DeltaV™ Virtualization – High Availability and Disaster Recovery

This document describes High Availiability and Disaster Recovery features supported by DeltaV Virtual Studio.

DELTAV

EMERSON
Process Management

Introduction

Virtualization is a standard practice in many industries and is rapidly being adopted within the manufacturing community as a way to reduce costs, increase productivity, extend system life, and improve overall reliability and system performance. The main concerns that manufacturers have with virtualization are system complexity and reliability. Manufacturers want their virtualized control platform to be easy to implement and maintain, and as reliable as their traditional control platform.

Emerson's DeltaV system is the first control system to make virtualization easy to implement and maintain. Our DeltaV Virtual Studio provides an integrated application environment designed for easy implementation and management of virtual DeltaV systems for both off-line and on-line operation. And starting with DeltaV Virtual Studio 2.3, Emerson delivers High Availability and Disaster Recovery for unprecendented control system reliability.

Why High Availability and Disaster Recovery is Important

Virtualization is appealing because it enables you to optimize your computer resources. You can run multiple virtual computers (i.e., virtual machines or VMs) within a single host server, and run multiple host servers in a server cluster using shared network storage. See ***Figure 1*** for an illustration of a typical virtualization environment. This capability delivers attractive benefits, but it also introduces new risks. What if a host server fails? What if your shared network storage fails? Instead of losing just one workstation, you can lose all of the VMs running in the host or using shared storage. So if you put all your eggs in one basket, how do you protect the basket?

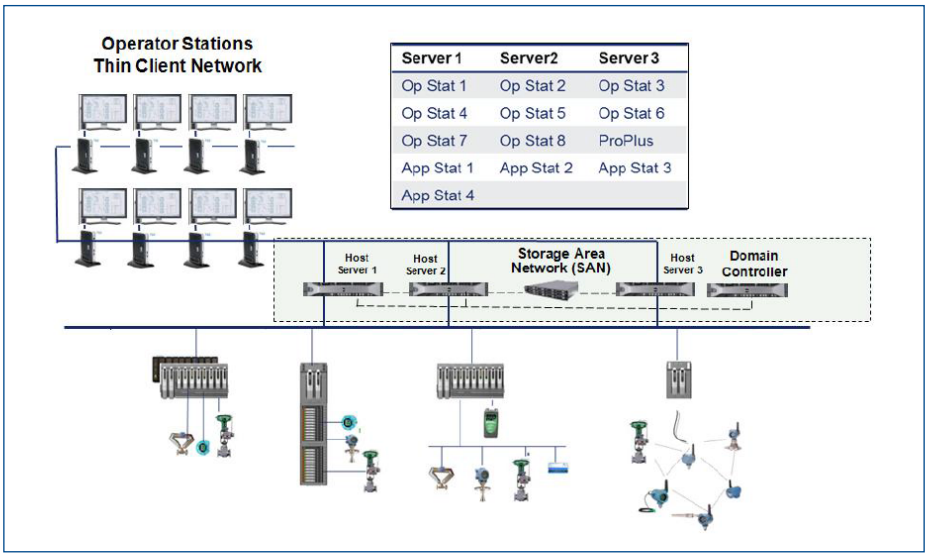

Figure 1 – Example DeltaV Virtualization System

Protecting your virtualized control system consists of three basic components:

1. ***Data Backup*** – Protects against data loss due to individual file loss, data corruption, or human error. To be protected, data must be backed up regularly. Having multiple recovery points available guards against data loss that is not detected immediately.

2. ***Image Backup*** – Protects against disaster scenarios where your primary host computer or shared storage fails. Image backup and recovery are often referred to as "disaster recovery". To minimize recovery time, images must be captured periodically and a secondary host must be available to restart image backups.

3. ***Automatic Failover*** – Protects against host system failure by automatically moving virtual machines from a primary host that has failed to a secondary host. This capability is commonly referred to as "High Availability" and is essential when your virtual machines are critical.

How to best protect your control system from computer hardware and software failures can depend on the size of your system, and the critical importance of the applications. This whitepaper explores the different levels of availability provided with DeltaV Virtual Studio to protect your system and insure your required level of reliability.

Levels of Availability

Control systems are a critical component of a manufacturing facility. Their availability is paramount to a plant's performance and profitability. There are several means to protect a control system from downtime caused by data loss or hardware / software failures. This section describes the levels of availability provided for different control configurations.

1. Traditional Backup and Recovery Systems

Traditional backup systems provide file and/or image backup for individual workstations or servers. The backup files are created by backup agents that run on the managed computers, and are stored on a separate backup storage server. Individual data files can be restored in the case of data corruption. In the event of a machine failure, backup images can be loaded on a spare recovery machine. ***Figure 2*** illustrates a traditional backup and recovery system.

DeltaV's Backup and Recovery product for traditional system backup provides an easy-to-use backup system with predefined DeltaV templates and automated installation. DeltaV Backup and Recovery is based on industry-leading Acronis Backup & Recovery software. See *DeltaV Backup and Recovery Product Data Sheet* for more information.

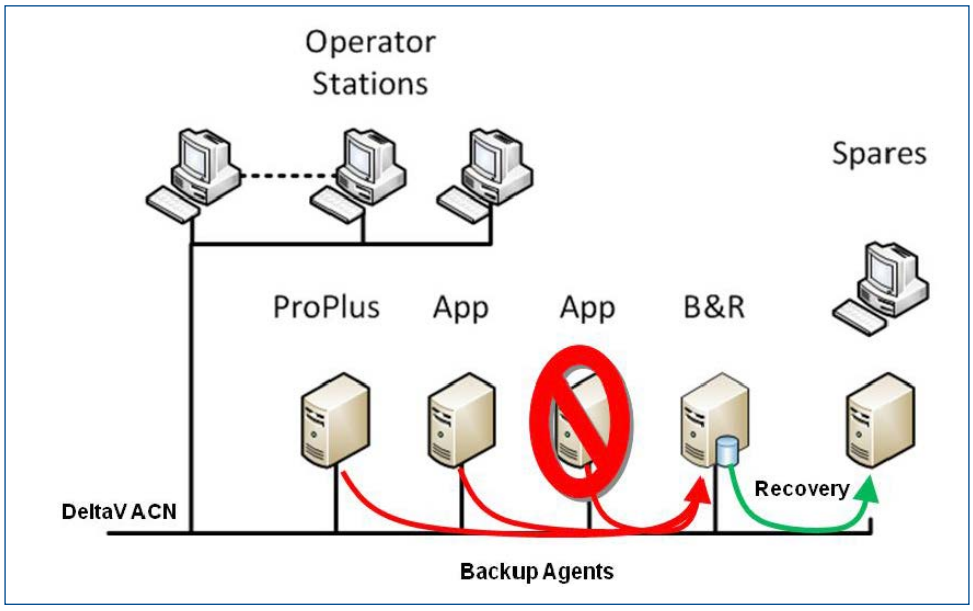

Figure 2 – Traditional Backup and Recovery System

2. Traditional Backup with Recovery to Virtual Machine

Traditional image backup of workstations and servers can also be recovered to a virtual machine running in a recovery host server. The advantage of this option is multiple workstations can be recovered to the same host computer. You can also use the host computer for running other applications (e.g., maintenance or engineering stations) so that you no longer have to keep idle spare machines around. ***Figure 3*** illustrates traditional backup with recovery to a virtual machine.

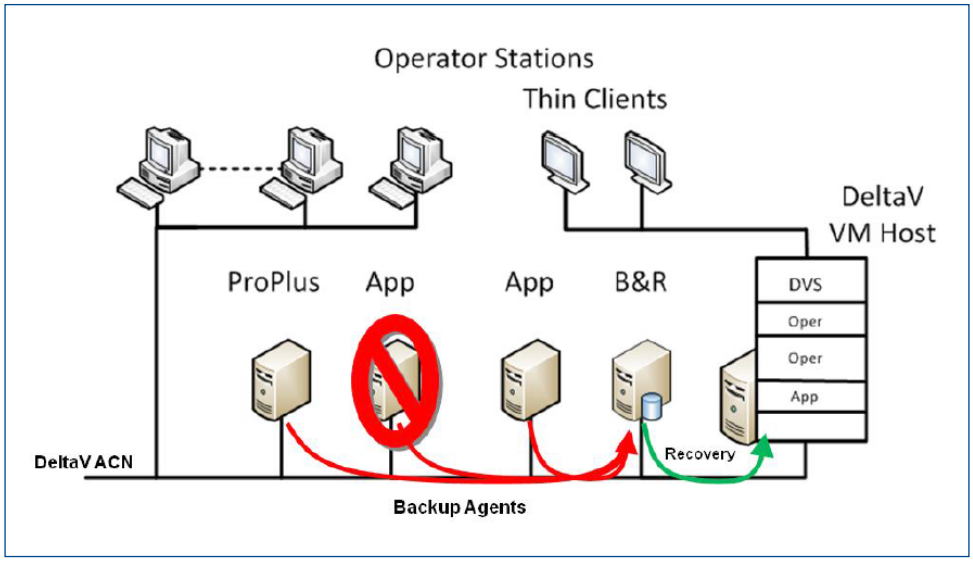

Figure 3 – Traditional backup with recovery to virtual machine

3. Redundant Application Servers

The DeltaV system provides redundancy for several critical applications including Batch Executive, OPC Server, Zone Servers, and advanced control modules executing in an Application Station. The standard redundancy configuration is to have two application servers that continuously synchronize over Ethernet. If the primary server fails, the secondary server automatically picks up operations without any application downtime.

These redundant DeltaV applications can be implemented in a virtual environment using two virtual machines in separate host servers. No shared network storage or high availability virtualization capabilities are required. ***Figure 4*** shows redundant applications running in a pair of virtual application stations in two host servers.

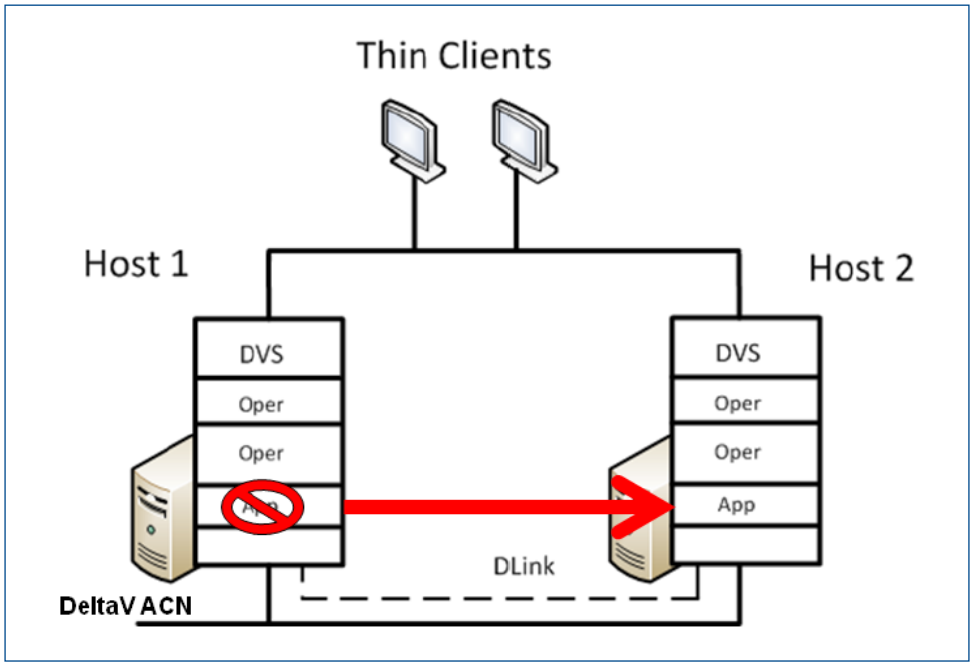

Figure 4 – Redundant applications running in virtual Application Stations

4. Redundant Network Storage and Switches

Shared network storage, such as Storage Area Network devices (SANs), provides the ability to easily move virtual machines between host computers. Shared storage can be configured as part of server failover clusters to allow VMs to be automatically moved and restarted in the case of a host server failure (i.e., High Availability). However, the more VMs running on a shared network storage device, the more VMs are at risk of downtime in case of SAN failure. To protect against failure, SANs are built with multiple redundant components including:

- Redundant Disk Drives – RAID 10 guarantees uninterrupted data access in the event of a disk drive failure.
- Redundant Disk Drive Controllers – Disk Drive Controllers are responsible for the read / write operations performed on the disks. Redundant controllers ensure there is no single point of failure in the disk IO operations.
- Redundant Power – Power supplies are a common source of failure, therefore redundant, independent power supplies are mandatory.

Figures 5 & 6 illustrate a SAN configuration and redundant storage components.

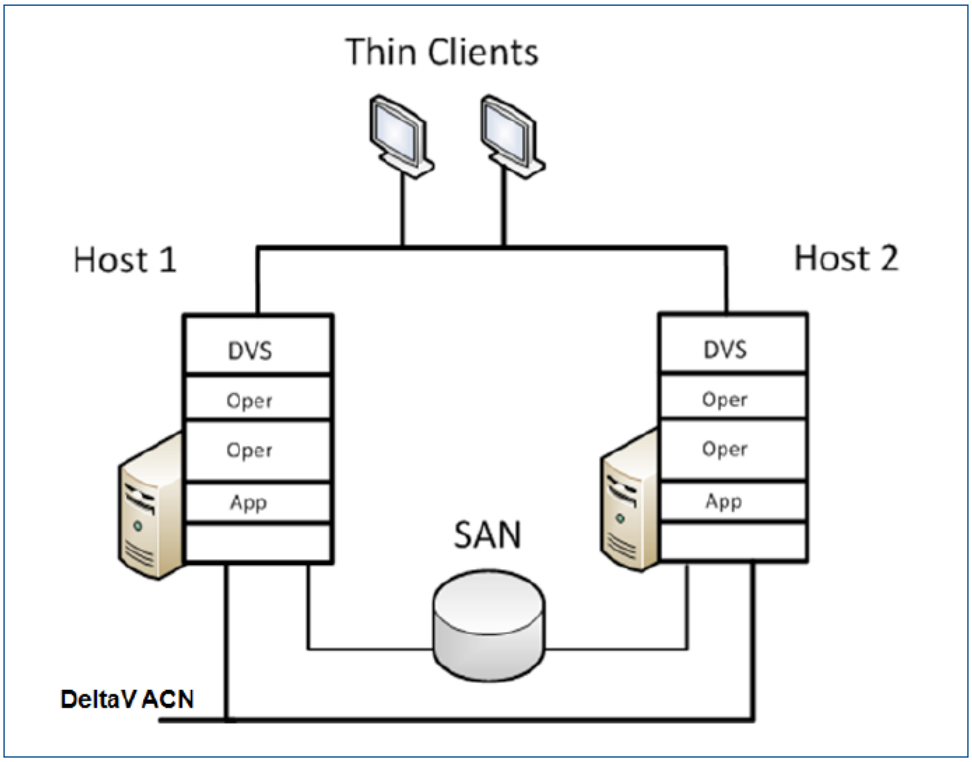

Figure 5 – Virtual Machines running on Storage Area Network (SAN) device

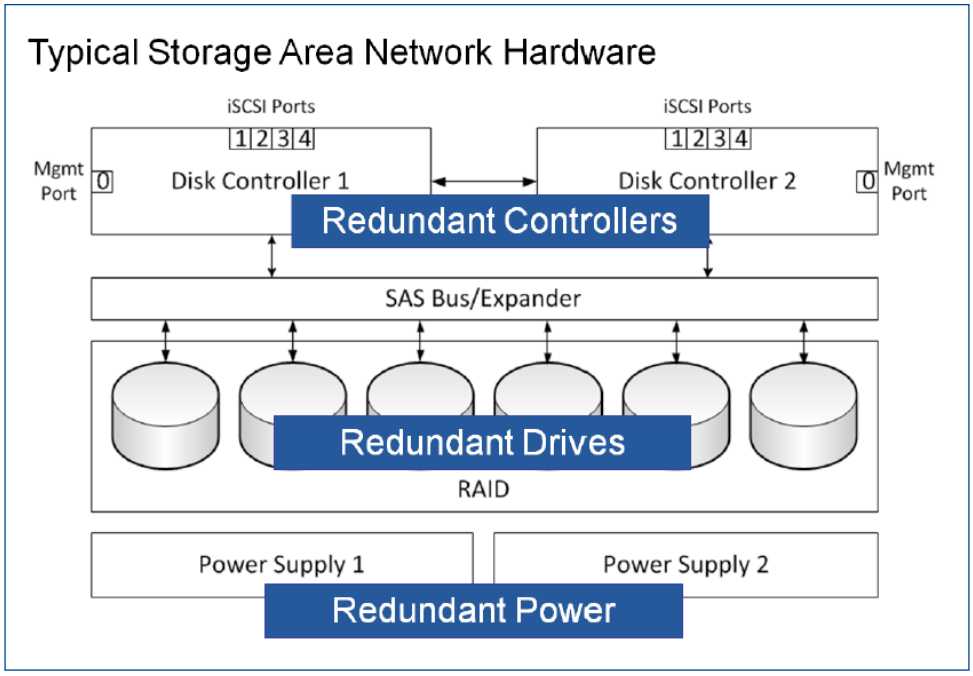

Figure 6 – Redundant components in a Storage Area Network Device

Redundant network switches are also important for critical applications. Redundant network switches are recommended for the thin client network, storage area network (SAN), host management / domain controller network, and DeltaV system ACN.

5. Virtual Machine Replication

VM Replication is the automatic image copy of the virtual machine across a standard network connection using standard virtualization hypervisor software. VM Replication is commonly used as a Disaster Recovery (DR) strategy whereby virtual machines can be replicated to a remote server, and quickly restored in the case of primary component failure such as host server, shared network storage, or cluster failure. VM Replication is very attractive for DR because it's included at no extra cost with Windows Server 2012 software, and it does not require an expensive SAN device.

An example of VM Replication is shown in ***Figure 7***. In this example, there are two host servers and a remote replication server. The replication server is connected to the ACN and thin client network so that the DeltaV workstations can operate on the replication server in case of disaster scenario.

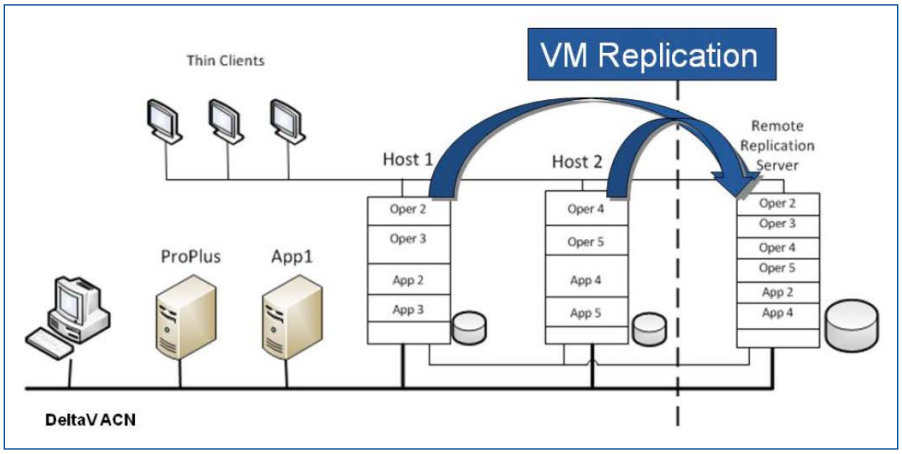

Figure 7 – VM Replication to Remote Server

How VM Replication Works

VM Replication is a standard feature of Windows Server 2012 (or Hyper-V Server 2012). The Windows VM Replication engine has a module called “Change Tracking” that captures every “write” within the virtual machine hard disk files and stores each in a log file. This replication happens at the operating system level for the Virtual Hard Drive (VHD), making it easy and efficient to replicate VMs. The replication using these logs occurs periodically (e.g., every 5 minutes) and asynchronously through an HTTP or HTTPS connection. All the data that must be replicated to the replication server uses the network module, which optimizes the workload such that it works in slow network connections (e.g. WANs). All you need are two physical servers running Windows Server 2012 (or Hyper-V Server 2012) and a network connection between them. It does not need any third-party hardware or software applications.

Windows Server 2012 (and Hyper-V Server 2012) also includes other special features that support disaster recovery including support for Failover Testing, synchronized Fail Back capability, and live storage migration.

A limitation with DeltaV VM Replication is that a manual restart of the replicated virtual machine is required when the primary VM has failed. Although it is relatively easy to restart the VM, it requires administrative privileges that most operators don’t have. For critical virtual machines that require an automatic restart, High Availability Clusters are recommended.

6. High Availability Cluster

For on-line production control systems, high availability with automatic VM failover capabilities is important to insure maximum uptime. DeltaV Virtual Studio supports automatic failover of virtual machines from a primary host to a secondary host in the event of primary host failure. To take advantage of this capability, a failover cluster must be configured using a storage area network (SAN) device with a host domain controller computer to manage the failover switch. ***Figure 8*** illustrates two host computers with shared storage and automatic failover enabled.

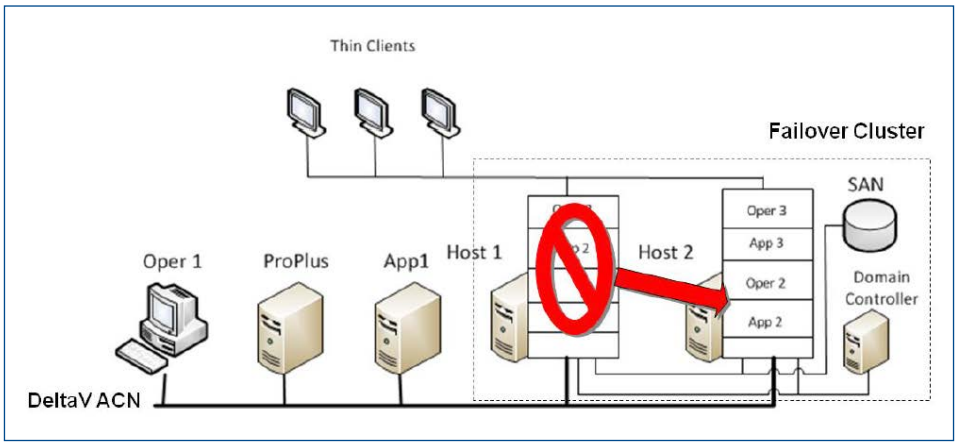

Figure 8 – Two host Computers with SAN and Automatic Failover

When a host failure occurs, DeltaV Virtual Studio automatically moves the virtual machines that were running on the failed host, to a specified secondary host. With High Availability (HA) enabled, the virtual machines are moved and automatically restarted.

How High Availability Clusters Work

High Availability is enabled using Windows Failover Clustering. A failover cluster is a group of independent host servers that work together to increase the availability and scalability of clustered applications and services. If one or more of the cluster host servers fail, other servers begin to provide service (a process known as failover). In addition, the clustered servers are proactively monitored to verify that they are working properly. If they are not working properly, their applications are restarted or moved to another server. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed storage area for all clustered servers. With the Failover Clustering feature, users experience a minimum of disruptions in service.

Live Migration

Another key component of High Availability is Live Migration. Live Migration allows VMs to be migrated from one host to another, while the VM is running. Live Migration enables you to shutdown a host server for maintenance without shutting down the VMs on that host and potentially disrupting operations.

High Availability provides increased uptime in the case of host server failure. But because it relies on shared network storage, there is a remote chance of storage failure which would not be protected by High Availability Clusters. Although the shared network storage SANs are extremely reliable with multiple layers of redundancy, we recommend additional disaster recovery provisions be considered for critical applications. Specifically, we recommend a combination of High Availability with VM Replication as described in the next section.

7. High Availability Cluster with Integrated VM Replication

The ultimate in High Availability and Disaster Recovery is to take advantage of both failover clusters and VM replication. In this case, High Availability provides very fast recovery from a host failure, and VM Replication protects against a disaster in which the shared network storage or failover cluster fails.

To illustrate, consider the example in ***Figure 8*** with High Availability consisting of a single failover cluster and a remote replication server to store backup VM images.

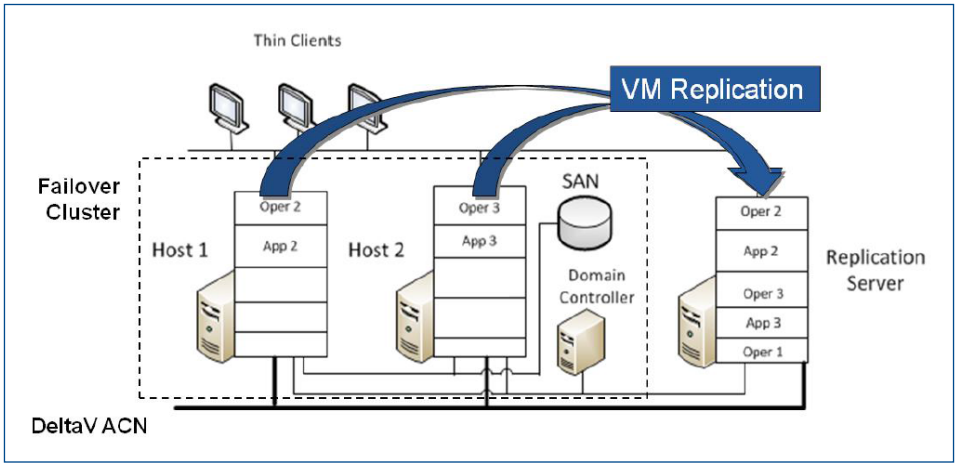

Figure 8 – High Availability with VM Replication

In this case, HA addresses host failures, and VM Replication addresses storage or cluster failure. One other precaution to ensure system visibility after a disaster involves having one or more workstations available outside the shared storage or failover cluster. You could have one or more Operator or Professional stations implemented in traditional hardware, or implemented as virtual machines in the replication server (not on the SAN and not part of the failover cluster).

Another advantage of a disaster recovery system is easier hardware upgrades and planned equipment repair. For example, in the system shown in Figure 7, you can switch to your replica VMs to perform a SAN upgrade or managed repair. DeltaV Virtual Studio also supports Live Storage Migration to your replication server for zero downtime of your VMs.

For larger systems, two or more clusters can be set-up with "cross replication" between the clusters. This configuration is shown in ***Figure 9***. The advantage of this configuration is that you never lose complete visibility of the plant because half of the virtual machines remain running in the case of a SAN or cluster failure.

High Availability with Disaster Recovery are recommended for all on-line production systems.

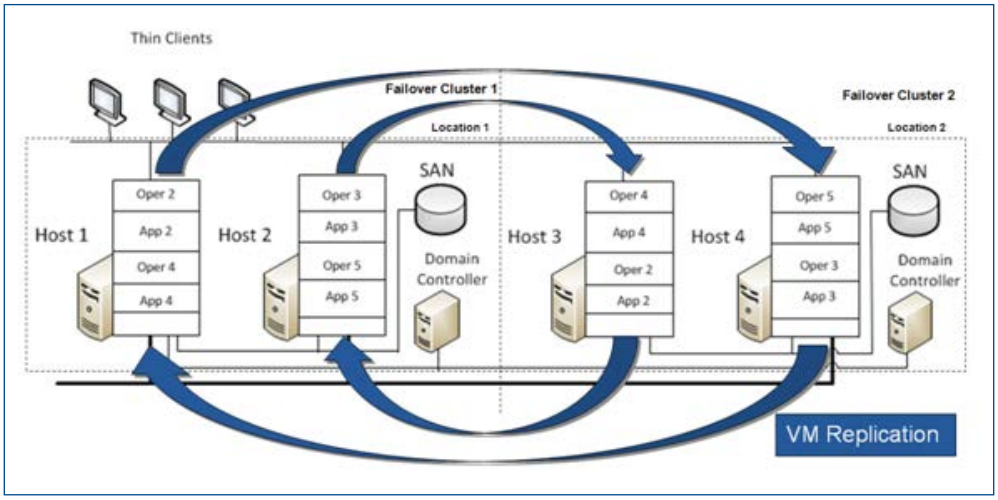

Figure 9 – Multiple HA Failover Clusters with VM Replication

Integrated Platform for High Availability and Disaster Recovery

An attractive solution for High Availability and VM Replication uses an integrated platform consisting of blade servers, shared network storage and switches. Shown in **Figure 10** is the new Dell PowerEdge VRTX integrated platform which supportes DeltaV Virtual Studio 2.3.1 with High Availibility and Disaster Recovery. Each PowerEdge VRTX can support up to 40 DeltaV virtual machines and is available in a tower or rack-mount format. With two Dell PowerEdge VRTX units, the High Availability and Disaster Recovery configuration in **Figure 9** is easily configured. All VRTX on-line solutions are required to have disaster recovery capabilities to support managed repair. For more information on the Dell PowerEdge VRTX, please see *DeltaV Virtualization Hardware* product data sheet.

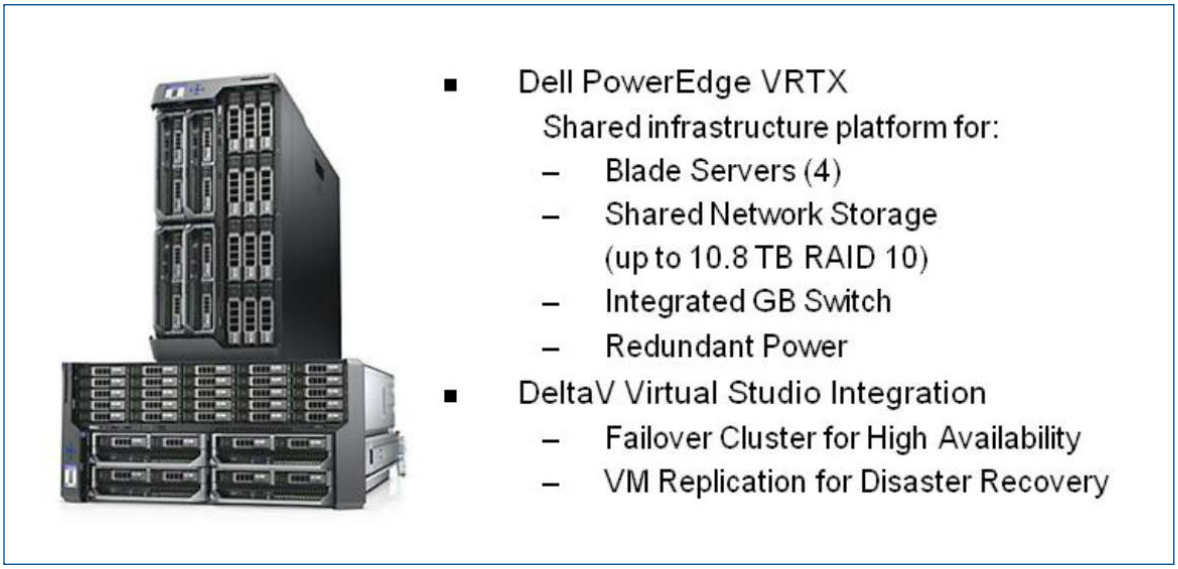

Figure 10 – Dell PowerEdge VRTX

Summary

As virtualization is adopted for distributed control systems, high availability and disaster recovery are necessary to ensure a plant's performance and profitability. The level of availability needed depends on the importance of the virtualized applications and the scope of the virtualized environment. DeltaV Virtual Studio's high availability and disaster recovery features provide multiple layers of protection to minimize downtime and ensure availability. DeltaV Virtual Studio bundled with an integrated virtualization hardware environment like Dell's PowerEdge VRTX blade servers make HA and DR easy to implement at an attractive price.

Emerson Process Management

Asia Pacific: 65.6777.8211
Europe, Middle East: 41.41.768.6111
North America, Latin America:
+1 800.833.8314 or
+1 512.832.3774

www.EmersonProcess.com/DeltaV

DELTAV

EMERSON
Process Management