuring, replicating, managing, and monitoring thin devices and thin pools involves the same tools and the same or very similar functions as those used for years with standard provisioning. Symmetrix Virtual Provisioning will operate on typical Fibre Channel or new Flash drives, and thin pools can coexist with standard storage in the same system.

Tiered storage

Symmetrix Virtual Provisioning adds a new dimension to tiered storage in the array, most notably (when used appropriately) by improving the utilization of storage tiers that can tolerate performance variability. Thin devices and storage pools can be used in a tiered storage configuration in the same manner as standard devices. Thin devices will adhere to priorities assigned to them by Symmetrix Priority Controls. Data devices within a thin pool, meanwhile, can be assigned to cache partitions created by Dynamic Cache Partitioning. In addition, different RAID groups can be assigned to different thin pools, in accordance with their designated tier.

Performance

One of the direct consequences of Virtual Provisioning is the wide striping that occurs. Data volumes can be striped across a large number of drives, which can balance the workload and improve performance. In some cases, this

improvement can allow you to use larger, more efficient drives, thus saving energy for the same reasons shown in the tiering discussion.

Configuration summary

The DMX-3 and the DMX-4 are the most scalable, and most flexible, storage arrays available today, and are the logical choice for situations where managing and minimizing power is important. In addition, EMC has tools that will provide detailed information about specific configurations, making it simple to plan and budget.

Following the guidelines below will result in the best combination of performance and efficiency:

- Minimize DA pairs required.
- Use more daisy chain storage bays to obtain needed capacity with less DA pairs.
- Fully load drive enclosures with drives (15) to reduce excess power overhead from cooling, logic, and power supply load efficiency.
- Fully populate your DA pairs before adding additional pairs.
- Order storage bays in increments of half to fully utilize enclosures.
- Use larger capacity drives to reduce spindles.
- Use tiered storage to reduce the number of higher speed drives when requirements allow.
- Use Flash drives for maximum performance with latency sensitive applications.
- Use RAID 5 or other as opposed to RAID 1 full mirroring.
- Add incremental storage bays and DA pairs as demand changes.
- Use shorter loops for high performance drives, longer loops for lower performance drives.
- Use Virtual Provisioning to allocate less storage and therefore reduce hardware requirements.

Power density and cooling

The power consumption, kVA (kilovolts x amperes), of fully configured DMX-3 and DMX-4 machines with 10k rpm drives is used to determine the heat dissipation, BTU/hr (British thermal units per hour). It is assumed that all the input power of the system is converted into heat energy and no mechanical work is being done. This provides a worse-case number for facilities designers in sizing the HVAC system. The results are tabulated in Table 2. The typical power of the fully configured system bay is 5.8 kVA and is factored into the total of each storage bay value.

Table 2. Symmetrix DMX power consumption and heat dissipation

DMX-3 and DMX-4								
	1 Storage bay	2 Storage bay	3 Storage bay	4 Storage bay	5 Storage bay	6 Storage bay	7 Storage bay	8 Storage bay
Drive count, max	240	480	720	960	1200	1440	1680	1920
Power consumption (kVA)	12.5	18.6	27.7	30.8	36.9	43.0	49.1	55.2
Heat dissipation (BTU/hr)	41,380	61,180	80,980	100,780	120,580	140,380	160,180	179,980

The power consumed to remove the waste heat from the building can typically range from 40 percent to 60 percent of the input power to the cabinet. This also includes the 40 percent to 50 percent of the waste heat in the data center from the voltage transformers, UPS, and electrical distribution. The DMX-3 and DMX-4 have a patent pending cooling design that allows better flow of hot exhaust air to help improve efficiency of the HVAC cooling system. A simplified picture is shown in Figure 7. The hot exhaust air is forced to the ceiling, away from the front intake of the machine and toward the air return duct. This helps keep a higher air temperature entering the chiller and refrigeration unit, which improves its cooling efficiency. Less air needs to be moved around the room and refrigerated, humidified, and returned. It also reduces hot spots around the equipment that can permit for an overall increase in data center room temperatures, which can further reduce power consumption. The DMX-3 and DMX-4 cabinets also have the flexibility to be placed in a hot-aisle/cold-aisle or traditional bottom-to-top cooling environment.

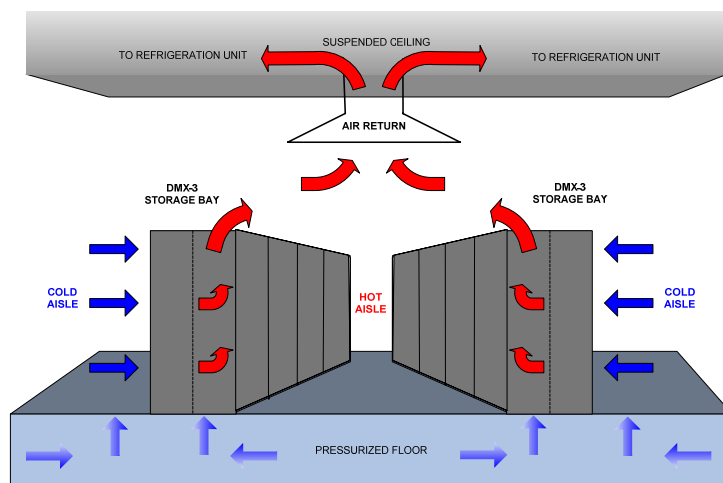


Figure 7. DMX-4 air flow diagram in data center

Power Calculator tool: Getting the exact information

The figures in this white paper are for “typical systems,” and may be useful for general planning, but your configuration probably isn’t exactly like the examples. EMC has developed an effective tool that can be used to project the power consumption of various subsystems in a DMX-3 or DMX-4 storage system. The EMC Power Calculator is a flexible tool that estimates power consumption of a specific configuration. This tool enables users to derive an estimate of AC and DC power of various EMC storage system configurations. The power estimate ranges from typical to maximum and includes a recharge power value that represents the time after the battery discharges and then recharges.

The purpose of the tool is to provide a means for customers to estimate power and cooling requirements, and then be able to interactively vary the configurations to meet specific power requirements. This can be combined with the configuration guidelines listed above to allow customers to meet their power and cooling targets. Contact your local EMC Sales Representative to obtain additional information on the EMC Power Calculator.

Conclusion

Controlling energy consumption in a data center environment is becoming more important with every passing day. As the need for processing power grows, along with the need for more storage, so does power consumption and the concentration of that consumption into smaller and smaller areas.

The DMX-3 and DMX-4 products were designed with efficiency as a high priority and offer many configuration options to balance energy usage with performance requirements and capacity. These options make it possible to deliver the capacity, performance, and features required with the lowest possible energy usage.

EMC provides the information and tools needed to take advantage of this flexibility and configure the most efficient storage arrays available.

References

The following documents provide additional sources of Symmetrix information. They are available on EMC.com and Powerlink, EMC’s customer- and partner-only extranet, except where noted.

- *Power Efficiency and Storage Arrays* white paper
- *EMC Symmetrix DMX-3 Best Practices Technical Note* (Powerlink only)
- *EMC Symmetrix Virtual Provisioning* white paper
- *EMC Symmetrix DMX-4 Ultra-Performance Tier 0 Using Flash Drives* white paper