



Redundancy in a virtual environment

How an appropriate design can mitigate risk and increase availability

Christopher Di Biase, Inbar Blankenship

Availability in an industrial control system is critical to production efficiency, quality and safety. IT infrastructure for ICS is a critical piece to ensuring overall system availability.

Why is redundancy needed?

Hardware and software failures can and do happen, even in well-designed and well-built systems. The server and network foundation for the industrial control system (ICS) needs to consider common failure modes and be able to appropriately respond to a failure. Whether that failure is hardware or software, several questions may arise for a plant manager or lead engineer:

- Does production continue?
- Does application performance suffer?
- Is data lost?
- What is the regulatory or product quality impact?

Designing a redundancy virtual infrastructure, in conjunction with application-level redundancy, can minimize the risks of production or data loss due to IT component failures. It is part of an overall system design strategy focused on availability.

This paper will outline the types of redundancy available in a virtualized environment. Selecting methods of redundancy is a balancing act between cost, administrative complexity and risk. When evaluating a complete solution, it is important to understand the cost of a single production disruption in terms of lost product, damaged equipment, environmental and safety impact as this will help define what is practical and what is not.

Once the high-level impact is understood, each component of the system should be evaluated as a failure point and its individual impact assessed. For example, a reboot of a FactoryTalk Historian SE server would disrupt the ability to generate reports and HMI trends. For some systems this may cause an unacceptable loss of process visibility, but for others it may be nothing more than a nuisance if the data collection notes are able to buffer data and avoid permanent data loss. For some process systems a batch server restart is catastrophic; for others it may cause a small delay. Understanding these impacts helps drive decisions on which redundancy methods should be applied and what costs can be tolerated to ensure application availability.

Definitions and technologies

Virtualization: A solution where the application software and its underlying operating system is decoupled from its hardware by a hypervisor. This decoupling provides the host multiple OS instances independent of the hardware. The hardware can be changed or upgraded without OS reinstallation or the OS and application can be migrated to new hardware altogether.

Hypervisor: A piece of software that manages the allocation of the physical resources to the virtual machines.

Host or Host Server: Physical server hardware which runs a hypervisor.

Virtual Machine (VM): A computer object that can host a standard operating system and its applications, backed by hardware resources provided by a hypervisor.

Guest OS: Operating system running within a VM. Commonly Microsoft Windows or Linux.

Network Attached Storage (NAS): A hardware appliance that uses a file-based approach to shared storage on a network using protocols such as SMB (Windows based file servers) or NFS (Linux based file servers). Network storage is presented as a folder with file to the hypervisor.

Storage Area Network (SAN): One or more hardware appliances that use a block-based approach to shared storage on a network using protocols such as iSCSI or FCoE. Network storage is presented as a disk that can be formatted and used by the hypervisor.

Hyper-Converged Infrastructure (HCI): A virtual infrastructure solution that eliminates the need for a traditional NAS or SAN storage appliance by utilizing internal host disks as a virtualized SAN. An example HCI solution is VMware vSAN.

Cluster: Group of 2 or more host servers that operate as a single pool of compute resources, for both management simplicity and system redundancy.

VMware vSphere: VMware's hypervisor product. Also commonly referred to as ESXi

VMware vMotion: a feature of VMware vSphere that enables a running VM to migrate from one host server to another without disrupting either Guest OS or applications.

VMware High Availability (HA): a feature of VMware vSphere that enables virtual machines to restart automatically in the event of a catastrophic failure of a host server.

VMware vSAN: an add-on feature for VMware vSphere that enables a cluster of host servers to share their internal disks as a virtualized SAN. This enables all the advanced management and resiliency features available in the vSphere suites without requiring a dedicated storage appliance. vSAN is a type of HCI solution.

Typical architecture overview

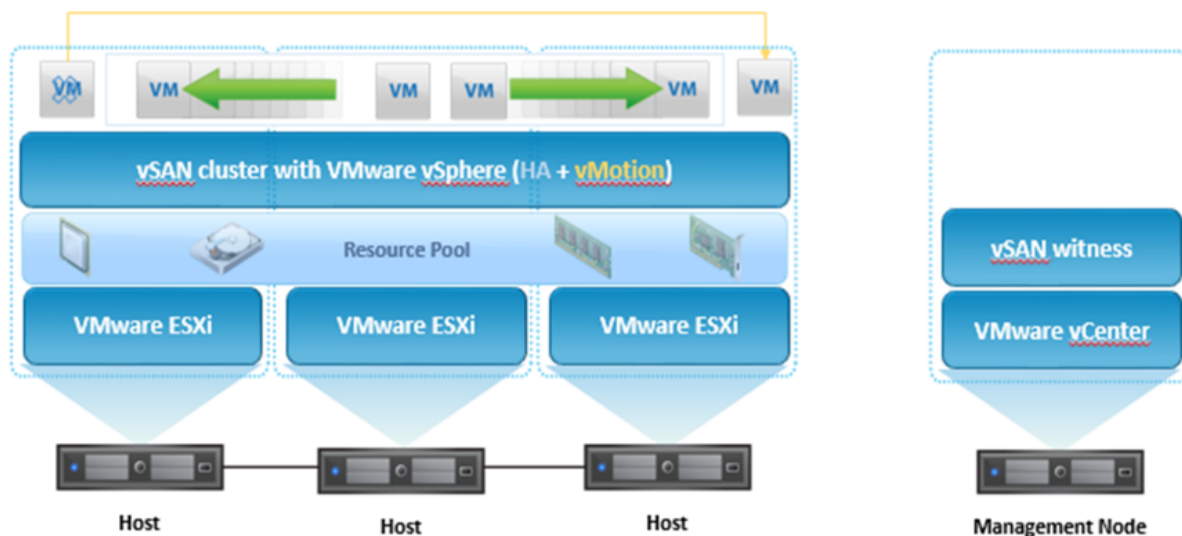


FIG 1. Typical 3 node vSAN cluster.

Types of redundancy

Redundancy can be designed into the system at every layer of the virtual infrastructure, from hardware to hypervisor, to the applications themselves. A well-designed system combines the following types of redundancy to maximize application availability and minimize disruption from a component failure.

Hardware redundancy

Hardware-level redundancy should be used to protect the system against common component failures. Such solutions should mitigate the loss of a hardware component without impact to the hypervisor or applications hosted on it. Hardware redundancy is best applied to the most common component failures, including moving parts such as fans and spinning disks as well as components vulnerable to outside forces such as power supplies and network connections.

Disk drive redundancy

RAID (redundant array of independent disk) combines multiple disk drives into a single logical unit (LUN) to mitigate the risk of data loss for a hard drive failure. In a RAID group, data is stored on at least two disks, preventing data loss in the event of a disk failure. There are different levels of RAID that allow for tradeoff between performance and resiliency:

RAID 0 (Striping): Not an actual redundancy solution, multiple disks are combined together and presented as one larger disk to the hypervisor or operating system. This is done to for administrative simplicity and to improve overall performance as reads and writes are spread over multiple physical disks. When combined with RAID 1 to add redundancy it is often referred to as RAID 10.

RAID 1: Also referred to as disk mirroring. Data is written onto or mirrored across two or more disks. When used to add redundancy to RAID 0 groups this is often referred to as RAID 10.

RAID 5 (Erasure Coding): RAID 5 is also referred to as disk striping with parity. Like RAID 0, data is spread (striped) across multiple disks. But unlike RAID 0, each disk stores an encoded copy of the data for another disk in the group which takes up much less space than the source drive. In the event that any one disk fails, its data can be recovered from the remaining disks in the system. RAID 5 provides better space efficiency than RAID 1 but has slower write performance.

RAID 6 (Erasure Coding): RAID 6 is also referred to as disk striping with double parity. Similar to RAID 5, except that each drive stores the parity stripe from two drives instead of one, allowing the RAID group to tolerate two failures. RAID 6 is often used for very large capacity drives (>4TB) that could take days to rebuild after a failure.

Power infrastructure redundancy

Ensuring that servers remain powered through component, facilities and utility failures is a key facet of system availability. Redundancy or resiliency can be designed into the system at each level.

Redundant power supplies (PSU): Each device in the solution should have at least two power supplies so that the device can continue uninterrupted should one fail.

Redundant Power Distribution Units (PDU): While PDU solutions do not fail often, they can fail. Furthermore having split PDUs allows for each PSU to be fed from a different circuit.

Separate power circuits: The PDU should be connected to separate power feeds so that a failure that causes a breaker to trip does not cause a complete system failure. For very critical systems these independent power circuits could be provided by separate utilities.

Uninterruptible power supply (UPS): As a final protection from a utility failure, a local battery backup can provide continuous power to a system until it can be safely shutdown, or until local temporary generators can be started.

Cooling redundancy

Proper cooling for IT equipment is critical. Fans are responsible for keeping the equipment cool. When selecting a server or other components such as switches ensure that system can tolerate the loss of one fan without thermal failure. Fans should also be field replicable while the component is running to minimize disruptions during repairs.

Heating, Ventilation and Air Conditioning (HVAC) is also a critical component of application availability. While loss of HVAC will typically not result in an immediate failure, thermal excursions above manufacturer rated maximums have been proven to drastically reduce the service life of servers and other IT components.

Network connectivity redundancy

Network connectivity is critical to application availability. If a host server were to lose its network connection, it would lose access to the storage behind the virtual machines and the applications themselves would lose connectivity to the underlying control system equipment and operator terminals.

Server network interface cards (NIC) should have at least two ports per connection required. When properly configured in the hypervisor, this allows for continued operation should a network cable become damaged or removed for some reason.

Server access switches should likewise be redundant. This allows for both maintenance tasks and catastrophic switch failures without production disruption.

Network availability beyond server access switching is beyond the scope of this white paper but equally as important to keeping applications available to system operators.

Memory redundancy

An unhandled failure of an internal memory module can cause data corruption or other unexpected behaviors in the applications. Server hardware should be able to detect when data stored in RAM has been corrupted and either fault the server or fault the DIMM responsible to prevent data corruption. Two common solutions are:

1. Error-correcting code memory (ECC) is used in most servers and is typically a requirement for servers used in a virtual environment. ECC memory ensures that any errors caused by either electro-magnetic interference or hardware degradation are corrected. Many server manufacturers also include a parity function so the system can detect an error in a specific memory module and mark it for replacement.
2. Mirrored Mode can be described as RAID 1 for RAM and is supported by server and workstation class CPU and motherboard designs. This is typically combined with ECC memory to provide maximum reliability of active data on a server system.

Hypervisor redundancy

Modern hypervisors support a number of redundancy and management tools to enable the system to both recovery from failures and allow preventative maintenance before an issue causes an application failure.

- **Virtual machine restart**

Despite hardware level redundancy, host servers will occasionally experience faults that either cannot be mitigated with hardware or are cost prohibitive to mitigate with hardware. In the event of such a failure, a hypervisor should be able to automatically restart impacted virtual machines on remaining host servers in the system. Features, such as VMware HA, enable host servers in a system to monitor each other for failure and act should a failure occur.

Automatic VM restarts can also be applied to correct from a VM level failure that causes the guest OS to become unresponsive.

- **Virtual machine mirroring**

For some applications, a VM restart would still cause unacceptable disruption to the system. In these cases, VM mirroring with hypervisor level features, such as VMware Fault Tolerance, can be used to ensure that the virtual machine remains running without disruption. In this case the virtual machine executes on two host servers in lockstep with each other over the network. Virtual machine mirroring will cause applications to run more slowly on identical hardware when enabled, due to the latencies introduced by the network synchronization, but it can be a useful tool to protect applications that otherwise lack redundancy.

- **Storage clustering (HCI)**

Hyper-converged infrastructure (HCI) solutions, can apply data redundancy and availability at the virtual machine object level instead of the virtual disk level. VMware vSAN can manage multiple copies of a virtual machine across the hosts in a cluster and emulate RAID1, RAID10, RAID5 or RAID6 levels of data protection depending on the size of the cluster and requirements of the application. This provides a software alternative to traditional hardware redundancy for disk drives.

- **Virtual machine migration**

While not a true redundancy solution, migrating running virtual machines between physical hosts can be used as a means of increasing application availability. Virtual machine migrations, using VMware vMotion, can be used to evacuate a host server that has experienced a non-catastrophic failure. This allows for maintenance and break/fix activities without adding risk to the running applications.

- **Health alerts and alarms**

The hypervisor and its management solution should be able to provide system administrators with notification when a hardware or infrastructure software component has failed or is trending towards a failure. Redundancy alone will not ensure application availability forever. Without proper monitoring and alerting, a system administrator may not know that a hardware or infrastructure software component has failed until the applications go offline at the second or third failure.

Application redundancy

The last line of defense for ensuring application availability is application level redundancy. Many applications that are considered mission critical and will disrupt production during a reboot will include built in redundancy or clustering methods to ensure that plant operators maintain visibility and control of the system. Application level redundancy also serves to protect against software faults that do not cause a VM hang and are otherwise undetectable by the hardware or hypervisor. Examples of these applications are:

- **FactoryTalk[®] redundancy**

The FactoryTalk Services Platform provides redundancy for many of the production critical servers that are part of a FactoryTalk or PlantPAx[®] solution. This includes, OPC and LiveData data servers, HMI servers and Alarm servers. This ensures that operator workstations remain connected to a valid source of data should a hardware or software event occur. When combined with VMware HA, FactoryTalk redundancy allows operators to continue to have access to critical process control and information, while VMware HA restarts lost servers, healing the software redundant state of the system.

- **FactoryTalk[®] Historian**

FactoryTalk Historian SE has two software redundant modes that protect against different failures. Data collection interface failover is included with all licenses and enables a pair of interface nodes to work in parallel and ensure that no software or hardware faults cause a loss of data collection. Historian server collectives ensure data accessibility during failure events.

- **ThinManager[®]**

ThinManager offers two means to improve operator terminal availability. The ThinManager server itself can be made redundant, ensuring that terminals can always find a session when brought online. ThinManager also allows multiple RDS hosts to be assigned to provide sessions to a given terminal, ensuring that a session host is available for a terminal, even during maintenance of the remote desktop service environment or during a hardware or software failure.

Choosing a redundancy solution

In general, a solution will have a mix of redundancy options. The scale to which a solution is implemented depends on the criticality of the applications that are running in the environment. The greater the impact to production due to application disruption, the more redundancy options should be implemented.

Is there such a thing as too much redundancy? Redundancy adds cost and complexity to any system. There comes a point of diminished returns for any system, where the cost of higher availability cannot be justified by the impact of failure. This point will vary depending on the cost of downtime, regulatory requirements and safety implications.

Should a system ever be deployed with no redundancy? There are situations where the cost of applying a redundant system could outweigh the need for it. If an application is not critical to ensure continued production, regulatory compliance or safe system operation, it may be acceptable to avoid the complexity of redundancy. It is critical to understand the criticality of every application before determining what level of redundancy is required.

Other factors to consider

Network

When considering the level of redundancy in your virtualization environment be sure to also take into account the rest of the infrastructure – particularly the network. The resiliency of the network is key in ensuring that your virtual infrastructure can communicate. If the network is not properly set up or has bandwidth issues, your redundancy measures will not be able to take over when needed.

Data protection

While not usually considered in a discussion of redundancy, data protection is an important part of avoiding extended production outages. While hardware and software designs strive to avoid data loss due to failure or corruption, having a backup copy of production critical workloads is necessary to avoid reengineering HMI projects and lost production or recalls due to loss of regulatory data.

At a minimum, any data protection solution should run automatically on a periodic basis and store the backup data on a secondary storage device or medium. For example, if using a VMware vSAN based system, backup data should be stored on an external appliance, such as a DataDomain, or on a separate server that does not participate in the cluster. Best practices would include archiving backup data off site, either to another corporate location, to the cloud or to tape stored offsite.

Often this evaluate must include components outside the server environment and include an evaluation of networking and control system.

Management

An important consideration is who will support and maintain the virtual environment? If the virtual environment is not maintained and managed, a redundancy solution can become a heavy burden. From managing the virtualization environment (creating the templates, add/remove users, etc.) to actively monitoring the hardware to ensure that redundancy is up and running, management of the system is critical to maintain high availability.

Use cases

Let's review a few situations and discuss the type of redundancy solution best implemented

- **Updates and/or patching for the hypervisor.**

To minimize downtime associated with providing updates to the hypervisor or patching to it, implement redundancy combination of hypervisor and server hardware. This combination will allow you to move the images running on a server to another server while the first server is being updated/patched. Without this combination, you'll experience downtime to shutdown the images as the updates are made and the system is restarted.

- **Application/OS software failure – application**

To minimize downtime associated with the result of application software or operating system failure, implement application redundancy. If the primary application server fails, the secondary server will take over.

- **N=# hardware failure**

To minimize the effects of the hardware failure, consider which type of hardware failure you're anticipating. If you are concerned regarding a specific hardware failure, having redundant hardware might suffice. If the concern is general, a combination of hardware and hypervisor redundancy solution can provide a layer on top of the hardware failure – restarting the image at hardware failure.

- **Hardware lifecycle – hypervisor.**

To minimize the downtime required for life cycling hardware implement a hypervisor redundancy solution. With this type of solution, if there are multiple servers, you can lifecycle the hardware in steps

Resources

You can read the following for further reading:

[VMware VMotion Product Datasheet](#)

[VMware High Availability Product Datasheet](#)

Rockwell Automation has a selection of products, product features, and services to enhance application availability in the compute space such as the VersaVirtual™ Appliance and Industrial Data Center. For more information please visit:

[VersaVirtual Appliance Product Page on ab.com](#)

[Pre-Engineered Network Solutions page on RockwellAutomation.com](#)

PLEASE CONTACT THE FOLLOWING PEOPLE FOR MORE INFORMATION.





Rockwell Automation

Christopher Di Biase
Architect - Pre-engineered Solutions
cdbiase@rockwellautomation.com

Rockwell Automation

Inbar Blankenship
Product Manager
iepstein@rockwellautomation.com



Connect with us.    

rockwellautomation.com

expanding **human possibility**[™]

AMERICAS: Rockwell Automation, 1201 South Second Street, Milwaukee, WI 53204-2496 USA, Tel: (1) 414.382.2000, Fax: (1) 414.382.4444
EUROPE/MIDDLE EAST/AFRICA: Rockwell Automation NV, Pegasus Park, De Kleetlaan 12a, 1831 Diegem, Belgium, Tel: (32) 2 663 0600, Fax: (32) 2 663 0640
ASIA PACIFIC: Rockwell Automation, Level 14, Core F, Cyberport 3, 100 Cyberport Road, Hong Kong, Tel: (852) 2887 4788, Fax: (852) 2508 1846

Expanding human possibility[™], FactoryTalk[®], PlantPAx[®], ThinManager[®], VersaVirtual[™] Appliance.

Publication GMSN-SP012A-EN-P-September 2019
Copyright © 2019 Rockwell Automation, inc. All Rights Reserved. Printed in USA.