Juku.it

Down to earth report: Microsoft Windows Server 2012 introduces new storage scenarios

Enrico Signoretti

July 2013

Index

Executive summary

Every end user is continuously looking for solutions aimed at reducing the complexity and to drive down the total cost of ownership (TCO) of his IT infrastructure. Especially in the SMB space, where resources are usually limited, there is a lot of attention payed to this issue.

Storage is a fundamental part of any modern IT infrastructure and it is constantly growing. Even in the smallest enterprises, we need to consider different storage solutions to solve different kind of problems: primary or secondary data, block or file access, backup, archiving, and so on. In many cases, traditional approaches are no longer effective and they need to be revised in all their aspects: ease of use, acquisition and operational costs, longevity and commoditization are on the top of users' requests.

Moreover, end users are looking with an increasing interest at new "software-defined" solutions that allow much more flexibility in designing next generation infrastructures. In this scenario hardware commoditization, coupled with smart software solutions, enable huge savings both in terms of CAPEX and OPEX.

Another key element of modern infrastructures is virtualization: Integration with the hypervisors should be a design principle in any next generation infrastructure.

Last but not least, next generation applications work better with modern infrastructures. Scale-out instead of scale-up with fast, and cheap, local storage instead of shared storage is only an example of what I mean.

With the introduction of Windows Server 2012, Microsoft added many new key features to this product. Plenty of them are related to storage and storage management. We are going to cover some of these features and how they can help the end user to implement modern and efficient storage solutions at a relatively low cost. We won't go deep on features but we will work with examples and use cases to explain the advantages of this approach and the benefits for the IT department and the end users. This document has the intent to show up some ideas on how to use Windows Server 2012 to effectively solve some common storage issues without spending a lot of money.

How to read this document

This paper has been written to talk about the new storage capabilities of Microsoft Windows Server 2012. Instead of writing about single features and how they work, I preferred to address each single issue from different points of view giving realistic examples and use cases on how Microsoft Windows Server 2012 can help you to brilliantly solve them. This is a “Down to earth report” and, as such, it wants to be educative and informative.

IT managers, Executives but also Architects and Infrastructure specialists who care about the IT transformation of their company will find information that will be useful to look at Windows Server from a different perspective.

The introduction was thought out to describe the scenario and the challenges while the rest of the document goes in depth on some of the most interesting use cases.

Every single topic was divided in four chapters: description of the problem, a more technical explanation, a practical example and why it is important for you. The idea is to give the scenario and the different reading perspectives. You can read all of them or choose the most interesting for your professional role.

Introduction

More data means more problems

It's nothing new that unstructured data are growing exponentially. Even smaller enterprises are seeing this trend.

All files are heavier, phones have high resolution cameras, you can't delete old data, retentions are getting longer and longer, you are experiencing first machine generated data collections, and you can't afford to lose any kind of information. Even more, you need to store and maintain data as the raw material to produce information, sometimes the same data produce different information for different enterprise businesses and needs: this causes a proliferation of data elaboration mechanisms and a perceptible boost in the number and in the quality of accesses. Data and information are a fundamental part of any modern competitive strategy: you need better ways to store, manage data and deal with its growth.

Up to now, for small and medium enterprises, it was very hard to maintain big amounts of data. The cost of traditional large shared storage solutions is very high for them and, in most cases, cheaper storage systems usually means slow access speeds that prevent further data elaborations.

Things are changing very quickly now. Enterprises are more confident on next generation solutions based on standard x86 hardware and off-the-shelf software, they have also understood that scale-out architectures can be a viable and flexible option to manage huge amounts of data at a lower cost. Applications need to be closer to data with the objective of shrinking latencies and the operating systems are evolving by showing new features allowing this.

New data means new challenges

Have you ever considered dealing with huge amounts of data? Are you ready to do that? and at which cost?

These questions could sound weird to many enterprises but, for example, even smaller manufacturers are developing new kinds of machinery that produces data and, sooner or later, they will want to gather sensors and logs from those machines. Data collected directly from the machines can be easily converted into useful information that help manufacturers to give better support services and speed up development of future products. This could result in a major improvement of the company agility and competitiveness.

Storage is a key component of the underlaying infrastructure of next generation analytics, even if you can't call it "Big Data" it's not hard to reach tens of Terabytes from extensive data collections. Those kind of data become fundamental to build new product, marketing and support strategy. They have a history, a long retention and they need to be accessed as quickly as possible from traditional tools. In fact, Small and Medium

Enterprises can't afford to develop applications based on new programming models at the moment (like Hadoop for example), that kind of developer is difficult to find and it is very costly. Small enterprises want to stick to well known products and technologies, like SQL for example.

What is Microsoft doing about storage?

The recently introduced Microsoft Windows Server 2012 added plenty of key features, many of them are related to storage and storage management.

It is important to notice that the experience Microsoft has made in the cloud space, thanks to its public cloud proposition (Azure), is at the base of most of these features. Cloud providers think about scalability, reliability, agility, performance and availability but always with an eye to the costs of their infrastructure. This is why they often choose to implement cheap and powerful commodity hardware with state of the art software features on top of that. The latest Microsoft Windows Server releases are well aligned with this philosophy and they are bringing new opportunities to a wider spectrum of end users:

- Storage Spaces: A new capability in Windows Server 2012 (also available in Windows 8) which enables the user to build highly available, resilient and cheap storage systems made of commodity components.
- Resilient FS: ReFS is a next generation FileSystem introduced with Windows Server 2012. It maximizes data integrity, reliability and scalability. it's the perfect companion to Storage Spaces enabling the creation of an end-to-end integrated architecture.
- Deduplication: NTFS is the first Microsoft file system to benefit of this feature. Deduplication is a data compression technique that is used to eliminate repeating data blocks and it can allow great space savings.
- SMB 3.0: the latest version of the most common network file sharing protocol brings a load of new features aimed to improve scalability, performance and availability. SMB 3.0 shares can also be used as a Hyper-V storage.

- NFS 4.1 and support for VMware: A new NFS server is available. Windows NFS server implements the latest version of the protocol, it's fully integrated with other Windows services (like SMB or VSS) and offers an affordable way to have a highly available storage resource to store VMware Virtual machines.

- iSCSI target: A cheap way to provision block storage from industry standard x86 servers.

- Offload Data Transfer (ODX): this feature, if supported by the underlaying storage system, minimizes the usage of host CPU and network by moving data between arrays without involving the hosts. Array-to-array data movements can also result in a huge performance improvement.

- VHDX and live storage migrations: A new, and improved, format for storing VMs. Now it's possible to grow further.

- Improved storage management.

In this paper we are going to cover all of these features with examples and ideas that should help end users and Microsoft partners to build next generation affordable and software-driven storage infrastructures.

You can find detailed technical information about these technology on Microsoft websites: http://aka.ms/storagews2012

Recycling old storage

The problem

Enterprises often manage operational and organizational changes: acquisitions, projects of any kind, business and process reorganizations, and so on. The whole technology stack follows these changes and the IT department is continuously involved in provisioning, management or retirement of applications and related resources, such as storage or computational power.

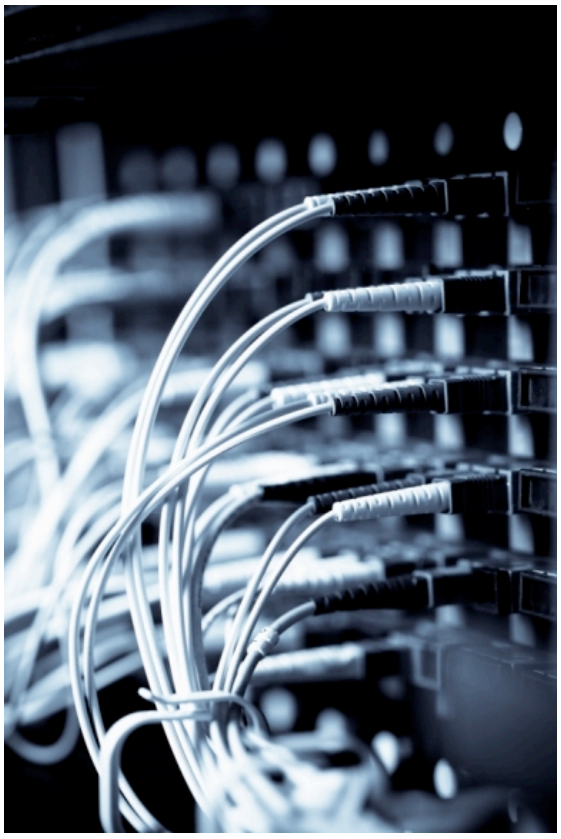

In time, those changes, strongly contribute to create heterogeneous, difficult to manage, environments. Many concurrent factors, such as organizational aspects or acquiring processes, often lead to totally separated technology stacks organized in silos! TCO grows and reclaimed or unused resources are frequently difficult to reallocate.

On the server side virtualization and commoditization helped to solve the problem: now you have common pools of resources and you can provision, dismiss or reuse them very easily. But, on the contrary, storage doesn't often offer the same elasticity because storage virtualization comes at a higher cost and through proprietary complex solutions.

The solution

Microsoft Windows Server 2012, especially now with the recent launch of version R2, allows to easily virtualize your old used resources, giving them a performance boost and sharing them again to your servers or clients.

Storage spaces, automated tiering, deduplication, VSS and others can concur all together to enable the creation of an efficient and full-of-features tier 2 storage system, made from old and reused different resources. Virtualized storage systems of this type, due to the differences in IOPS and latencies provided by different resources, can hardly replace ordinary enterprise storage systems. On the other hand, as already described in previous chapters, you can use a Windows Server Storage solution to provision huge amounts of space, and IOPS, at a very low cost: possible fields of application are countless in all those cases where space is more important then pure performance.

A practical example

A couple of clustered servers with FC HBAs can be easily seen as a high available powerful virtualization system. Connected to an old SAN, these servers can mount different LUNs and organize them through Storage Spaces. Once you have created LUNs on the old storage arrays all management can be demanded to windows Server. Storage Spaces is also useful to rationalize the data protection mechanisms used by different storage in the backend and add, at no cost, many useful new features like thin provisioning or deduplication, for example.

Automated tiering functionalities, introduced with Windows Server 2012 R2, enable to use a portion of Storage Spaces shared storage to maintain most accessed files on the fastest disks. If storage tiering is configured using a small amount of SSDs the performance improvement can be huge. SSDs can also be configured as a write-back cache to improve writes.

Why it is important for you

Especially in this case, through virtualization and recycling of old resources, Windows Server makes possible interesting savings on both the CAPEX and OPEX point of view. Storage spaces can also be considered a viable solution to build a next generation storage system for many types of workloads.

With a small investment it's possible to give new life to old storage systems and delay hardware investments. Future expansions could be done by buying cheap SAS JBODS instead of expensive proprietary disk trays. Intelligent SSD usage can boost performance and flexibility of the overall storage infrastructure. Features like deduplication can vastly improve storage space utilization efficiency and, again, help to delay further investments.

At the same time you can ease management doing all storage administration on Windows Server. Automated tiering and wide striping provided by Storage Spaces help to avoid hot spots. Virtualization can also be used to perform non disruptive migrations of old storage systems retirements.

Rethinking NAS

The problem

File servers are very relevant to almost all business activities. In recent years, due to growth of unstructured data, many end users have decided to implement NAS (Network Attached Storage) systems or NAS gateways to avoid some of the complexity of implementing and maintaining a clustered server environment. File servers weren't only complex but also less scalable, reliable and available due to file systems limitations. NAS storage systems are still a great option but they are very costly and lately have become more complex, especially when you need to scale a lot. The advantages of using NAS are less evident now.

In fact, these new types of configurations aren't only complex but, many times, they struggle to keep the pace with enhancements introduced by new protocols, new operating systems and new technology in general. Powerful commodity x86 servers and new software capabilities can easily change the scenario.

The solution

Microsoft Windows Server 2012 has introduced a lot of improvements to SMB and NFS servers and to the clustering system (CSV2), and now, it's

possible to achieve better performance, scalability and reliability at a low cost. Storage Spaces enables to build a very powerful flexible, and resilient pool of storage on the backend that, coupled with a next generation filesystem like ReFS, can easily manage huge amounts of data.

How it works

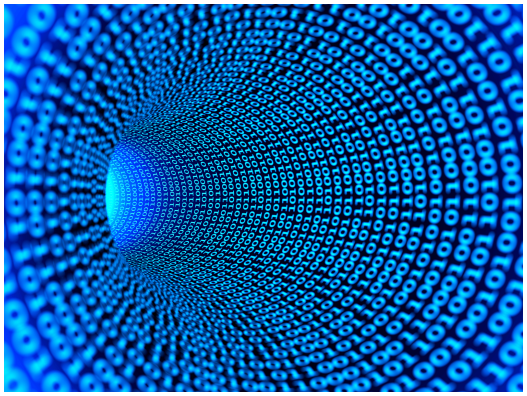

ReFS is a brand-new and robust next generation file system: it's optimized for data integrity, availability and scalability. Despite its recent launch it is already clear that it will soon become the primary choice for all new installations. This FS will receive a lot of improvements in future releases and it also maintains NTFS compatibility to ease migrations from older Windows Server versions.

Microsoft Windows Server 2012 has also introduced a few state of the art enhancements to its SMB protocol (now at version 3). The new SMB features include:

- the ability to create transparent-failover clusters, that means no service disruption in case of a failure;
- support for Remote Direct Memory Access (RDMA) network adapters. As for local DMA it helps to extremely speed up communications while freeing CPU resources;
- SMB multichannel, a way to bond different ethernet links together to get better resiliency, speed and load balancing;

All these new features can dramatically change the performance of the storage system and allow to confidently store Microsoft SQL Databases and Hyper-V VMs. With the Microsoft NFS service it is even possible (and supported) to use Microsoft Windows based storage as a repository for VMware VM files (VMDKs).

The usage of Microsoft iSCSI target, now a built-in option, can provide Storage Area Network functionalities. This means that it is possible to share

a block device to other hosts and even boot servers from it, particularly interesting in labs and in all those environments that change configuration very often (like training rooms for example).

A practical example

There are many reference architectures and best practices available ranging from smaller 2-nodes failover clusters to multi-node scale-out topologies. Regarding back-end storage, it is possible to design different configurations starting from ordinary FC (Fiber Channel) or iSCSI Storage Area Networks to the most radical and complex Storage Spaces based configurations.

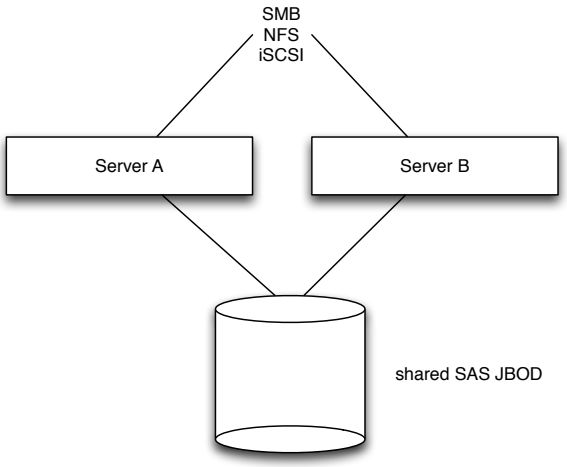

Windows Server 2012 HA Cluster

Thanks to Storage Spaces, every Microsoft Windows system administrator can easily build a simple 2-node cluster that accesses one or more shared SAS JBODs while avoiding the cost of a traditional shared expensive storage (Fiber Channel or iSCSI SAN). Modern CPUs can easily manage all storage related operations like parity, compression or deduplication calculations. At the same time the efficiency introduced by SMB 3, NFS 4 and iSCSI services enables to achieve the best efficiency in terms of data communication.

Why it is important for you

This is a way to build a NAS at the cost of servers, disks and Windows Server licenses. It is not only cheap but also powerful, scalable and reliable. These kind of commodity storage solutions have many real life applications and can be widely used in many different environments, a few examples follow:

- Internet and cloud service providers: where price/performance ratio of storage is always a key element of their architecture.
- Media/graphic: users that need to store big amounts of photos and/or videos are currently looking for cheap, fast commodity storage. They also need to access it as quickly as possible from their workstations to manage all the copies and editing processes during their day-to-day work.
- R&D departments with loads of data to gather, manage and delete every now and then. They usually can't afford to spend a lot in storage but the more storage they have the more online data they can manage to get better results faster.

Shrinking backup costs

The problem

The more data you have the more backup becomes a problem. Tapes are difficult to manage. In most of the cases, obtaining a high efficiency during the write operations automatically means poor restore performance. This is no longer acceptable in business terms. When your company needs to restore data it means that it is necessary as soon as possible: every hour spent waiting for a restore is an hour wasted in terms of productivity. An efficient backup/restore policy coupled with the right hardware solution could lower the overall TCO of the whole infrastructure but the acquisition costs and its maintenance is a barrier to its adoption.

How it works

The latest common type of tape drives are fast but you need to push and mix many different data streams together on a single drive to be sure to maintain a sustained throughput. This technique is called multiplexing. A good solution for writing backups but, on the flip side, it's really painful when you need to restore them. In fact mixing different streams results in scattered data on the tape and the restoring process goes very slowly due to the fragmentation of data on the tape. This impedes fast sequential reads, and consequently the restores are slow. Furthermore, multiplexing, like many

other similar techniques can severely compromise compression and deduplication efficiency.

VTLs (virtual tape libraries) are a possible solution: these appliances usually expose NFS or SMB shared volumes, if not emulating a physical tape library. Some types of VTL also have data footprint reduction mechanisms like compression or deduplication. In many cases, they brilliantly solve part of the problem but they are also expensive.

The solution

Microsoft Windows Server 2012 with storage spaces, NTFS, deduplication, SMB 3 and NFS 4 could be a great alternative to a VTL for price sensitive customers and it can also be better in certain circumstances.

In fact, thanks to the characteristics of storage spaces you can create a big and resilient pool of storage resources starting from commodity hard disks and build compressed/deduplicated NTFS data volumes on top of it. NFS and SMB protocols allow you to easily share these volumes with the backup/media server and get the maximum in terms of throughput. Expansions on the storage side comes at the cost of the disk trays.

An industry standard x86 server with 24 disk slots

Microsoft Windows Server can also be configured to manage all the task: in this case you can use the same server to mange the disk space and the tape library taking full advantage of modern CPUs and avoiding useless network traffic going back and forth between the VTL and the tape library.

And it's not all, in case of small companies, where the backup/media server is a single system, you can configure the backup software, the VTL-like and tape management in the same server. This is the most affordable and easiest to manage configuration for this kind of environment.

A practical example

New Windows Server 2012 features, like storage spaces and deduplication, are free of charge and they enable great availability and space savings.

These two, coupled with a next generation x86 server capable of accommodating up to 24TB of raw storage, can offer a good backup repository: it actually means more than 18TB usable space and much more when using deduplication. These numbers are more than enough for maintaing the backup of a SMB end user and they remain quite large for many enterprises looking at a short term, but fast, backup on disks while maintaining long term backups on tapes.

Why it is important for you

Using Microsoft Windows Server 2012 to build a backup appliance can dramatically improve your backup strategy, giving your company faster backup and almost instant restores.

This kind of solution could easily increase your backup retention and shrink backup windows at a very low cost without needing to adopt dedicated hardware.

These improvements enable to achieve better efficiency and agility in data management that results in productivity enhancements and big savings for the company. Also TCO is affected positively because the end user can continue to use standard tools and procedures currently used for other Windows Server.

Disaster Recovery made easier

The problem

Implementing a Disaster Recovery plan for an IT environment is often costly and, usually, it's hard to maintain efficient.

In traditional IT environments, when you have RPO (Recovery Point Objective) and RTO (Recovery Time Objective) constraints, most of the DR projects are based on storage hardware remote replication. This means that the SAN/NAS on the primary site must have a remote copy functionality enabled and there is a similar hardware on the disaster recovery site.

The cost of doubling the whole infrastructure, especially if the second site is in a hot stand-by mode, is not affordable for the majority of SMB companies. Moreover, things got more complicated if virtualization is involved:

- To properly replicate data on a secondary site, it's necessary to have a strong integration between the hypervisor and the storage system, that's because you need to be sure of having consistent data sets at any given time. This could also mean more costs in terms of software licenses.
- In most cases you are obliged to remotely copy full volumes of data, even if the VMs involved in the DR plan are few. This could result in waste of bandwidth and space or in more complex management of storage.
- Some FC (Fiber Channel) protocol based storage systems can't replicate data via TCP/IP based protocols, such as iSCSI. This means that the end user needs very costly FC switches capable of encapsulating/decapsulating FC into IP (FCIP).

because of all this requirements and constraints, the cost of implementing and maintaining a DR plan rises, and this is why many small and medium enterprises struggle to look at it with confidence.

The solution

Microsoft, with the introduction of Windows Server 2012, has added important new features to Hyper-V and the underneath storage layer, allowing the end user to avoid most hardware and software costs in simple Disaster Recovery designs for virtualized environments.

Windows Server 2012, coupled with System Center 2012, can also take advantage of Azure's Hyper-V Recovery Manager: a new service that can help you to coordinate replication and recovery from within the cloud. Hyper-V Recovery manager takes full advantage of Hyper-V replica and doesn't need third party products to operate. This Service helps to automate and coordinate recovery operations on the secondary site in case of an outage of the primary site. At the same time the service can also be used to test DR procedures and all the activities can be automated with Powershell.

A practical example

Hyper-V VM replica enables customers to replicate the content of single VMs using standard IP protocols and without specialized storage. The replication can happen on unprotected networks, and without a VPN, because the user can choose to encrypt VM Replica traffic also. It is managed on a per single VM basis with a very fine granularity regarding the frequency of the snapshots taken for the replica.

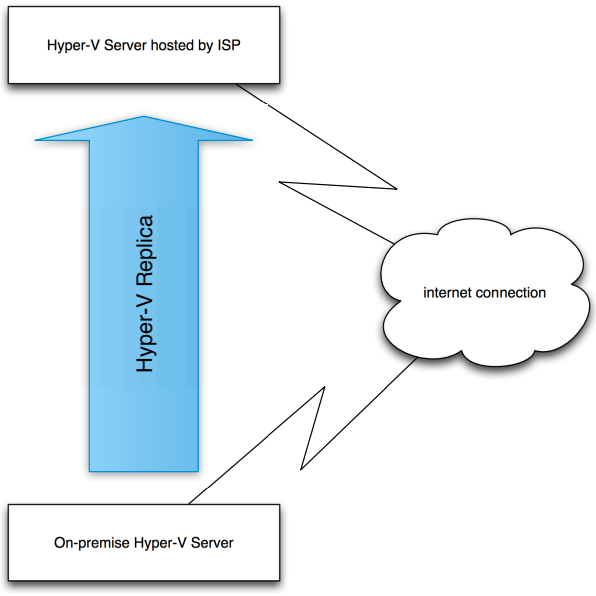

The secondary site can be built on standard servers with local disks. In this case, once again,

Storage Spaces and ReFS could easily manage availability and reliability of commodity hard disks while giving the best uptime.

Using internal disks also helps to save rack space and power, allowing the customers to get better pricing if the secondary site is hosted by an internet service provider. In fact, many small enterprises don't have a second secure remote site, but if they do it's usually not far enough.

In more complex infrastructures, it's possible to think about a shared storage built on the basis of the indications written in the above chapter.

Why it is important for you

Disaster Recovery and Business Continuity plans are quite complex to realize and maintain but, they are insurance to protect the business in some critical circumstances.

Many small companies are still using tape vaulting to implement their DR because they can't afford a replication solution. That's the cheaper approach but the RTO is often incompatible with current business needs. Nonetheless, testing DR is often a big problem and a very costly procedure.

Now, with these new Windows Server 2012 features it is possible to rethink most of these DR processes and dramatically reduce the RTO. Realizing a simple DR plan is slightly easier now, especially if Hyper-V is involved. The cost of the secondary infrastructure can be easily slashed down due to the use of Storage Spaces and commodity storage, while orchestration and automation on the cloud can help to maintain the highest level of efficiency and agility at a minimum cost.

The overall savings make the implementation of a simple DR very attractive even for smaller enterprises.

Reacting to business peaks

The problem

In recent years enterprises have been seeing an incredible data growth, all the previous chapter talked about this issue and how to solve it with your on-premise infrastructure but, in some circumstances, it's not enough. Organizations are experiencing many new ways to do business and economic scenarios can suddenly change. Explosive growths are often accompanied with stagnation or even periods of recession. Opportunities are to be grabbed as they pop up, the agility of the IT is part of the company competitiveness.

You must be ready to react at business changes: rapid resource allocation and deallocation are fundamental to keep up with business needs.

Traditional on-premise infrastructures aren't quick enough to respond to these business needs: procurement, implementation and even deallocation of the resources are non compatible with the immediacy of the demand.

The solution

In order to provide the most flexible solutions, Microsoft offers many different integrations with Azure cloud services. In fact, most of the features that we can now find on Windows Server were previously developed to implement Azure cloud. The integration between Azure, System Center and Windows Server is a consequence of those efforts.

A practical example

The most useful example comes from backup, Windows Azure Backup can be used through the standard backup tools of Windows Server or System Center Data Protection Manager. It provides:

- Incremental backups to save bandwidth;
- Encryption;
- Integrations with other Microsoft tools;
- Flexible retention policies;
- Its cost is only related to the effective occupied space.

Microsoft has also integrated Windows Server with Azure using other technologies, like StorSimple for example. The StorSimple appliance acts as an ordinary local storage but, actually it is only a cache that relies on the Azure cloud storage service. It offers the performance of a local storage with a cost-effective and infinite space. The tight integration between Windows Server and StorSimple enables very interesting scenarios for all major tier 2 storage applications. In fact, StoreSimple integration allows many applications (like Exchange and SharePoint) to take the maximum advantage in terms of data protection and recovery policies in case of need.

Why it is important for you

The IT landscape is rapidly changing and new challenges are arising. A modern IT department of any type or size of enterprise must be ready for any requests coming from end-users and management.

The public cloud is the only way to immediately provide the service at a minimum risk, but Microsoft, thanks to Azure with system Center and Windows, could do much more than that. It provides a way to transparently migrate data and workloads back and forth between your enterprise and the cloud at a low cost. The IT becomes more flexible and agile while the Enterprise saves money and has the advantage of a better service.

Gathering and analyzing data

The problem

Terms like “internet of things” is often used to explain how next generation devices work: it refers to the ability of objects to be connected to the internet, record and then share data with other devices or humans. It can be a huge opportunity if you are able to take advantage of it but, on the flip side, it could become a major weakness for your enterprise if you are not ready to properly manage it.

The best way to understand the size of the problem is to work with examples: if your company produces machinery, and your product is capable of sending its status, or sensors and logs, to a central internet repository, you need something that stays up and runs all the time, listening and saving all the data for future elaboration. Take for example a company that produces 1000 machines every year and each one of those produces a database line every hour. It’s easy to do some calculations on how many data will be produced after the first year of production:

- 1000 * 24 hours a day = 24.000 DB lines per day
- 24.000 becomes 720.000 in a month
- 720.000 in a year: 262.800.000 DB lines per year!!!

This is a quite an impressive example, you won’t probably get there but your company will have more data to manage in the near future. It means a lot of

storage space (even if this is a DB), hard to backup and maintain with ordinary, expensive, shared storage solutions.

Moreover, collecting and storing is only the first part of the problem, then you need to transform data in information as fast as possible.

The solution

Microsoft offers all the tools to collect, store and transform data in information, now with Windows Server 2012 and Microsoft SQL Server 2012 it's possible to mange much more data than in previous versions.

Internet Information Services can host the listener/collecting applications and can easily store data directly on files or in a DB. Storage spaces provides the cheap and powerful storage layer to sustain the potentially massive growth of data.

In practice the solution lies on Windows Server 2012 and well known standard tools that every developer already knows or can easily learn.

Why it is important for you

All modern analytics and Big Data architectures show scale-out designs with applications and storage resources distributed on the same nodes of the clusters. Even if an ordinary enterprise doesn't target Petabyte scale scenarios and this kind of architecture, Windows Server 2012 and Windows SQL Server 2012 show similar scaled-down characteristics and can be very helpful to manage big amounts of data if compared to previous versions.

Latest Microsoft SQL versions follow the principle of scale-out, has columnar compression and other performance optimizations while Windows Server 2012 is optimized to efficiently manage local storage.

The recently announced Windows Server 2012 R2 also adds automatic tiering functionalities to enable transparent usage of flash memory available on the servers. This feature has a tremendous impact in performance especially in those environments that have access to big amounts of data with only small datasets active at one time.

Bottom line

Microsoft Windows Server 2012 opens new interesting storage scenarios. It allows to build many different affordable and flexible solutions that can easily fit small and medium enterprise business needs. For some use cases it can also be considered as a valid option in a large enterprise.

End users and system integrators have new opportunities to address unstructured storage growth with next generation solutions based on commodity hardware and off-the-shelf easy to use software.

Storage Spaces, ReFS and other recently introduced features show some limitations but, fortunately, most of them were already been addressed by Microsoft with the introduction of Windows Server 2012 R2. In fact the modern design of Windows Server 2012 storage layer will allow end users to receive enhancements as soon as they will be available. The new version also add more impressive features, protocols enhancements and an improvements to all the Hyper-V storage layer.

Most of the enhancements introduced to the latest Windows Server versions come from the experience that Microsoft has matured with Azure cloud services, and this is also why it is very well integrated with it. This integration brings a great freedom of choice to the end user. Now one can decide where to deploy data and services and move them between on-premise infrastructure and Azure cloud in very simple and cost-effective ways.

Appendix

In this short section the reader will find some data storage specific technical terms and acronyms used it this paper:

- **DR**: is the abbreviation for Disaster Recovery
- **FC**: Fiber channel is a high speed networking technology primarily used to connect data storage devices.
- **HBA**: Host Bus Adapter is an expansion (PCI) card used to connect the host to a network or a storage device.
- **IOPS**: Input Output Operations per Second, is a measurement used to benchmark data storage devices.
- **iSCSI**: Internet Small Computer System Interface, an IP-based storage networking protocol.
- **JBOD**: (Just a Bunch Of Drives), normally refers to a group of disks that are not organized in a RAID group.
- **LUN**: Logical Unit Number is a number used to address a logic device (volume) in a storage device.
- **NAS**: Network Attached Storage is a file-level network attached storage device that acts as a file server.
- **NFS**: Network File System, it's a network file sharing protocol. An open standard commonly used on Unix and Linux platforms.
- **Remote Direct Memory Access (RDMA)**: is a communication protocol, like for local DMA, that helps to extremely speed up communications while freeing CPU resources.
- **RPO**: Recovery Point Objective is the maximum period of time in which data may be lost due to a disaster.
- **RTO**: Recovery Time Objective is the time in which a service should be restored after a disaster
- **SAN**: Storage Area Network is a dedicated network for storage used to connect arrays, tape libraries and other storage devices to computer hosts.
- **SAS**: Serial Attached SCSI is an evolution of SCSI protocol that uses serial links instead of a parallel bus.
- **SMB**: Server Message Block (also known as CIFS) is a network sharing protocol developed by Microsoft.

- **Scale-out**: the ability of an IT systems to scale horizontally by adding more nodes to a cluster.
- **Scale-up**: the ability of an IT system to scale vertically in the same box.
- **SSD**: Solid State Disk, is a device that acts as a traditional mechanical disk but is made only of electronic parts like flash memory.

Juku

Why Juku

Jukus are Japanese specialized cram schools and our philosophy is the same. Not to replace the traditional information channels, but to help those who make decisions for their IT environments, to inform and discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in IT should look around themselves: things are changing rapidly and there is the need to stay informed, learn quickly and to support important decisions, but how? Through our support, our ideas, the result of our daily interaction that we have globally on the web and social networking with vendors, analysts, bloggers, journalists and consultants. But our work doesn't stop there, the comparison and the search is global, but the sharing and application of our ideas must be local and that is where our daily experience, with companies rooted in local areas, becomes essential to provide a sincere and helpful vision. That's why we have chosen: "think global, act local" as a payoff for Juku.

Author

Enrico Signoretti, Consultant and passionate blogger (not necessarily in that order). Having immersed into IT environments for over 20 years, his career began with Assembler in the second half of the 80's before moving on to UNIX platforms (but always with the Mac at heart) until now when he joined the "Cloudland". During these years his job has changed from deep technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm higly focused on supporting end users, vendors and third parties in the development of their IT infrastructure strategies. He is constantly keeping an eye on how market evolves and continuously looking for new ideas and innovative solutions. You can find Enrico's social profiles here: http://about.me/esignoretti