

What is the business case for Windows-Server virtualization?

Windows-Server virtualization, the deployment of a virtual version of a Windows-Server operating environment, is used to reduce hardware costs, gain efficiencies, and improve the availability of computing resources. It refers to installing a virtual environment onto one or more “virtualized” hardware servers (termed Physical Hosts) and deploying multiple virtual Windows-Server operating systems (termed Virtual Guests) onto this virtual environment.

In small to medium-sized businesses, we typically see three levels of Windows-Server virtualization with these increasing benefits:

- Single Physical Host – Cost savings (energy and hardware) with some flexibility
- Multiple hosts with Storage Area Network (SAN) – Highly available environment with minimal downtime
- Multiple hosts with Site-to-Site Failover – Disaster recovery to separate location

We review each of these levels below.

Single Physical Host

This virtualization level has these components:

- Single hardware server with onboard storage – This hardware server is the platform for the Physical Host; it could be a HP ML350/ML370 tower server or equivalent with multiple disk drives.
- Virtualizing software – The operating environment for virtualization; typically the free versions of either VMware’s VSphere or Microsoft’s Hyper-V. (These products are available as free downloads from the manufacturer.) Installing the virtualizing software onto the hardware server creates the Physical Host.
- Multiple Virtual Guests – The virtual operating systems installed onto the Physical Host; usually one or more instances of Microsoft’s Windows Server. (These instances must each be licensed copies of Windows Server and any associated, server-based applications.)

This environment consolidates several Windows Server instances onto a single hardware server with sufficient processing capability, Random Access Memory (RAM), and on-board disk storage. It introduces cost savings in hardware, energy, and support and provides some flexibility in the transfer of a virtualized instance to a new hardware platform (although this transfer is manual and requires a second hardware server).

Some caveats:

- The hardware server (and its components) is the primary point of failure; if it is down, all of the installed Virtual Guests are unavailable.

- Ports on the Physical Host are handled differently in a virtual environment; attached backup devices and UPS equipment might need special setup.

Primary business benefits:

- Less up-front acquisition cost (capital expenditure or CapEx) since a single hardware server can be used rather than two or more hardware servers. Plus, the virtualizing software at this level is basically free.
- Less energy required to power a single hardware server than multiple hardware servers; leads to reduced operating expenses (OpEx).
- Fewer components to support; could lead to lower support costs.
- Increased flexibility and scalability when migrating to a new hardware server.

This virtualizing environment works well in a business with a couple of Windows Servers that is looking to capital and operating reduce costs.

Multiple Physical Hosts with a Storage Area Network

At this level, we separate the storage (disk-drives) from the Physical Host and move them to a separate Storage Area Network (SAN)¹. We also add sophisticated virtualizing software capable of automatically managing the location of Virtual Guests.

A major benefit of this approach is termed: "High availability".

High availability refers to "A system design approach and associated service implementation that ensures a prearranged level of operational performance will be met..." (from Wikipedia under "High availability"). Basically, if designed properly, this level provides complete redundancy of all critical components within the equipment stack such that any single component can fail without compromising system reliability.

Improved performance is also likely since the virtualizing software can automatically balance available resources against Virtual Guest needs.

This virtualization level has these primary hardware components:

- Storage Area Network (SAN), preferably with redundant disk chassis and network switching²
- Two or more Physical Hosts, preferably with N+1 redundancy³
- Two or more VLAN-capable Ethernet switches⁴

Each item is a critical of the overall design:

- All data and Virtual Guests reside on the SAN
- Virtual Guests are balanced among the Physical Hosts
- Ethernet switches route all the traffic between the SAN and the Physical Hosts

If any item fails, the system fails. So, each item must be redundant (to increase reliability) and must be properly maintained.

Notes:

¹ Technically, the Storage Area Network consists of disk arrays and the inter-connecting fabric, which is TCP/IP over Ethernet over UDP in the case of an iSCSI SAN.

² The SAN is the data storage; it should have redundant components capable of automatic failover. A single-chassis SAN (like the HP P2000 series) has redundant controllers and power supplies, but fails if its disk backplane fails; a redundant-chassis SAN (like the HP P4000 series) consists of two or more separate storage arrays. The chance of a failure in a redundant-chassis SAN affecting all arrays at once is extremely small.

³ Physical Host N+1 redundancy refers to adding one more Physical Host than required to meet performance standards. The additional Physical Host permits performance standards to be retained, even if a Physical Host fails.

⁴ In addition to providing the SAN connectivity, the Ethernet switches provide redundant network links between the Physical Hosts and the remainder of the network.

Multiple Hosts with Site-to-Site Failover

Our highest level of Windows Server virtualization, Multiple Hosts with Site-to-Site Failover, addresses the issue of a single-site failure; how long does it take to recover to a new location if your primary site fails (as in a building catastrophe such as long-term power outage, flooding, fire, theft, etc.).

Like most data-center-uptime strategies, redundancy is the core concept; in this case, a second site is equipped with comparable equipment and the data is synchronized between the primary and secondary site. Done properly, the secondary site can be brought up either automatically or, when budget is a constraint, within a short interval of an hour or less.

Configuring for automatic failover can be considerably more expensive than allowing a short interval of an hour or less to recover since you essentially need to duplicate the primary site at the remote location, have sufficient bandwidth between the locations to permit real-time replication, and deploy some additional equipment and software to manage the automatic failover. While automatic failover is feasible, we structure the failover interval (automatic or short) to meet the client's requirements and budget.

When configuring for a short delay, we use HP Proliant servers with VMware's vSphere virtualization platform. Storage is provided through an HP P4500-series SAN (Storage Area Network), which offers complete redundancy within the SAN (redundant-chassis, dual power supplies per chassis, redundant array controllers, and a Network-RAID array to spread the data across the

P4500) as well as block-by-block transfer of data to a storage device at one or more remote locations. (This replication is not real-time; it is based on snapshots taken and copied to the remote location. These snapshots can be taken no more frequently than every 15 minutes, but this time period often needs to be lengthened to accommodate bandwidth constraints.)

The P4500 is setup at the primary site with a lower-cost HP P2000 deployed at the secondary site(s). The P4500 is configured to provide synchronization aligned with the circuit bandwidth between sites, allowing the P2000 to retain the same data and configuration without compromising performance.

The secondary site(s) would also have HP Proliant servers and two (or more) VLAN-capable Ethernet switches. The Proliant servers run the VMware virtualizing software, but are basically dormant until needed.

When configuring for automatic failover, several items must be adjusted:

- P4500 SANs must be deployed at the primary and remote site(s) and must be configured in a multi-site cluster
- VMware vSphere Enterprise or better is required and must be licensed for both the primary and remote (recovery) site(s)
- Windows Server licensing at the primary site must be duplicated for the recovery site(s)
- Sufficient bandwidth must exist for real-time disk-writes since this configuration cannot fall behind and catch-up during slack periods
- Additional VMware utilities and enhanced licensing for applications may be required to enable true automatic failover

For more information, see the Bryley Systems case study on the virtualization of RTA Transit Services, Inc.; the company operating the Worcester Regional Transit Authority at <http://www.bryley.com/documents/2012/Bryley%20--%20WRTA%20Case%20Study%20--%202012.pdf>.