ICONZ | Hosting a world of opportunities

Whitepaper

The disk classes defined

By Sean McDonald –ICONZ CEO

Table of Contents

1. Introduction 2
- 1.1 Choosing the right computing solution 2
- 1.2 When is ICONZ VERSA the answer and when is it not? 2

2. Disk performance factors 3
- 2.1 Offline Storage 4
- 2.2 Nearline storage 5
- 2.3 Online storage 5

3. Typical workloads of applications 6

4. Disk types 6
- 4.1 PATA (IDE) drives 6
- 4.2 SATA (IDE) drives 6
- 4.3 FATA drives 7
- 4.4 SAS drives 7
- 4.5 Fibre Channel (FC) drives 7

5. The impact of 'RAID' on performance 8
- 5.1 RAID 0 9
- 5.2 RAID 1 9
- 5.3 RAID 0+1 10
- 5.4 RAID 10 (1+0) 10
- 5.5 RAID 5 11
- 5.6 Virtual RAID (VRAID) 12

6. Conclusion 12

1. Introduction

SATA, FATA, SAS, Fibre Channel....the choices for disk storage in servers are endless. With today's technology solution ó*du jour*ô (Cloud Computing) where performance and reliability is critical, the operating disk class and raid type configuration of your cloud hosting provider is crucial to the performance of your hosted application.

With the hype around ó*The Cloud*ô there is enough confusion around what this actually is, let alone the level of detail most important, the hardware environments and infrastructure powering these clouds.

The following paper is not so much original thought but information from a variety of sources that explain the various disk storage classes on offer from hosting providers that present themselves as ó*Cloud Providers*ô Like most things, there are ó*horses for courses*ôand the author does not attempt to belittle providers on the basis of their infrastructure offering. Rather, according to levels of performance and reliability required for business applications in ó*The Cloud*ô, there are considerations, not least of which being price, that should be considered.

1.1 Choosing the right computing solution

IT administrators are continuously looking to maximize the value of their storage resources and are constantly challenged with expanding capacity requirements. Moreover, these challenges become even greater with the need to provide timely backups and quick restoration as the data sets continue to grow. Additionally, IT managers are eager to acquire storage that is easily managed and cost effective.

Many applications do not access data frequently, for example, document imaging and static websites, but require relatively fast access to the data when the need arises. In the case of high transactional ecommerce sites, fast access is required to data sources 24 x 7 365 with high levels of availability ranging from 99.9% right through to 99.999% in some cases.

Technology being what it is today, there is no shortage of choice as to computing power (CPU, RAM and DISK). The question is therefore, what is the right fit for purpose and what of the many choices makes sense from a price, reliability and performance perspective?

What then does a business do with their server of choice once the decision has been made? Host the server internally, host it with data centre, or is there another option?

The hosting choices are many and varied. One of these hosting options is the ICONZ VERSA Cloud hosting service. So why and when should ICONZ VERSA be considered?

1.2 When is ICONZ VERSA the answer and when is it not?

ICONZ VERSA has been designed on true enterprise class infrastructure with redundancy, reliability and performance in mind. Managed by the ICONZ Managed Hosting Data Centre team, its place in the market is designed to deliver the highest levels of availability, and performance at a price point that makes practical sense.

That said, we are often asked by clients, ñWhy does the VERSA Cloud Service option appear more expensive compared to a simple server that I can buy off the shelf and then, host with ICONZ.
The answer to that question is complicated but in the same breath, easy to articulate. Simply put, if your business application is of the ó*never fail variety*ô then VERSA may well be the answer.

Performance and reliability will be of utmost importance and as a base point, this will be predetermined by the business application/ solution that you are looking to host.

Ultimately and conversely, the server compute (CPU, RAM) and storage (FATA, SATA, Fibre Channel, RAID etc...) metrics chosen for the business application/solution will predetermine the level of service or user experience received. Diagram A (following page) illustrates this metric and demonstrates the potential outcome of an imbalance in this formula when sizing server platforms for business applications/solutions.

The proposition offered by ICONZ VERSA is one of total flexibility, void of the constraints of old (price and technology ólock inô) and a platform that infinitely scales to meet the demands and requirements of business applications whilst also matching these against the budgetary demands of business today.

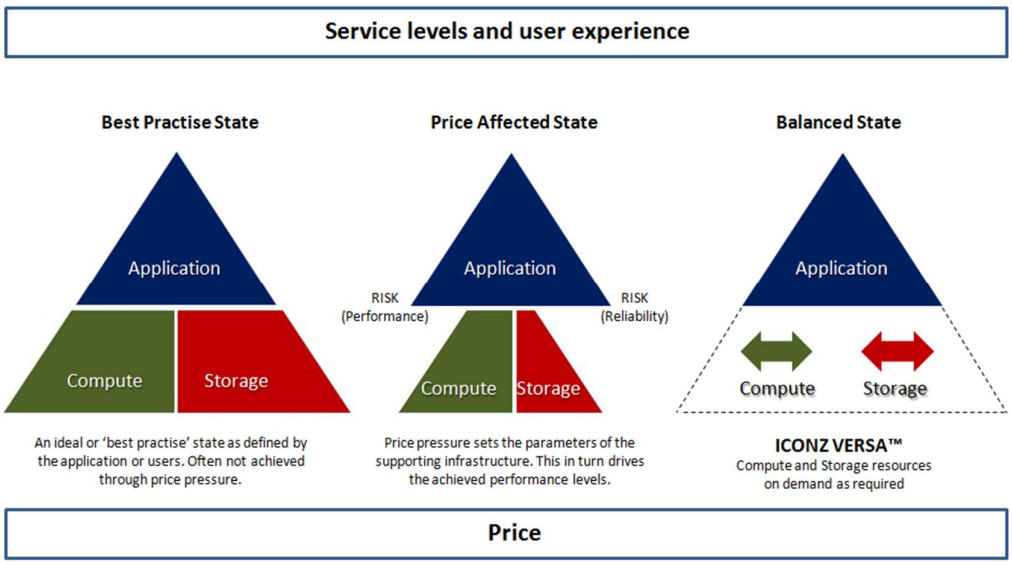

VERSA is a fit right across the hosting spectrum but when the shopping list demands lowest price and ósometimes performanceôthen VERSA may not be the right choice.
This paper is not intended as marketing and sales collateral, rather, a guide for executives and business owners when assessing the various server and hosting technologies that will care for their business application or website.

2. Disk performance factors

In the server world, disk is a critical factor in the performance and reliability stakes. One of the highest cost items when looking at server solutions is disk storage. For ICONZ VERSA, the same consideration is true. From the PC and Server from your local supplier, to the vendor mainframe and ICONZ VERSA, the decision process should always be the same, what are my requirements and therefore, which storage (disk) option in right for my business?
PATA, SATA, FATA, SAS, Fibre Channel....the choices are endless! What then is the right fit for purpose?

To start with, we need an explanation of what is what. Firstly, each disk class is defined by a set off ôuse parametersô known as Online, Offline and Nearline.
The table below defines these disk classes.

Storage class	Disk Type	Application Profile	Application types
Offline	PATA Desktop SATA	Departmental file/print/application serving; also small Web or e-mail servers. Economical single-user storage; least demanding storage environment	• Fixed content • Websites • Email • Desktop applications
Nearline	SATA II FATA	Storing infrequently accessed information, including backup/restore data and reference data, such as medical Records. Low-cost bulk storage; requires higher reliability, seamless integration	• File recovery • Archiving • Medium traffic websites • Data migration • Email arching • Medical imaging • File and print • Application testing • Low end video imaging
Online	SAS Fibre Channel	Transactional databases for companies of all sizes, contact management, ERP, e-commerce platforms other large-scale applications. Demanding storage environments that require the highest levels of performance and reliability	• All application types above • Databases • ERP / CRM • Financial applications • GPS processing • Geospatial applications • E-commerce • High definition video • Data mining

2.1 Offline Storage

Offline storage can be seen as the economical jack-of-all-trades in enterprise computing. Easy to configure and deploy, theyôre general-use disk solutions that often support a relatively small group or subset of users (for example, a specific department within a large corporation). In such environments, overall user counts are relatively small and the number of concurrent users modest. In some less demanding low-cost server environments, desktop-class SATA disc drives can be economical alternatives to high-performance SCSI drives, particularly when Native Command Queuing (NCQ) capabilities are incorporated. So equipped, moderate-capacity desktop drives can deliver respectable throughput, typically a higher priority than capacity in low-cost server applications.

2.2 Nearline storage

Nearline storage is the linchpin of tiered storage in midrange and high-end network environments, filling the gap between online and offline (archival) enterprise storage.

Nearline applications can be divided into two basic categories: data protection and reference data. Both entail data that doesnôt justify the cost of high-availability, mission-critical storage, but must still be readily accessible at any time.

Nearline applications work in very rigorous disk drive environments. In order to ensure uninterrupted data access for a broad variety of users, Nearline systems are always on. By contrast, desktop-class drives store data for only a single user, and are typically powered on only during normal business hours (for instance, eight hours a day, five days a week). Being an ñalways-onò repository for enterprise data significantly raises Nearline-ready drivesô reliability requirements and explains why desktop-class SATA drives are unsuitable for Nearline duty.

2.3 Online storage

Online storage application requirements can be neatly summed up with an age-old aphorism: Time is money. Online data-powering everything from highly critical transactional databases to local e-mail serversð is the very lifeblood of many enterprises, and online storage is all about keeping that data flowing. In the online world, diminished throughput and bouts with downtime are far more than inconvenient; they often translate into substantial lost revenues. Thatôs why the enterprise demands high-performance Fibre Channel and Serial Attached SCSI (SAS) disc drives.

Specifically engineered for the rigors of constant access by a multitude of concurrent users, these drives employ sophisticated (and expensive) technology to deliver maximum performance and reliability. In this unforgiving environment, SATA drives are simply out of place.

The diagram below) illustrates the various disk classes and some of the cross over that exists within the classes. Of these, ICONZ VERSA utilises Fibre Channel exclusively. We have chosen this disk platform over all others for the reasons of high availability, I/O performance and data integrity, no other platform (SATA, SAS, FATA) comes close. With customer demand though for a choice of disk performance levels, ICONZ will be adding a secondary platform offering to VERSA for the use of lower demand applications.

The diagram below illustrates the performance and price point comparisons against the ôDisk Classesô)

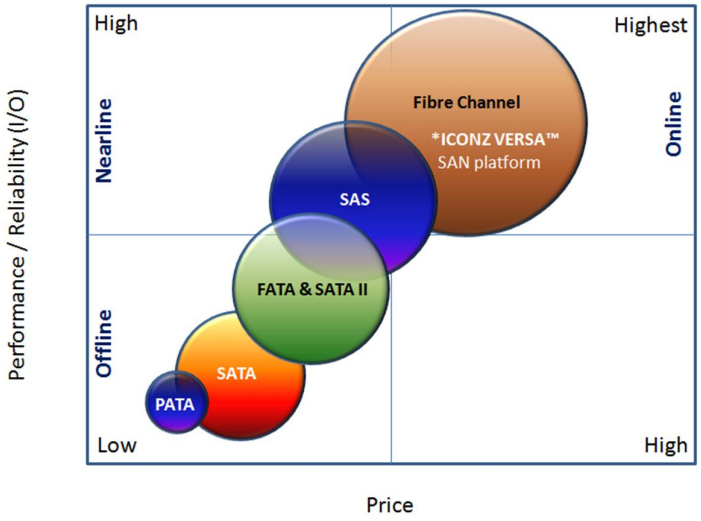

3. Typical workloads of applications

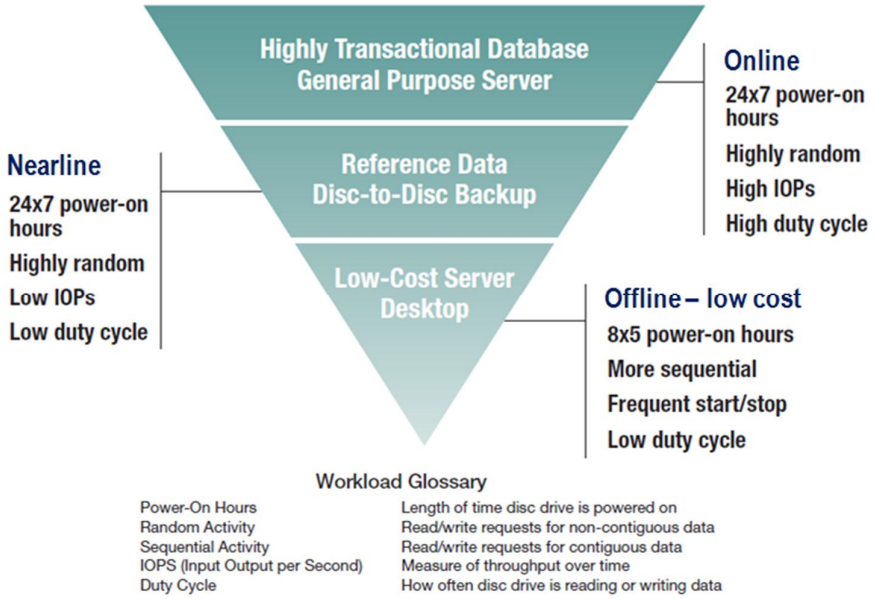

4. Disk types

4.1 PATA (IDE) drives

(Designed for desktop PC) – Offline storage

The PATA /IDE (Parallel Attached Technology Adapted) disk drive has been the storage interface for the desktop/PC disk drive for nearly 20 years. With its limited, maximum interface speed of 133 MB/s and large 40-connection parallel ribbon cable, the drive is seldom considered for new designs.

PATA Disk is designed for 8 ï 12 power on hours (POH) per day, single users and low workloads. It is designed only for minimal sequential and random I/O activity and should not be considered for enterprise class servers.

4.2 SATA (IDE) drives

(Designed for desktop and small servers) – Offline storage

The SATA (Serial Attached Technology Adapted) disk drive is a single port interface on an ATA disk. Of the disk family (excluding PATA), SATA has the lowest price per GB but performance levels are poor compared to its cousins. The SATA interface allows a point-to-point connection enabling multiple ports (disk drive connections) aggregated on a single controller.

The SATA interface supports performance capabilities of approximately 150 MB/s yet is still only suitable for single use applications per server. Performance on SATA can be improved by adding additional spindles but this does not change the úseôqualities of the disk. Its primary use should only be for single use applications per server.

SATA drives lack of enterprise-class reliability makes them unsuitable for 24x7 high availability applications, but are good option for secondary storage applications where constant access to the drives is not required.

4.3 FATA drives

(Designed for enterprise server arrays) – Nearline storage

FATA (Fibre Attached Technology Adapted) disk is a hybrid hard drive first introduced by HP in 2004 that combines both Fibre Channel and ATA technologies. FATA drives use an ATA drive mechanism, offering the same performance and capacity as a standard ATA drive, but also features a Fibre Channel connector, which enables the FATA drive to be used where conventional Fibre Channel drives are currently connected.

FATA is essentially FC attached disk drive with dual porting similar to a regular FC disk drive. The big difference is in the price and lower performance -- putting FATA on par as an alternative to Serial Advanced Technology Attachment (SATA) disk drives.
FATA provides better performance levels over PATA and SATA and is designed for multi-user arrays and sequential or random I/O. It also provides good levels of reliability.

4.4 SAS drives

(Designed for enterprise server arrays) – Online Storage

Serial Attached SCSI (SAS), an evolution of parallel SCSI into a point-to-point serial peripheral interface in which controllers are linked directly to disk drives.

SAS is a performance improvement over traditional SCSI because SAS enables multiple devices (up to 128) of different sizes and types to be connected simultaneously with thinner and longer cables; its full-duplex signal transmission supports 3.0Gb/s. In addition, SAS drives can be hot-plugged.

Serial-attached SCSI (SAS) hard drives are ideal for servers and external storage running mission-critical, high I/O workloads applications requiring maximum reliability (about 3.5 times that of Entry Drives) and performance, e.g. high transaction applications such as email, ERP, and CRM.

SAS offers data transfer rates in excess of 3 gigabits per second (Gbps) with potential rates of 10 Gbps or more. SAS is said to offer an ideal solution for businesses with substantial storage, backup, and archiving demands.

4.5 Fibre Channel (FC) drives

(Designed for mission critical and high transactional environments) – Online Storage

Fibre Channel is a highly-reliable, gigabit interconnect technology that allows concurrent communications among workstations, mainframes, servers, data storage systems, and other peripherals using SCSI and IP protocols. It provides interconnect systems for multiple topologies that can scale to a total system bandwidth on the order of a terabit per second.

Fibre Channel delivers a new level of reliability and throughput. Switches, hubs, storage systems, storage devices, and adapters are among the products that are on the market today, providing the ability to implement a total system solution.

Fibre Channel has been adopted by the major computer systems and storage manufacturers as the next technology for enterprise storage. It eliminates distance, bandwidth, scalability, and reliability issues of iSCSI and SAS.
Fibre Channel is simply the easiest, most reliable solution for information storage and retrieval.
Fibre Channel, a powerful ANSI standard, economically and practically meets the challenge with these advantages:

- **Price Performance Leadership** - Fibre Channel delivers cost-effective solutions for storage and networks.
- **Solutions Leadership** - Fibre Channel provides versatile connectivity with scalable performance.
- **Reliable** - Fibre Channel, a most reliable form of communications, sustains an enterprise with assured information delivery.
- **Gigabit Bandwidth** - multi gigabit bandwidth now (currently at 4GBps) with 8GBps and 10GBps and greater on the horizon with converged network and storage adapters..
- **Multiple Topologies** - Dedicated point-to-point, shared loops, and scaled switched topologies meet application requirements.
- **Multiple Protocols** - Fibre Channel delivers data. SCSI, TCP/IP, video, or raw data can all take advantage of high-performance, reliable Fibre Channel technology.
- **Scalable** - From single point-to-point gigabit links to integrated enterprises with hundreds of servers, Fibre Channel delivers unmatched performance.
- **Congestion Free** - Fibre Channel's credit-based flow control delivers data as fast as the destination buffer is able to receive it.
- **High Efficiency** - Real price performance is directly correlated to the efficiency of the technology. Fibre Channel has very little transmission overhead. Most important, the Fibre Channel protocol is specifically designed for highly efficient operation using hardware.

Fibre Channel has the highest power on hour (POH) measure of all the disk classes. Running at 24 hours continuously, Fibre Channel is purpose built for multi user, multi application and array use. Fibre Channel is designed for sequential I/O and high random occurrence with 100% activity. For these reason, it is the storage platform of choice for the ICONZ VERSA platform.

5. The impact of 'RAID' on performance

RAID stands for Redundant Array of Inexpensive Disks. It enables a large number of low cost hard drives could be linked together to form a single large capacity storage device that offered superior performance, storage capacity and reliability over older storage solutions.

There are three primary reasons that RAID was implemented:

- Redundancy
- Increased Performance
- Lower Costs

Redundancy is the most important factor in the development of RAID for server environments. It allows for a form of backup of the data in the storage array in the event of a failure. If one of the drives in the array fails, it can either be swapped out for a new drive without turning the systems off

(referred to as hot swappable) or the redundant drive can be used. The method of redundancy depends on which version of RAID is used.
The increased performance is only found when specific versions of the RAID are used. Performance is also dependent upon the number of drives used in the array and the controller.

Unquestionably, IT managers like low costs. When the RAID standards were being developed, cost was also a key issue. The point of a RAID array is to provide the same or greater storage capacity for a system compared to using individual high capacity hard drives. A good example of this can be seen in the price differences between the highest capacity hard drives and lower capacity drives as shown in Diagram A on page 5. Three drives of a smaller size could cost less than an individual high-capacity drive but provide more capacity.

There are typically three forms of RAID used for desktop computer systems: RAID 0, RAID 1 and RAID 5. In most cases, only the first two of these versions is available and one of the two technically is not a form of RAID.

5.1 RAID 0

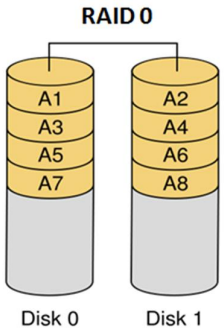

The lowest designated level of RAID, level 0, is actually not a valid type of RAID. It was given the designation of level 0 because it fails to provide any level of redundancy for the data stored in the array. Thus, if one of the drives fails, all the data is damaged.

RAID 0 uses a method called stripping. Stripping takes a single chunk of data like a graphic image, and spreads that data across multiple drives. The advantage that stripping has is in improved performance. Twice the amount of data can be written in a given time frame to the two drives compared to that same data being written to a single drive.

5.2 RAID 1

RAID 1 provides a simple form of redundancy for data through a process called mirroring. This form typically requires two individual drives of similar capacity. One drive is the active drive and the secondary drive is the mirror. When data is written to the active drive, the same data is written to the mirror drive.

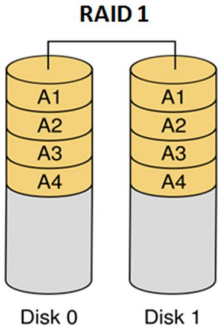

This provides a full level of redundancy for the data on the system. If one of the drives fails, the other drive still has all the data that existed in the system. The big drawback of course is that the capacity of the RAID will only be as big as the smaller of the two drives, effectively halving the amount of storage capacity if the two drives were used independently.

Advantages:

- Provides full redundancy of data

Disadvantages

- Storage capacity is only as large as the smallest drive
- No performance increases
- Some downtime to change active drive during a failure

5.3 RAID 0+1

This is a hybrid form of RAID (not to be confused with RAID 10) that some manufacturers have implemented to try and give the advantages of each of the two versions combined. Typically this can only be done on a system with a minimum of 4 hard drives. It then combines the methods of mirroring and stripping to provide the performance and redundancy. The first set of drives will be active and have the data stripped across them while the second set of drives will be a mirror of the data on the first two.

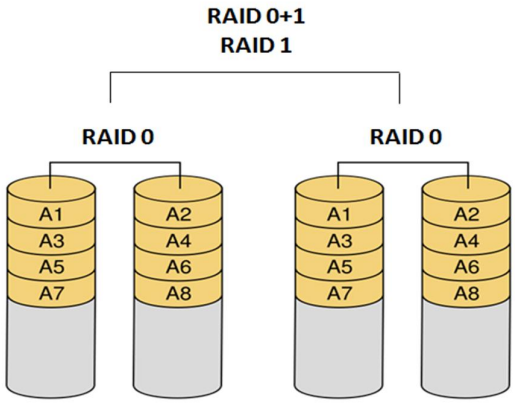

Advantages:

- Increased performance
- Data is fully redundant

Disadvantages:

- Large number of drives required
- Effective data capacity is halved

5.4 RAID 10 (1+0)

RAID 10 is effectively a similar version to RAID 0+1. Rather than stripping data between the disk sets and then mirroring them, the first two drives in the set are mirrored together. The second two drives form another set of disks that are mirror of one another but store stripped data with the first pair. This is a form of nested RAID setup. Drives 1 and 2 are a RAID 1 mirror and drives 3 and 4 are also a mirror. These two sets are then setup as stripped array.

Just like the RAID 0+1 setup, RAID 10 requires a minimum of four hard drives to function. Performance is pretty much the same but the data is a bit more protected than the RAID 0+1 setup.

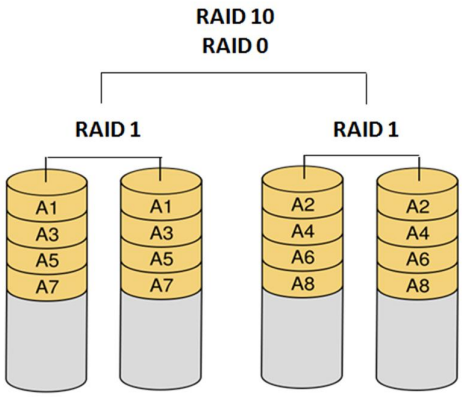

Advantages:

- Increased performance
- Data is fully redundant

Disadvantages:

- Large number of drives required
- Effective data capacity is halved

5.5 RAID 5

This is the most powerful form of RAID that can be found in a desktop computer system. Typically it requires the form of a hardware controller card to manage the array, but some desktop operating systems can create these via software.

This method uses a form of stripping with parity to maintain data redundancy. A minimum of three drives is required to build a RAID 5 array and they should be identical drives for the best performance.

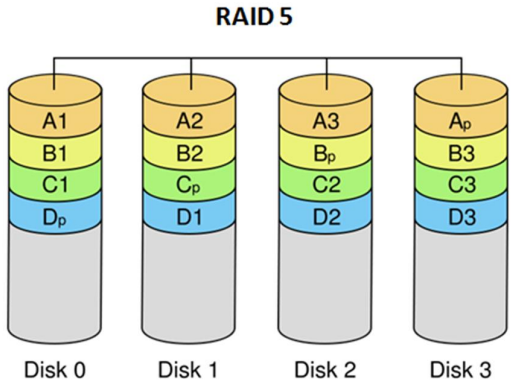

Parity is essentially a form of binary math that compares two blocks of data and forms a third data block based upon the sum of the first two. The easiest way to explain it is even and odd. If the sum of the two data blocks is even, then the parity bit is even. If the sum of the two data blocks is odd, the parity bit is odd. So 0+0 and 1+1 both equal 0 while 0+1 or 1+0 will equal 1. Based on this form of binary math, a failure in one drive in the array will allow the parity bit to reconstruct the data when the drive is replaced.

The parity bit shifts between the drives to increase the performance and reliability of the data. The drive array will still have increased performance over a single drive because the multiple drives can write the data faster than a single drive. The data is also fully redundant because of the parity bits. In the case of drive 2 failing, the data can be rebuilt based on the data and parity bits on the two remaining drives. Data capacity is reduced due to the parity data blocks. In practice the capacity of the array is based on the following equation where n is the number of drives and z is the capacity:

(n-1)z = Array Capacity

In the case of three 500 gigabyte hard drives, the total capacity would be (3-1) x 500GB or 1000 gigabytes.

Hardware RAID 5 implementations can also have a function called hot swap. This allows for drives to be replaced while the array is still functioning to either increase the drives capacity or to replace a damaged drive. The drive controller then takes time while the array is running to rebuild the data array across the drives. This is a valuable feature for systems that require 24x7 operations.

Advantages:

- Increased storage array performance
- Full data redundancy
- Ability to run 24x7 with hot swap

Disadvantages

- High costs to implement
- Performance degrades during rebuilding

5.6 Virtual RAID (VRAID)

VRAID is a term used to describe virtual RAID arrays. In the HP enterprise server world this is defined as, VRAID 0, VRAID 1 and VRAID 5

VRAID0 (none): data is striped across all physical disks in the disk group.

VRAID5 (moderate): data is striped with parity across all physical disks in the disk group. Always 5 (4+1) physical disks per stripe are used.

VRAID1 (high): data is stripe mirrored across all physical disks (even number of them) in the disk group. established pairs of physical disks mirror each other.

On the EVA SAN used by the ICONZ VERSA platform, VRAID 5 is used on all disks with the load spread virtually and evenly over the entire disk group, this delivers the highest levels of efficiency over any of the RAID categories.

Conventional RAID	Virtual RAID
Performance limited by # of disk drives in the storage set.	Performance limited by # of disk drives in disk group
Possible but difficult to find customer data unless? one knows the LBN and chunk size.	Customer data distributed across all disks in a group
Load balancing required of application and databases over available backend (SCSI) busses	Eliminate load balancing procedures for applications and databases.
I/Os balanced across storage set	I/Os balanced across disk group

6. Conclusion

IT executives realize that all data does not have the same value and does not require the same performance. For this reason, there are many solutions in the market that have a ôright-fitô for the purpose intended.

Whether this is a standalone 1RU hosted server through to the high performance ICONZ VERSA platform, the same math applies;

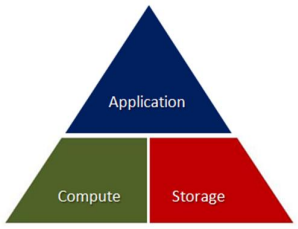

APPLICATION = COMPUTE POWER x STORAGE

It is no longer, nor some would suggest, was it ever, about the cost of the solution first. In the virtual hosting environment where you entrust your business application (mission critical or otherwise) to the provider to manage, the level of investment and care that this provider brings to this equation is key to your decision making.

If you are considering a hosted (collocation or virtual) solution for your business application, your ICONZ account manager and technical consultant will be able to advise you of the many solutions available from ICONZ that will deliver a solution against your mathematical equation.